

MQSeries

SC33-1873-00

System Administration



MQSeries

SC33-1873-00

System Administration

Note!

Before using this information and the product it supports, be sure to read the general information under Appendix G, "Notices" on page 321.

First edition (September 1997)

This edition applies to the following products:

- MQSeries for AIX Version 5
- MQSeries for HP-UX Version 5
- MQSeries for OS/2 Warp Version 5
- MQSeries for Sun Solaris Version 5
- MQSeries for Windows NT Version 5

and to any subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

At the back of this publication is a page titled "Sending your comments to IBM". If you want to make comments, but the methods described are not available to you, please address them to:

IBM United Kingdom Laboratories,
Information Development,
Mail Point 095,
Hursley Park,
Winchester,
Hampshire,
England,
SO21 2JN

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

© **Copyright International Business Machines Corporation 1994,1997. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

About this book	ix
Who this book is for	ix
What you need to know to understand this book	ix
Terms used in this book	ix
Using MQSeries for UNIX systems	x
Using MQSeries for OS/2 Warp and Windows NT	x
The calls MQCONN and MQCONNX	xi
MQSeries publications	xi
MQSeries cross-platform publications	xi
MQSeries platform-specific publications	xiii
MQSeries Level 1 product publications	xv
Softcopy books	xv
MQSeries information available on the Internet	xvi
Related publications	xvi

Part 1. Guidance	1
Chapter 1. Introduction to MQSeries	7
MQSeries and message queuing	7
Messages and queues	8
Objects	9
System default objects	14
Administration	14
Clients and servers	15
Extending queue manager facilities	15
Security	17
Transactional support	18
Chapter 2. Understanding administration command sets	19
Control commands	19
MQSeries commands (MQSC)	20
PCF commands	21
Comparing command sets	22
Chapter 3. Managing queue managers	27
Getting started	27
Guidelines for creating queue managers	27
Understanding MQSeries file names	30
Working with queue managers	32
Managing the command server for remote administration	37
Chapter 4. Administering local MQSeries objects	39
Supporting application programs that use the MQI	39
Issuing MQSC commands for administration	40
Running MQSC commands from text files	44
Resolving problems with MQSC	47
Working with local queues	49
Working with alias queues	55
Working with model queues	57

Managing objects for triggering	59
Chapter 5. Administering remote MQSeries objects	63
Understanding channels and remote queuing	63
Remote administration	64
Creating a local definition of a remote queue	71
Using remote queue definitions as aliases	74
Data conversion	75
Chapter 6. Protecting MQSeries objects	77
Before you begin (UNIX systems)	77
Before you begin (Windows NT)	78
Why you need to protect MQSeries resources	79
Understanding the Object Authority Manager	79
Using the Object Authority Manager commands	82
Object Authority Manager guidelines	85
Understanding the authorization specification tables	88
Understanding authorization files	94
Chapter 7. Configuration files	99
What configuration files are	99
MQSeries configuration file	99
Queue manager configuration file	103
Editing configuration files	108
Configuring the logs	108
Specifying log file sizes	112
Configuring XA_-compliant databases	112
Chapter 8. The MQSeries dead-letter queue handler	115
Invoking the DLQ handler	115
The DLQ handler rules table	116
How the rules table is processed	123
An example DLQ handler rules table	125
Chapter 9. Instrumentation events	127
What instrumentation events are	127
Why use events?	128
Chapter 10. Linking to Lotus Notes	133
What is Lotus Notes?	133
Linking applications	133
Server or client?	135
Disconnected requests	137
Setting up your system	137
Starting the server add-in task	138
Verifying that Lotus Notes can link to MQSeries	139
Stopping the link server task	140
User notification	140

Chapter 11. Transactional support	141
Database coordination	142
DB2 configuration	146
Oracle configuration	152
Multiple database configurations	157
Administration tasks	159
External syncpoint coordination	164
Using CICS	166
Chapter 12. Recovery and restart	171
Making sure that messages are not lost (logging)	171
Checkpointing – ensuring complete recovery	174
Dumping the contents of the log using the dmpmqlog command	176
Managing logs	194
Using the log for recovery	197
Backup and restore	200
Recovery scenarios	201
Chapter 13. Problem determination	203
Preliminary checks	203
What to do next	207
Application design considerations	211
Incorrect output	212
Error logs	215
Dead-letter queues	219
Configuration files and problem determination	219
Tracing	219
First-failure support technology (FFST)	226
Problem determination with clients	230
Part 2. Reference	233
Chapter 14. MQSeries control commands	235
Names of MQSeries objects	235
How to read syntax diagrams	236
Syntax help	237
crtmqcvx (Data conversion)	238
crtmqm (Create queue manager)	240
dlmqm (Delete queue manager)	244
dmpmqlog (Dump log)	246
dspmqaut (Display authority)	248
dspmqcsv (Display command server)	252
dspmqfls (Display MQSeries files)	253
dspmqtrc (Display MQSeries formatted trace output)	255
dspmqtrn (Display MQSeries transactions)	256
endmqcsv (End command server)	258
endmqlsr (End listener)	260
endmqm (End queue manager)	261
endmqtrc (End MQSeries trace)	263
rcdmqimg (Record media image)	265
rcrmqobj (Recreate object)	267
rsvmqtrn (Resolve MQSeries transactions)	269
runmqchi (Run channel initiator)	271

Contents

runmqchl (Run channel)	272
runmqdlq (Run dead-letter queue handler)	273
runmqlsr (Run listener)	275
runmqsc (Run MQSeries commands)	277
runmqtmc (Start client trigger monitor)	280
runmqtrm (Start trigger monitor)	281
scmmqm (Add the queue manager to the Windows NT Service Control Manager)	282
setmqaut (Set/reset authority)	284
setmqprd (Enroll production license)	290
setmqtry (Start trial period)	291
strmqcsv (Start command server)	292
strmqm (Start queue manager)	293
strmqtrc (Start MQSeries trace)	295

Part 3. Appendixes 299

Appendix A. System and default objects	301
---	-----

Appendix B. Directory structure (UNIX systems)	303
Queue manager log directory structure	305

Appendix C. Directory structure (OS/2)	307
Queue manager log directory structure	309

Appendix D. Directory structure (Windows NT)	311
Queue manager log directory structure	313

Appendix E. Stopping and removing queue managers manually	315
Stopping a queue manager manually	315
Removing queue managers manually	316

Appendix F. User identifier service	319
--	-----

Appendix G. Notices	321
Trademarks	322

Part 4. Glossary and index 323

Glossary of terms and abbreviations	325
--	-----

Index	335
--------------	-----

Figures

1. Queues, messages, and applications	39
2. Extract from the MQSC command file, myprog.in	45
3. Extract from the MQSC report file, myprog.out	46
4. Typical results from queue browser	54
5. Remote administration	65
6. Setting up channels and queues for remote administration	66
7. Example MQSeries configuration file for UNIX systems	101
8. Example MQSeries configuration file for MQSeries for OS/2 Warp and Windows NT	102
9. Example queue manager configuration file for MQSeries for UNIX systems	105
10. Example queue manager configuration file for OS/2	106
11. Example queue manager configuration file for Windows NT	107
12. An example rule from a DLQ handler rules table	118
13. Understanding instrumentation events	128
14. Monitoring queue managers across different platforms, on a single node	129
15. Lotus Notes and MQSeries	134
16. Configuration with the add-in task on the MQSeries client	136
17. Source code for db2swit.c for platforms other than Windows NT	146
18. Source code for db2swit.c on Windows NT (Microsoft Visual C++ specific)	146
19. Source code for db2swit.def on OS/2	147
20. Makefile for DB2 switch on OS/2	147
21. Source code for db2swit.def on Windows NT	148
22. Makefile for DB2 switch on Windows NT	148
23. Makefile for DB2 switch on AIX	149
24. Makefile for DB2 switch on Sun Solaris	149
25. Makefile for DB2 switch on HP-UX	149
26. Sample XAResourceManager entry for DB2 on OS/2 and Windows NT	150
27. Sample XAResourceManager entry for DB2 on UNIX platforms	150
28. Source code for oraswit.c	153
29. Makefile for Oracle switch on AIX	153
30. Makefile for Oracle switch on Sun Solaris	153
31. Makefile for Oracle switch on HP-UX	154
32. Sample XAResourceManager entry for Oracle on UNIX platforms	155
33. Sample XAResourceManager entries for multiple DB2 databases	157
34. Sample XAResourceManager entries for a DB2 and Oracle database	157
35. Sample dspmqtrn output	160
36. Sample dspmqtrn output for a transaction in error	162
37. Commented out XAResourceManager stanza	163
38. Checkpointing	174
39. Checkpointing with a long-running transaction	175
40. Example dmpmqlog output	180
41. Extract from an MQSeries error log	218
42. Sample AIX trace	221
43. Sample HP-UX trace	223
44. Sample MQSeries for Sun Solaris trace	224
45. Sample MQSeries for Windows NT trace	225
46. Sample MQSeries for Sun Solaris First Failure Symptom Report	226
47. Sample MQSeries for Windows NT First Failure Symptom Report	228

Tables

48.	Default directory structure (UNIX systems) after a queue manager has been started	303
49.	Default file tree (OS/2) after a queue manager has been started	307
50.	Default file tree (Windows NT) after a queue manager has been started	311

Tables

1.	Commands for queue manager administration	23
2.	Commands for command server administration	23
3.	Commands for queue administration	24
4.	Commands for process administration	24
5.	Commands for channel administration	25
6.	Other control commands	26
7.	Security authorization needed for MQI calls	89
8.	MQSC commands and security authorization needed	92
9.	PCF commands and security authorization needed	93
10.	Log overhead sizes	112
11.	XA-compliant relational databases	143
12.	XA-compliant external syncpoint coordinators	164
13.	CICS task termination exits	168
14.	Sample exits	169
15.	MQS_TRACE_OPTIONS settings	221
16.	How to read syntax diagrams	236
17.	Security authorities from the dspmqaut command	249
18.	Specifying authorizations for different object types	287
19.	The system and default objects	301

About this book

This book applies to the MQSeries Version 5 products, which are:

- MQSeries for AIX V5.0
- MQSeries for HP-UX V5.0
- MQSeries for OS/2 Warp V5.0
- MQSeries for Sun Solaris V5.0
- MQSeries for Windows NT V5.0

These products provide application programming services that enable application programs to communicate with each other using *message queues*. This form of communication is referred to as *commercial messaging*. The applications involved can exist on different nodes on a wide variety of machine and operating system types. They use a common application programming interface, called the Message Queuing Interface or MQI, so that programs developed on one platform can readily be transferred to another.

This book describes the system administration aspects of the MQSeries Version 5 products, and the services they provide to support commercial messaging. This includes managing the queues that applications use to receive their messages, and ensuring that applications have access to the queues that they require.

Installation of MQSeries is described in the *MQSeries Quick Beginnings* booklet for your platform. Post-installation configuration of a distributed queuing network is described in the *MQSeries Intercommunication* book.

Who this book is for

This book is intended for system administrators, and system programmers who manage the configuration and administration tasks for MQSeries. It is also useful to application programmers who must have some understanding of MQSeries administration tasks.

What you need to know to understand this book

To use this book, you should have a good understanding of the operating systems described here, and of the utilities associated with them. You do not need to have worked with message queuing products before, but you should have an understanding of the basic concepts of message queuing.

Terms used in this book

In this book, the term “the **MQSeries Version 5 products**” means:

- MQSeries for AIX V5.0
- MQSeries for HP-UX V5.0
- MQSeries for OS/2 Warp V5.0
- MQSeries for Sun Solaris V5.0
- MQSeries for Windows NT V5.0

The term “MQSeries for UNIX systems” means:

MQSeries for AIX V5.0
MQSeries for HP-UX V5.0
MQSeries for Sun Solaris V5.0

Using MQSeries for UNIX systems

The following restrictions apply to the use of UNIX operating-system facilities with the MQSeries product:

1. MQSeries for AIX and MQSeries for HP-UX use the UNIX subroutine `ftok` to generate standard interprocess communication keys. Using `ftok` exclusively within a node ensures that these keys are unique, which is a requirement of MQSeries. Therefore, do not use any code that generates interprocess keys in a different way.
2. MQCONN sets up its own signal handler for the signals:

SIGSEGV
SIGBUS

User handlers for these signals are restored after every MQI call.

The remaining signals are handled differently.

SIGINT
SIGQUIT
SIGFPE
SIGTERM
SIGHUP

If any handler for this second group of signals receives an interrupt within an MQI call, the application handler must exit the application. MQI may not be called.

3. For each MQI call, MQSeries uses the UNIX interval timer `ITIMER_REAL` to generate `SIGALRM` signals. Any previous `SIGALRM` handler and timer interval is saved on entry to MQI and restored on exit. Any timer interval set is therefore frozen while within MQI.

The installation directory

Throughout this book, the name **mqmtop** has been used to represent the name of the installation directory in the UNIX environments.

- For MQSeries for AIX, **mqmtop** represents the directory `/usr/lpp/mqm`.
- For MQSeries for HP-UX, **mqmtop** represents the directory `/opt/mqm/`.
- For MQSeries for Sun Solaris, **mqmtop** represents the directory `/opt/mqm/`.

Using MQSeries for OS/2 Warp and Windows NT

Examples in this book relevant to MQSeries for Windows NT may use New Technology file system (NTFS), high performance file system (HPFS), or file allocation table (FAT) file names. Examples relevant to MQSeries for OS/2 Warp may use HPFS or FAT file names.

The examples are valid for all file-naming systems, the name being transformed if necessary when the FAT system is in use. Name transformation is described in “Understanding MQSeries file names” on page 30.

The calls MQCONN and MQCONNX

References in this book to the call MQCONN - Connect queue manager can be replaced by references to the call MQCONNX - Connect queue manager (extended); MQCONNX requires an additional parameter. The syntax of both calls is described in the *MQSeries Application Programming Reference* manual.

MQSeries publications

This section describes the documentation available for all current MQSeries products.

MQSeries cross-platform publications

Most of these publications, which are sometimes referred to as the MQSeries “family” books, apply to all MQSeries Level 2 products. The latest MQSeries Level 2 products are:

- MQSeries for AIX V5.0
- MQSeries for AT&T GIS UNIX V2.2
- MQSeries for Digital OpenVMS V2.2
- MQSeries for HP-UX V5.0
- MQSeries for MVS/ESA V1.2
- MQSeries for OS/2 Warp V5.0
- MQSeries for OS/400 V3R2
- MQSeries for OS/400 V3R7
- MQSeries for SINIX and DC/OSx V2.2
- MQSeries for SunOS V2.2
- MQSeries for Sun Solaris V5.0
- MQSeries Three Tier
- MQSeries for Windows V2.0
- MQSeries for Windows V2.1
- MQSeries for Windows NT V5.0

Any exceptions to this general rule are indicated. (Publications that support the MQSeries Level 1 products are listed in “MQSeries Level 1 product publications” on page xv. For a functional comparison of the Level 1 and Level 2 MQSeries products, see the *MQSeries Planning Guide*.)

MQSeries Brochure

The *MQSeries Brochure*, G511-1908, gives a brief introduction to the benefits of MQSeries. It is intended to support the purchasing decision, and describes some authentic customer use of MQSeries.

MQSeries: An Introduction to Messaging and Queuing

MQSeries: An Introduction to Messaging and Queuing, GC33-0805, describes briefly what MQSeries is, how it works, and how it can solve some classic interoperability problems. This book is intended for a more technical audience than the *MQSeries Brochure*.

MQSeries Planning Guide

The *MQSeries Planning Guide*, GC33-1349, describes some key MQSeries concepts, identifies items that need to be considered before MQSeries is installed, including storage requirements, backup and recovery, security, and migration from earlier releases, and specifies hardware and software requirements for every MQSeries platform.

MQSeries Intercommunication

The *MQSeries Intercommunication* book, SC33-1872, defines the concepts of distributed queuing and explains how to set up a distributed queuing network in a variety of MQSeries environments. In particular, it demonstrates how to (1) configure communications to and from a representative sample of MQSeries products, (2) create required MQSeries objects, and (3) create and configure MQSeries channels. The use of channel exits is also described.

MQSeries Clients

The *MQSeries Clients* book, GC33-1632, describes how to install, configure, use, and manage MQSeries client systems.

MQSeries System Administration

The *MQSeries System Administration* book, SC33-1873, supports day-to-day management of local and remote MQSeries objects. It includes topics such as security, recovery and restart, transactional support, problem determination, the dead-letter queue handler, and the MQSeries links for Lotus Notes**. It also includes the syntax of the MQSeries control commands.

This book applies to the following MQSeries products only:

- MQSeries for AIX V5.0
- MQSeries for HP-UX V5.0
- MQSeries for OS/2 Warp V5.0
- MQSeries for Sun Solaris V5.0
- MQSeries for Windows NT V5.0

MQSeries Command Reference

The *MQSeries Command Reference*, SC33-1369, contains the syntax of the MQSC commands, which are used by MQSeries system operators and administrators to manage MQSeries objects.

MQSeries Programmable System Management

The *MQSeries Programmable System Management* book, SC33-1482, provides both reference and guidance information for users of MQSeries events, programmable command formats (PCFs), and installable services.

MQSeries Messages

The *MQSeries Messages* book, GC33-1876, which describes “AMQ” messages issued by MQSeries, applies to these MQSeries products only:

- MQSeries for AIX V5.0
- MQSeries for HP-UX V5.0
- MQSeries for OS/2 Warp V5.0
- MQSeries for Sun Solaris V5.0
- MQSeries for Windows NT V5.0
- MQSeries for Windows V2.0
- MQSeries for Windows V2.1

This book is available in softcopy only.

MQSeries Application Programming Guide

The *MQSeries Application Programming Guide*, SC33-0807, provides guidance information for users of the message queue interface (MQI). It describes how to design, write, and build an MQSeries application. It also includes full descriptions of the sample programs supplied with MQSeries.

MQSeries Application Programming Reference

The *MQSeries Application Programming Reference*, SC33-1673, provides comprehensive reference information for users of the MQI. It includes: data-type descriptions; MQI call syntax; attributes of MQSeries objects; return codes; constants; and code-page conversion tables.

MQSeries Application Programming Reference Summary

The *MQSeries Application Programming Reference Summary*, SX33-6095, summarizes the information in the *MQSeries Application Programming Reference* manual.

MQSeries Using C++

MQSeries Using C++, SC33-1877, provides both guidance and reference information for users of the MQSeries C++ programming-language binding to the MQI. MQSeries C++ is supported by V5.0 of MQSeries for AIX, HP-UX, OS/2 Warp, Sun Solaris, and Windows NT, and by MQSeries clients supplied with those products and installed in the following environments:

- AIX
- HP-UX
- OS/2
- Sun Solaris
- Windows NT
- Windows 3.1
- Windows 95

MQSeries platform-specific publications

Each MQSeries product is documented in at least one platform-specific publication, in addition to the MQSeries family books.

MQSeries for AIX

MQSeries for AIX V5.0 Quick Beginnings, GC33-1867

MQSeries for AT&T GIS UNIX

MQSeries for AT&T GIS UNIX Version 2.2 System Management Guide, SC33-1642

MQSeries for Digital OpenVMS

MQSeries for Digital OpenVMS Version 2.2 System Management Guide, GC33-1791

MQSeries for HP-UX

MQSeries for HP-UX V5.0 Quick Beginnings, GC33-1869

MQSeries for MVS/ESA

MQSeries for MVS/ESA Version 1 Release 2 Licensed Program Specifications, GC33-1350

MQSeries for MVS/ESA Version 1 Release 2 Program Directory

MQSeries for MVS/ESA Version 1 Release 2 System Management Guide, SC33-0806

MQSeries for MVS/ESA Version 1 Release 2 Messages and Codes, GC33-0819

MQSeries publications

MQSeries for MVS/ESA Version 1 Release 2 Problem Determination Guide, GC33-0808

MQSeries for OS/2 Warp

MQSeries for OS/2 Warp V5.0 Quick Beginnings, GC33-1868

MQSeries for OS/400

MQSeries for OS/400 Version 3 Release 2 Licensed Program Specifications, GC33-1360 (softcopy only)

MQSeries for OS/400 Version 3 Release 2 Administration Guide, GC33-1361

MQSeries for OS/400 Version 3 Release 2 Application Programming Reference (RPG), SC33-1362

Note: The MQSeries for OS/400 Version 3 Release 2 publications apply also to MQSeries for OS/400 Version 3 Release 7.

MQSeries link for R/3

MQSeries link for R/3 Version 1.0 User's Guide, GC33-1934

MQSeries for SINIX and DC/OSx

MQSeries for SINIX and DC/OSx Version 2.2 System Management Guide, GC33-1768

MQSeries for SunOS

MQSeries for SunOS Version 2.2 System Management Guide, GC33-1772

MQSeries for Sun Solaris

MQSeries for Sun Solaris V5.0 Quick Beginnings, GC33-1870

MQSeries Three Tier

MQSeries Three Tier Administration Guide, SC33-1451

MQSeries Three Tier Reference Summary, SX33-6098

MQSeries Three Tier Application Design, SC33-1636

MQSeries Three Tier Application Programming, SC33-1452

MQSeries for Windows

MQSeries for Windows Version 2.0 User's Guide, GC33-1822

MQSeries for Windows Version 2.1 User's Guide, GC33-1965

MQSeries for Windows NT

MQSeries for Windows NT V5.0 Quick Beginnings, GC33-1871

MQSeries Level 1 product publications

For information about the MQSeries Level 1 products, see the following publications:

MQSeries: Concepts and Architecture, GC33-1141

MQSeries Version 1 Products for UNIX Operating Systems Messages and Codes, SC33-1754

MQSeries for Digital VMS VAX Version 1.5 User's Guide, SC33-1144

MQSeries for SCO UNIX Version 1.4 User's Guide, SC33-1378

MQSeries for Tandem NonStop Kernel Version 1.5.1 User's Guide, SC33-1755

MQSeries for UnixWare Version 1.4.1 User's Guide, SC33-1379

MQSeries for VSE/ESA Version 1 Release 4 Licensed Program Specifications, GC33-1483

MQSeries for VSE/ESA Version 1 Release 4 User's Guide, SC33-1142

Softcopy books

Most of the MQSeries books are supplied in both hardcopy and softcopy formats.

BookManager format

The MQSeries library is supplied in IBM BookManager format on a variety of online library collection kits, including the *Transaction Processing and Data* collection kit, SK2T-0730. You can view the softcopy books in IBM BookManager format using the following IBM licensed programs:

BookManager READ/2
 BookManager READ/6000
 BookManager READ/DOS
 BookManager READ/MVS
 BookManager READ/VM
 BookManager READ for Windows

PostScript format

The MQSeries library is provided in PostScript (.PS) format with many MQSeries products, including all MQSeries V5.0 products. Books in PostScript format can be printed on a PostScript printer or viewed with a suitable viewer.

HTML format

The MQSeries documentation is provided in HTML format with these MQSeries products:

- MQSeries for AIX V5.0
- MQSeries for HP-UX V5.0
- MQSeries for OS/2 Warp V5.0
- MQSeries for Sun Solaris V5.0
- MQSeries for Windows NT V5.0

The MQSeries books are also available from the MQSeries software-server home page at URL:

<http://www.software.ibm.com/mqseries/>

Information Presentation Facility (IPF) format

In the OS/2 environment, the MQSeries documentation is supplied in IBM IPF format on the MQSeries product CD-ROM.

Windows Help format

The *MQSeries for Windows User's Guide* is provided in Windows Help format with MQSeries for Windows Version 2.0 and MQSeries for Windows Version 2.1.

MQSeries information available on the Internet

MQSeries URL

The URL of the MQSeries product family home page is:

<http://www.software.ibm.com/mqseries/>

Related publications

This section lists other documentation referred to in this book.

CICS in Open Systems Administration Reference, SC33-1533

Transaction Server for Windows NT, Version 4: Administration Guide (CICS), SC33-1881

Part 1. Guidance

Chapter 1. Introduction to MQSeries	7
MQSeries and message queuing	7
Time-independent applications	7
Message-driven processing	7
Messages and queues	8
What messages are	8
What queues are	8
Objects	9
Object names	9
Managing objects	10
MQSeries queue managers	10
MQSeries queues	11
Process definitions	13
Channels	14
System default objects	14
Administration	14
Local and remote administration	14
Clients and servers	15
MQSeries applications in a client-server environment	15
Extending queue manager facilities	15
User exits	16
Installable services	16
Security	17
Object Authority Management (OAM)	17
DCE Security	17
Transactional support	18
Chapter 2. Understanding administration command sets	19
Control commands	19
Using control commands	20
MQSeries commands (MQSC)	20
Running MQSC commands	21
PCF commands	21
Attributes in MQSC and PCFs	22
Escape PCFs	22
Comparing command sets	22
Chapter 3. Managing queue managers	27
Getting started	27
Guidelines for creating queue managers	27
Specifying a unique queue manager name	28
Limiting the number of queue managers	28
Specifying the default queue manager	28
Specifying a dead-letter queue	29
Specifying a default transmission queue	29
Specifying the required logging parameters	30
Backing up configuration files after creating a queue manager	30
Understanding MQSeries file names	30
Queue manager name transformation	31
Object name transformation	32

Working with queue managers	32
Creating a default queue manager	32
Starting a queue manager	33
Stopping a queue manager	34
Restarting a queue manager	35
Making an existing queue manager the default	36
Deleting a queue manager	36
Managing the command server for remote administration	37
Starting the command server	37
Displaying the status of the command server	37
Stopping a command server	38
Chapter 4. Administering local MQSeries objects	39
Supporting application programs that use the MQI	39
Issuing MQSC commands for administration	40
Before you start	40
Using the MQSC facility interactively	41
Feedback from MQSC commands	42
Ending interactive input to MQSC	42
Displaying queue manager attributes	42
Using a queue manager that is not the default	44
Altering queue manager attributes	44
Running MQSC commands from text files	44
MQSC command files	45
MQSC reports	46
Running the supplied MQSC command files	47
Using runmqsc to verify commands	47
Resolving problems with MQSC	47
Working with local queues	49
Defining a local queue	49
Defining a dead-letter queue	50
Displaying default object attributes	51
Copying a local queue definition	51
Changing local queue attributes	52
Clearing a local queue	53
Deleting a local queue	53
Browsing queues	53
Working with alias queues	55
Defining an alias queue	55
Using other commands with alias queues	57
Working with model queues	57
Defining a model queue	58
Using other commands with model queues	58
Managing objects for triggering	59
Defining an application queue for triggering	59
Defining an initiation queue	60
Creating a process definition	60
Displaying your process definition	61
Fastpath binding	61

Chapter 5. Administering remote MQSeries objects	63
Understanding channels and remote queuing	63
Remote administration	64
Preparing queue managers for remote administration	64
Preparing channels and transmission queues for remote administration	65
Defining channels and transmission queues	66
Start the channels	67
Issuing MQSC commands remotely	69
Working with queue managers on MVS/ESA	69
If you have problems using MQSC remotely	70
Creating a local definition of a remote queue	71
Understanding how local definitions of remote queues work	71
An alternative way of putting messages on a remote queue	72
Using other commands with remote queues	73
Creating a transmission queue	73
Using remote queue definitions as aliases	74
Queue manager aliases	74
Reply-to queue aliases	74
Data conversion	75
File ccsid.tbl	75
Conversion of messages in user-defined formats	76
Chapter 6. Protecting MQSeries objects	77
Before you begin (UNIX systems)	77
User IDs in user group mqm (UNIX systems)	77
Before you begin (Windows NT)	78
User IDs (Windows NT systems)	78
Restricted access NT objects	79
Why you need to protect MQSeries resources	79
Understanding the Object Authority Manager	79
How the OAM works	80
Managing access through user groups	80
Default user group	81
Resources you can protect with the OAM	81
Using groups for authorizations	81
Disabling the object authority manager	82
Using the Object Authority Manager commands	82
What you specify when you use the OAM commands	83
Using the setmqaut command	83
Access authorizations	84
Display authority command	84
Object Authority Manager guidelines	85
User IDs	85
Queue manager directories	85
Queues	85
Alternate-user authority	85
Context authority	86
Remote security considerations	87
Channel command security	87
Understanding the authorization specification tables	88
MQI authorizations	89
Administration authorizations	91
Authorizations for MQSC commands in escape PCFs	92

Understanding authorization files	94
Authorization file paths	94
What the authorization files contain	96
Managing authorization files	97
Chapter 7. Configuration files	99
What configuration files are	99
MQSeries configuration file	99
What the MQSeries configuration file contains	99
Queue manager configuration file	103
What the queue manager configuration file contains	103
Editing configuration files	108
Changing the default prefix	108
Implementing changes to configuration files	108
Recommendations for configuration files	108
Configuring the logs	108
Log configuration stanzas	109
Specifying log file sizes	112
Configuring XA_-compliant databases	112
Chapter 8. The MQSeries dead-letter queue handler	115
Invoking the DLQ handler	115
The sample DLQ handler, amqsdliq	116
The DLQ handler rules table	116
Control data	117
Rules (patterns and actions)	118
Rules table conventions	121
How the rules table is processed	123
Ensuring that all DLQ messages are processed	124
An example DLQ handler rules table	125
Chapter 9. Instrumentation events	127
What instrumentation events are	127
Why use events?	128
Types of event	129
Event notification through event queues	130
Enabling and disabling events	130
Event messages	131
Chapter 10. Linking to Lotus Notes	133
What is Lotus Notes?	133
Linking applications	133
Add-in task requirements	135
Server or client?	135
Disconnected requests	137
Setting up your system	137
Lotus Notes setup	137
Setting up the server add-in task	137
Starting the server add-in task	138
Verifying that Lotus Notes can link to MQSeries	139
Stopping the link server task	140
User notification	140

Chapter 11. Transactional support	141
Database coordination	142
Restrictions	143
Database connections	143
Configuring database managers	144
DB2 configuration	146
Checking the environment variable settings	146
Creating the DB2 switch load file	146
Adding the XAResourceManager stanza to the qm.ini file	150
Changing DB2 configuration parameters	151
Oracle configuration	152
Checking Oracle level and applying patches	152
Checking the environment variable settings	152
Enabling Oracle XA support	152
Creating the Oracle switch load file	153
Adding the XAResourceManager stanza to the qm.ini file	154
Changing the Oracle configuration parameters	156
Multiple database configurations	157
Security considerations	157
Administration tasks	159
In-doubt units of work	159
Using the dspmqtrn command	160
Using the rsvmqtrn command	161
Mixed outcomes and errors	162
Changing the qm.ini configuration file	163
External syncpoint coordination	164
The MQSeries XA switch structure	165
Using CICS	166
The CICS two-phase commit process	166
The CICS single-phase commit process	168
Chapter 12. Recovery and restart	171
Making sure that messages are not lost (logging)	171
What logs look like	171
Types of logging	172
Checkpointing – ensuring complete recovery	174
Dumping the contents of the log using the dmpmqlog command	176
Managing logs	194
What happens when a disk gets full	194
Managing log files	195
Using the log for recovery	197
Recovering from problems	197
Media recovery	197
Recovering damaged objects during startup	199
Recovering damaged objects at other times	199
Backup and restore	200
Backing up MQSeries	200
Restoring MQSeries	200
Recovery scenarios	201
Disk drive failures	201
Damaged queue manager object	202
Damaged single object	202
Automatic media recovery failure	202

Chapter 13. Problem determination	203
Preliminary checks	203
Has MQSeries run successfully before?	204
Are there any error messages?	204
Are there any return codes explaining the problem?	204
Can you reproduce the problem?	204
Have any changes been made since the last successful run?	204
Has the application run successfully before?	205
Problems with commands	206
Does the problem affect specific parts of the network?	206
Does the problem occur at specific times of the day?	206
Is the problem intermittent?	207
Have you applied any service updates?	207
What to do next	207
Have you obtained incorrect output?	208
Have you failed to receive a response from a PCF command?	208
Are some of your queues failing?	209
Does the problem affect only remote queues?	210
Is your application or system running slowly?	210
Application design considerations	211
Effect of message length	211
Effect of message persistence	211
Searching for a particular message	211
Queues that contain messages of different lengths	211
Frequency of syncpoints	212
Use of the MQPUT1 call	212
Number of threads in use	212
Incorrect output	212
Messages that do not appear on the queue	212
Messages that contain unexpected or corrupted information	214
Problems with incorrect output when using distributed queues	214
Error logs	215
Log files	216
Early errors	217
Operator messages	217
An example error log	217
The MQSeries log-dump utility	218
Dead-letter queues	219
Configuration files and problem determination	219
Tracing	219
Tracing MQSeries for AIX	219
Tracing MQSeries for HP-UX and MQSeries for Sun Solaris	222
Tracing MQSeries for OS/2 Warp and MQSeries for Windows NT	224
First-failure support technology (FFST)	226
FFST: MQSeries for UNIX systems	226
FFST: MQSeries for OS/2 Warp and Windows NT	227
FFST: MQSeries for OS/2 Warp	228
Problem determination with clients	230
Terminating clients	230
Error messages with clients	230

Chapter 1. Introduction to MQSeries

This chapter introduces the MQSeries Version 5 products from an administrator's perspective, and describes the basic concepts of MQSeries and messaging. It contains these sections:

- "MQSeries and message queuing"
- "Messages and queues" on page 8
- "Objects" on page 9
- "System default objects" on page 14
- "Administration" on page 14
- "Clients and servers" on page 15
- "Extending queue manager facilities" on page 15
- "Security" on page 17
- "Transactional support" on page 18

MQSeries and message queuing

MQSeries lets applications use message queuing to participate in message-driven processing. Applications can communicate across different platforms by using the appropriate message queuing software products. For example, HP-UX and MVS/ESA applications can communicate through MQSeries for HP-UX and MQSeries for MVS/ESA respectively. The applications are shielded from the mechanics of the underlying communications.

MQSeries products implement a common application programming interface (message queue interface or MQI) whatever platform the applications are run on. This makes it easier to port applications from one platform to another.

The MQI is described in detail in the *MQSeries Application Programming Reference manual*.

Time-independent applications

With message queuing, the exchange of messages between the sending and receiving programs is time independent. This means that the sending and receiving applications are decoupled so that the sender can continue processing without having to wait for the receiver to acknowledge the receipt of the message. In fact, the target application does not even have to be running when the message is sent. It can retrieve the message after it is started.

Message-driven processing

On arriving on a queue, messages can automatically start an application using a mechanism known as *triggering*. If necessary, the applications can be stopped when the message or messages have been processed.

Messages and queues

Messages and queues are the basic components of a message queuing system.

What messages are

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

MQSeries messages have two parts; the *application data* and a *message descriptor*. The content and structure of the application data is defined by the application programs that use them. The message descriptor identifies the message and contains other control information, such as the type of message and the priority assigned to the message by the sending application.

The format of the message descriptor is defined by MQSeries. For a complete description of the message descriptor, see the *MQSeries Application Programming Reference* manual.

Message lengths

The maximum message length is 100 MB (where 1 MB equals 1 048 576 bytes). In practice, the message length may be limited by:

- The maximum message length defined for the receiving queue
- The maximum message length defined for the queue manager
- The maximum message length defined by either the sending or receiving application
- The amount of storage available for the message

It may take several messages to send all the information that an application requires.

What queues are

A *queue* is a data structure that stores zero or more messages. The messages may be put on the queue by applications or by a queue manager as part of its normal operation.

Each queue belongs to a *queue manager*, which is responsible for maintaining it. The queue manager puts the messages it receives on the appropriate queues.

Applications send and receive messages using MQI calls. For example, one application can put a message on a queue, and another application can retrieve the message from the same queue. The sending application opens the queue for put operations by making an MQOPEN call. Then it issues an MQPUT call to put the message onto that queue. When the receiving application opens the same queue for gets, it can retrieve the message from the queue by issuing an MQGET call.

For more information about MQI calls, see the *MQSeries Application Programming Reference* manual.

Predefined and dynamic queues

Queues can be characterized by the way they are created:

- *Predefined queues* are created by an administrator using the appropriate command set. For example, the MQSC command DEFINE QLOCAL creates a predefined local queue. Predefined queues are permanent; they exist independently of the applications that use them and survive MQSeries restarts.
- *Dynamic queues* are created when an application issues an open request specifying the name of a model queue. The queue created is based on a template queue definition, which is the model queue. You can create a model queue using the MQSC command DEFINE QMODEL. The attributes of a model queue, for example the maximum number of messages that can be stored on it, are inherited by any dynamic queue that is created from it.

Model queues have an attribute that specifies whether the dynamic queue is to be permanent or temporary. Permanent queues survive application and queue manager restarts; temporary queues are lost on restart.

Retrieving messages from queues

Suitably authorized applications can retrieve messages from a queue according to these retrieval algorithms:

- First-in-first-out (FIFO).
- Message priority, as defined in the message descriptor. Messages that have the same priority are retrieved on a FIFO basis.
- A program request for a specific message.

The MQGET request from the application determines the method used.

Objects

Many of the tasks described in this book involve manipulating MQSeries *objects*. In the MQSeries Version 5 products, there are four different types of object:

- Queue managers; see “MQSeries queue managers” on page 10.
- Queues; see “MQSeries queues” on page 11.
- Process definitions; see “Process definitions” on page 13.
- Channels; see “Channels” on page 14.

Object names

Each instance of a queue manager is known by its name. This name must be unique within the network of interconnected queue managers, so that one queue manager can unambiguously identify the target queue manager to which any given message should be sent.

For the other types of object, each object has a name associated with it and can be referenced by that name. These names must be unique within one queue manager and object type. For example, you can have a queue and a process with the same name, but you cannot have two queues with the same name.

In MQSeries, names can have a maximum of 48 characters, with the exception of *channels*, which have a maximum of 20 characters. For more information about names see “Names of MQSeries objects” on page 235.

Managing objects

MQSeries provides commands for creating, altering, displaying, and deleting objects. These include:

- MQSeries commands (MQSC), which can be typed in from a keyboard or read from a file.
- Programmable Command Format (PCF) commands, which can be used in a program.

For more information, see Chapter 2, “Understanding administration command sets” on page 19.

Object attributes

The properties of an object are defined by its attributes. Some you can specify, others you can only view. For example, the maximum message length that a queue can accommodate is defined by its *MaxMsgLength* attribute; you can specify this attribute when you create a queue. The *DefinitionType* attribute specifies how the queue was created; you can only display this attribute.

In MQSeries, there are two ways of referring to an attribute:

- Using its PCF name, for example, *MaxMsgLength*.
- Using its MQSC name, for example, MAXMSGL.

The formal name of an attribute is its PCF name. Because using the MQSC facility is an important part of this book, you are more likely to see the MQSC name in examples than the PCF name of a given attribute.

MQSeries queue managers

A queue manager provides queuing services to applications, and manages the queues that belong to it. It ensures that:

- Object attributes are changed according to the commands received.
- Special events such as trigger events or instrumentation events are generated when the appropriate conditions are met.
- Messages are put on the correct queue, as requested by the application making the MQPUT call. The application is informed if this cannot be done, and an appropriate reason code is given.

Each queue belongs to a single queue manager and is said to be a *local queue* to that queue manager. The queue manager to which an application is connected is said to be the local queue manager for that application. For the application, the queues that belong to its local queue manager are local queues. A *remote queue* is simply a queue that belongs to another queue manager. A *remote queue manager* is any queue manager other than the local queue manager. A remote queue manager may exist on a remote machine across the network or it may exist on the same machine as the local queue manager. MQSeries supports multiple queue managers on the same machine.

MQI calls

A queue manager object may be used in some MQI calls. For example, you can inquire about the attributes of the queue manager object using the MQI call MQINQ.

Note: You cannot put messages on a queue manager object; messages are always put on queue objects, not on queue manager objects.

MQSeries queues

Queues are defined to MQSeries using the appropriate MQSC DEFINE command or the PCF Create Queue command. The command specifies the type of queue and its attributes. For example, a local queue object has attributes that specify what happens when applications reference that queue in MQI calls. Examples of attributes are:

- Whether applications can retrieve messages from the queue (GET enabled).
- Whether applications can put messages on the queue (PUT enabled).
- Whether access to the queue is exclusive to one application or shared between applications.
- The maximum number of messages that can be stored on the queue at the same time (maximum queue depth).
- The maximum length of messages that can be put on the queue.

For further details about defining queue objects, see the *MQSeries Command Reference* or the *MQSeries Programmable System Management* manual.

Using queue objects

In MQSeries, there are four types of queue object. Each type of object can be manipulated by the product commands and is associated with real queues in different ways:

1. A *local queue* object identifies a local queue belonging to the queue manager to which the application is connected. All queues are local queues in the sense that each queue belongs to a queue manager and, for that queue manager, the queue is a local queue.
2. A *remote queue object* identifies a queue belonging to another queue manager. This queue must be defined as a local queue to that queue manager. The information you specify when you define a remote queue object allows the local queue manager to find the remote queue manager, so that any messages destined for the remote queue go to the correct queue manager.

You must also define a transmission queue and channels between the queue managers, before applications can send messages to a queue on another queue manager.

3. An *alias queue object* allows applications to access a queue by referring to it indirectly in MQI calls. When an alias queue name is used in an MQI call, the name is resolved to the name of either a local or a remote queue at run time. This allows you to change the queues that applications use without changing the application in any way—you merely change the alias queue definition to reflect the name of the new queue to which the alias resolves.

An alias queue is not a queue, but an object that you can use to access another queue.

4. A *model queue object* defines a set of queue attributes that are used as a template for creating a dynamic queue. Dynamic queues are created by the queue manager when an application issues an MQOPEN request specifying a queue name that is the name of a model queue. The dynamic queue that is created in this way is a local queue whose attributes are taken from the model queue definition. The dynamic queue name can be specified by the application or the queue manager can generate the name and return it to the application.

Dynamic queues defined in this way may be temporary queues, which do not survive product restarts, or permanent queues, which do.

Specific local queues used by MQSeries

MQSeries uses some local queues for specific purposes related to its operation. You *must* define them before MQSeries can use them.

Application queues: A queue that is used by an application (through the MQI) is referred to as an *application queue*. This can be a local queue on the queue manager to which an application is linked, or it can be a remote queue that is owned by another queue manager.

Applications can put messages on local or remote queues. However, they can only get messages from a local queue.

Initiation queues: *Initiation queues* are queues that are used in triggering. A queue manager puts a trigger message on an initiation queue when a trigger event occurs. A trigger event is a logical combination of conditions that is detected by a queue manager. For example, a trigger event may be generated when the number of messages on a queue reaches a predefined depth. This event causes the queue manager to put a trigger message on a specified initiation queue. This trigger message is retrieved by a *trigger monitor*, a special application that monitors an initiation queue. The trigger monitor then starts up the application program that was specified in the trigger message.

If a queue manager is to use triggering, at least one initiation queue must be defined for that queue manager.

See “Managing objects for triggering” on page 59, and “runmqtrm (Start trigger monitor)” on page 281. For more information about triggering, see the *MQSeries Application Programming Guide*.

Transmission queues: A *transmission queue* temporarily stores messages that are destined for a remote queue manager. You must define at least one transmission queue for each remote queue manager to which the local queue manager is to send messages directly. These queues are also used in remote administration; see “Remote administration” on page 64. For information about the use of transmission queues in distributed queuing, see the *MQSeries Intercommunication* book.

Dead-letter queues: A *dead-letter queue* stores messages that cannot be routed to their correct destinations. This occurs when, for example, the destination queue is full. The supplied dead-letter queue is called SYSTEM.DEAD.LETTER.QUEUE. These queues are sometimes referred to as undelivered-message queues.

For distributed queuing, you should define a dead-letter queue on each queue manager involved.

Command queues: The command queue, named SYSTEM.ADMIN.COMMAND.QUEUE, is a local queue to which suitably authorized applications can send MQSeries commands for processing. These commands are then retrieved by an MQSeries component called the command server. The command server validates the commands, passes the valid ones on for processing by the queue manager, and returns any responses to the appropriate reply-to queue.

A command queue is created automatically for each queue manager when that queue manager is created.

Reply-to queues: When an application sends a request message, the application that receives the message can send back a reply message to the sending application. This message is put on a queue, called a reply-to queue, which is normally a local queue to the sending application. The name of the reply-to queue is specified by the sending application as part of the message descriptor.

Event queues: The MQSeries Version 5 products support instrumentation events, which can be used to monitor queue managers independently of MQI applications. Instrumentation events can be generated in several ways, for example:

- An application attempting to put a message on a queue that is not available or does not exist.
- A queue becoming full.
- A channel being started.

When an instrumentation event occurs, the queue manager puts an event message on an event queue. This message can then be read by a monitoring application which may inform an administrator or initiate some remedial action if the event indicates a problem.

Note: Trigger events are quite different from instrumentation events in that trigger events are not caused by the same conditions, and do not generate event messages.

For more information about instrumentation events, see the *MQSeries Programmable System Management* manual.

Process definitions

A *process definition object* defines an application that is to be started in response to a trigger event on an MQSeries queue manager. See “Initiation queues” on page 12 for more information.

The process definition attributes include the application ID, the application type, and data specific to the application.

Use the MQSC command DEFINE PROCESS or the PCF command Create Process to create a process definition.

Channels

Channels are objects that provide a communication path from one queue manager to another. Channels are used in distributed message queuing to move messages from one queue manager to another. They shield applications from the underlying communications protocols. The queue managers may exist on the same, or different, platforms. For queue managers to communicate with one another, you must define one channel object at the queue manager that is to send messages, and another, complementary one, at the queue manager that is to receive them.

For information on channels and how to use them, see the *MQSeries Intercommunication* book, and also “Preparing channels and transmission queues for remote administration” on page 65.

System default objects

The *system default objects* are a set of object definitions that are created automatically whenever a queue manager is created. You can copy and modify any of these object definitions for use in applications at your installation. Default object names have the stem SYSTEM.DEF; for example, the default local queue is SYSTEM.DEFAULT.LOCAL.QUEUE, and the default receiver channel is SYSTEM.DEF.RECEIVER. You cannot rename these objects; default objects of these names are required.

When you define an object, any attributes that you do not specify explicitly are copied from the appropriate default object. For example, if you define a local queue, those attributes you do not specify are taken from the default queue SYSTEM.DEFAULT.LOCAL.QUEUE.

Administration

In MQSeries, you carry out administration tasks by issuing *commands*. Three command sets are provided, depending on which tasks you want to perform and how you want to perform them. The command sets are described in Chapter 2, “Understanding administration command sets” on page 19. Administration tasks include:

- 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. This is described in detail in the *MQSeries Intercommunication* book.

Local and remote administration

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

MQSeries supports administration from a single point through what is known as *remote administration*. This allows you to issue commands from your local system

that are processed on another system. You do not have to log on to that system, although you do need to have the appropriate channels defined. The queue manager and command server on the target system must be running. For example, you can issue a remote command to change a queue definition on a remote queue manager.

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

Clients and servers

MQSeries supports client-server configurations for MQSeries applications.

An *MQSeries client* is a part of the MQSeries product that is installed on a machine to accept MQI calls from applications and pass them to an *MQI server* machine. There they are processed by a queue manager. Typically, the client and server reside on different machines but they can also exist on the same machine.

An *MQI server* is a queue manager that provides queuing services to one or more clients. All the MQSeries objects, for example queues, exist only on the queue manager machine, that is, on the MQI server machine. A server can support normal local MQSeries applications as well.

The difference between an MQI server and an ordinary queue manager is that a server has a dedicated communications link with each client. For more information about creating channels for clients and servers, see the *MQSeries Intercommunication* book. For information about client support in general, see the *MQSeries Clients* book.

MQSeries applications in a client-server environment

When linked to a server, client MQSeries applications can issue most MQI calls in the same way as local applications. The client application issues an MQCONN call to connect to a specified queue manager. Any additional MQI calls that specify the connection handle returned from the connect request are then processed by this queue manager. You must link your applications to the appropriate client libraries. See the *MQSeries Application Programming Guide* for further information.

Extending queue manager facilities

The facilities provided by a queue manager can be extended by:

- User exits
- Installable services

User exits

User exits provide a mechanism for users to insert their own code into a queue manager function. Two types of user exit are supported:

- *Channel exits*, which change the way that channels operate. Channel exits are described in the *MQSeries Intercommunication* book.
- *Data conversion exits*, which create source code fragments that can be put into application programs to convert data from one format to another. Data conversion exits are described in the *MQSeries Application Programming Guide*.

Installable services

Installable services are more extensive than exits in that they have formalized interfaces (an API) with multiple entry points.

An implementation of an installable service is called a *service component*. You can use the components supplied with the MQSeries product, or you can write your own component to perform the functions that you require. Currently, the following installable services are provided:

Authorization service

The authorization service allows you to build your own security facility.

The default service component that implements the service is the Object Authority Manager (OAM), which is supplied with MQSeries for UNIX systems and the MQSeries for Windows NT product. (The OAM is not supplied with MQSeries for OS/2 Warp.) By default, the OAM is active, and you do not have to do anything to configure it. You can use the authorization service interface to create other components to replace or augment the OAM. For more information about the OAM, see Chapter 6, "Protecting MQSeries objects" on page 77.

Under MQSeries for OS/2 Warp, you must write your own service component if you want to implement the authorization service. For example, you can create your own security features based on a third-party security product.

Name service

The name service enables the sharing of queues by allowing applications to identify remote queues as though they were local queues.

A default service component that implements the name service is provided with the MQSeries Version 5 products. It uses the Open Software Foundation (OSF) Distributed Computing Environment (DCE). You can also write your own name service component. (You might want to do this if you do not have DCE installed, for example.) By default, the name service is inactive.

For more information, see the *MQSeries Programmable System Management* book.

User identifier service

The user identifier service is supported by MQSeries for OS/2 Warp only. It allows MQI applications in an OS/2 environment to associate a user ID (other than the default user ID, OS2) with MQSeries messages. The receiving applications are then able to identify the source of the messages. A sample user identifier service component is supplied.

Note that this is not intended to provide a **secure** service. There is no mechanism to prevent applications from copying this user ID.

For more information, see Appendix F, "User identifier service" on page 319.

See the *MQSeries Programmable System Management* manual for more information about the installable services.

Security

In the MQSeries Version 5 products, there are two methods of providing security.

Object Authority Management (OAM)

In MQSeries for UNIX systems and MQSeries for Windows NT, authorization for using MQI calls, commands, and access to objects is provided by the Object Authority Manager (OAM), which by default is enabled. Access to MQSeries entities is controlled through MQSeries user groups and the OAM. A command line interface is provided to enable administrators to grant or revoke authorizations as required.

No OAM security features are provided either by MQSeries for OS/2 Warp or by OS/2 itself. You should consider what your security requirements are, and design your system to provide these facilities or, in their absence, to ensure that your applications are aware of the lack of security and are not therefore compromised.

Note: The authorization service is available in MQSeries for OS/2 Warp, but no authorization service component is supplied. If security is essential to your enterprise, you could consider writing your own authorization service component. This component would use the supplied interface to access the facilities provided by a third-party security manager.

For more information about creating authorization service components, see the *MQSeries Programmable System Management* book.

DCE Security

Channel exits that use the DCE Generic Security Service (GSS) are provided by MQSeries. For more information, see the *MQSeries Intercommunication* book.

Transactional support

In the MQSeries Version 5 products, there are three ways of updating resources under syncpoint control. They are:

- Local unit of work

The MQSeries queue manager is the only participant in a local unit of work (UOW). Changes made to MQSeries resources are committed using MQCMIT or backed out using MQBACK. This is a single-phase commit process.

- Global unit of work

A global UOW is coordinated by an MQSeries queue manager, but can include calls to XA-compliant external database managers, such as DB2 and Oracle**. The global UOW uses a two-phase commit process.

- Internal coordination

Global units of work are started using the MQBEGIN verb and then committed using MQCMIT or backed out using MQBACK. A two-phase commit process is used whereby XA-compliant database managers such as DB2 and Oracle are firstly all asked to prepare to commit. Only if all are prepared will they then be asked to commit. If any resource manager signals that it cannot commit, each will be asked to back out instead.

- External coordination

Here the coordination is performed by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo. Units of work are started and committed under control of the transaction manager. The MQBEGIN, MQCMIT and MQBACK verbs are unavailable.

For more information, see Chapter 11, “Transactional support” on page 141.

Chapter 2. Understanding administration command sets

Read this chapter for an overview of the different methods that you can use to perform system administration tasks on MQSeries objects. This chapter also helps you to understand the different methods, and when each should be used.

Administration tasks include creating, starting, altering, viewing, stopping, and deleting MQSeries objects, that is, queue managers, queues, processes, and channels. To perform these tasks, you must select the appropriate command from one of the supplied command sets.

MQSeries provides three command sets for invoking administration tasks:

- Control commands
- MQSC commands
- PCF commands

This chapter describes the command sets that are available and provides a summary of the different commands in “Comparing command sets” on page 22.

Control commands

Control commands fall into three categories:

- *Queue manager commands*, including commands for creating, starting, stopping, and deleting queue managers and command servers.
- *Channel commands*, including commands for starting and ending channels and channel initiators.
- *Utility commands*, including commands associated with:
 - Running MQSC commands
 - Conversion exits
 - Authority management
 - Recording and recovering media images of queue manager resources
 - Displaying and resolving transactions
 - Trigger monitors
 - Displaying the file names of MQSeries objects

Using control commands

In MQSeries for UNIX systems, you enter control commands in a shell window. In these environments, control commands, including the command name itself, the flags, and any arguments, are case sensitive. For example, in the command:

```
crtmqm -u SYSTEM.DEAD.LETTER.QUEUE jupiter.queue.manager
```

- The command name must be `crtmqm`, not `CRTMQM`.
- The flag must be `-u`, not `-U`.
- The dead-letter queue is `SYSTEM.DEAD.LETTER.QUEUE`.
- The argument is specified as `jupiter.queue.manager`, which is different from `JUPITER.queue.manager`.

Therefore, take care to type the commands exactly as you see them in the examples.

In MQSeries for OS/2 Warp and Windows NT, you enter control commands at a command prompt. In these environments, control commands and their flags are not case sensitive, but arguments to those commands (such as queue names and queue-manager names) are case sensitive. For example, in the command:

```
crtmqm /u SYSTEM.DEAD.LETTER.QUEUE jupiter.queue.manager
```

- The command name can be entered in uppercase or lowercase, or a mixture of the two. These are all valid: `crtmqm`, `CRTMQM`, and `CRTmqm`.
- The flag can be entered as `-u`, `-U`, `/u`, or `/U`.
- The arguments `SYSTEM.DEAD.LETTER.QUEUE` and `jupiter.queue.manager` must be entered exactly as shown.

Chapter 14, “MQSeries control commands” on page 235 describes the syntax and purpose of each command.

MQSeries commands (MQSC)

You use the MQSeries (MQSC) commands to manage queue manager objects, including the queue manager itself, channels, queues, and process definitions. For example, there are commands to define, alter, display, and delete a specified queue.

When you display a queue, using the `DISPLAY QUEUE` command, you display the queue *attributes*. For example, the `MAXMSGL` attribute specifies the maximum length of a message that can be put on the queue. The command does not show you the messages on the queue.

MQSC commands are available on other platforms, including AS/400, and MVS/ESA.

These commands are summarized in “Comparing command sets” on page 22. For detailed information about each MQSC command, see the *MQSeries Command Reference*.

Running MQSC commands

You run MQSC commands by invoking the control command **runmqsc**. You can run MQSC commands:

- Interactively by typing them at the keyboard. See “Using the MQSC facility interactively” on page 41.
- As a sequence of commands from an ASCII text file. See “Running MQSC commands from text files” on page 44.

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

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

For more information about using the MQSC facility and text files, see “Using the MQSC facility interactively” on page 41. For more information about the **runmqsc** command, see “runmqsc (Run MQSeries commands)” on page 277.

PCF commands

The purpose of the MQSeries programmable command format (PCF) commands is to allow administration tasks to be programmed into an administration program. In this way you can create queues and process definitions, and change queue managers, from a program. In fact, PCF commands cover the same range of functions that are provided by the MQSC facility. You can therefore write a program to issue PCF commands to any queue manager in the network from a single node. In this way, you can both centralize and automate administration tasks.

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

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

Comparing command sets

Briefly, these are some of the things the application programmer must specify to create a PCF command message:

Message descriptor

This is a standard MQSeries message descriptor, in which:

Message type (*MsgType*) is MQMT_REQUEST.

Message format (*Format*) is MQFMT_ADMIN.

Application data

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

The PCF message type (*Type*) specifies MQCFT_COMMAND.

The command identifier specifies the command, for example, *Change Queue* (MQCMD_CHANGE_Q).

For a complete description of the PCF data structures and how to implement them, see the *MQSeries Programmable System Management* manual.

Attributes in MQSC and PCFs

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

Object attributes in PCF, which are not limited to eight characters, are shown in this book in italics. For example, the PCF equivalent of RQMNAME is *RemoteQMgrName*.

Escape PCFs

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

Comparing command sets

The following tables compare the facilities available from the different administration command sets.

Note: Only those MQSC commands that are supported by the MQSeries Version 5 products are shown.

<i>Table 1. Commands for queue manager administration</i>		
PCF	MQSC	Control
Change Queue Manager	ALTER QMGR	–
(Create queue manager) ¹	–	crtmqm
(Delete queue manager) ¹	–	dltmqm
Inquire Queue Manager	DISPLAY QMGR	–
(Stop queue manager) ¹	–	endmqm
Ping Queue Manager	PING QMGR	–
(Start queue manager) ¹	–	strmqm
(Service Control Manager) ¹	–	scmmqm ²
Notes: 1. Not available as PCF commands 2. Supported by MQSeries for Windows NT only		

<i>Table 2. Commands for command server administration</i>	
Description	Control
Display command server	dspmqcsv
Start command server	strmqcsv
Stop command server	endmqcsv
Note: Functions in this group are available only as control commands. There are no MQSC or PCF equivalents.	

Comparing command sets

<i>Table 3. Commands for queue administration</i>	
PCF	MQSC
Change Queue	ALTER QLOCAL ALTER QALIAS ALTER QMODEL ALTER QREMOTE
Clear Queue	CLEAR QUEUE
Copy Queue	DEFINE QLOCAL(x) LIKE(y) DEFINE QALIAS(x) LIKE(y) DEFINE QMODEL(x) LIKE(y) DEFINE QREMOTE(x) LIKE(y)
Create Queue	DEFINE QLOCAL DEFINE QALIAS DEFINE QMODEL DEFINE QREMOTE
Delete Queue	DELETE QLOCAL DELETE QALIAS DELETE QMODEL DELETE QREMOTE
Inquire Queue	DISPLAY QUEUE
Inquire Queue Names	DISPLAY QUEUE
<p>Note:</p> <p>There are no control-command equivalents for these MQSC and PCF commands.</p>	

<i>Table 4. Commands for process administration</i>	
PCF	MQSC
Change Process	ALTER PROCESS
Copy Process	DEFINE PROCESS(x) LIKE(y)
Create Process	DEFINE PROCESS
Delete Process	DELETE PROCESS
Inquire Process	DISPLAY PROCESS
Inquire Process Names	DISPLAY PROCESS
<p>Note:</p> <p>There are no control-command equivalents for these MQSC and PCF commands.</p>	

<i>Table 5. Commands for channel administration</i>		
PCF	MQSC	Control
Change Channel	ALTER CHANNEL	–
Copy Channel	DEFINE CHANNEL(x) LIKE(y)	–
Create Channel	DEFINE CHANNEL	–
Delete Channel	DELETE CHANNEL	–
Inquire Channel	DISPLAY CHANNEL	–
Inquire Channel Names	DISPLAY CHANNEL	–
Ping Channel	PING CHANNEL	–
Reset Channel	RESET CHANNEL	–
Resolve Channel	RESOLVE CHANNEL	–
Start Channel	START CHANNEL	runmqchl
Start Channel Initiator	START CHINIT	runmqchi
Start Channel Listener 1	–	runmqlsr 2
Stop Channel	STOP CHANNEL	–
<p>Notes:</p> <ol style="list-style-type: none"> 1. This PCF command is not supported by MQSeries for UNIX systems. You set up the inetd daemon or SNA product for this task, as described in the <i>MQSeries Intercommunication</i> book. 2. Supported by MQSeries for OS/2 Warp and MQSeries for Windows NT only. 		

Comparing command sets

<i>Table 6. Other control commands. (See note 1.)</i>	
Description	Control
Create MQSeries conversion exit	crtmqcvx
Dump MQSeries log	dmpmqlog
Display authority	dspmqaut
Display files used by objects	dspmqfls
Display MQSeries formatted trace	dspmqtrc ³
Display MQSeries transactions	dspmqtrn
End MQSeries trace	endmqtrc ²
Record media image	rcdmqimg
Recreate media object	rcrmqobj
Resolve MQSeries transactions	rsvmqtrn
Run dead-letter queue handler	runmqdlq
Run MQSC commands	runmqsc
Run trigger monitor	runmqtrm
Run client trigger monitor	runmqtrmc
Set or reset authority	setmqaut
Enroll production license	setmqprd
Start trial period	setmqtry
Start MQSeries trace	strmqtrc ²
<p>Notes:</p> <ol style="list-style-type: none"> 1. Functions in this group are available only as control commands. There are no direct PCF or MQSC equivalents. 2. Not supported by MQSeries for AIX. 3. Supported by MQSeries for HP-UX and MQSeries for Sun Solaris only. 	

Chapter 3. Managing queue managers

This chapter describes how you can perform operations on queue managers and command servers. It contains these sections:

- “Getting started”
- “Guidelines for creating queue managers”
- “Understanding MQSeries file names” on page 30
- “Working with queue managers” on page 32
- “Managing the command server for remote administration” on page 37

Getting started

Before you can do anything with messages and queues, you must create at least one queue manager and its associated objects. To create a queue manager, you use the MQSeries control command **crtmqm**. The **crtmqm** command automatically creates the required default objects and system objects. Default objects form the basis of any object definitions that you make; system objects are required for queue manager operation. When a queue manager and its objects have been created, you use the **strmqm** command to start the queue manager.

See Chapter 2, “Understanding administration command sets” on page 19 for more information about commands that can be used with the MQSeries Version 5 products, and the different methods of invoking them.

Guidelines for creating queue managers

A queue manager manages the resources associated with it, in particular the queues that it owns. It provides queueing services to applications for Message Queuing Interface (MQI) calls and commands to create, modify, display, and delete MQSeries objects. You create a queue manager using the **crtmqm** command. However, before you try this, especially in a production environment, work through this checklist:

- Specify a unique queue manager name.
- Limit the number of queue managers.
- Specify a default queue manager.
- Specify a dead-letter queue.
- Specify a default transmission queue.
- Specify the required logging parameters.
- Back up configuration files after creating a queue manager.

The tasks in this list are explained in the sections that follow.

Specifying a unique queue manager name

When you create a queue manager, you must ensure that no other queue manager has the same name, anywhere in your network. Queue manager names are not checked at creation time, and nonunique names prevent you from creating channels for distributed queuing.

One method of ensuring uniqueness is to prefix each queue manager name with its own (unique) node name. For example, if a node is called `accounts`, you could name your queue manager `accounts.saturn.queue.manager`, where `saturn` identifies a particular queue manager and `queue.manager` is an extension you can give to all queue managers. Alternatively, you can omit this, but note that `accounts.saturn` and `accounts.saturn.queue.manager` are *different* queue manager names.

If you are using MQSeries for communication with other enterprises, you can also include your own enterprise as a prefix. We do not actually do this in the examples, because it makes them more difficult to follow.

Note: Queue manager names in control commands are case-sensitive. This means that you could create two queue managers with the names `jupiter.queue.manager` and `JUPITER.queue.manager`. Such complications are best avoided.

Limiting the number of queue managers

You can create as many queue managers as resources allow. However, because each queue manager requires its own resources, it is generally better to have one queue manager with 100 queues than ten queue managers with ten queues each. In production systems, many nodes will be run with a single queue manager, but larger server machines may run with multiple queue managers.

Specifying the default queue manager

Each node should have a default queue manager, though it is possible to configure MQSeries on a node without one.

To create a default queue manager, specify the `-q` flag on the `crtmqm` command. For a detailed description of this command and its parameters, see “`crtmqm` (Create queue manager)” on page 240.

What is a default queue manager?

The default queue manager is the queue manager that applications connect to if they do not specify a queue manager name in an MQCONN call. It is also the queue manager that processes MQSC commands when you invoke the `runmqsc` command without specifying a queue manager name.

How do you specify a default queue manager?

You include the `-q` flag on the `crtmqm` command to specify that the queue manager you are creating is the default queue manager. Omit this flag if you do not want to create a default queue manager.

Specifying a queue manager as the default *replaces* any existing default queue manager specification for the node.

What happens if I make another queue manager the default?

If you change the default queue manager, this can affect other users or applications. The change has no effect on currently-connected applications, because they can use the handle from their original connect call in any further MQI calls. This handle ensures that the calls are directed to the same queue manager. Any applications connecting after the change connect to the new default queue manager.

This may be what you intend, but you should take this into account before you change the default.

Specifying a dead-letter queue

The dead-letter queue is a local queue where messages are put if they cannot be routed to their correct destination.

Attention: It is vitally important to have a dead-letter queue on each queue manager in your network. Failure to do so may mean that errors in application programs cause channels to be closed or that replies to administration commands are not received.

For example, if an application attempts to put a message on a queue on another queue manager, but the wrong queue name is given, the channel is stopped, and the message remains on the transmission queue. Other applications cannot then use this channel for their messages.

The channels are not affected if the queue managers have dead-letter queues. The undelivered message is simply put on the dead-letter queue at the receiving end, leaving the channel and its transmission queue available.

Therefore, when you create a queue manager you should use the `-u` flag to specify the name of the dead-letter queue. You can also use an MQSC command to alter the attributes of a queue manager and specify the dead-letter queue to be used. See “Altering queue manager attributes” on page 44 for an example of an MQSC ALTER command.

When you find messages on a dead-letter queue, you can use the dead-letter queue handler, supplied with MQSeries, to process these messages. See Chapter 8, “The MQSeries dead-letter queue handler” on page 115 for further information about the dead-letter queue handler.

Specifying a default transmission queue

A transmission queue is a local queue on which messages in transit to a remote queue manager are queued pending transmission. The default transmission queue is the queue that is used when no transmission queue is explicitly defined. Each queue manager can be assigned a default transmission queue.

When you create a queue manager you should use the `-d` flag to specify the name of the default transmission queue. This does not actually create the queue; you have to do this explicitly later on. See “Working with local queues” on page 49 for more information.

Specifying the required logging parameters

You can specify logging parameters on the **crtmqm** command, including the type of logging, and the path and size of the log files. In a development environment, the default logging parameters should be adequate. However, you can change the defaults if, for example:

- You have a low-end system configuration that cannot support large logs.
- You anticipate a large number of long messages being on your queues at the same time.

For more information about specifying logging parameters:

- Using the **crtmqm** command, see “crtmqm (Create queue manager)” on page 240.
- Using configuration files, see “Log configuration stanzas” on page 109.

Backing up configuration files after creating a queue manager

There are two configuration files to consider:

1. When you install the product, the MQSeries configuration file (mqs.ini) is created. It contains a list of queue managers, which is updated each time you create or delete a queue manager. There is one mqs.ini file per node.
2. When you create a new queue manager, a new queue manager configuration file (qm.ini) is automatically created. This contains configuration parameters for the queue manager.

You should make a backup of these files. If, later on, you create another queue manager that causes you problems, you can reinstate the backups when you have removed the source of the problem. As a general rule, you should back up your configuration files each time you create a new queue manager.

For more information about configuration files, see Chapter 7, “Configuration files” on page 99.

Understanding MQSeries file names

Each MQSeries queue, queue manager, and process object is represented by a file. Because object names are not necessarily valid file names, the queue manager converts the object name into a valid file name, where necessary.

The path to a queue manager directory is formed from the following:

- A prefix, which is defined in the queue manager configuration file.

In MQSeries for UNIX systems, the default prefix is:

`/var/mqm`

In MQSeries for OS/2 Warp and Windows NT, the default prefix is:

`c:\mqm`

- A literal:

`qmgrs`

- A coded queue manager name, which is the queue manager name transformed into a valid directory name. For example, the queue manager:

`queue.manager`

would be represented as:

`queue!manager`

This process is referred to as *name transformation*.

Queue manager name transformation

In MQSeries you can give a queue manager a name containing up to 48 characters. For example, you could name a queue manager:

`QUEUE.MANAGER.ACCOUNTING.SERVICES`

However, each queue manager is represented by a file and there are limitations to the maximum length a file name can be, and to the characters that can be used in the name. As a result, the names of files representing objects are automatically transformed to meet the requirements of the file system.

The rules governing the transformation of a queue manager name, using the example of a queue manager with the name `queue.manager`, are as follows:

1. Transform individual characters:

. becomes !

/ becomes &

2. If the name is still not valid:

- a. Truncate it to eight characters

- b. Append a three-character numeric suffix

For example, assuming the default prefix, the queue manager name in MQSeries for UNIX systems becomes:

`/var/mqm/qmgrs/queue!manager`

In MQSeries for OS/2 Warp and Windows NT with HPFS (or NTFS), the queue manager name becomes:

`c:\mqm\qmgrs\queue!manager`

In MQSeries for OS/2 Warp and Windows NT with FAT, the queue manager name becomes:

`c:\mqm\qmgrs\queue!ma`

The transformation algorithm also allows distinction between names that differ only in case, on file systems that are not case sensitive.

Object name transformation

Object names are not necessarily valid file system names. Therefore the object names may need to be transformed. The method used is different from that for queue manager names because, although there only a few queue manager names per machine, there can be a large number of other objects for each queue manager. Only process definitions and queues are represented in the file system; channels are not affected by these considerations.

When a new name is generated by the transformation process there is no simple relationship with the original object name. You can use the **dspmqfls** command to convert between real and transformed object names.

Working with queue managers

MQSeries provides control commands for creating, starting, ending, and deleting queue managers. You can also display a queue manager's attributes using the MQSC command DISPLAY QMGR and change them using ALTER QMGR. See "Displaying queue manager attributes" on page 42 and "Altering queue manager attributes" on page 44.

Creating a default queue manager

The following command: creates a default queue manager called saturn.queue.manager; creates the default and system objects; and specifies the names of both a default transmission queue and a dead-letter queue:

```
crtmqm -q -d MY.DEFAULT.XMIT.QUEUE -u SYSTEM.DEAD.LETTER.QUEUE saturn.queue.manager
```

where:

- | | |
|-----------------------------|--|
| -q | Indicates that this queue manager is the default queue manager. |
| -d MY.DEFAULT.XMIT.QUEUE | Is the name of the default transmission queue. |
| -u SYSTEM.DEAD.LETTER.QUEUE | Is the name of the dead-letter queue. |
| saturn.queue.manager | Is the name of this queue manager. This must be the last parameter specified on the crtmqm command. |

The system and default objects are listed in Appendix A, "System and default objects" on page 301.

Note: In MQSeries for UNIX systems only, you can create the queue manager directory `/var/mqm/qmgrs/<qmgr>`, even on a separate local file system, before you use the **crtmqm** command. When you use **crtmqm**, if the `/var/mqm/qmgrs/<qmgr>` directory exists, is empty, and is owned by mqm, it is used for the queue manager data. If the directory is not owned by mqm, the creation fails with an FFST. If the directory is not empty, then a new directory is created.

Starting a queue manager

Although you have created a queue manager, it cannot process commands or MQI calls until it has been started. Start the queue manager by typing in this command:

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

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

Starting a queue manager automatically

In MQSeries for Windows NT only, a queue manager can be invoked automatically when the system starts. To request automatic startup of a queue manager, enter:

```
scmmqm -a -s c:\mqm\startup.cmd saturn.queue.manager
```

where:

-a	Specifies that the queue manager is to be added to the list of those that will be automatically started.
-s	References a command file, of any name and location, that describes what actions are taken when this queue manager starts.
saturn.queue.manager	Is the name of this queue manager. This parameter must be the final parameter specified on the scmmqm command.

The command file, referenced -s, is a text file containing a series of commands. The commands are executed in order, and must start in column one. Valid commands are:

- **runmqchi**
- **runmqchl**
- **runmqlsr**
- **strmqcsv**
- **strmqm**
- **tpstart** (SNA command)

The commands must be followed by their correct parameters and each queue manager included in the automatic startup list needs its own command file.

For example, the following command file:

```
strmqm saturn.queue.manager
runmqlsr -t TCP -m saturn.queue.manager -n server
```

starts a queue manager and starts a listener for that queue manager.

Working with queue managers

All queue managers in the automatic start-up list are started and stopped by one service named IBM MQSeries. This service can be modified from the Control Panel or from the DOS prompt and it accepts two commands only: start and stop.

You can also prevent a queue manager from starting automatically. See “Removing a queue manager from the automatic start-up list” for more information.

Removing a queue manager from the automatic start-up list

In order to prevent Windows NT from starting a queue manager automatically, use the following command:

```
scmmqm -d saturn.queue.manager
```

See “scmmqm (Add the queue manager to the Windows NT Service Control Manager)” on page 282 for more information about the **scmmqm** command.

Stopping a queue manager

To stop a queue manager, use the **endmqm** command. For example, to stop a queue manager called `saturn.queue.manager` use this command:

```
endmqm saturn.queue.manager
```

Quiesced shutdown

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

Use this type of shutdown to notify applications to stop; you are not told when they have stopped.

Immediate shutdown

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

Use this as the normal way to stop the queue manager, optionally after a quiesce period. For an immediate shutdown, the command is:

```
endmqm -i saturn.queue.manager
```

Preemptive shutdown

Preemptive shutdown

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

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

```
endmqm -p saturn.queue.manager
```

This stops all queue manager code immediately.

If this method still does not work, see “Stopping a queue manager manually” on page 315 for an alternative.

For a detailed description of the **endmqm** command and its options, see “endmqm (End queue manager)” on page 261.

If you have problems

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

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

If a problem does occur while stopping the queue manager, break out of the **endmqm** command using Ctrl-C.

You can then issue another **endmqm** command, but this time with a flag that specifies the type of shutdown that you require.

Restarting a queue manager

To restart a queue manager, use the command:

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

Making an existing queue manager the default

When you create a default queue manager, the name of the default queue manager is inserted in the *DefaultQueueManager* stanza in the MQSeries configuration file (mqs.ini). The stanza and its contents are automatically created if they do not exist.

You may need to edit this stanza:

- To make an existing queue manager the default. To do this you have to change the queue manager name in this stanza to the name of the new default queue manager. You must do this manually, using a text editor.
- If you do not have a default queue manager on the node, and you want to make an existing queue manager the default. To do this you must create the *DefaultQueueManager* stanza—with the required name—yourself.
- If you accidentally make another queue manager the default and wish to revert to the original default queue manager. To do this, edit the *DefaultQueueManager* stanza in the MQSeries configuration file, replacing the name of the unwanted default queue manager with that of the one you do want.

See Chapter 7, “Configuration files” on page 99 for information about configuration files. When the stanza contains the required information, stop the queue manager and restart it.

Deleting a queue manager

To delete a queue manager, first stop it, then use the following command:

```
dltmqm saturn.queue.manager
```

Notes:

1. Deleting a queue manager is a drastic step, because you also delete all resources associated with that queue manager, including all queues and their messages, and all object definitions.
2. In MQSeries for Windows NT, the **dltmqm** command also removes a queue manager from the automatic start-up list (described in “Starting a queue manager automatically” on page 33).

For a description of the **dltmqm** command and its options, see “dltmqm (Delete queue manager)” on page 244. You should ensure that only trusted administrators have the authority to use this command.

If the usual methods for deleting a queue manager do not work, see “Removing queue managers manually” on page 316 for an alternative.

Managing the command server for remote administration

Each queue manager can have a command server associated with it. A command server processes any incoming commands from remote queue managers, or PCF commands from applications. It presents the commands to the queue manager for processing and returns a completion code or operator message depending on the origin of the command. There are separate control commands for starting and stopping the command server.

Note: For remote administration, you must ensure that the target queue manager is running. Otherwise, the messages containing commands cannot leave the queue manager from which they are issued. Instead, these messages are queued in the local transmission queue that serves the remote queue manager. This situation should be avoided, if at all possible.

Starting the command server

To start the command server use this command:

```
strmqcsv saturn.queue.manager
```

where `saturn.queue.manager` is the queue manager for which the command server is being started.

Displaying the status of the command server

For remote administration, you must ensure that the command server on the target queue manager is running. If it is not running, no remote commands can be processed. Any messages containing commands are queued in the target queue manager's command queue.

To display the status of the command server for a queue manager, called here `saturn.queue.manager`, the command is:

```
dspmqcsv saturn.queue.manager
```

You must issue this command on the target machine. If the command server is running, the following message is returned:

```
AMQ8027    MQSeries Command Server Status ...: Running
```

Stopping a command server

To end a command server, the command, using the previous example is:

```
endmqcsv saturn.queue.manager
```

You can stop the command server in two different ways:

- For a controlled stop, use the **endmqcsv** command with the **-c** flag, which is the default.
- For an immediate stop, use the **endmqcsv** command with the **-i** flag.

Note: Stopping a queue manager also ends the command server associated with it (if one has been started).

Chapter 4. Administering local MQSeries objects

This chapter describes how to administer local MQSeries objects to support application programs that use the Message Queuing Interface (MQI). In this context, local administration means creating, displaying, changing, copying, and deleting MQSeries objects.

This chapter contains these sections:

- “Supporting application programs that use the MQI”
- “Issuing MQSC commands for administration” on page 40
- “Running MQSC commands from text files” on page 44
- “Resolving problems with MQSC” on page 47
- “Working with local queues” on page 49
- “Working with alias queues” on page 55
- “Working with model queues” on page 57
- “Managing objects for triggering” on page 59

Supporting application programs that use the MQI

MQSeries application programs need certain objects before they can run successfully. For example, Figure 1 shows an application that removes messages from a queue, processes them, and then sends some results to another queue on the same queue manager.




Figure 1. Queues, messages, and applications

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

Issuing MQSC commands

Before this application can be run, these conditions must be satisfied:

- The queue manager must exist and be running.
- The first application queue, from which the messages are to be removed, must be defined.
- The second queue, on which the application puts the messages, must also be defined.
- The application must be able to connect to the queue manager. To do this it must be linked to the product code. See the *MQSeries Application Programming Guide* for more information.
- The applications that put the messages on the first queue must also connect to a queue manager. If they are remote, they must also be set up with transmission queues and channels. This part of the system is not shown in Figure 1 on page 39.

Issuing MQSC commands for administration

In this section, we assume that you will be issuing commands using the **runmqsc** command. You can do this interactively—entering the commands at the keyboard—or you can redirect the standard input device (stdin) to run a sequence of commands from an ASCII text file. In both cases, the format of the commands is the same.

The *MQSeries Command Reference* manual contains a description of each MQSC command and its syntax.

Before you start

Before you can run MQSC commands, you must have created and started the queue manager that is going to run the commands. For more information, see “Creating a default queue manager” on page 32.

MQSeries object names

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

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

```
ORANGE.LOCAL.QUEUE
```

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

We also use the name saturn.queue.manager as a queue manager name.

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

You do not have to use these names, but if you do not, you must modify any commands in examples that specify them.

Case-sensitivity in MQSC commands

MQSC commands, including their attributes, can be written in uppercase or lowercase. Object names in MQSC commands are folded (that is, QUEUE and queue are not differentiated), unless the names are within single quotation marks. If quotation marks are not used, the object is processed with a name in uppercase. See the *MQSeries Command Reference* manual for more information.

However, the **runmqsc** command invocation, in common with all MQSeries control commands, is case sensitive in some MQSeries environments. See “Using control commands” on page 20 for more information.

Standard input and output

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

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

Using the MQSC facility interactively

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

```
runmqsc
```

In this command, a queue manager name has not been specified, therefore the MQSC commands will be processed by the default queue manager. Now you can type in any MQSC commands, as required. For example, try this one:

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

Continuation characters must be used to indicate that a command is continued on the following line:

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

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

Feedback from MQSC commands

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

```
AMQ8006: MQSeries queue created.  
.  
.  
.  
AMQ8405: Syntax error detected at or near end of command segment below:-  
Z  
AMQ8426: Valid MQSC commands are:  
  
ALTER  
CLEAR  
DEFINE  
DELETE  
DISPLAY  
END  
PING  
RESET  
RESOLVE  
START  
STOP  
4 : end
```

The first message confirms that a queue has been created; the second indicates that you have made a syntax error.

These messages are sent to the standard output device. If you have not entered the command correctly, refer to the *MQSeries Command Reference* manual for the correct syntax.

Ending interactive input to MQSC

To end interactive input of MQSC commands, enter the MQSC END command:

```
END
```

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

If you are redirecting input from other sources, such as a text file, you do not have to do this.

Displaying queue manager attributes

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

```
DISPLAY QMGR
```

Typical output from this command is:

```

1 : display qmgr all
AMQ8408: Display Queue Manager details.
DESCR( )                                DEADQ( )
DEFXMITQ( )                              CHADEXIT( )
COMMANDQ(SYSTEM.ADMIN.COMMAND.QUEUE)    QMNAME(saturn.queue.manager)
TRIGINT(999999999)                      MAXHANDS(256)
MAXUMSGS(10000)                          AUTHOREV(DISABLED)
INHIBTEV(DISABLED)                      LOCALEV(DISABLED)
REMOTEEV(DISABLED)                      PERFMEEV(DISABLED)
STRSTPEV(ENABLED)                       CHAD(DISABLED)
CHADEV(DISABLED)                        MAXMSGL(4194304)
MAXPRTY(9)                               CCSID(850)
CMDLEVEL(230)                            PLATFORM(OS2)
SYNCPT                                    DISTL(YES)

2 : display qmgr
AMQ8408: Display Queue Manager details.
DESCR( )                                DEADQ( )
DEFXMITQ( )                              CHADEXIT( )
COMMANDQ(SYSTEM.ADMIN.COMMAND.QUEUE)    QMNAME(saturn.queue.manager)
TRIGINT(999999999)                      MAXHANDS(256)
MAXUMSGS(10000)                          AUTHOREV(DISABLED)
INHIBTEV(DISABLED)                      LOCALEV(DISABLED)
REMOTEEV(DISABLED)                      PERFMEEV(DISABLED)
STRSTPEV(ENABLED)                       CHAD(DISABLED)
CHADEV(DISABLED)                        MAXMSGL(4194304)
MAXPRTY(9)                               CCSID(850)
CMDLEVEL(230)                            PLATFORM(OS2)
SYNCPT                                    DISTL(YES)

3 : end

```

The ALL parameter on the DISPLAY QMGR command causes all the queue manager attributes to be displayed. In particular, the output tells us the default queue manager name (saturn.queue.manager), and the names of the dead-letter queue (SYSTEM.DEAD.LETTER.QUEUE) and the command queue (SYSTEM.ADMIN.COMMAND.QUEUE).

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

```

DISPLAY QUEUE (SYSTEM.*)

```

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

Using a queue manager that is not the default

To run MQSC commands on a local queue manager other than the default queue manager, you specify the name of the queue manager on input to the **runmqsc** command. For example, to run MQSC commands on queue manager `jupiter.queue.manager`, use the command:

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

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

You can also run MQSC commands on a remote queue manager; see “Issuing MQSC commands remotely” on page 69.

Altering queue manager attributes

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

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

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

Running MQSC commands from text files

Running MQSC commands interactively is suitable for quick tests, but if you have very long commands, or are using a particular sequence of commands repeatedly, consider redirecting `stdin` from a text file. (See “Standard input and output” on page 41 for information about `stdin` and `stdout`.) To do this, first create a text file containing the MQSC commands using your usual text editor. When you use the **runmqsc** command, use the redirection operators. For example, the following command runs a sequence of commands contained in the text file `myprog.in`:

```
runmqsc < myprog.in
```

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

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

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

This command invokes the MQSC commands contained in the MQSC command file `myprog.in`. Because we have not specified a queue manager name, the MQSC commands are run against the default queue manager. The output is sent to the report file `myprog.out`. Figure 2 shows an extract from the MQSC command file `myprog.in` and Figure 3 on page 46 shows the corresponding extract of the output in `myprog.out`.

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

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

MQSC command files

MQSC commands are written in human-readable form, that is, in ASCII text. Figure 2 is an extract from an MQSC command file showing an MQSC command (`DEFINE QLOCAL`) with its attributes. The *MQSeries Command Reference* manual contains a description of each MQSC command and its syntax.

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

Figure 2. Extract from the MQSC command file, `myprog.in`

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

MQSC reports

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

- A header identifying MQSC as the source of the report:
Starting MQSeries Commands.
- An optional numbered listing of the MQSC commands issued. By default, the text of the input is echoed to the output. Within this output, each command is prefixed by a sequence number, as shown in Figure 3. However, you can use the `-e` flag on the `runmqsc` command to suppress the output.
- A syntax error message for any commands found to be in error.
- An *operator message* indicating the outcome of running each command. For example, the operator message for the successful completion of a `DEFINE QLOCAL` command is:
AMQ8006: MQSeries queue created.
- Other messages resulting from general errors when running the script file.
- A brief statistical summary of the report indicating the number of commands read, the number of commands with syntax errors, and the number of commands that could not be processed.

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

```
Starting MQSeries Commands.
.
.
12:    DEFINE QLOCAL('RED.LOCAL.QUEUE') REPLACE +
:      DESCR(' ') +
:      PUT(ENABLED) +
:      DEFPRTY(0) +
:      DEFPSIST(NO) +
:      GET(ENABLED) +
:      MAXDEPTH(5000) +
:      MAXMSGL(1024) +
:      DEFSOPT(SHARED) +
:      USAGE(NORMAL) +
:      NOTRIGGER;
AMQ8006: MQSeries queue created.
:
.
.
```

Figure 3. Extract from the MQSC report file, `myprog.out`

Running the supplied MQSC command files

These MQSC command files are supplied with MQSeries:

amqscos0.tst	Definitions of objects used by sample programs.
amqscic0.tst	Definitions of queues for CICS transactions.
amqslnk0.tst	Definitions of queues used by the Lotus Notes sample application.

In MQSeries for UNIX systems, these files are located in the directory `mqmtop/samp`; see “The installation directory” on page x for details of the installation directory `mqmtop`.

In MQSeries for OS/2 Warp and Windows NT, these files are located in the directory `c:\mqm\tools\mqsc\samples`.

Using runmqsc to verify commands

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

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

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

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

This report is similar to that shown in Figure 3 on page 46.

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

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

the `-w` flag, which you use to indicate that the queue manager is remote, is ignored, and the command is run locally in verification mode.

Resolving problems with MQSC

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

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

- Use the indirection operator `<` when redirecting input from a file. If you omit the indirection operator, the queue manager interprets the file name as a queue manager name, and issues the following error message:

```
AMQ8118: MQSeries queue manager does not exist.
```

- If you redirect output to a file, use the `>` indirection operator. By default, the file is put in the current working directory at the time **runmqsc** is invoked. Specify a fully-qualified file name to send your output to a specific file and directory.
- Check that you have created the queue manager that is going to run the commands.

To do this, look in the configuration file `mq5.ini`. This file contains the names of the queue managers and the name of the default queue manager, if you have one.

- The queue manager should already be started, if it is not, start it; see “Starting a queue manager” on page 33. You get an error message if it is already started.
- Specify a queue manager name on the **runmqsc** command if you have not defined a default queue manager, otherwise you get this error:

```
AMQ8146: MQSeries queue manager not available.
```

To correct this type of problem, see “Making an existing queue manager the default” on page 36.

- You cannot specify an MQSC command as parameter of the **runmqsc** command. For example, this is invalid:

```
runmqsc DEFINE QLOCAL(FRED)
```

- You cannot enter MQSC commands before you issue the **runmqsc** command.
- You cannot run control commands from **runmqsc**. For example, you cannot issue the **strmqm** command to start a queue manager while you are running MQSC interactively.

```

runmqsc
.
.
Starting MQSeries Commands.

strmqm saturn.queue.manager
  1 : strmqm saturn.queue.manager
AMQ8405: Syntax error detected at or near end of command segment below:-
s
AMQ8426: Valid MQSC commands are:

    ALTER
    CLEAR
    DEFINE
    DELETE
    DISPLAY
    END
    PING
    RESET
    RESOLVE
    START
    STOP
  4 : end

```

See also “If you have problems using MQSC remotely” on page 70.

Working with local queues

This section contains examples of some of the MQSC commands that you can use to manage local, model, and alias queues. Refer to the *MQSeries Command Reference* for a complete description of these commands.

Defining a local queue

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

Use the MQSC command `DEFINE QLOCAL` to create a definition of a local queue and also to create the data structure that is called a queue. You can also modify the queue characteristics from those of the default local queue.

In this example, the queue we define, `ORANGE.LOCAL.QUEUE`, is specified to have these characteristics:

- It is enabled for gets, disabled for puts, and operates on a first-in-first-out (FIFO) basis.
- It is an ‘ordinary’ queue, that is, it is not an initiation queue or a transmission queue, and it does not generate trigger messages.
- The maximum queue depth is 1000 messages; the maximum message length is 2000 bytes.

The following MQSC command does this:

```
DEFINE QLOCAL (ORANGE.LOCAL.QUEUE) +
  DESCR('Queue for messages from other systems') +
  PUT (DISABLED) +
  GET (ENABLED) +
  NOTRIGGER +
  MSGDLVSQ (FIFO) +
  MAXDEPTH (1000) +
  MAXMSGL (2000) +
  USAGE (NORMAL);
```

Notes:

1. Most of these attributes are the defaults as supplied with the product. However, they are shown here for purposes of illustration. You can omit them if you are sure that the defaults are what you want or have not been changed. See also “Displaying default object attributes” on page 51.
2. USAGE (NORMAL) indicates that this queue is not a transmission queue.
3. If you already have a local queue on the same queue manager with the name ORANGE.LOCAL.QUEUE, this command fails. Use the REPLACE attribute, if you want to overwrite the existing definition of a queue, but see also “Changing local queue attributes” on page 52.

Defining a dead-letter queue

Each queue manager should have a local queue to be used as a dead-letter queue so that messages that cannot be delivered to their correct destination can be stored for later retrieval. You must explicitly tell the queue manager about the dead-letter queue. You can do this by specifying a dead-letter queue on the **crtmqm** command, or you can use the ALTER QMGR command to specify one later. You must also define the dead-letter queue before it can be used.

A sample dead-letter queue called SYSTEM.DEAD.LETTER.QUEUE is supplied with the product. This queue is automatically created when you run the sample. You can modify this definition if required. There is no need to rename it, although you can if you like.

A dead-letter queue has no special requirements except that it must be a local queue and its MAXMSGL (maximum message length) attribute must enable the queue to accommodate the largest messages that the queue manager has to handle.

MQSeries provides a dead-letter queue handler that allows you to specify how messages found on a dead-letter queue are to be processed or removed. For further information, see Chapter 8, “The MQSeries dead-letter queue handler” on page 115.

Displaying default object attributes

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

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

Note: The syntax of this command is different from that of the corresponding `DEFINE` command.

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

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

This command displays the three specified attributes as follows:

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

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

Copying a local queue definition

You can copy a queue definition using the `LIKE` attribute on the `DEFINE` command. For example:

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

This command creates a queue with the same attributes as our original queue `ORANGE.LOCAL.QUEUE`, rather than those of the system default local queue.

Working with local queues

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

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

This command copies the attributes of the queue ORANGE.LOCAL.QUEUE to the queue THIRD.QUEUE, but specifies that the maximum message length on the new queue is to be 1024 bytes, rather than 2000.

Notes:

1. When you use the LIKE attribute on a DEFINE command, you are copying the queue attributes only. You are not copying the messages on the queue.
2. If you define a local queue, without specifying LIKE, it is the same as DEFINE LIKE(SYSTEM.DEFAULT.LOCAL.QUEUE).

Changing local queue attributes

You can change queue attributes in two ways, using either the ALTER QLOCAL command or the DEFINE QLOCAL command with the REPLACE attribute. In “Defining a local queue” on page 49, we defined the queue ORANGE.LOCAL.QUEUE. Suppose, for example, you wanted to increase the maximum message length on this queue to 10 000 bytes.

- Using the ALTER command:

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

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

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

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

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

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

Clearing a local queue

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

```
CLEAR QLOCAL (MAGENTA.QUEUE)
```

You cannot clear a queue if:

- There are uncommitted messages that have been put on the queue under syncpoint.
- An application currently has the queue open.

Deleting a local queue

Use the MQSC command DELETE QLOCAL to delete a local queue. A queue cannot be deleted if it has uncommitted messages on it. However, if the queue has one or more committed messages, and no uncommitted messages, it can only be deleted if you specify the PURGE option. For example:

```
DELETE QLOCAL (PINK.QUEUE) PURGE
```

Specifying NOPURGE instead of PURGE ensures that the queue is not deleted if it contains any committed messages.

Browsing queues

MQSeries provides a sample queue browser that you can use to look at the contents of the messages on a queue. The browser is supplied in both source and executable formats.

In MQSeries for UNIX systems, the default file names and paths are:

```
Source      mqmtop/samp/amqsbcg0.c
Executable  mqmtop/samp/bin/amqsbcg
```

In MQSeries for OS/2 Warp and Windows NT, the default file names and paths are:

```
Source      c:\mqm\tools\c\samples\amqsbcg0.c
Executable  c:\mqm\tools\c\samples\bin\amqsbcg.exe
```

The sample requires two input parameters, the queue manager name and the queue name. For example:

```
amqsbcg SYSTEM.ADMIN.QMGREVENT.tpp01 saturn.queue.manager
```

There are no defaults; both parameters are required.

Typical results from this command are:

```
AMQSBCG0 - starts here
*****

MQOPEN - 'SYSTEM.ADMIN.QMGR.EVENT'

MQGET of message number 1
****Message descriptor****

StrucId : 'MD ' Version : 2
Report  : 0 MsgType : 8
Expiry  : -1 Feedback : 0
Encoding : 546 CodedCharSetId : 850
Format  : 'MQEVENT '
Priority : 0 Persistence : 0
MsgId   : X'414D512073617475726E2E71756575650005D30033563DB8'
CorrelId : X'0000000000000000000000000000000000000000000000000000'
BackoutCount : 0
ReplyToQ      : ' '
ReplyToQMgr   : 'saturn.queue.manager '
** Identity Context
UserIdentifier : ' '
AccountingToken :
X'0000000000000000000000000000000000000000000000000000000000000000'
ApplIdentityData : ' '
** Origin Context
PutApplType      : '7'
PutApplName      : 'saturn.queue.manager '
PutDate          : '19970417' PutTime : '15115208'
ApplOriginData   : ' '

GroupId : X'0000000000000000000000000000000000000000000000000000'
MsgSeqNumber : '1'
Offset       : '0'
MsgFlags     : '0'
OriginalLength : '104'
```

Figure 4 (Part 1 of 2). Typical results from queue browser


```

****   Message   ****

length - 104 bytes

00000000:  0700 0000 2400 0000 0100 0000 2C00 0000 '....¢.....'
00000010:  0100 0000 0100 0000 0100 0000 AE08 0000 '.....'
00000020:  0100 0000 0400 0000 4400 0000 DF07 0000 '.....D.....'
00000030:  0000 0000 3000 0000 7361 7475 726E 2E71 '....θ...saturn.q'
00000040:  7565 7565 2E6D 616E 6167 6572 2020 2020 'ueue.manager'
00000050:  2020 2020 2020 2020 2020 2020 2020 2020 '
00000060:  2020 2020 2020 2020 '

No more messages
MQCLOSE
MQDISC

```

Figure 4 (Part 2 of 2). Typical results from queue browser

Working with alias queues

An alias queue (also known as a queue alias) provides a method of redirecting MQI calls. An alias queue is not a real queue but a definition that resolves to a real queue. The alias queue definition contains a target queue name which is specified by the TARGQ attribute (*BaseQName* in PCF). When an application specifies an alias queue in an MQI call, the queue manager resolves the real queue name at run time.

For example, an application has been developed to put messages on a queue called MY.ALIAS.QUEUE. It specifies the name of this queue when it makes an MQOPEN request and, indirectly, if it puts a message on this queue. The application is not aware that the queue is an alias queue. For each MQI call using this alias, the queue manager resolves the real queue name, which could be either a local queue or a remote queue defined at this queue manager.

By changing the value of the TARGQ attribute, you can redirect MQI calls to another queue, possibly on another queue manager. This is useful for maintenance, migration, and load-balancing.

Defining an alias queue

The following command creates an alias queue:

```
DEFINE QALIAS (MY.ALIAS.QUEUE) TARGQ (YELLOW.QUEUE)
```

This command redirects MQI calls that specify MY.ALIAS.QUEUE to the queue YELLOW.QUEUE. The command does not create the target queue; the MQI calls fail if the queue YELLOW.QUEUE does not exist at run time.

Working with alias queues

If you change the alias definition, you can redirect the MQI calls to another queue. For example:

```
ALTER QALIAS (MY.ALIAS.QUEUE) TARGQ (MAGENTA.QUEUE)
```

This command redirects MQI calls to another queue, MAGENTA.QUEUE.

You can also use alias queues to make a single queue (the target queue) appear to have different attributes for different applications. You do this by defining two aliases, one for each application. Suppose there are two applications:

- Application ALPHA can put messages on YELLOW.QUEUE, but is not allowed to get messages from it.
- Application BETA can get messages from YELLOW.QUEUE, but is not allowed to put messages on it.

You can do this using the following commands:

```
* This alias is put enabled and get disabled for application ALPHA  
  
DEFINE QALIAS (ALPHAS.ALIAS.QUEUE) +  
  TARGQ (YELLOW.QUEUE) +  
  PUT (ENABLED) +  
  GET (DISABLED)  
  
* This alias is put disabled and get enabled for application BETA  
  
DEFINE QALIAS (BETAS.ALIAS.QUEUE) +  
  TARGQ (YELLOW.QUEUE) +  
  PUT (DISABLED) +  
  GET (ENABLED)
```

ALPHA uses the queue name ALPHAS.ALIAS.QUEUE in its MQI calls; BETA uses the queue name BETAS.ALIAS.QUEUE. They both access the same queue, but in different ways.

You can use the LIKE and REPLACE attributes when you define queue aliases, in the same way that you use these attributes with local queues.

Using other commands with alias queues

You can use the appropriate MQSC commands to display or alter queue alias attributes, or delete the queue alias object. For example:

```
* Display the queue alias's attributes

DISPLAY QUEUE (ALPHAS.ALIAS.QUEUE)

* ALTER the base queue name, to which the alias resolves.
* FORCE = Force the change even if the queue is open.

ALTER QALIAS (ALPHAS.ALIAS.QUEUE) TARGQ(ORANGE.LOCAL.QUEUE) FORCE

* Delete this queue alias, if you can.

DELETE QALIAS (ALPHAS.ALIAS.QUEUE)
```

You cannot delete a queue alias if, for example, an application currently has the queue open or has a queue open that resolves to this queue. See the *MQSeries Command Reference* manual for more information about this and other queue alias commands.

Working with model queues

A queue manager creates a *dynamic queue* if it receives an MQI call from an application specifying a queue name that has been defined as a model queue. The name of the new dynamic queue is generated by the queue manager when the queue is created. A *model queue* is a template that specifies the attributes of any dynamic queues created from it.

Model queues provide a convenient method for applications to create queues as they are required.

Defining a model queue

You define a model queue with a set of attributes in the same way that you define a local queue. Model queues and local queues have the same set of attributes except that on model queues you can specify whether the dynamic queues created are temporary or permanent. (Permanent queues are maintained across queue manager restarts, temporary ones are not). For example:

```
DEFINE QMODEL (GREEN.MODEL.QUEUE) +
  DESCR('Queue for messages from application X') +
  PUT (DISABLED) +
  GET (ENABLED) +
  NOTRIGGER +
  MSGDLVSQ (FIFO) +
  MAXDEPTH (1000) +
  MAXMSGL (2000) +
  USAGE (NORMAL) +
  DEFTYPE (PERMDYN)
```

This command creates a model queue definition. From the DEFTYPE attribute, the actual queues created from this template are permanent dynamic queues.

Note: The attributes not specified are automatically copied from the `SYSYSTEM.DEFAULT.MODEL.QUEUE` default queue.

You can use the `LIKE` and `REPLACE` attributes when you define model queues, in the same way that you use them with local queues.

Using other commands with model queues

You can use the appropriate MQSC commands to display or alter a model queue's attributes, or delete the model queue object. For example:

```
* Display the model queue's attributes

DISPLAY QUEUE (GREEN.MODEL.QUEUE)

* ALTER the model to enable puts on any
* dynamic queue created from this model.

ALTER QMODEL (BLUE.MODEL.QUEUE) PUT(ENABLED)

* Delete this model queue:

DELETE QMODEL (RED.MODEL.QUEUE)
```

Managing objects for triggering

MQSeries provides a facility for starting an application automatically when certain conditions on a queue are met. One example of the conditions is when the number of messages on a queue reaches a specified number. This facility is called *triggering* and is described in detail in the *MQSeries Application Programming Guide*. This section describes how to set up the required objects to support triggering on MQSeries.

Defining an application queue for triggering

An application queue is a local queue that is used by applications for messaging, through the MQI. Triggering requires a number of queue attributes to be defined on the application queue. Triggering itself is enabled by the *Trigger* attribute (TRIGGER in MQSC).

In this example, a trigger event is to be generated when there are 100 messages of priority 5 or greater on the local queue MOTOR.INSURANCE.QUEUE, as follows:

```
DEFINE QLOCAL (MOTOR.INSURANCE.QUEUE) +
        PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS) +
        MAXMSGL (2000) +
        DEFPSIST (YES) +
        INITQ (MOTOR.INS.INIT.QUEUE) +
        TRIGGER +
        TRIGTYPE (DEPTH) +
        TRIGDPTH (100)+
        TRIGMPRI (5)
```

where:

QLOCAL (MOTOR.INSURANCE.QUEUE)

Specifies the name of the application queue being defined.

PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS)

Specifies the name of the application to be started by a trigger monitor program.

MAXMSGL (2000)

Specifies the maximum length of messages on the queue.

DEFPSIST (YES)

Specifies that messages on this queue are persistent by default.

INITQ (MOTOR.INS.INIT.QUEUE)

Is the name of the initiation queue on which the queue manager is to put the trigger message.

TRIGGER

Is the trigger attribute value.

TRIGTYPE (DEPTH)

Specifies that a trigger event is generated when the number of messages of the required priority (TRIMPRI) reaches the number specified in TRIGDPTH.

TRIGDPTH (100)

Specifies the number of messages required to generate a trigger event.

Managing objects for triggering

TRIGMPRI (5)

Is the priority of messages that are to be counted by the queue manager in deciding whether to generate a trigger event. Only messages with priority 5 or higher are counted.

Defining an initiation queue

When a trigger event occurs, the queue manager puts a trigger message on the initiation queue specified in the application queue definition. Initiation queues have no special settings, but you can use the following definition of the local queue MOTOR.INS.INIT.QUEUE for guidance:

```
DEFINE QLOCAL(MOTOR.INS.INIT.QUEUE) +
  GET (ENABLED) +
  NOSHARE +
  NOTRIGGER +
  MAXMSGL (2000) +
  MAXDEPTH (1000)
```

Creating a process definition

Use the DEFINE PROCESS command to create a process definition. A process definition associates an application queue with the application that is to process messages from the queue. This is done through the PROCESS attribute on the application queue MOTOR.INSURANCE.QUEUE. The following MQSC command defines the required process, MOTOR.INSURANCE.QUOTE.PROCESS, identified in this example:

```
DEFINE PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS) +
  DESCR ('Insurance request message processing') +
  APPLTYPE (UNIX) +
  APPLICID ('/u/admin/test/IRMP01') +
  USERDATA ('open, close, 235')
```

where:

MOTOR.INSURANCE.QUOTE.PROCESS

Is the name of the process definition.

DESCR ('Insurance request message processing')

Is a description of the application program to which this definition relates.

This text is displayed when you use the DISPLAY PROCESS command.

This can help you to identify what the process does. If you use spaces in the string, you must enclose the string in single quotation marks.

APPLTYPE (UNIX)

Is the type of application to be started.

APPLICID ('/u/admin/test/IRMP01')

Is the name of the application executable file, specified as a fully qualified file name. In MQSeries for OS/2 Warp and Windows NT, a typical APPLICID value would be c:\app1\test\irmp01.exe.

USERDATA ('open, close, 235')

Is user-defined data, which can be used by the application.

Displaying your process definition

Use the DISPLAY PROCESS command to examine the results of your definition. For example:

```

DISPLAY PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS)

      24 : DISPLAY PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS) ALL
AMQ8407: Display Process details.
DESCR ('Insurance request message processing')  APPLICID ('/u/admin/test/IRMP01')
USERDATA (open, close, 235)                    PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS)
APPLTYPE (UNIX)
    
```

You can also use the MQSC command ALTER PROCESS to alter an existing process definition, and the DELETE PROCESS command to delete a process definition.

Fastpath binding

The normal (default) method of binding causes an application and the local-queue-manager agent to run in separate units of execution. Using a fastpath binding causes the application and the local-queue-manager agent to be part of the same unit of execution. For information about situations where the use of fastpath bindings may be useful, and situations where it may not be recommended, see the *MQSeries Application Programming Guide*.

You can specify the environment variable MQ_CONNECT_TYPE=FASTPATH or MQ_CONNECT_TYPE=STANDARD, which specify a fastpath binding or a standard binding respectively. (Note that the environment variable is case sensitive.)

An application that uses the fastpath binding is known as a *trusted* application. There are only two ways of making a program run as trusted:

1. By specifying MQCNO_FASTPATH_BINDING on the MQCONNX call and not specifying the environment variable MQ_CONNECT_TYPE.
2. By specifying MQCNO_FASTPATH_BINDING on the MQCONNX call and setting the environment variable MQ_CONNECT_TYPE=FASTPATH.

Managing objects for triggering

Chapter 5. Administering remote MQSeries objects

This chapter describes how to administer MQSeries objects on another queue manager. It also describes how you can use remote queue objects to control the destination of messages and reply messages.

It contains these sections:

- “Understanding channels and remote queuing”
- “Remote administration” on page 64
- “Creating a local definition of a remote queue” on page 71
- “Using remote queue definitions as aliases” on page 74
- “Data conversion” on page 75

For more information about channels, their attributes, and how to set them up, refer to the *MQSeries Intercommunication* book.

Understanding channels and remote queuing

Queue managers communicate with each other using channels. For example, if an application is to put a message on a queue managed by a remote queue manager, a channel must be set up between the two queue managers. The channel is defined to the queue managers at each end of the connection. Each channel is named and has a number of attributes that define, for example, the type of channel and the protocol to be used for communication.

Channels are used for sending messages between queue managers. These messages may originate from:

- User-written application programs that transfer data from one node to another.
- User-written administration applications that use PCFs.
- Queue managers sending:
 - Instrumentation event messages to another queue manager.
 - MQSC commands issued from a **runmqsc** command in indirect mode (where the commands are run on another queue manager).

Channels are unidirectional. That is, messages can be sent in one direction only. Channel definitions are made in complementary pairs, one at each end of the connection. For example, if one end is a sender, the other must be a receiver.

Channels are ‘linked’ to queue managers (and therefore the applications they serve) by transmission queues and remote queue definitions. A transmission queue is used to forward messages (through a channel) to another queue manager. A remote queue definition identifies a queue on another queue manager. To give you an idea of how these things can fit together:

- A remote queue definition specifies a transmission queue.
- A channel serves a transmission queue, which is specified when the channel is defined.

“Preparing channels and transmission queues for remote administration” on page 65 shows how to use these definitions to set up remote administration.

Remote administration

You define a channel using the MQSC command `DEFINE CHANNEL`. Channels, their attributes, and how you use them in distributed queuing, are discussed at length in the *MQSeries Intercommunication* book. In this section, the examples concerned with channels use the default channel attributes unless otherwise specified.

Remote administration

This section tells you how to administer a remote queue manager from a local queue manager. You can implement remote administration from a local node using:

- MQSC commands
- PCF commands

Preparing the queues and channels is essentially the same for both methods. In this book, the examples show MQSC commands, because they are easier to understand. However, you can convert the examples to PCFs if you wish. For more information about writing administration programs using PCFs, see the *MQSeries Programmable System Management* book.

In remote administration you send MQSC commands to a remote queue manager—either interactively or from a text file containing the commands. The remote queue manager may be on the same machine or, more typically, on a different machine. You can remotely administer queue managers in other MQSeries environments, including UNIX systems, AS/400, MVS/ESA, OS/2, and Windows NT.

To implement remote administration, you must create specific objects. Unless you have specialized requirements, you should find that the default values (for example, for message length) are sufficient.

Preparing queue managers for remote administration

Figure 5 on page 65 shows the configuration of queue managers and channels that are required for remote administration. The object `source.queue.manager` is the *source* queue manager from which you can issue MQSC commands and to which the results of these commands (operator messages) are returned. The object `target.queue.manager` is the destination queue manager, which processes the commands and generates any operator messages.

Note: If you are using MQSC with the `-w` option, `source.queue.manager` *must* be the default queue manager. For further information on creating a queue manager, see “`crtmqm` (Create queue manager)” on page 240.




Figure 5. Remote administration

On both systems, if you have not already done so, you must:

- Create the queue manager and the default objects using the **crtmqm** command.
- Start the queue manager, using the **strmqm** command.

You have to run these commands locally or over a network facility such as Telnet.

On the destination queue manager:

- The command queue, SYSTEM.ADMIN.COMMAND.QUEUE, must be present. This queue is created by default when a queue manager is created.
- The command server must be started, using the **strmqcsv** command.

Preparing channels and transmission queues for remote administration

To run MQSC commands remotely, you must set up two channels, one for each direction, and their associated transmission queues. This example assumes that TCP/IP is being used as the transport type and that you know the TCP/IP address involved.

The channel `source.to.target` is for sending MQSC commands from the source queue manager to the destination. Its sender is at `source.queue.manager` and its receiver is at queue manager `target.queue.manager`. The channel `target.to.source` is for returning the output from commands and any operator messages that are generated to the source queue manager. You must also define a transmission queue for each sender. This queue is a local queue that is given the name of the receiving queue manager. Figure 6 on page 66 summarizes this configuration.

Remote administration

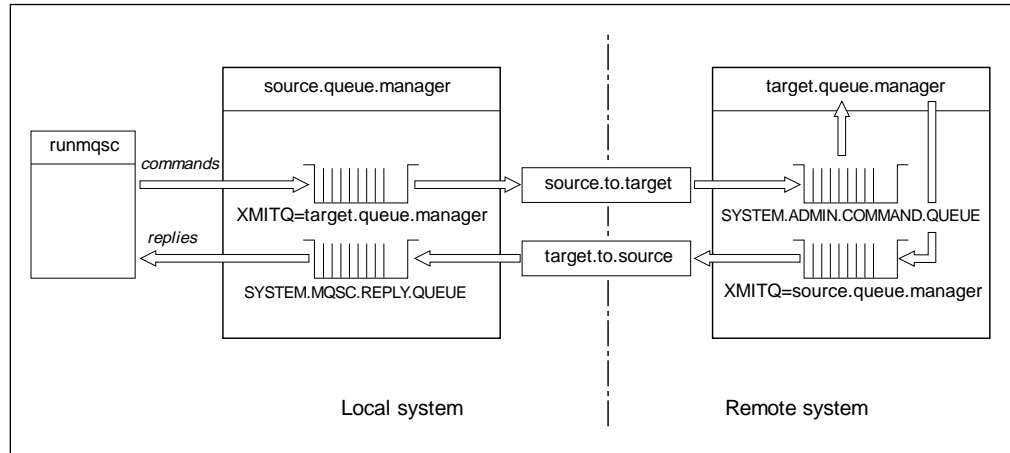


Figure 6. Setting up channels and queues for remote administration

See the *MQSeries Intercommunication* book for more information about setting up remote channels.

Defining channels and transmission queues

On the source queue manager, issue these MQSC commands to define the channels and the transmission queue:

* Define the sender channel at the source queue manager

```
DEFINE CHANNEL ('source.to.target') +  
  CHLTYPE(SDR) +  
  CONNAME (RH5498) +  
  XMITQ ('target.queue.manager') +  
  TRPTYPE(TCP)
```

* Define the receiver channel at the source queue manager

```
DEFINE CHANNEL ('target.to.source') +  
  CHLTYPE(RCVR) +  
  TRPTYPE(TCP)
```

* Define the transmission queue on the source

```
DEFINE QLOCAL ('target.queue.manager') +  
  USAGE (XMITQ)
```

Issue these commands on the destination queue manager (`target.queue.manager`), to create the channels and the transmission queue there:

```
* Define the sender channel on the destination queue manager

DEFINE CHANNEL ('target.to.source') +
  CHLTYPE(SDR) +
  CONNAME (RHX7721) +
  XMITQ ('source.queue.manager') +
  TRPTYPE(TCP)

* Define the receiver channel on the destination queue manager

DEFINE CHANNEL ('source.to.target') +
  CHLTYPE(RCVR) +
  TRPTYPE(TCP)

* Define the transmission queue on the destination queue manager

DEFINE QLOCAL ('source.queue.manager') +
  USAGE (XMITQ)
```

Note: The TCP/IP connection names specified for the `CONNAME` attribute in the sender channel definitions are for illustration only. This is the network name of the machine at the *other* end of the connection. Use the values appropriate for your network.

Start the channels

The way in which you start the channels depends on the environments in which MQSeries is running.

In MQSeries for UNIX systems, ensure that the `inetd` daemons have been configured for MQSeries and are running. Then start the channels as background processes:

- On the source queue manager, type:

```
runmqchl -c source.to.target &
```

- On the destination queue manager, type:

```
runmqchl -c target.to.source &
```

Remote administration

In MQSeries for OS/2 Warp and Windows NT, start a listener as a background process at the receiver end of each channel.

- On the source queue manager, type:

```
START runmqlsr -t TCP -m source.queue.manager
```

- On the destination queue.manager, type:

```
START runmqlsr -t TCP -m target.queue.manager
```

Then start the channels, again as background processes:

- On the source queue manager, type:

```
START runmqchl -c source.to.target
```

- On the destination queue manager, type:

```
START runmqchl -c target.to.source
```

The **runmqlsr** and **runmqchl** commands are MQSeries control commands. They cannot be issued using **runmqsc**. Channels can however be started using **runmqsc** commands or scripts (start channel).

Automatic definition of channels

Automatic definition of channels applies only if the target queue manager is running on of the MQSeries Version 5 products. If an inbound attach request is received and an appropriate receiver or server-connection definition cannot be found in the channel definition file (CDF), MQSeries creates a definition automatically and adds it to the CDF. Automatic definitions are based on two default definitions supplied with MQSeries, SYSTEM.AUTO.RECEIVER and SYSTEM.AUTO.SVRCONN. You enable automatic definition of receiver and server-connection definitions by updating the queue manager object using MQSC ALTER QMGR (or the PCF command Change Queue Manager).

For more information about the automatic creation of channel definitions, see the *MQSeries Intercommunication* book.

Issuing MQSC commands remotely

The command server *must* be running on the destination queue manager, if it is going to process MQSC commands remotely. (This is not necessary on the source queue manager.)

- On the destination queue manager, type:

```
strmqcsv target.queue.manager
```

- On the source queue manager, you can then run MQSC interactively in indirect mode by typing:

```
runmqsc -w 30 target.queue.manager
```

This form of the **runmqsc** command—with the **-w** flag—runs the MQSC commands in indirect mode, where commands are put (in a modified form) on the command-server input queue and executed in order.

When you type in an MQSC command, it is redirected to the remote queue manager, in this case, `target.queue.manager`. The timeout is set to 30 seconds; if a reply is not received within 30 seconds, the following message is generated on the local (source) queue manager:

```
AMQ8416: MQSC timed out waiting for a response from the command server.
```

At the end of the MQSC session, the local queue manager displays any timed-out responses that have arrived. When the MQSC session is finished, any further responses are discarded.

In indirect mode, you can also run an MQSC command file on a remote queue manager. For example:

```
runmqsc -w 60 target.queue.manager < mycomds.in > report.out
```

where `mycomds.in` is a file containing MQSC commands and `report.out` is the report file.

Working with queue managers on MVS/ESA

You can issue MQSC commands to an MVS/ESA queue manager from an MQSeries Version 5 queue manager. However, to do this, you must modify the **runmqsc** command and the channel definitions at the sender.

In particular, you add the **-x** flag to the **runmqsc** command on the source node:

```
runmqsc -w 30 -x target.queue.manager
```

Remote administration

On the sender channel, set the CONVERT attribute to YES. This specifies that the required data conversion between the systems is performed at the non-MVS/ESA end. The channel definition command now becomes:

```
* Define the sender channel at the source queue manager

DEFINE CHANNEL (source.to.target) +
    CHLTYPE(SDR) +
    CONNAME (RHX5498) +
    XMITQ (target.queue.manager) +
    TRPTYPE(TCP) +
    CONVERT (YES)
```

You must also define the receiver channel and the transmission queue at the source queue manager. Again, this example assumes that TCP/IP is the transmission protocol being used.

Recommendations for remote queuing

When you are implementing remote queuing:

1. Put the MQSC commands to be run on the remote system in a command file.
2. Verify your MQSC commands locally, by specifying the `-v` flag on the **runmqsc** command.

You cannot use **runmqsc** to verify MQSC commands on another queue manager.

3. Check that the command file runs locally without error.
4. Finally, run the command file against the remote system.

If you have problems using MQSC remotely

If you have difficulty in running MQSC commands remotely, use the following checklist to see if you have:

- Started the command server on the destination queue manager.
- Defined a valid transmission queue.
- Defined the two ends of the message channels for both:
 - The channel along which the commands are being sent.
 - The channel along which the replies are to be returned.
- Specified the correct connection name (CONNAME) in the channel definition.
- Started the listeners before you started the message channels.
- Checked that the disconnect interval has not expired, for example, if a channel started but then shut down after some time. This is especially important if you start the channels manually.
- Ensure that you are not trying to use MQSC from an MQSeries V2.2 source queue manager to control aspects of a later version of an MQSeries queue manager that cannot be controlled by an MQSeries V2.2 queue manager. Similarly, ensure that you are not trying to use MQSC from an MQSeries V2.2 source queue manager to control aspects of an MQSeries V5.0 queue manager that cannot be controlled by an MQSeries V2.2 queue manager.

- Ensure that you are not sending requests from an MQSeries V5.0 source queue manager that do not make sense to an MQSeries V2.2 queue manager (for example, requests that include new parameters).

See also “Resolving problems with MQSC” on page 47.

Creating a local definition of a remote queue

You can use a remote queue definition as a local definition of a remote queue. You create a remote queue definition on your local queue manager to identify a local queue on another queue manager.

Understanding how local definitions of remote queues work

An application connects to a local queue manager and then issues an MQOPEN call. In the open call, the queue name specified is that of a remote queue definition on the local queue manager. The remote queue definition supplies the names of the destination queue, the destination queue manager, and optionally, a transmission queue. To put a message on the remote queue, the application issues an MQPUT call, specifying the handle returned from the MQOPEN call. The queue manager uses the remote queue name and the remote queue manager name in a transmission header prepended to the message. This information is used to route the message to its correct destination in the network.

As administrator, you can control the destination of the message by altering the remote queue definition.

Example

Purpose: An application is required to put a message on a queue owned by a remote queue manager.

How it works: The application connects to a queue manager, for example, saturn.queue.manager. The destination queue is owned by another queue manager.

On the MQOPEN call, the application specifies these fields:

Field value	Description
<i>ObjectName</i> CYAN.REMOTE.QUEUE	Specifies the local name of the remote queue object. This defines the destination queue and the destination queue manager.
<i>ObjectType</i> (Queue)	Identifies this object as a queue.
<i>ObjectQmgrName</i> Blank or saturn.queue.manager	This field is optional. If blank, the name of the local queue manager is assumed. (This is the queue manager on which the remote queue definition exists.)

After this, the application issues an MQPUT call to put a message on to this queue.

Local definition of remote queue

On the local queue manager, you can create a local definition of a remote queue using the following MQSC commands:

```
DEFINE QREMOTE (CYAN.REMOTE.QUEUE) +
  DESCR ('Queue for auto insurance requests from the branches') +
  RNAME (AUTOMOBILE.INSURANCE.QUOTE.QUEUE) +
  RQMNAME (jupiter.queue.manager) +
  XMITQ (INQUOTE.XMIT.QUEUE)
```

where:

QREMOTE (CYAN.REMOTE.QUEUE)

Specifies the local name of the remote queue object. This is the name that applications connected to this queue manager must specify in the MQOPEN call to open the queue AUTOMOBILE.INSURANCE.QUOTE.QUEUE on the remote queue manager jupiter.queue.manager.

DESCR ('Queue for auto insurance requests from the branches')

Additional text that describes the use of the queue.

RNAME (AUTOMOBILE.INSURANCE.QUOTE.QUEUE)

Specifies the name of the destination queue on the remote queue manager. This is the real destination queue for messages that are sent by applications that specify the queue name CYAN.REMOTE.QUEUE. The queue AUTOMOBILE.INSURANCE.QUOTE.QUEUE must be defined as a local queue on the remote queue manager.

RQMNAME (jupiter.queue.manager)

Specifies the name of the remote queue manager that owns the destination queue AUTOMOBILE.INSURANCE.QUOTE.QUEUE.

XMITQ (INQUOTE.XMIT.QUEUE)

Specifies the name of the transmission queue. This is optional; if the name of a transmission queue is not specified, a queue with the same name as the remote queue manager is used.

In either case, the appropriate transmission queue must be defined as a local queue with a *Usage* attribute specifying that it is a transmission queue (USAGE(XMIT) in MQSC).

An alternative way of putting messages on a remote queue

Using a local definition of a remote queue is not the only way of putting messages on a remote queue. Applications can specify the full queue name, which includes the remote queue manager name, as part of the MQOPEN call. In this case, a local definition of a remote queue is not required. However, this alternative means that applications must either know or have access to the name of the remote queue manager at run time.

Using other commands with remote queues

You can use the appropriate MQSC commands to display or alter the attributes of a remote queue object, or you can delete the remote queue object. For example:

```
* Display the remote queue's attributes.

DISPLAY QUEUE (CYAN.REMOTE.QUEUE)

* ALTER the remote queue to enable puts.
* This does not affect the destination queue,
* only applications that specify this remote queue.

ALTER QREMOTE (CYAN.REMOTE.QUEUE) PUT(ENABLED)

* Delete this remote queue
* This does not affect the destination queue
* only its local definition

DELETE QREMOTE (CYAN.REMOTE.QUEUE)
```

Note: When you delete a remote queue, you delete only the local representation of the remote queue. You do not delete the remote queue itself or any messages on it.

Creating a transmission queue

A transmission queue is a local queue that is used when a queue manager forwards messages to a remote queue manager through a message channel. The channel provides a one-way link to the remote queue manager. Messages are queued at the transmission queue until the channel can accept them. When you define a channel, you must specify a transmission queue name at the sending end of the message channel.

The *Usage* attribute (USAGE in MQSC) defines whether a queue is a transmission queue or a normal queue.

Default transmission queues

Optionally, you can specify a transmission queue in a remote queue object, using the *XmitQName* attribute (XMITQ in MQSC). If no transmission queue is defined, a default is used. When applications put messages on a remote queue, if a transmission queue with the same name as the destination queue manager exists, that queue is used. If this queue does not exist, the queue specified by the *DefaultXmitQ* attribute (DEFXMITQ in MQSC) on the local queue manager is used.

Aliases

For example, the following MQSC command creates a default transmission queue on source.queue.manager for messages going to target.queue.manager:

```
DEFINE QLOCAL ('target.queue.manager') +  
  DESCR ('Default transmission queue for target qm') +  
  USAGE (XMITQ)
```

Applications can put messages directly on a transmission queue, or they can be put there indirectly, for example, through a remote queue definition. See also “Creating a local definition of a remote queue” on page 71.

Using remote queue definitions as aliases

In addition to locating a queue on another queue manager, you can also use a local definition of a remote queue for both:

- Queue manager aliases
- Reply-to queue aliases

Both types of aliases are resolved through the local definition of a remote queue.

As usual in remote queuing, the appropriate channels must be set up if the message is to arrive at its destination.

Queue manager aliases

An alias is the process by which the name of the destination queue manager—as specified in a message—is modified by a queue manager on the message route. Queue manager aliases are important because you can use them to control the destination of messages within a network of queue managers.

You do this by altering the remote queue definition on the queue manager at the point of control. The sending application is not aware that the queue manager name specified is an alias.

For more information about queue manager aliases, see the *MQSeries Intercommunication* book.

Reply-to queue aliases

Optionally, an application can specify the name of a reply-to queue when it puts a *request message* on a queue. If the application that processes the message extracts the name of the reply-to queue, it knows where to send the *reply message*, if required.

A reply-to queue alias is the process by which a reply-to queue – as specified in a request message – is altered by a queue manager on the message route. The sending application is not aware that the reply-to queue name specified is an alias.

A reply-to queue alias lets you alter the name of the reply-to queue and optionally its queue manager. This in turn lets you control which route is used for reply messages.

For more information about request messages, reply messages, and reply-to queues, see the *MQSeries Application Programming Reference*. For more information about reply-to queue aliases, see the *MQSeries Intercommunication* book.

Data conversion

Message data in MQSeries-defined formats (also known as *built-in formats*) can be converted by the queue manager from one coded character set to another, provided that both character sets relate to a single language or a group of similar languages. For example, conversion between coded character sets whose identifiers (CCSIDs) are 850 and 500 is supported, because both apply to Western European languages. Supported conversions are defined in the *MQSeries Application Programming Reference Manual*.

The queue manager cannot automatically convert messages in built-in formats if their CCSIDs represent different national-language groups. For example, conversion between CCSID 850 and CCSID 1025 (which is an EBCDIC coded character set for languages using Cyrillic script) is not supported, because many of the characters in one coded character set cannot be represented in the other. If you have a network of queue managers working in different national languages, and data conversion among some of the coded character sets is not supported, you can enable a default conversion. Default data conversion is described in “Default data conversion.”

File ccsid.tbl

The file `ccsid.tbl` is used for the following purposes:

- In MQSeries for Windows NT it records all the supported code sets. In MQSeries for OS/2 Warp and UNIX systems the supported code sets are held internally by the operating system.
- It specifies any additional code sets. To specify additional code sets, you need to edit `ccsid.tbl` (guidance on how to do this is provided in the file).
- It specifies any default data conversion.

You can update the information recorded in `ccsid.tbl`; you might want to do this if, for example, a future release of your operating system supports additional coded character sets.

In UNIX environments a sample `ccsid.tbl` file is provided as `mqmtop/samp/ccsid.tbl`.

In MQSeries for UNIX systems, `ccsid.tbl` is located in directory `/var/mqm/conv/table`. In MQSeries for OS/2 Warp and Windows NT, `ccsid.tbl` is located on the boot drive in directory `\MQM\CONV\TABLE`.

Default data conversion

To implement default data conversion, you edit `ccsid.tbl` to specify a default EBCDIC CCSID and a default ASCII CCSID, and also to specify the defaulting CCSIDs. Instructions for doing this are included in the file.

If you update `ccsid.tbl` to implement default data conversion, the queue manager must be restarted before the change can take effect.

Data conversion

The default data-conversion process is as follows:

- If conversion between the source and target CCSIDs is not supported, but the CCSIDs of the source and target environments are either both EBCDIC or both ASCII, the character data is passed to the target application without conversion.
- If one CCSID represents an ASCII coded character set, and the other represents an EBCDIC coded character set, MQSeries converts the data using the default data-conversion CCSIDS defined in `ccsid.tbl`.

Note: You should try to restrict the characters being converted to those that have the same code values in the coded character set specified for the message and in the default coded character set. If you use only that set of characters that is valid for MQSeries object names (as defined in “Names of MQSeries objects” on page 235) you will, in general, satisfy this requirement. Exceptions occur with EBCDIC CCSIDs 290, 930, 1279, and 5026 used in Japan, where the lowercase characters have different codes from those used in other EBCDIC CCSIDs.

Conversion of messages in user-defined formats

Messages in user-defined formats cannot be converted from one coded character set to another by the queue manager. If data in a user-defined format requires conversion, you must supply a data-conversion exit for each such format. The use of default CCSIDs for converting character data in user-defined formats is not recommended, though it is possible. For more information about converting data in user-defined formats and about writing data conversion exits, see the *MQSeries Application Programming Guide*.

Chapter 6. Protecting MQSeries objects

This chapter describes how to prevent unauthorized access to MQSeries objects in these environments:

- MQSeries for AIX
- MQSeries for HP-UX
- MQSeries for Sun Solaris
- MQSeries for Windows NT

This chapter does not apply to MQSeries for OS/2 Warp.

It contains these sections:

- “Before you begin (UNIX systems)”
- “Before you begin (Windows NT)” on page 78
- “Why you need to protect MQSeries resources” on page 79
- “Understanding the Object Authority Manager” on page 79
- “Using the Object Authority Manager commands” on page 82
- “Object Authority Manager guidelines” on page 85
- “Understanding the authorization specification tables” on page 88
- “Understanding authorization files” on page 94

Before you begin (UNIX systems)

In MQSeries for UNIX systems, UNIX restrictions mean that all user IDs must be defined in lowercase.

In MQSeries for UNIX systems, all queue manager processes run with these IDs:

User ID	mqm
Group	mqm

You must create this user ID and group before you install MQSeries. For more information about setting the IDs, see the appropriate *MQSeries Quick Beginnings* booklet.

User IDs in user group mqm (UNIX systems)

If your user ID belongs to group mqm, you have all authorities to all MQSeries resources. Your user ID **must** belong to group mqm to be able to use all the MQSeries control commands, except **crtmqcvx**. In particular, you need this authority to:

- Use the **runmqsc** command to run MQSC commands
- Administer authorities using the **setmqaut** command

If you are sending channel commands to remote queue managers, you must ensure that your user ID is a member of group mqm on the target system. For a list of PCF and MQSC channel commands, see “Channel command security” on page 87.

It is not essential for your user ID to belong to group mqm for issuing:

- PCF commands—including Escape PCFs—from an administration program
- MQI calls from an application program

Before you begin (Windows NT)

If the local mqm group does not already exist on the local computer, it is created automatically when MQSeries for Windows NT is installed.

User IDs (Windows NT systems)

If your user ID belongs to the local mqm or Administrators group, you can administer any queue manager on that system. The system-defined user ID 'SYSTEM' can also administer any queue manager.

The name of the local mqm group to be used for privileged MQSeries administration is fixed, and it can contain (directly, or indirectly by the inclusion of global groups) users who require MQSeries authority to any queue manager on the workstation or server.

In order to run all the MQSeries for Windows NT control commands, your user ID must belong to the local mqm or Administrators group. In particular, you need this authority to:

- Use the **runmqsc** command to run MQSC commands
- Administer authorities on MQSeries for Windows NT using the **setmqaut** command
- Create a queue manager using **crtmqm**

If you are sending channel commands to queue managers on a remote Windows NT system, you must ensure that your user ID is a member of the mqm or Administrators group on the target system. For a list of PCF and MQSC channel commands, see "Channel command security" on page 87.

Some control commands, for example, **crtmqm**, manipulate authorities on MQSeries objects using the Object Authority Manager (OAM). As described in "Understanding the Object Authority Manager" on page 79, the OAM uses a predefined search order to determine the authority rights for a given user ID. Consequently the authorities granted to your user ID may differ from those determined by the OAM. For example, if you issue **crtmqm** from a user ID authenticated by a domain controller that has membership of the local mqm group through a global group, the command fails if the system has a local user of the same name who is not in the local mqm group.

Your user ID does not have to belong to group mqm in order to issue:

- PCF commands—including Escape PCFs—from an administration program
- MQI calls from an application program

Notes:

1. For MQSeries authorizations, names of user IDs and groups are limited to a maximum of 12 characters, and no spaces are allowed. This means that the Windows NT system-defined 'Administrator' user ID cannot issue MQSeries control commands.
2. When you use a Domain user ID defined on a remote machine, you must be a member of the local mqm or Administrators group to (1) issue commands (such as create queue manager), and (2) grant MQSeries authorities.

Restricted access NT objects

When MQSeries creates restricted access NT objects, full control permission is given to the following entities:

- The local mqm group on the local computer
- The local Administrators group on the local computer
- The SYSTEM user ID

Why you need to protect MQSeries resources

Because MQSeries queue managers handle the transfer of information that is potentially valuable, you need the safeguard of an authority system. This ensures that the resources that a queue manager owns and manages are protected from unauthorized access, which could lead to the loss or disclosure of the information. In a secure system, it is essential that none of the following is accessed or changed by any unauthorized user or application:

- Connections to a queue manager
- Access to MQSeries objects such as queues, channels, and processes
- Commands for queue manager administration, including MQSC commands and PCF commands
- Access to MQSeries messages
- Context information associated with messages

You should develop your own policy with respect to which users have access to which resources.

Understanding the Object Authority Manager

By default, access to queue-manager resources is controlled through an authorization service installable component. This component is formally called the Object Authority Manager (OAM) for MQSeries. It is supplied with MQSeries, and is automatically installed and enabled for each queue manager you create, unless you specify otherwise. In this chapter, the term OAM is used to denote the Object Authority Manager supplied with MQSeries.

The OAM is an *installable component* of the authorization service. Providing the OAM as an installable service gives you the flexibility to:

- Replace the supplied OAM with your own authorization service component using the interface provided.
- Augment the facilities supplied by the OAM with those of your own authorization service component, again using the interface provided.
- Remove or disable the OAM, and run with no authorization service at all.

For more information on installable services, see the *MQSeries Programmable System Management* manual.

Object authority manager

The OAM manages users' authorizations to manipulate MQSeries objects, including queues and process definitions. It also provides a command interface through which you can grant or revoke access authority to an object for a specific group of users. The decision to allow access to a resource is made by the OAM, and the queue manager follows that decision. If the OAM cannot make a decision, the queue manager prevents access to that resource.

How the OAM works

The OAM works by exploiting the security features of the underlying operating system. In particular, the OAM uses operating system user and group IDs. Users can access queue manager objects only if they have the required authority.

Managing access through user groups

In the command interface, we use the term *principal* rather than user ID. The reason for this is that authorities granted to a user ID can also be granted to other entities, for example, an application program that issues MQI calls, or an administration program that issues PCF commands. In these cases, the principal associated with the program is not necessarily the user ID that was used when the program was started. However, in this discussion, principals are always user IDs.

Group sets and the primary group

Managing access permissions to MQSeries resources is based on user groups (that is, on groups of principals). A principal can belong to one or more groups. If it belongs to more than one group, the groups to which it belongs are known as its *group set*.

In MQSeries for UNIX systems only, one of the groups in the group set is the *primary group*. For MQSeries for Windows NT, the role of the primary group is fulfilled by the user ID. The Windows NT primary group associated with a user ID is given no special treatment by MQSeries; it is handled in the same way as any other group.

MQSeries for Windows NT only

The OAM searches for the specified user in the following order:

1. In the security database of the local computer
2. In the security database of the domain controller for the local computer
3. In the security database of the domain controller for the user ID under which the OAM is running

The first user ID encountered is used when checking for group membership.

Note that each of these user IDs may have different group memberships on a particular computer.

End of MQSeries for Windows NT only

The OAM maintains authorizations at the level of groups rather than individual principals. The mapping of principals to group names is carried out within the OAM, and operations are carried out at the group level. You can, however, display the authorizations of an individual principal.

When a principal belongs to more than one group

The authorizations that a principal has are derived from the union of the authorizations of its group set. Whenever a principal requests access to a resource, the OAM computes this union, and uses the resultant authorization to check the principal's access to the resource. You can use the control command **setmqaut** to set the authorizations for a specific principal. However, for MQSeries for UNIX systems, this also gives the same authorizations to the principal's primary group.

The group set associated with a principal is cached when the group authorizations are computed by the OAM. Any changes made to a group's authorizations after the group set has been cached are not recognized until the queue manager is restarted.

Default user group

The OAM recognizes a default user group to which all users are nominally assigned. This group has a group ID of 'nobody'. By default, no authorizations are given to this group. Users without specific authorizations can be granted access to MQSeries resources through this group ID.

Resources you can protect with the OAM

Through OAM you can control:

- Access to MQSeries objects through the MQI. When an application program attempts to access an object, the OAM checks that the user ID making the request has the authorization for the operation requested.

In particular, this means that queues, and the messages on queues, can be protected from unauthorized access.

- Permission to use PCF commands.

Different groups of users may be granted different kinds of access authority to the same object. For example, for a specific queue, one group may be allowed to perform both put and get operations; another group may be allowed only to browse the queue (MQGET with browse option). Similarly, some groups may have get and put authority to a queue, but are not allowed to alter or delete the queue.

Using groups for authorizations

Using groups, rather than individual principals, for authorization reduces the amount of administration required. Typically, a particular kind of access is required by more than one principal. For example, you might define a group consisting of end users who want to run a particular application. New users can be given access simply by adding their user ID to the appropriate group.

Try to keep the number of groups as small as possible. For example, dividing principals into one group for application users and one for administrators is a good place to start.

Notes:

1. In MQSeries for UNIX systems, when you change the authorization of a principal, you also change the authorization of its primary group. This makes it especially important to ensure that you do not change the authorization of a principal inadvertently, simply because it belongs to the same primary group as the principal you specified when you changed an authorization.
2. MQSeries for Windows NT treats the local Administrators group in a special manner. Members of this group are always granted full access rights which cannot be removed.

Disabling the object authority manager

By default, the OAM is enabled. You can disable it by setting the operating system variable MQSNOAUT before the queue manager is created.

In MQSeries for UNIX systems, you set MQSNOAUT as follows:

```
export MQSNOAUT=yes
```

For MQSeries for Windows NT, you set MQSNOAUT as follows:

```
SET MQSNOAUT=yes
```

However, if you do this you cannot, in general, restart the OAM later. A better approach is to have the OAM enabled and ensure that all users and applications have access through an appropriate group or user ID.

You can also disable the OAM, for testing purposes only, by removing the authorization service stanza in the queue manager configuration file (qm.ini).

Using the Object Authority Manager commands

The OAM provides a command interface for granting and revoking authority. Before you can use these commands, you must be suitably authorized:

- In MQSeries for UNIX systems, your user ID must belong to the group mqm, which you define when you install MQSeries.
- In MQSeries for Windows NT, your user ID must belong to either the local mqm group or the local Administrators group.

If your user ID is a member of mqm, or, for MQSeries for Windows NT, of either mqm or the local Administrators group, you have a 'super user' authority to the queue manager, which means that you are authorized to issue any MQI request or command from your user ID.

The OAM provides two control commands that allow you to manage the authorizations of users. These are:

- **setmqaut** (Set or reset authority)
- **dspmqaut** (Display authority)

Authority checking occurs in the following calls: MQCONN, MQOPEN, MQPUT1, and MQCLOSE. Therefore, any changes made to the authority of an object using **setmqaut** do not take effect until you reset the object.

What you specify when you use the OAM commands

The authority commands **setmqaut** and **dspmqaut** apply to the specified queue manager; if you do not specify the name of a queue manager, the default queue manager is assumed. On these commands, you must also identify the object uniquely (that is, you must specify the object name and its type). You also have to specify the principal or group name to which the authority applies.

Authorization lists

On the **setmqaut** command you specify a list of authorizations. This is simply a shorthand way of specifying whether authorization is to be granted or revoked, and of identifying the resources to which the change in authorization applies. Each authorization in the list is specified as a lowercase keyword, prefixed with a plus sign (+) or a minus sign (-). Use a plus sign to add the specified authorization, and a minus sign to remove the authorization. You can specify any number of authorizations in a single command. For example:

```
+browse -get +put
```

Using the setmqaut command

Provided you have the required authorization, you can use the **setmqaut** command to grant or revoke authorization of a principal or user group to access a particular object. The following example shows how the **setmqaut** command is used:

```
setmqaut -m saturn.queue.manager -t queue -n RED.LOCAL.QUEUE -g groupa +browse -get +put
```

In this example:

- saturn.queue.manager is the queue manager name.
- queue is the object type.
- RED.LOCAL.QUEUE is the object name.
- groupa is the ID of the group whose the authorizations are to change.
- +browse -get +put is the authorization list for the specified queue. There must be no spaces between the '+' or '-' signs and the keyword.
 - +browse adds authorization to browse messages on the queue (to issue MQGET with the browse option).
 - -get removes authorization to get (MQGET) messages from the queue.
 - +put adds authorization to put (MQPUT) messages on the queue.

In summary, applications started with user IDs that belong to user group groupa have at least these authorizations.

Using OAM commands

You can specify one or more principals and, at the same time, one or more groups. For example, the following command revokes put authority on the queue MyQueue from the principal fvuser and from groups groupa and groupb.

```
setmqaut -m saturn.queue.manager -t queue -n MyQueue -p fvuser -g groupa -g groupb -put
```

Note: For MQSeries for UNIX systems, this command also revokes put authority for all principals in the primary group of FvUser.

For a formal definition of the command and its syntax, see “setmqaut (Set/reset authority)” on page 284.

Authority commands and installable services

The **setmqaut** command takes an additional parameter that specifies the name of the installable service component to which the update applies. You must specify this parameter if you have multiple installable components running at the same time. By default, this is not the case. If the parameter is omitted, the update is made to the first installable service of that type, if one exists. By default, this is the supplied OAM.

Access authorizations

Authorizations defined by the authorization list associated with the **setmqaut** command can be categorized as follows:

- Authorizations related to MQI calls
- Authorization related administration commands
- Context authorizations
- General authorizations, that is, for MQI calls, for commands, or both

Each authorization is specified by a keyword used with the **setmqaut** and **dspmqa** commands. These are described in “setmqaut (Set/reset authority)” on page 284.

Display authority command

You can use the command **dspmqa** to view the authorizations that a specific principal or group has for a particular object. The flags have the same meaning as those in the **setmqaut** command. Authorization can be displayed for only one group or principal at a time. See “dspmqa (Display authority)” on page 248 for a formal specification of this command.

For example, the following command displays the authorizations that the group GpAdmin has to a process definition named Annuities on queue manager QueueMan1.

```
dspmqa -m QueueMan1 -t process -n Annuities -g GpAdmin
```

The keywords displayed as a result of this command identify the authorizations that are active.

Object Authority Manager guidelines

Some operations are particularly sensitive and should be limited to privileged users. For example,

- Accessing some special queues, such as transmission queues or the command queue SYSTEM.ADMIN.COMMAND.QUEUE
- Running programs that use full MQI context options
- Creating and copying application queues

User IDs

This information applies to MQSeries for UNIX systems only.

The special user ID mqm that you create is intended for use by the product only. It should never be available to nonprivileged users.

If an MQ process is associated with a login session, then the authorization routines check the real (logged-in) user ID.

If an MQ process is not associated with a login session (for example, if the process is invoked from a daemon such as inetd), the effective user ID is used for authorization. In a CICS environment, the CICS user ID associated with the transaction is used.

All objects are owned by user ID mqm.

Queue manager directories

The directory containing queues and other queue manager data is private to the product. Do not use standard operating system commands to grant or revoke authorizations to MQI resources.

Queues

The authority to a dynamic queue is based on, but is not necessarily the same as, that of the model queue from which it is derived. See note 11 on page 91 for more information.

For alias queues and remote queues, the authorization is that of the object itself, not the queue to which the alias or remote queue resolves. It is, therefore, possible to authorize a user ID to access an alias queue that resolves to a local queue to which the user ID has no access permissions.

You should limit the authority to create queues to privileged users. If you do not, users may bypass the normal access control simply by creating an alias.

Alternate-user authority

Alternate-user authority controls whether one user ID can use the authority of another user ID when accessing an MQSeries object. This is essential where a server receives requests from a program and the server wishes to ensure that the program has the required authority for the request. The server may have the required authority, but it needs to know whether the program has the authority for the actions it has requested.

For example:

- A server program running under user 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, you can use alternate-user authority to control whether PAYSERV is allowed to specify USER1 as an alternate-user ID when it opens the reply-to queue.

The alternate-user ID is specified on the *AlternateUserId* field of the object descriptor.

Note: You can use alternate-user IDs on any MQSeries object. Use of an alternate-user ID does not affect the user ID used by any other resource managers.

Context authority

Context is information that applies to a particular message and is contained in the message descriptor, MQMD, which is part of the message. The context information comes in two sections:

Identity section This part specifies who the message came from. It consists of the following fields:

- *UserIdentifier*
- *AccountingToken*
- *ApplIdentityData*

Origin section This section specifies where the message came from, and when it was put onto the queue. It consists of the following fields:

- *PutApplType*
- *PutApplName*
- *PutDate*
- *PutTime*
- *ApplOriginData*

Applications can specify the context data when either an MQOPEN or an MQPUT call is made. This data may be generated by the application, it may be passed on from another message, or it may 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, testing whether the message came from an application, running under an authorized user ID.

A server program can use the *UserIdentifier* to determine the user ID of an alternate user.

You use context authorization 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 *MQSeries Application Programming Guide*. For descriptions of the message descriptor fields relating to context, see the *MQSeries Application Programming Reference* manual.

Remote security considerations

For remote security, you should consider:

Put authority For security across queue managers you can specify the put authority that is used when a channel receives a message sent from another queue manager.

Specify the channel attribute PUTAUT as follows:

DEF Default user ID. This is the user ID that the message channel agent is running under.

CTX The user ID in the message context.

Transmission queues

Queue managers automatically put remote messages on a transmission queue; no special authority is required for this. However, putting a message directly on a transmission queue requires special authorization; see Table 7 on page 89.

Channel exits Channel exits can be used for added security.

For more information about remote security, see the *MQSeries Intercommunication* book.

Channel command security

Channel commands can be issued as PCF commands, MQSC commands, and control commands.

PCF commands

You can issue PCF channel commands by sending a PCF message to the SYSTEM.ADMIN.COMMAND.QUEUE on a remote MQSeries system. The user ID, as specified in the message descriptor of the PCF message, must belong to group mqm (or the Administrator's group in the MQSeries for Windows NT) on the target system. These commands are:

- *ChangeChannel*
- *CopyChannel*
- *CreateChannel*
- *DeleteChannel*
- *PingChannel*
- *ResetChannel*
- *StartChannel*
- *StartChannelInitiator*
- *StopChannel*
- *ResolveChannel*

See the *MQSeries Programmable System Management* manual for the PCF security requirements.

MQSC channel commands

You can issue MQSC channel commands to a remote MQSeries system either by sending the command directly in a PCF escape message or by issuing the command using **runmqsc** in indirect mode. The user ID as specified in the message descriptor of the associated PCF message must belong to group mqm (or the Administrator's group in MQSeries for Windows NT) on the target system.

Authorization specification tables

(PCF commands are implicit in MQSC commands issued from **runmqsc** in indirect mode.) These commands are:

- ALTER CHANNEL
- DEFINE CHANNEL
- DELETE CHANNEL
- PING CHANNEL
- RESET CHANNEL
- START CHANNEL
- START CHINIT
- STOP CHANNEL
- RESOLVE CHANNEL

For MQSC commands issued from the **runmqsc** command, the user ID in the PCF message is normally that of the current user.

Control commands for channels

For the control commands for channels, the user ID that issues them must belong to user group mqm (or the Administrator's group in MQSeries for Windows NT). These commands are:

- **runmqchi** (Run channel initiator)
- **runmqchl** (Run channel)
- **runmqlsr** (Run listener). The **runmqlsr** command is supported by MQSeries for Windows NT only.

Understanding the authorization specification tables

The authorization specification tables starting on page 89 define precisely how the authorizations work and the restrictions that apply. The tables apply to these situations:

- Applications that issue MQI calls
- Administration programs that issue MQSC commands as escape PCFs
- Administration programs that issue PCF commands

In this section, the information is presented as a set of tables that specify the following:

Action to be performed MQI option, MQSC command, or PCF command.

Access control object Queue, process, or queue manager.

Authorization required Expressed as an 'MQZAO_' constant.

In the tables, the constants prefixed by MQZAO_ correspond to the keywords in the authorization list for the **setmqaut** command for the particular entity. For example, MQZAO_BROWSE corresponds to the keyword +browse; similarly, the keyword MQZAO_SET_ALL_CONTEXT corresponds to the keyword +setall and so on. These constants are defined in the header file cmqzc.h, which is supplied with the product. See "What the authorization files contain" on page 96 for more information.

MQI authorizations

An application is allowed to issue specific MQI calls and options only if the user identifier under which it is running (or whose authorizations it is able to assume) has been granted the relevant authorization.

Four MQI calls may require authorization checks: MQCONN, MQOPEN, MQPUT1, and MQCLOSE.

For MQOPEN and MQPUT1, the authority check is made on the name of the object being opened, and not on the name, or names, resulting after a name has been resolved. For example, an application may be granted authority to open an alias queue without having authority to open the base queue to which the alias resolves. The rule is that the check is carried out on the first definition encountered during the process of name resolution that is not a queue-manager alias, unless the queue-manager alias definition is opened directly; that is, its name appears in the *ObjectName* field of the object descriptor. Authority is always needed for the particular object being opened; in some cases additional queue-independent authority—which is obtained through an authorization for the queue-manager object—is required.

Table 7 summarizes the authorizations needed for each call.

<i>Table 7 (Page 1 of 2). Security authorization needed for MQI calls</i>			
Authorization required for:	Queue object (1)	Process object	Queue manager object
MQCONN option	Not applicable	Not applicable	MQZAO_CONNECT
MQOPEN Option			
MQOO_INQUIRE	MQZAO_INQUIRE (2)	MQZAO_INQUIRE (2)	MQZAO_INQUIRE (2)
MQOO_BROWSE	MQZAO_BROWSE	Not applicable	No check
MQOO_INPUT_*	MQZAO_INPUT	Not applicable	No check
MQOO_SAVE_ALL_CONTEXT (3)	MQZAO_INPUT	Not applicable	No check
MQOO_OUTPUT (Normal queue) (4)	MQZAO_OUTPUT	Not applicable	No check
MQOO_PASS_IDENTITY_CONTEXT (5)	MQZAO_PASS_IDENTITY_CONTEXT	Not applicable	No check
MQOO_PASS_ALL_CONTEXT (5, 6)	MQZAO_PASS_ALL_CONTEXT	Not applicable	No check
MQOO_SET_IDENTITY_CONTEXT (5, 6)	MQZAO_SET_IDENTITY_CONTEXT	Not applicable	MQZAO_SET_IDENTITY_CONTEXT (7)
MQOO_SET_ALL_CONTEXT (5, 8)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (7)
MQOO_OUTPUT (Transmission queue) (9)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (7)
MQOO_SET	MQZAO_SET	Not applicable	No check
MQOO_ALTERNATE_USER_AUTHORITY	(10)	(10)	MQZAO_ALTERNATE_USER_AUTHORITY (10, 11)
MQPUT1 Option			
MQPMO_PASS_IDENTITY_CONTEXT	MQZAO_PASS_IDENTITY_CONTEXT (12)	Not applicable	No check
MQPMO_PASS_ALL_CONTEXT	MQZAO_PASS_ALL_CONTEXT (12)	Not applicable	No check

Authorization specification tables

MQPMO_SET_IDENTITY_CONTEXT	MQZAO_SET_IDENTITY_CONTEXT (12)	Not applicable	MQZAO_SET_IDENTITY_CONTEXT (7)
MQPMO_SET_ALL_CONTEXT	MQZAO_SET_ALL_CONTEXT (12)	Not applicable	MQZAO_SET_ALL_CONTEXT (7)
(Transmission queue) (9)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (7)
MQPMO_ALTERNATE_USER_AUTHORITY	(13)	Not applicable	MQZAO_ALTERNATE_USER_AUTHORITY (11)
MQCLOSE Option			
MQCO_DELETE	MQZAO_DELETE (14)	Not applicable	Not applicable
MQCO_DELETE_PURGE	MQZAO_DELETE (14)	Not applicable	Not applicable

Specific notes:

1. If a model queue is being opened:
 - MQZAO_DISPLAY authority is needed for the model queue, in addition to the authority to open the model queue for the type of access for which you are opening.
 - MQZAO_CREATE authority is not needed to create the dynamic queue.
 - The user identifier used to open the model queue is automatically granted all of the queue-specific authorities (equivalent to MQZAO_ALL) for the dynamic queue created.
2. Either the queue, process, or queue manager object is checked, depending on the type of object being opened.
3. MQOO_INPUT_* must also be specified. This is valid for a local, model, or alias queue.
4. This check is performed for all output cases, except the case specified in note 9.
5. MQOO_OUTPUT must also be specified.
6. MQOO_PASS_IDENTITY_CONTEXT is also implied by this option.
7. This authority is required for both the queue manager object and the particular queue.
8. MQOO_PASS_IDENTITY_CONTEXT, MQOO_PASS_ALL_CONTEXT, and MQOO_SET_IDENTITY_CONTEXT are also implied by this option.
9. This check is performed for a local or model queue that has a *Usage* queue attribute of MQUS_TRANSMISSION, and is being opened directly for output. It does not apply if a remote queue is being opened (either by specifying the names of the remote queue manager and remote queue, or by specifying the name of a local definition of the remote queue).
10. At least one of MQOO_INQUIRE (for any object type), or (for queues) MQOO_BROWSE, MQOO_INPUT_*, MQOO_OUTPUT, or MQOO_SET must also be specified. The check carried out is as for the other options specified, using the supplied alternate-user identifier for the specific-named object authority, and the current application authority for the MQZAO_ALTERNATE_USER_IDENTIFIER check.

11. This authorization allows any *AlternateUserId* to be specified.
12. An MQZAO_OUTPUT check is also carried out, if the queue does not have a *Usage* queue attribute of MQUS_TRANSMISSION.
13. The check carried out is as for the other options specified, using the supplied alternate-user identifier for the specific-named queue authority, and the current application authority for the MQZAO_ALTERNATE_USER_IDENTIFIER check.
14. The check is carried out only if both of the following are true:
 - A permanent dynamic queue is being closed and deleted.
 - The queue was not created by the MQOPEN which returned the object handle being used.

Otherwise, there is no check.

General notes:

1. The special authorization MQZAO_ALL_MQI includes all of the following that are relevant to the object type:
 - MQZAO_CONNECT
 - MQZAO_INQUIRE
 - MQZAO_SET
 - MQZAO_BROWSE
 - MQZAO_INPUT
 - MQZAO_OUTPUT
 - MQZAO_PASS_IDENTITY_CONTEXT
 - MQZAO_PASS_ALL_CONTEXT
 - MQZAO_SET_IDENTITY_CONTEXT
 - MQZAO_SET_ALL_CONTEXT
 - MQZAO_ALTERNATE_USER_AUTHORITY
2. MQZAO_DELETE (see note 14) and MQZAO_DISPLAY are classed as administration authorizations. They are not therefore included in MQZAO_ALL_MQI.
3. 'No check' means that no authorization checking is carried out.
4. 'Not applicable' means that authorization checking is not relevant to this operation. For example, you cannot issue an MQPUT call to a process object.

Administration authorizations

These authorizations allow a user to issue administration commands. This can be an MQSC command as an escape PCF message or as a PCF command itself. These methods allow a program to send an administration command as a message to a queue manager, for execution on behalf of that user.

Authorizations for MQSC commands in escape PCFs

Table 8 summarizes the authorizations needed for each MQSC command that is contained in Escape PCF.

<i>Table 8. MQSC commands and security authorization needed</i>			
(2) Authorization required for:	Queue object	Process object	Queue manager object
MQSC command			
ALTER object	MQZAO_CHANGE	MQZAO_CHANGE	MQZAO_CHANGE
CLEAR QLOCAL	MQZAO_CLEAR	Not applicable	Not applicable
DEFINE object NOREPLACE (3)	MQZAO_CREATE (4)	MQZAO_CREATE (4)	Not applicable
DEFINE object REPLACE (3, 5)	MQZAO_CHANGE	MQZAO_CHANGE	Not applicable
DELETE object	MQZAO_DELETE	MQZAO_DELETE	Not applicable
DISPLAY object	MQZAO_DISPLAY	MQZAO_DISPLAY	MQZAO_DISPLAY

Specific notes:

1. The user identifier, under which the program (for example, **runmqsc**) which submits the command is running, must also have MQZAO_CONNECT authority to the queue manager.
2. Either the queue, process, or queue manager object is checked, depending on the type of object.
3. For DEFINE commands, MQZAO_DISPLAY authority is also needed for the LIKE object if one is specified, or on the appropriate SYSTEM.DEFAULT.xxx object if LIKE is omitted.
4. The MQZAO_CREATE authority is not specific to a particular object or object type. Create authority is granted for all objects, for a specified queue manager, by specifying an object type of QMGR on the **setmqaut** command.
5. This applies if the object to be replaced does in fact already exist. If it does not, the check is as for DEFINE object NOREPLACE.

General notes:

1. To perform any PCF command, you must have DISPLAY authority on the queue manager.
2. The authority to execute an escape PCF depends on the MQSC command within the text of the escape PCF message.
3. 'Not applicable' means that authorization checking is not relevant to this operation. For example, you cannot issue a CLEAR QLOCAL on a queue manager object.

Authorizations for PCF commands

Table 9 summarizes the authorizations needed for each PCF command.

<i>Table 9. PCF commands and security authorization needed</i>			
(2) Authorization required for:	Queue object	Process object	Queue manager object
PCF command			
Change object	MQZAO_CHANGE	MQZAO_CHANGE	MQZAO_CHANGE
Clear Queue	MQZAO_CLEAR	Not applicable	Not applicable
Copy object (without replace) (3)	MQZAO_CREATE (4)	MQZAO_CREATE (4)	Not applicable
Copy object (with replace) (3, 6)	MQZAO_CHANGE	MQZAO_CHANGE	Not applicable
Create object (without replace) (5)	MQZAO_CREATE (4)	MQZAO_CREATE (4)	Not applicable
Create object (with replace) (5, 6)	MQZAO_CHANGE	MQZAO_CHANGE	Not applicable
Delete object	MQZAO_DELETE	MQZAO_DELETE	Not applicable
Inquire object	MQZAO_DISPLAY	MQZAO_DISPLAY	MQZAO_DISPLAY
Inquire object names	No check	No check	No check
Reset queue statistics	MQZAO_DISPLAY and MQZAO_CHANGE	Not applicable	Not applicable

Specific notes:

1. The user identifier under which the program submitting the command is running must also have authority to connect to its local queue manager, and to open the command admin queue for output.
2. Either the queue, process, or queue-manager object is checked, depending on the type of object.
3. For Copy commands, MQZAO_DISPLAY authority is also needed for the From object.
4. The MQZAO_CREATE authority is not specific to a particular object or object type. Create authority is granted for all objects, for a specified queue manager, by specifying an object type of QMGR on the **setmqaut** command.
5. For Create commands, MQZAO_DISPLAY authority is also needed for the appropriate SYSTEM.DEFAULT.* object.
6. This applies if the object to be replaced already exists. If it does not, the check is as for Copy or Create without replace.

General notes:

1. To perform any PCF command, you must have DISPLAY authority on the queue manager.
2. The special authorization MQZAO_ALL_ADMIN includes all of the following that are relevant to the object type:
 - MQZAO_CHANGE
 - MQZAO_CLEAR
 - MQZAO_DELETE
 - MQZAO_DISPLAY

MQZAO_CREATE is not included, because it is not specific to a particular object or object type.
3. 'No check' means that no authorization checking is carried out.

4. 'Not applicable' means that authorization checking is not relevant to this operation. For example, you cannot use a Clear Queue command on a process object.

Understanding authorization files

Note: The information in this section is given for problem determination purposes. Under normal circumstances, use authorization commands to view and change authorization information.

MQSeries uses a specific file structure to implement security. You should not have to do anything with these files, except to ensure that all the authorization files are themselves secure.

Security is implemented by authorization files. From this perspective, there are three types of authorization:

- Authorizations applying to single object, for example, the authority to put a message on an queue.
- Authorizations applying to a class of objects, for example, the authority to create a queue.
- Authorizations applying across all classes of objects, for example, the authority to perform operations on behalf of different users.

Authorization file paths

The path to an authorization file depends on its type. When you specify an authorization for an object, for example, the queue manager creates the appropriate authorization files. It puts these files into a subdirectory, the path of which is defined by the queue manager name, the type of authorization, and where appropriate, the object name.

Not all authorizations apply directly to instances of objects. For example, the authorization to create an object applies to the class of objects rather than to an individual instance. Also, some authorizations apply across the entire queue manager, for example, alternate-user authority means that a user can assume the authorities associated with another user.

Authorization directories

In MQSeries for UNIX systems, the default authorization directories, for a queue manager called saturn, are:

<code>/var/mqm/qmgrs/saturn/auth/queues</code>	Authorization files for queues.
<code>/var/mqm/qmgrs/saturn/auth/procdef</code>	Authorization files for process definitions.
<code>/var/mqm/qmgrs/saturn/auth/qmanager</code>	Authorization files for the queue manager.
<code>/var/mqm/qmgrs/saturn/auth/@aclass</code>	Authorizations applying to all classes.

For MQSeries for Windows NT, the default authorization directories, for a queue manager called saturn, are:

<code>\mqm\qmgrs\saturn\auth\queues</code>	Authorization files for queues.
<code>\mqm\qmgrs\saturn\auth\procdef</code>	Authorization files for process definitions.
<code>\mqm\qmgrs\saturn\auth\qmanager</code>	Authorization files for the queue manager.
<code>\mqm\qmgrs\saturn\auth\@aclass</code>	Authorizations applying to all classes.

In the object directories, the @class files hold the authorizations related to the entire class.

Note: There is a difference between @class (the authorization file that specifies authorization for a particular class) and @aclass (the directory that contains a file that specifies authorizations to all classes).

The paths of the object authorization files are based on those of the object itself, where auth is inserted ahead of the object type directory. You can use the **dspmqls** command to display the path to a specified object.

For example, if the name and path of SYSTEM.DEFAULT.LOCAL.QUEUE is:

```
/var/mqm/qmgrs/saturn/queues/SYSTEM!DEFAULT!LOCAL!QUEUE
```

the name and path of the corresponding authorization file is:

```
/var/mqm/qmgrs/saturn/auth/queues/SYSTEM!DEFAULT!LOCAL!QUEUE
```

If the name and path of SYSTEM.DEFAULT.LOCAL.QUEUE is:

```
\mqm\qmgrs\saturn\queues\SYSTEM!DEFAULT!LOCAL!QUEUE
```

the name and path of the corresponding authorization file is:

```
\mqm\qmgrs\saturn\auth\queues\SYSTEM!DEFAULT!LOCAL!QUEUE
```

Note: In this case, the actual names of the files associated with the queue are not the same as the name of the queue itself. See “Understanding MQSeries file names” on page 30 for details.

What the authorization files contain

The authorizations of a particular group are defined by a set of stanzas in the authorization file. See “Understanding authorization files” on page 94 for more information. The authorizations apply to the object associated with this file. For example:

```
groupB:
  Authority=0x0040007
```

This stanza defines the authority for the group groupB. The authority specification is the union of the individual bit patterns based on the following assignments:

Authorization keyword	Formal name	Hexadecimal Value
connect	MQZAO_CONNECT	0x00000001
browse	MQZAO_BROWSE	0x00000002
get	MQZAO_INPUT	0x00000004
put	MQZAO_OUTPUT	0x00000008
inq	MQZAO_INQUIRE	0x00000010
set	MQZAO_SET	0x00000020
passid	MQZAO_PASS_IDENTITY_CONTEXT	0x00000040
passall	MQZAO_PASS_ALL_CONTEXT	0x00000080
setid	MQZAO_SET_IDENTITY_CONTEXT	0x00000100
setall	MQZAO_SET_ALL_CONTEXT	0x00000200
altusr	MQZAO_ALTERNATE_USER_AUTHORITY	0x00000400
allmqi	MQZAO_ALL_MQI	0x000007FF
crt	MQZAO_CREATE	0x00010000
dlt	MQZAO_DELETE	0x00020000
dsp	MQZAO_DISPLAY	0x00040000
chg	MQZAO_CHANGE	0x00080000
clr	MQZAO_CLEAR	0x00100000
chgaut	MQZAO_AUTHORIZE	0x00800000
alladm	MQZAO_ALL_ADMIN	0x009E0000
none	MQZAO_NONE	0x00000000
all	MQZAO_ALL	0x009E07FF

These definitions are made in the header file cmqzc.h. In the following example, groupB has been granted authorizations based on the hexadecimal number 0x40007. This corresponds to:

```
MQZAO_CONNECT          0x00000001
MQZAO_BROWSE           0x00000002
MQZAO_INPUT            0x00000004
MQZAO_DISPLAY          0x00040000
-----
Authority is:          0x00040007
```

These access rights mean that anyone in groupB can issue the MQI calls:

```
MQCONN
MQGET (with browse)
```

They also have DISPLAY authority for the object associated with this authorization file.

Class authorization files

The *class authorization files* hold authorizations that relate to the entire class. These files are called “@class” and exist in the same directory as the files for specific objects. The entry MQZAO_CRT in the @class file gives authorization to create an object in the class. This is the only class authority.

All class authorization files

The *all class authorization file* holds authorizations that apply to an entire queue manager. This file is called “@aclass” and exists in the auth subdirectory of the queue manager.

The following authorizations apply to the entire queue manager and are held in the all-class authorization file:

- The entry MQZAO_ALTERNATE_USER_AUTHORITY gives authorization to assume the identity of another user when interacting with MQSeries objects.
- The entry MQZAO_SET_ALL_CONTEXT gives authorization to set the context of a message when issuing MQPUT.
- The entry MQZAO_SET_IDENTITY_CONTEXT gives authorization to set the identity context of a message when issuing MQPUT.

Managing authorization files

Here are some items that you need consider when managing your authorization files:

1. You must ensure that the authorization files are secure and not write-accessible by non-trusted general users. See “Authorizations to authorization files.”
2. To be able to reproduce your file authorizations, ensure that you do at least one of the following:
 - Back up the auth subdirectory after any significant updates
 - Retain shell scripts or command files containing the commands used
3. You can copy and edit authorization files. However, you should not normally have to create or repair them manually. Should an emergency occur, you can use the information given here to recover lost or damaged authorization files.

Authorizations to authorization files

In MQSeries for UNIX systems, authorization files must be readable by any principal. However, only the mqm user ID and the mqm group should be allowed to update these files.

The permissions on authorization files, created by the OAM, are:

```
-rw-rw-r--      mqm      mqm
```

Authorization files

Do not alter these permissions without reviewing carefully whether there are any security exposures.

To alter authorizations using the command supplied with MQSeries, your user ID must either be mqm, or it must belong to the mqm group.

For MQSeries for Windows NT, authorization files must be readable by any principal. However, only the mqm or Administrator's group should be allowed to update these files.

To alter authorizations using the **setmqaut** command supplied with MQSeries for Windows NT, your Windows NT user ID must belong to the local mqm group or the local Administrators group.

Chapter 7. Configuration files

MQSeries uses *configuration files* to hold basic product configuration information. This chapter describes what they are and how you can use them to change the way that queue managers operate.

What configuration files are

Configuration files define optional values for individual queue managers and for MQSeries on the node as a whole. These files have file name extensions of “ini”, and are also referred to as *ini files* or *stanza files*.

A configuration file contains one or more *stanzas*; a stanza is simply a group of lines in the file that together have a common function or define part of a system. For example, there are stanzas associated with logs, with channels, and with installable services.

Configuration files can be modified automatically by commands that change the configuration of queue managers on the node. You can also edit configuration files manually.

There are two types of configuration file:

- The *MQSeries configuration file*, which specifies values for MQSeries on the node as a whole. There is one MQSeries configuration file per node.
- *Queue manager configuration files*, which specify values for specific queue managers. There is one queue manager configuration file for each queue manager on the node.

MQSeries configuration file

The MQSeries configuration file, `mqs.ini`, contains information relevant to all the queue managers on a node. It is created automatically during installation. In particular, the MQSeries configuration file is used to locate the data associated with each queue manager. The MQSeries configuration file is located in the `mqm` directory. In MQSeries for UNIX systems, the default `mqm` directory is `/var/mqm`. In MQSeries for OS/2 Warp and Windows NT, the default `mqm` directory is `C:\MQM`.

What the MQSeries configuration file contains

The `mqs.ini` file contains the names of the queue managers, the name of the default queue manager, and the location of the files associated with each of them. The following stanzas can appear in `mqs.ini`:

AllQueueManagers

This stanza is used in MQSeries for OS/2 Warp and Windows NT only. It specifies the path to the `qmgrs` directory where the files associated with a queue manager are stored, and the path to the executable and DLL libraries.

ClientExitPath

Specifies the exit path for a client. to be `/var/mqm/exi ts`.

DefaultQueueManager

Specifies the default queue manager for the node. This queue manager processes any commands for which a queue manager name is not explicitly specified. The stanza is automatically updated if you create a new default queue manager. If you inadvertently create a default queue manager and then wish to revert to the original, you must alter this stanza manually.

QueueManager

There is one such stanza for each queue manager. This specifies the queue manager name and the location of the files associated with that queue manager. The names of these files are based on the queue manager name but are transformed if the queue manager name is not a valid file name. See “Understanding MQSeries file names” on page 30 for more information about name transformation.

LogDefaults

Specifies the default log parameters for the node. The `DefaultPrefix` and `LogDefaultPath` entries allow for the queue manager and its log to be on different physical drives. This is recommended, although by default they are on the same drive. See “Configuring the logs” on page 108 for more information about the log file stanzas.

Figure 7 on page 101 shows an example of an MQSeries configuration file in MQSeries for UNIX systems. Figure 8 on page 102 shows an example of an MQSeries configuration file for OS/2 Warp and Windows NT.

```

#####
#* Module Name: mqs.ini                                     *#
#* Type       : MQSeries Configuration File               *#
#* Function   : Define MQSeries resources for the node    *#
#*                                                   *#
#####
#* Notes      :                                           *#
#* 1) This is an example MQSeries configuration file      *#
#*                                                   *#
#####
AllQueueManagers:
#####
#* The path to the qmgrs directory, below which queue manager data *#
#* is stored                                               *#
#####
DefaultPrefix=/var/mqm

LogDefaults:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=1024
  LogType=CIRCULAR
  LogBufferPages=17
  LogDefaultPath=/var/mqm/log

QueueManager:
  Name=saturn.queue.manager
  Prefix=/var/mqm
  Directory=saturn!queue!manager

QueueManager:
  Name=pluto.queue.manager
  Prefix=/var/mqm
  Directory=pluto!queue!manager

DefaultQueueManager:
  Name=saturn.queue.manager

```

Figure 7. Example MQSeries configuration file for UNIX systems

MQSeries configuration file

```
#####  
#* Module Name: mqs.ini *#  
#* Type      : MQSeries Machine-wide Configuration File *#  
#* Function   : Define MQSeries resources for an entire machine *#  
#* *#  
#####  
#* Notes      : *#  
#* 1) This is the installation time default configuration *#  
#* *#  
#####  
AllQueueManagers:  
#####  
#* The path to the qmgrs directory, below which queue manager data *#  
#* is stored *#  
#####  
DefaultPrefix=c:\mqm  
  
DefaultFilePrefix=c:\mqm  
  
LogDefaults:  
  LogPrimaryFiles=3  
  LogSecondaryFiles=2  
  LogFilePages=256  
  LogType=CIRCULAR  
  LogBufferPages=17  
  LogDefaultPath=c:\mqm\log  
  
QueueManager:  
  Name=saturn.queue.manager  
  Prefix=c:\mqm  
  Directory=saturn!queue!manager  
  
QueueManager:  
  Name=venus  
  Prefix=c:\mqm  
  Directory=venus  
  
DefaultQueueManager:  
  Name=saturn.queue.manager
```

Figure 8. Example MQSeries configuration file for MQSeries for OS/2 Warp and Windows NT

In Figure 7 on page 101 and Figure 8, MQSeries on the node is using the default locations for queue managers and for the logs.

The queue manager saturn.queue.manager is the default queue manager for the node. The directory for files associated with this queue manager has been automatically transformed into a valid file name for the file system.

Note: Because the MQSeries configuration file is used to locate the data associated with queue managers, a nonexistent or incorrect configuration file can cause some or all MQSeries commands to fail. Also, applications cannot connect to a queue manager that is not defined in the MQSeries configuration file.

Queue manager configuration file

A queue manager configuration file, `qm.ini`, contains information relevant to a specific queue manager. There is one queue manager configuration file for each queue manager. It is created automatically when the queue manager with which it is associated is created.

The file is held in the root of the directory tree occupied by the queue manager. For example, in the MQSeries for UNIX systems, the path and name for a configuration file for a queue manager called QMNAME is `/var/mqm/qmgrs/QMNAME/qm.ini`. In MQSeries for OS/2 Warp and Windows NT, the path and name for a configuration file for a queue manager called QMNAME is `C:\MQM\QMGRS\QMNAME\QM.INI`.

Note: The queue manager name can be up to 48 characters in length. However, this does not guarantee that the name is valid or unique. Therefore, a directory name is generated based on the queue manager name. This process is known as name transformation; for a description, see “Understanding MQSeries file names” on page 30.

What the queue manager configuration file contains

The stanzas that can appear in a queue manager configuration file, `qm.ini`, are as follows:

Service

This stanza specifies the name of one of the installable services, and the number of entry points to that service. There is one stanza for each service. The following services are available:

Authorization service

In MQSeries for UNIX systems and MQSeries for Windows NT, the `AuthorizationService` stanza and its associated `ServiceComponent` stanza are added automatically when the queue manager is created. When the Object Authority Manager (OAM) is enabled, you can disable it by:

1. Deleting the queue manager (using the **dltmqm** command)
2. Creating the queue manager again (using the **crtmqm** command) with the `MQSNOAUT` variable

In MQSeries for OS/2 Warp, you must add the `AuthorizationService` stanza to `qm.ini` manually to enable the service. For the text of this stanza, see Figure 10 on page 106. To disable the service, you delete the stanza.

Name service

The `NameService` stanza must be added to `qm.ini` manually to enable the supplied name service; for the text of this stanza, see Figure 9 on page 105, Figure 10 on page 106, or Figure 11 on page 107 as appropriate.

User ID service (MQSeries for OS/2 Warp only)

The user identifier service stanza must be added to `qm.ini` manually to enable the service. For the text of this stanza, see Figure 10 on page 106.

ServiceComponent

This stanza identifies the service component associated with a particular service. There can be more than one service component stanza for each service, but each service component stanza must match the corresponding service stanza. See the *MQSeries Programmable System Management* manual for more information.

In MQSeries for Windows NT and UNIX systems, the authorization service stanza is present by default, and the associated component, the OAM, is active.

Log

This stanza specifies the default log parameters for this queue manager. The fields in this stanza are same as those in the LogDefaults stanza in the mqs.ini file. The values can be changed, if required. See “Configuring the logs” on page 108 for more information about the log file stanzas.

XAResourceManager

This stanza identifies resource managers to be involved in global units of work. See “Database coordination” on page 142 for more information about adding an XAResourceManager stanza to a qm.ini file.

Channels

This stanza contains information about the channels. It defines the maximum number of channels (MaxChannels) that can be defined for the queue manager, and also limits the number of channels that can be active at any time (MaxActiveChannels). This stanza also defines whether a channel is to run as trusted (MQIBindType=FASTPATH) or not (MQIBindType=STANDARD, which is the default).

See the *MQSeries Intercommunication* book for more information about channels.

LU 6.2, NETBIOS, TCP, and SPX

These stanzas specify network protocol configuration parameters. They override the default parameters for channels. Only stanzas representing changed default values are actually present.

KeepAlive, if specified, causes TCP/IP or SPX periodically to check that the other end of the connection is still available. If it is not, the channel is closed.

See the *MQSeries Intercommunication* book for more information.

Figure 9 on page 105 shows how the stanzas might be arranged in a queue manager configuration file in MQSeries for UNIX systems; Figure 10 on page 106 shows a similar example in MQSeries for OS/2 Warp; and Figure 11 on page 107 shows an example in MQSeries for Windows NT.

```

#####
** Module Name: qm.ini                                **
** Type      : MQSeries queue manager configuration file **
** Function  : Define the configuration of a single queue manager **
**                                     **
#####
** Notes      :                                       **
** 1) This file defines the configuration of the queue manager **
**                                     **
#####
ExitPath:
  ExitsDefaultPath=/var/mqm/exits

Service:
  Name=AuthorizationService
  EntryPoints=9

ServiceComponent:
  Service=AuthorizationService
  Name=MQSeries.UNIX.auth.service
  Module=mqmtop/bin/amqzfu.o
  ComponentDataSize=0

Service:
  Name=NameService
  EntryPoints=5

ServiceComponent:
  Service=NameService
  Name=MQSeries.DCE.name.service
  Module=mqmtop/lib/amqzfa
  ComponentDataSize=0

Log:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=1024
  LogType=CIRCULAR
  LogBufferPages=17
  LogPath=/var/mqm/log/saturn!queue!manager/

XAResourceManager:
  Name=DB2 Resource Manager Bank
  SwitchFile=/usr/bin/db2swit
  XAOpenString=MQBankDB
  XACloseString=
  ThreadOfControl=PROCESS

CHANNELS:
  MaxChannels = 20           ; Maximum number of Channels allowed.
                             ; Default is 100.
  MaxActiveChannels = 10    ; Maximum number of Channels allowed to be
                             ; active at any time. The default is the
                             ; value of MaxChannels.

TCP:
                             ; TCP/IP entries.
  Port = 1800                ; Use port 1800 instead of the default 1414
  KeepAlive = Yes            ; Switch KeepAlive on

```

Figure 9. Example queue manager configuration file for MQSeries for UNIX systems

Queue manager configuration file

```
#####  
** Module Name: qm.ini *#  
** Type : MQSeries queue manager configuration file *#  
# Function : Define the configuration of a single queue manager *#  
** *#  
#####  
** Notes : *#  
** 1) This file defines the configuration of the queue manager *#  
** *#  
#####  
ExitPath:  
  ExitsDefaultPath=c:\mqm\exits  
  
Service:  
  Name=NameService  
  EntryPoints=5  
  
ServiceComponent:  
  Service=NameService  
  Name=MQSeries.DCE.name.service  
  Module=c:\mqm\amqzfa  
  ComponentDataSize=0  
  
Log:  
  LogPrimaryFiles=3  
  LogSecondaryFiles=2  
  LogFilePages=1024  
  LogType=CIRCULAR  
  LogBufferPages=17  
  LogPath=c:\mqm\log\venus\  
  
XAResourceManager:  
  Name=DB2 Resource Manager Bank  
  SwitchFile=\usr\bin\db2swit  
  XAOpenString=MQBankDB  
  XACloseString=  
  ThreadOfControl=PROCESS  
  
Channels:  
  MaxChannels=10 ; Maximum number of Channels allowed.  
                  ; Default is 100.  
  MaxActiveChannels=5 ; Maximum number of channels allowed to be active  
                      ; at any time. The default is the value of  
                      ; MaxChannels.  
  
TCP: ; TCP/IP entries  
  Port = 1800 ; Use port 1800 instead of the default 1414  
  Library1=DLLName1 ; Name of TCP/IP Sockets DLL  
  Library2=DLLName2 ; Same as above if code is in two libraries  
  KeepAlive=Yes ; Switch KeepAlive on
```

Figure 10. Example queue manager configuration file for OS/2

```

#####
#* Module Name: qm.ini                                *#
#* Type       : MQSeries queue manager configuration file *#
# Function    : Define the configuration of a single queue manager *#
#*                                                   *#
#####
#* Notes      :                                       *#
#* 1) This file defines the configuration of the queue manager *#
#*                                                   *#
#####
ExitPath:
  ExitsDefaultPath=c:\mqm\exits

Service:
  Name=NameService
  EntryPoints=5

ServiceComponent:
  Service=NameService
  Name=MQSeries.DCE.name.service
  Module=c:\mqm\amqzfa
  ComponentDataSize=0

Log:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=1024
  LogType=CIRCULAR
  LogBufferPages=17
  LogPath=c:\mqm\log\venus\

XAResourceManager:
  Name=DB2 Resource Manager Bank
  SwitchFile=\usr\bin\db2swit
  XAOpenString=MQBankDB
  XACloseString=
  ThreadOfControl=PROCESS

Channels:
  MaxChannels=10      ; Maximum number of Channels allowed, the
                      ; default value is 100
  MaxActiveChannels=5 ; Maximum number of channels allowed to be active
                      ; at any time. The default is the value of
                      ; MaxChannels.

LU62:                ; LU 6.2 entries
  TPName= RECV        ; TP Name to start on remote site
  Library1=DLLName1   ; Name of APPC DLL
  Library2=DLLName2   ; Same as above if code is in two libraries
  LocalLU=MyLocalLU   ; LU to use on local systems

```

Figure 11. Example queue manager configuration file for Windows NT

Editing configuration files

You can edit the default configuration files to alter the system defaults. However, before editing any configuration file, make sure that you have a backup that you can revert to.

In some circumstances, you may have to edit your configuration files. For example:

- If you lose a configuration file; recover from backup if possible.
- If you need to move one or more queue managers to a new directory.
- If you need to change your default queue manager; this could happen if you accidentally delete the existing queue manager.
- When advised to do so by your IBM Support Center.

Changing the default prefix

If you change the default prefix, `DefaultPrefix`, for the message queue manager, you must replicate the directory structure that was created at installation time (see Figure 48 on page 303). In particular, the `qmgrs` structure must be created. You must stop MQSeries before changing the default prefix, and restart MQSeries only after the structures have been moved to the new location and the default prefix has been changed.

As an alternative to changing the `DefaultPrefix`, you can use the environment variable `MQSPREFIX` to override the `DefaultPrefix` in the `mqs.ini` file for the `crtmqm` command.

Implementing changes to configuration files

Do not edit configuration files while the queue manager is running.

Recommendations for configuration files

When you create a new queue manager, you should:

- Back up the MQSeries configuration file
- Back up the new queue manager configuration file

Configuring the logs

The log parameters in the MQSeries configuration file are used as default values when you create a queue manager. These defaults can be overridden if you specify the log parameters on the `crtmqm` command. See “`crtmqm` (Create queue manager)” on page 240 for details of this command.

The values specified in the queue manager configuration file are read when the queue manager is started. The file is created when the queue manager is created.

The values in a configuration file are set according to these priorities:

1. Parameters entered on the command line override the MQSeries configuration file.
2. The MQSeries configuration file contains the supplied default values.

Note: In MQSeries for UNIX systems, user ID mqm and group mqm must have full authorities to the log files. If you change the locations of these files, you must give these authorities yourself. This is not required if the logs files are in the default locations supplied with the product.

If you use an invalid value in a configuration file, it is ignored. The effect is the same as missing out the value entirely. An operator message is issued to indicate the problem.

You can edit the MQSeries configuration file after installation and change the default values to your own requirements.

Log configuration stanzas

The size and location of the log is configured by stanzas in the MQSeries and queue manager configuration files. These stanzas specify the type of logging to be used, the log file size, and the log path.

The MQSeries configuration file contains a stanza called `LogDefaults`. In MQSeries for UNIX systems, `LogDefaults` is as follows:

```
LogDefaults:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=1024
  LogType=CIRCULAR
  LogBufferPages=17
  LogDefaultPath=/var/mqm/log
```

In MQSeries for OS/2 Warp and Windows NT, `LogDefaults` is as follows:

```
LogDefaults:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=256
  LogType=CIRCULAR
  LogBufferPages=17
  LogDefaultPath=c:\mqm\log
```

The values specified in the MQSeries configuration file are read whenever a queue manager is created, started, or deleted.

Each queue manager configuration file has a stanza called `Log`. In MQSeries for UNIX systems, `Log` is as follows:

```
Log:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=1024
  LogType=CIRCULAR
  LogBufferPages=17
  LogPath=/var/mqm/log/<QM_Dir_Name>/
```

In MQSeries for OS/2 Warp and Windows NT, Log is as follows:

```
Log:
  LogPrimaryFiles=3
  LogSecondaryFiles=2
  LogFilePages=256
  LogType=CIRCULAR
  LogBufferPages=17
  LogPath=c:\mqm\log\<QM_Dir_Name>/
```

<QM_Dir_Name> is the subdirectory name for this queue manager, providing a unique path to the logs. This is the queue manager name if it is valid for the file system; otherwise, it is a transformed name. (See “Understanding MQSeries file names” on page 30 for more information about name transformation.)

LogPrimaryFiles

Primary log files are the log files allocated during creation for future use.

The default number is 3. The default can be overridden by editing the LogPrimaryFiles value in mqs.ini and qm.ini.

The value is examined when the queue manager is created or started. You can change it after the queue manager has been created. However, a change in the value is not effective until the queue manager is restarted, and the effect may not be immediate.

The minimum number of primary log files is 2 and the maximum is 62. The total number of primary and secondary log files must not exceed 63, and must not be less than 3.

LogSecondaryFiles

Secondary log files are the log files allocated when the primary files are exhausted.

The default number is 2. The default can be overridden using the LogSecondaryFiles value in mqs.ini and qm.ini.

The value is examined when the queue manager is created or started. You can change this value, but changes are not effective until the queue manager is restarted, and the effect may not be immediate.

The minimum number of secondary log files is 1 and the maximum is 61. The total number of primary and secondary log files must not exceed 63, and must not be less than 3.

LogFilePages

The log data is held in a series of files called log files. The log file size is specified in units of 4 KB pages.

In MQSeries for UNIX systems, the default number of log file pages is 1024, giving a log file size of 4 MB. The minimum number of log file pages is 64 and the maximum is 384.

In MQSeries for OS/2 Warp and Windows NT, the default number of log file pages is 256, giving a log file size of 1 MB. The minimum number of log file pages is 32 and the maximum is 4095.

The log file size can be specified only during queue manager creation. The value used is obtained by taking the default (1024 or 256) and overriding it with the value in the `LogFilePages` attribute in the MQSeries configuration file, or by overriding with the value specified on the `crtmqm` command using the `-lf` flag.

Note: The size of the log files is specified during queue manager creation and cannot be changed for an existing queue manager.

LogType

The `LogType` parameter is used to define the type to be used, either CIRCULAR or LINEAR. The default is CIRCULAR.

If you want to change the default, you can either edit the MQSeries configuration file or specify linear logging with the `crtmqm` command. You cannot change the logging method after a queue manager has been created.

LogBufferPages

The amount of memory allocated to buffer records for writing is configurable. The size of the buffers is specified in units of 4 KB pages.

The default number of buffer pages is 17, equating to 68 KB.

The default can be overridden using the `LogBufferPages` value in `mqs.ini` and `qm.ini`.

The value is examined when the queue manager is created or started, and may be increased or decreased at either of these times. However, a change in the value is not effective until the queue manager is restarted.

The minimum number of buffer pages is 4 and the maximum is 32. Larger buffers lead to higher throughput, especially for larger messages.

LogPath

You can specify the directory in which the log files for a queue manager reside. The directory should exist on a local device to which the queue manager can write and, preferably, should be on a different drive from the message queues. Specifying a different drive gives added protection in case of system failure.

The default is `/var/mqm/log` in MQSeries for UNIX systems. The default is `c:\mqm\log` in MQSeries for OS/2 Warp and MQSeries for Windows NT.

You can specify the name of a directory on the `crtmqm` command using the `-ld` flag. When a queue manager is created, a directory is also created under the queue manager directory, and this is used to hold the log files. The name of this directory is based on the queue manager name. This ensures that the Log File Path is unique, and also that it conforms to any limitations on directory name lengths.

If you do not specify `-ld` on the `crtmqm` command, the value of the `LogDefaultPath` attribute in the MQSeries configuration file `mqs.ini` is used. If this attribute is missing, the directory specified on `LogPath` is used. The queue manager name is appended to the directory name to ensure that multiple queue managers use different log directories.

When the queue manager has been created, a LogPath value is created in the log stanza in the queue manager configuration file giving the complete directory name for the queue manager's log. This value is used to locate the log when the queue manager is started or deleted.

Specifying log file sizes

The size of the log file that you require depends on the number and size of messages that are to be handled by your system. Each operation adds an overhead to the size of the log. For example, when a persistent message is put to a queue, the message data must be written to the log to make recovery of the message possible. The message descriptor is also logged, together with some internal information that describes the effect of putting the message on the queue.

There is a trade-off between the size of your log files and the number of files that you have. Larger files are more difficult to handle but are more efficient.

Table 10 shows approximate values for the header information required for various types of operation.

Operation	Size
Put persistent message	750 bytes + message length If the message is large, it is divided into segments of 15700 bytes, each with a 300-byte overhead.
Get message	260 bytes
Syncpoint, commit	750 bytes
Syncpoint, roll-back	1000 bytes + 12 bytes for each get or put to be rolled back
Create object	1500 bytes
Delete object	300 bytes
Alter attributes	1024 bytes
Record media image	800 bytes + image The image is divided into segments of 15700 bytes, each having a 300-byte overhead.
Checkpoint	750 bytes + 200 bytes for each active unit of work. Additional data may be logged for any uncommitted puts or gets that have been buffered for performance reasons.

Configuring XA_-compliant databases

One XAResourceManager stanza is required in qm.ini for each instance of a resource manager participating in global units of work; no default values for this stanza are supplied via mqs.ini. See "Database coordination" on page 142 for more information about adding an XAResourceManager stanza to qm.ini.

Name Identifies the resource manager instance.

The Name value, which is required, can be up to 31 characters in length and must be unique within qm.ini. You can use the name of

the resource manager as defined in its XA switch structure. However, if you are using more than one instance of the same resource manager, you must construct a unique name for each instance. You could ensure uniqueness by including the name of the database in the Name string, for example.

MQSeries uses the Name value in messages and in output from the **dspmqrn** command.

You are recommended not to change the name of a resource manager instance, nor to delete its entry from `qm.ini`, once the associated queue manager has started and the resource-manager name is in effect.

SwitchFile Is the fully-qualified name of the load file containing the resource manager's XA switch structure. The `SwitchFile` value is required.

XAOpenString

Is the string of data to be passed to the resource manager's `xa_open` entry point. The contents of the string depend on the resource manager itself. For example, the string could identify the database that this instance of the resource manager is to access. Consult your resource manager documentation for the appropriate string.

`XAOpenString` is optional.

XACloseString

Is the string of data to be passed to the resource manager's `xa_close` entry point. The contents of the string depend on the resource manager itself. Consult your database documentation for the appropriate string. `XACloseString` is optional.

ThreadOfControl

Is used by the queue manager for serialization purposes when it needs to call the resource manager from one of its own multithreaded processes. The `ThreadOfControl` value can be `THREAD` or `PROCESS`.

`ThreadOfControl=THREAD` means that the resource manager is fully "thread aware." In a multithreaded MQSeries process, XA function calls can be made to the external resource manager from multiple threads at the same time.

`ThreadOfControl=PROCESS` means that the resource manager is not thread safe. In a multithreaded MQSeries process, only one XA function call at a time can be made to the resource manager.

The `ThreadOfControl` entry does not apply to XA function calls issued by the queue manager in a multithreaded application process. In general, an application that has concurrent units of work on different threads requires this mode of operation to be supported by each of the resource managers.

`ThreadOfControl` is mandatory in MQSeries for OS/2 Warp and Windows NT.

Chapter 8. The MQSeries dead-letter queue handler

A *dead-letter queue* (DLQ), sometimes referred to as an *undelivered-message queue*, is a holding queue for messages that cannot be delivered to their destination queues. Every queue manager in a network should have an associated DLQ.¹

Messages can be put on the DLQ by queue managers, by message channel agents (MCAs), and by applications. All messages on the DLQ should be prefixed with a *dead-letter header* structure, MQDLH. Messages put on the DLQ by a queue manager or by a message channel agent always have an MQDLH; applications putting messages on the DLQ are strongly recommended to supply an MQDLH. The *Reason* field of the MQDLH structure contains a reason code that identifies why the message is on the DLQ.

In all MQSeries environments, there should be a routine that runs regularly to process messages on the DLQ. MQSeries supplies a default routine, called the *dead-letter queue handler* (the DLQ handler), which you invoke using the **runmqdlq** command. Instructions for processing messages on the DLQ are supplied to the DLQ handler by means of a user-written *rules table*. That is, the DLQ handler matches messages on the DLQ against entries in the rules table: when a DLQ message matches an entry in the rules table, the DLQ handler performs the action associated with that entry.

This chapter contains the following sections:

- “Invoking the DLQ handler”
- “The DLQ handler rules table” on page 116
- “How the rules table is processed” on page 123
- “An example DLQ handler rules table” on page 125

Invoking the DLQ handler

You invoke the DLQ handler using the **runmqdlq** command. You can name the DLQ you want to process and the queue manager you want to use in two ways:

- As parameters to **runmqdlq** from the command prompt. For example:

```
runmqdlq ABC1.DEAD.LETTER.QUEUE ABC1.QUEUE.MANAGER <rule.ru1
```

- In the rules table. For example:

```
INPUTQ(ABC1.DEAD.LETTER.QUEUE) INPUTQM(ABC1.QUEUE.MANAGER)
```

The above examples apply to the DLQ called ABC1.DEAD.LETTER.QUEUE, owned by the queue manager ABC1.QUEUE.MANAGER.

¹ It is often preferable to avoid placing messages on a DLQ. For information about the use and avoidance of DLQs, see the *MQSeries Application Programming Guide*.

If you do not specify the DLQ or the queue manager as shown above, the default queue manager for the installation is used along with the DLQ belonging to that queue manager.

The **runmqdlq** command takes its input from `stdin`; you associate the rules table with **runmqdlq** by redirecting `stdin` from the rules table.

In order to run the DLQ handler, you must be authorized to access both the DLQ itself and any message queues to which messages on the DLQ are forwarded. Furthermore, if the DLQ handler is to be able to put messages on queues with the authority of the user ID in the message context, you must be authorized to assume the identity of other users.

For more information about the **runmqdlq** command, see “runmqdlq (Run dead-letter queue handler)” on page 273.

The sample DLQ handler, amqsdq

In addition to the DLQ handler invoked using the **runmqdlq** command, MQSeries provides the source of a sample DLQ handler, `amqsdq`, whose function is similar to that provided via **runmqdlq**. You can customize `amqsdq` to provide a DLQ handler that meets specific, local requirements. For example, you might decide that you want a DLQ handler that can process messages without dead-letter headers. (Both the default DLQ handler and the sample, `amqsdq`, process only those messages on the DLQ that begin with a dead-letter header, `MQDLH`. Messages that do not begin with an `MQDLH` are identified as being in error, and remain on the DLQ indefinitely.)

In MQSeries for UNIX systems, the source of `amqsdq` is supplied in the directory:

```
mqmtop/samp/dlq
```

and the compiled version is supplied in the directory:

```
mqmtop/samp/bin
```

In MQSeries for OS/2 Warp and Windows NT, the source of `amqsdq` is supplied in the directory:

```
c:\mqm\tools\c\samples\dlq
```

and the compiled version is supplied in the directory:

```
c:\mqm\tools\c\samples\bin
```

The DLQ handler rules table

The DLQ handler rules table defines how the DLQ handler is to process messages that arrive on the DLQ. There are two types of entry in a rules table:

- The first entry in the table, which is optional, contains *control data*.
- All other entries in the table are *rules* for the DLQ handler to follow. Each rule consists of a *pattern* (a set of message characteristics) that a message is matched against, and an *action* to be taken when a message on the DLQ matches the specified pattern. There must be at least one rule in a rules table.

Each entry in the rules table comprises one or more keywords.

Control data

This section describes the keywords that you can include in a control-data entry in a DLQ handler rules table. Please note the following:

- The default value for a keyword, if any, is underlined>.
- The vertical line (|) separates alternatives, only one of which can be specified.
- All keywords are optional.

INPUTQ (*QueueName*|' ')

Allows you to name the DLQ you want to process:

1. If you specify an INPUTQ value as a parameter to the **runmqdlq** command, this overrides any INPUTQ value in the rules table.
2. If you do not specify an INPUTQ value as a parameter to the **runmqdlq** command but you *do* specify a value in the rules table, the INPUTQ value in the rules table is used.
3. If no DLQ is specified or you specify INPUTQ(' ') in the rules table, the name of the DLQ belonging to the queue manager whose name is supplied as a parameter to the **runmqdlq** command is used.
4. If you do not specify an INPUTQ value as a parameter to the **runmqdlq** command or as a value in the rules table, the DLQ belonging to the queue manager named on the INPUTQM keyword in the rules table is used.

INPUTQM (*QueueManagerName*|' ')

Allows you to name the queue manager that owns the DLQ named on the INPUTQ keyword:

1. If you specify an INPUTQM value as a parameter to the **runmqdlq** command, this overrides any INPUTQM value in the rules table.
2. If you do not specify an INPUTQM value as a parameter to the **runmqdlq** command, the INPUTQM value in the rules table is used.
3. If no queue manager is specified or you specify INPUTQM(' ') in the rules table, the default queue manager for the installation is used.

RETRYINT (*Interval*|60)

Is the interval, in seconds, at which the DLQ handler should attempt to reprocess messages on the DLQ that could not be processed at the first attempt, and for which repeated attempts have been requested. By default, the retry interval is 60 seconds.

WAIT (YES|NO|*nnn*)

Indicates whether the DLQ handler should wait for further messages to arrive on the DLQ when it detects that there are no further messages that it can process.

YES Causes the DLQ handler to wait indefinitely.

NO Causes the DLQ handler to terminate when it detects that the DLQ is either empty or contains no messages that it can process.

nnn Causes the DLQ handler to wait for *nnn* seconds for new work to arrive before terminating, after it detects that the queue is either empty or contains no messages that it can process.

You are recommended to specify WAIT (YES) for busy DLQs, and WAIT (NO) or WAIT (*nnn*) for DLQs that have a low level of activity. If the DLQ handler is allowed to terminate, you are recommended to reinvoke it by means of triggering.

As an alternative to including control data in the rules table, you can supply the names of the DLQ and its queue manager as input parameters of the **runmqdlq** command. If any value is specified both in the rules table and on input to the **runmqdlq** command, the value specified on the **runmqdlq** command takes precedence.

Note: If a control-data entry is included in the rules table, it **must** be the first entry in the table.

Rules (patterns and actions)

Figure 12 shows an example rule from a DLQ handler rules table.

```
PERSIST(MQPER_PERSISTENT) REASON (MQRC_PUT_INHIBITED) +  
ACTION (RETRY) RETRY (3)
```

Figure 12. An example rule from a DLQ handler rules table. This rule instructs the DLQ handler to make 3 attempts to deliver to its destination queue any persistent message that was put on the DLQ because MQPUT and MQPUT1 were inhibited.

All keywords that you can use on a rule are described in the remainder of this section. Please note the following:

- The default value for a keyword, if any, is underlined. For most keywords, the default value is * (asterisk), which matches any value.
- The vertical line (|) separates alternatives, only one of which can be specified.
- All keywords except ACTION are optional.

This section begins with a description of the pattern-matching keywords (those against which messages on the DLQ are matched), and then describes the action keywords (those that determine how the DLQ handler is to process a matching message).

The pattern-matching keywords

The pattern-matching keywords, which you use to specify values against which messages on the DLQ are matched, are described below. All pattern-matching keywords are optional.

APPLIDAT (*ApplIdentityData**)

Is the *ApplIdentityData* value specified in the message descriptor, MQMD, of the message on the DLQ.

APPLNAME (*PutApplName**)

Is the name of the application that issued the MQPUT or MQPUT1 call, as specified in the *PutApplName* field of the message descriptor, MQMD, of the message on the DLQ.

APPLTYPE (*PutApplType**)

Is the *PutApplType* value specified in the message descriptor, MQMD, of the message on the DLQ.

DESTQ (*QueueName**)

Is the name of the message queue for which the message is destined.

DESTQM (*QueueManagerName**)

Is the name of the queue manager of the message queue for which the message is destined.

FEEDBACK (*Feedback**)

When the *MsgType* value is MQFB_REPORT, *Feedback* describes the nature of the report.

Symbolic names can be used. For example, you can use the symbolic name MQFB_COA to identify those messages on the DLQ that require confirmation of their arrival on their destination queues.

FORMAT (*Format**)

Is the name that the sender of the message uses to describe the format of the message data.

MSGTYPE (*MsgType**)

Is the message type of the message on the DLQ.

Symbolic names can be used. For example, you can use the symbolic name MQMT_REQUEST to identify those messages on the DLQ that require replies.

PERSIST (*Persistence**)

Is the persistence value of the message. (The persistence of a message determines whether it survives restarts of the queue manager.)

Symbolic names can be used. For example, you can use the symbolic name MQPER_PERSISTENT to identify those messages on the DLQ that are persistent.

REASON (*ReasonCode**)

Is the reason code that describes why the message was put to the DLQ.

Symbolic names can be used. For example, you can use the symbolic name MQRC_Q_FULL to identify those messages placed on the DLQ because their destination queues were full.

REPLYQ (*QueueName**)

Is the name of the reply-to queue specified in the message descriptor, MQMD, of the message on the DLQ.

REPLYQM (*QueueManagerName**)

Is the name of the queue manager of the reply-to queue, as specified in the message descriptor, MQMD, of the message on the DLQ.

USERID (*UserIdentifier**)

Is the user ID of the user who originated the message on the DLQ, as specified in the message descriptor, MQMD.

The action keywords

The action keywords, which you use to describe how a matching message is to be processed, are described below.

ACTION (DISCARD|IGNORE|RETRY|FWD)

Is the action to be taken for any message on the DLQ that matches the pattern defined in this rule.

- DISCARD** Causes the message to be deleted from the DLQ.
- IGNORE** Causes the message to be left on the DLQ.
- RETRY** Causes the DLQ handler to try again to put the message on its destination queue.
- FWD** Causes the DLQ handler to forward the message to the queue named on the FWDQ keyword.

The ACTION keyword must be specified. The number of attempts made to implement an action is governed by the RETRY keyword. The interval between attempts is controlled by the RETRYINT keyword of the control data.

FWDQ (*QueueName*&DESTQ|&REPLYQ)

Is the name of the message queue to which the message should be forwarded when ACTION (FWD) is requested.

QueueName

Is the name of a message queue. FWDQ(' ') is not valid.

- &DESTQ** Causes the queue name to be taken from the *DestQName* field in the MQDLH structure.
- &REPLYQ** Causes the name to be taken from the *ReplyToQ* field in the message descriptor, MQMD.

To avoid error messages when a rule specifying FWDQ (&REPLYQ) matches a message with a blank *ReplyToQ* field, you can specify REPLYQ (?*) in the message pattern.

FWDQM (*QueueManagerName*&DESTQM|&REPLYQM|'_')

Identifies the queue manager of the queue to which a message is to be forwarded.

QueueManagerName

Is the name of the queue manager of the queue to which a message is to be forwarded when ACTION (FWD) is requested.

&DESTQM

Causes the queue manager name to be taken from the *DestQMGrName* field in the MQDLH structure.

&REPLYQM

Causes the name to be taken from the *ReplyToQMGr* field in the message descriptor, MQMD.

' ' FWDQM(' '), which is the default value, identifies the local queue manager.

HEADER (YES|NO)

Specifies whether the MQDLH should remain on a message for which ACTION (FWD) is requested. By default, the MQDLH remains on the message. The HEADER keyword is not valid for actions other than FWD.

PUTAUT (DEF|CTX)

Defines the authority with which messages should be put by the DLQ handler:

DEF Causes messages to be put with the authority of the DLQ handler itself.

CTX Causes the messages to be put with the authority of the user ID in the message context. If you specify PUTAUT (CTX), you must be authorized to assume the identity of other users.

RETRY (*RetryCount*{1})

Is the number of times, in the range 1–999 999 999, that an action should be attempted (at the interval specified on the RETRYINT keyword of the control data).

Note: The count of attempts made by the DLQ handler to implement any particular rule is specific to the current instance of the DLQ handler; the count does not persist across restarts. If the DLQ handler is restarted, the count of attempts made to apply a rule is reset to zero.

Rules table conventions

The rules table must adhere to the following conventions regarding its syntax, structure, and contents:

- A rules table must contain at least one rule.
- Keywords can occur in any order.
- A keyword can be included once only in any rule.
- Keywords are not case sensitive.
- A keyword and its parameter value must be separated from other keywords by at least one blank or comma.
- Any number of blanks can occur at the beginning or end of a rule, and between keywords, punctuation, and values.
- Each rule must begin on a new line.
- For reasons of portability, the significant length of a line should not be greater than 72 characters.

- Use the plus sign (+) as the last nonblank character on a line to indicate that the rule continues from the first nonblank character in the next line. Use the minus sign (-) as the last nonblank character on a line to indicate that the rule continues from the start of the next line. Continuation characters can occur within keywords and parameters.

For example:

```
APPLNAME('ABC+  
D')
```

results in 'ABCD', and

```
APPLNAME('ABC-  
D')
```

results in 'ABC D'.

- Comment lines, which begin with an asterisk (*), can occur anywhere in the rules table.
- Blank lines are ignored.
- Each entry in the DLQ handler rules table comprises one or more keywords and their associated parameters. The parameters must follow these syntax rules:

- Each parameter value must include at least one significant character. The delimiting quotation marks in quoted values are not considered significant. For example, these parameters are valid:

```
FORMAT('ABC')    3 significant characters  
FORMAT(ABC)      3 significant characters  
FORMAT('A')      1 significant character  
FORMAT(A)        1 significant character  
FORMAT(' ')      1 significant character
```

These parameters are invalid because they contain no significant characters:

```
FORMAT(' ')  
FORMAT( )  
FORMAT()  
FORMAT
```

- Wildcard characters are supported: you can use the question mark (?) in place of any single character, except a trailing blank; you can use the asterisk (*) in place of zero or more adjacent characters. The asterisk (*) and the question mark (?) are **always** interpreted as wildcard characters in parameter values.
- Wildcard characters cannot be included in the parameters of these keywords: ACTION, HEADER, RETRY, FWDQ, FWDQM, and PUTAUT.
- Trailing blanks in parameter values, and in the corresponding fields in the message on the DLQ, are not significant when performing wildcard matches. However, leading and embedded blanks within strings in quotation marks are significant to wildcard matches.

- Numeric parameters cannot include the question mark (?) wildcard character. The asterisk (*) can be used in place of an entire numeric parameter, but cannot be included as part of a numeric parameter. For example, these are valid numeric parameters:

MSGTYPE(2)	Only reply messages are eligible
MSGTYPE(*)	Any message type is eligible
MSGTYPE('*')	Any message type is eligible

However, MSGTYPE('2*') is not valid, because it includes an asterisk (*) as part of a numeric parameter.
- Numeric parameters must be in the range 0–999 999 999. If the parameter value is in this range, it is accepted, even if it is not currently valid in the field to which the keyword relates. Symbolic names can be used for numeric parameters.
- If a string value is shorter than the field in the MQDLH or MQMD to which the keyword relates, the value is padded with blanks to the length of the field. If the value, excluding asterisks, is longer than the field, an error is diagnosed. For example, these are all valid string values for an 8-character field:

'ABCDEFGH'	8 characters
'A*C*E*G*I'	5 characters excluding asterisks
'*A*C*E*G*I*K*M*O*'	8 characters excluding asterisks
- Strings that contain blanks, lowercase characters, or special characters other than period (.), forward slash (/), underscore (_), and percent sign (%) must be enclosed in single quotation marks. Lowercase characters not enclosed in quotation marks are folded to uppercase. If the string includes a quotation, two single quotation marks must be used to denote both the beginning and the end of the quotation. When the length of the string is calculated, each occurrence of double quotation marks is counted as a single character.

How the rules table is processed

The DLQ handler searches the rules table for a rule whose pattern matches a message on the DLQ. The search begins with the first rule in the table, and continues sequentially through the table. When a rule with a matching pattern is found, the action from that rule is attempted. The DLQ handler increments the retry count for a rule by 1 whenever it attempts to apply that rule. If the first attempt fails, the attempt is repeated until the count of attempts made matches the number specified on the RETRY keyword. If all attempts fail, the DLQ handler searches for the next matching rule in the table.

This process is repeated for subsequent matching rules until an action is successful. When each matching rule has been attempted the number of times specified on its RETRY keyword, and all attempts have failed, ACTION (IGNORE) is assumed. ACTION (IGNORE) is also assumed if no matching rule is found.

Notes:

1. Matching rule patterns are sought only for messages on the DLQ that begin with an MQDLH. Messages that do not begin with an MQDLH are reported periodically as being in error, and remain on the DLQ indefinitely.
2. All pattern keywords can be allowed to default, such that a rule may consist of an action only. Note, however, that action-only rules are applied to all messages on the queue that have MQDLHs and that have not already been processed in accordance with other rules in the table.
3. The rules table is validated when the DLQ handler is started, and errors are flagged at that time. (Error messages issued by the DLQ handler are described in the *MQSeries Messages* book.) You can make changes to the rules table at any time, but those changes do not come into effect until the DLQ handler is restarted.
4. The DLQ handler does not alter the content of messages, of the MQDLH, or of the message descriptor. The DLQ handler always puts messages to other queues with the message option MQPMO_PASS_ALL_CONTEXT.
5. Consecutive syntax errors in the rules table may not be recognized because the implementation of the validation of the rules table is designed to eliminate the generation of repetitive errors.
6. The DLQ handler opens the DLQ with the MQOO_INPUT_AS_Q_DEF option.
7. Multiple instances of the DLQ handler could run concurrently against the same queue, using the same rules table. However, it is more usual for there to be a one-to-one relationship between a DLQ and a DLQ handler.

Ensuring that all DLQ messages are processed

The DLQ handler keeps a record of all messages on the DLQ that have been seen but not removed. If you use the DLQ handler as a filter to extract a small subset of the messages from the DLQ, the DLQ handler still has to keep a record of those messages on the DLQ that it did not process. Also, the DLQ handler cannot guarantee that new messages arriving on the DLQ will be seen, even if the DLQ is defined as first-in-first-out (FIFO). Therefore, if the queue is not empty, a periodic rescan of the DLQ is performed to check all messages. For these reasons, you should try to ensure that the DLQ contains as few messages as possible; if messages that cannot be discarded or forwarded to other queues (for whatever reason) are allowed to accumulate on the queue, the workload of the DLQ handler increases and the DLQ itself is in danger of filling up.

You can take specific measures to enable the DLQ handler to empty the DLQ. For example, try not to use ACTION (IGNORE), which simply leaves messages on the DLQ. (Remember that ACTION (IGNORE) is assumed for messages that are not explicitly addressed by other rules in the table.) Instead, for those messages that you would otherwise ignore, use an action that moves the messages to another queue. For example:

```
ACTION (FWD) FWDQ (IGNORED.DEAD.QUEUE) HEADER (YES)
```

Similarly, the final rule in the table should be a catchall to process messages that have not been addressed by earlier rules in the table. For example, the final rule in the table could be something like this:

```
ACTION (FWD) FWDQ (REALLY.DEAD.QUEUE) HEADER (YES)
```

This action causes messages that fall through to the final rule in the table to be forwarded to the queue REALLY.DEAD.QUEUE, where they can be processed manually. If you do not have such a rule, messages are likely to remain on the DLQ indefinitely.

An example DLQ handler rules table

Here is an example rules table that contains a single control-data entry and several rules:

```
*****
*           An example rules table for the runmqdlq command           *
*****
* Control data entry
* -----
* If no queue manager name is supplied as an explicit parameter to
* runmqdlq, use the default queue manager for the machine.
* If no queue name is supplied as an explicit parameter to runmqdlq,
* use the DLQ defined for the local queue manager.
*
inputqm(' ') inputq(' ')

* Rules
* ----
* We include rules with ACTION (RETRY) first to try to
* deliver the message to the intended destination.

* If a message is placed on the DLQ because its destination
* queue is full, attempt to forward the message to its
* destination queue. Make 5 attempts at approximately
* 60-second intervals (the default value for RETRYINT).

REASON(MQRC_Q_FULL) ACTION(RETRY) RETRY(5)

* If a message is placed on the DLQ because of a put inhibited
* condition, attempt to forward the message to its
* destination queue. Make 5 attempts at approximately
* 60-second intervals (the default value for RETRYINT).

REASON(MQRC_PUT_INHIBITED) ACTION(RETRY) RETRY(5)

* The AAAA corporation are always sending messages with incorrect
* addresses. When we find a request from the AAAA corporation,
* we return it to the DLQ (DEADQ) of the reply-to queue manager
* (&REPLYQM).
* The AAAA DLQ handler attempts to redirect the message.

MSGTYPE(MQMT_REQUEST) REPLYQM(AAAA.*) +
ACTION(FWD) FWDQ(DEADQ) FWDQM(&REPLYQM)
```

DLQ handler

- * The BBBB corporation never do things by half measures. If
- * the queue manager BBBB.1 is unavailable, try to
- * send the message to BBBB.2

```
DESTQM(bbbb.1) +  
  action(fwd) fwdq(&DESTQ) fwdqm(bbbb.2) header(no)
```

- * The CCCC corporation considers itself very security
- * conscious, and believes that none of its messages
- * will ever end up on one of our DLQs.
- * Whenever we see a message from a CCCC queue manager on our
- * DLQ, we send it to a special destination in the CCCC organization
- * where the problem is investigated.

```
REPLYQM(CCCC.*) +  
  ACTION(FWD) FWDQ(ALARM) FWDQM(CCCC.SYSTEM)
```

- * Messages that are not persistent run the risk of being
- * lost when a queue manager terminates. If an application
- * is sending nonpersistent messages, it should be able
- * to cope with the message being lost, so we can afford to
- * discard the message.

```
PERSIST(MQPER_NOT_PERSISTENT) ACTION(DISCARD)
```

- * For performance and efficiency reasons, we like to keep
- * the number of messages on the DLQ small.
- * If we receive a message that has not been processed by
- * an earlier rule in the table, we assume that it
- * requires manual intervention to resolve the problem.
- * Some problems are best solved at the node where the
- * problem was detected, and others are best solved where
- * the message originated. We don't have the message origin,
- * but we can use the REPLYQM to identify a node that has
- * some interest in this message.
- * Attempt to put the message onto a manual intervention
- * queue at the appropriate node. If this fails,
- * put the message on the manual intervention queue at
- * this node.

```
REPLYQM('?*') +  
  ACTION(FWD) FWDQ(DEADQ.MANUAL.INTERVENTION) FWDQM(&REPLYQM)
```

```
ACTION(FWD) FWDQ(DEADQ.MANUAL.INTERVENTION)
```

Chapter 9. Instrumentation events

You can use MQSeries instrumentation events to monitor the operation of queue managers. This chapter provides a short introduction to instrumentation events. For a more complete description, see the *MQSeries Programmable System Management* book.

What instrumentation events are

Instrumentation events cause special messages, called *event messages*, to be generated whenever the queue manager detects a predefined set of conditions. For example, the following conditions give rise to a *Queue Full* event:

- Queue Full events are enabled for a specified queue, and
- An application issues an MQPUT call to put a message on that queue, but the call fails because the queue is full.

Other conditions that can give rise to instrumentation events include:

- A predefined limit for the number of messages on a queue being reached
- A queue not being serviced within a specified time
- A channel instance being started or stopped
- In MQSeries for UNIX systems, an application attempting to open a queue and specifying a user ID that is not authorized

With the exception of channel events, all instrumentation events must be enabled before they can be generated.

Figure 13 on page 128 summarizes the production of an event message.

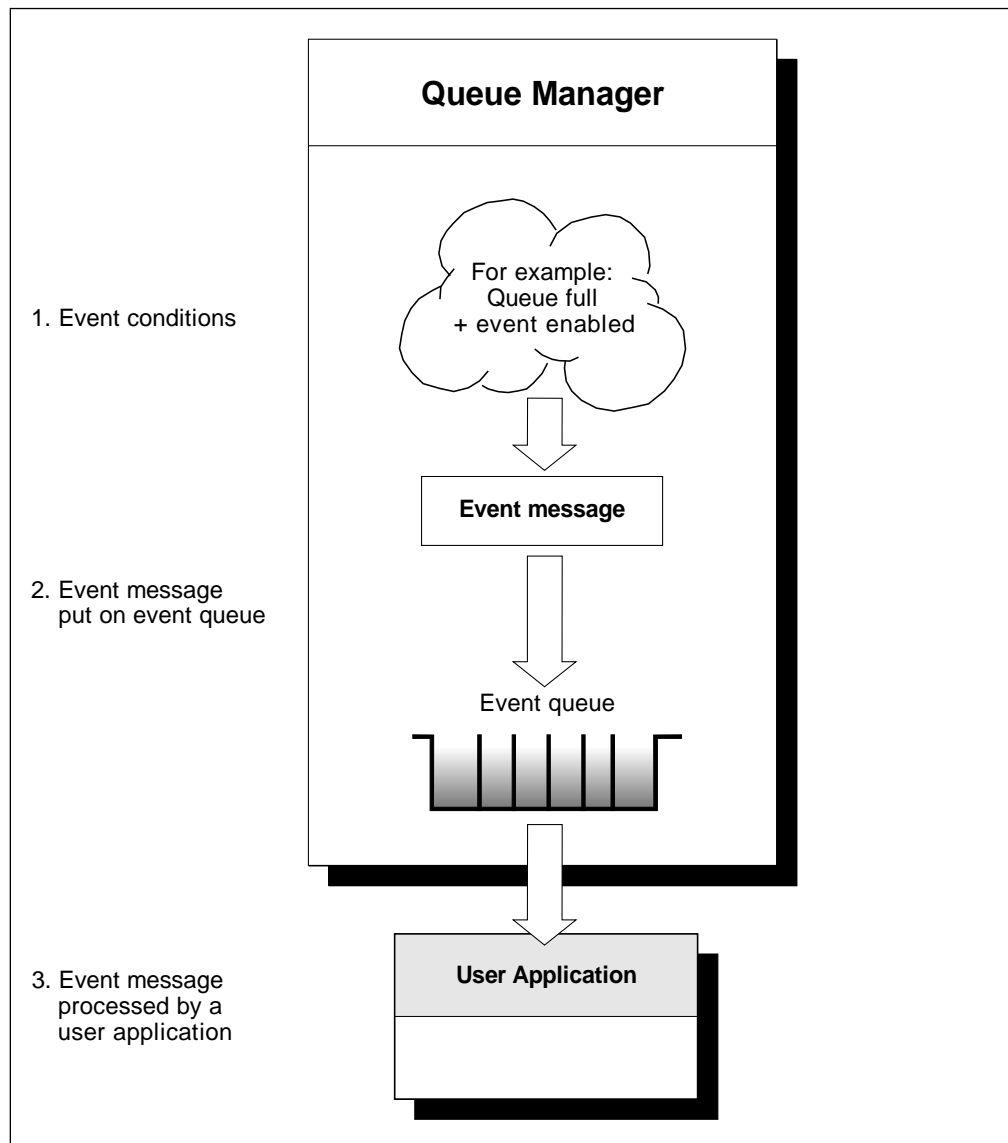


Figure 13. Understanding instrumentation events. When a queue manager detects that the conditions for an event have been met, it puts an event message on the appropriate event queue.

The event message contains information about the conditions giving rise to the event. An application can retrieve the event message from the event queue for analysis.

Why use events?

If you define your event queues as remote queues, you can put all the event queues on a single queue manager (for those nodes that support instrumentation events). You can then use the events generated to monitor a network of queue managers from a single node. Figure 14 on page 129 illustrates this.

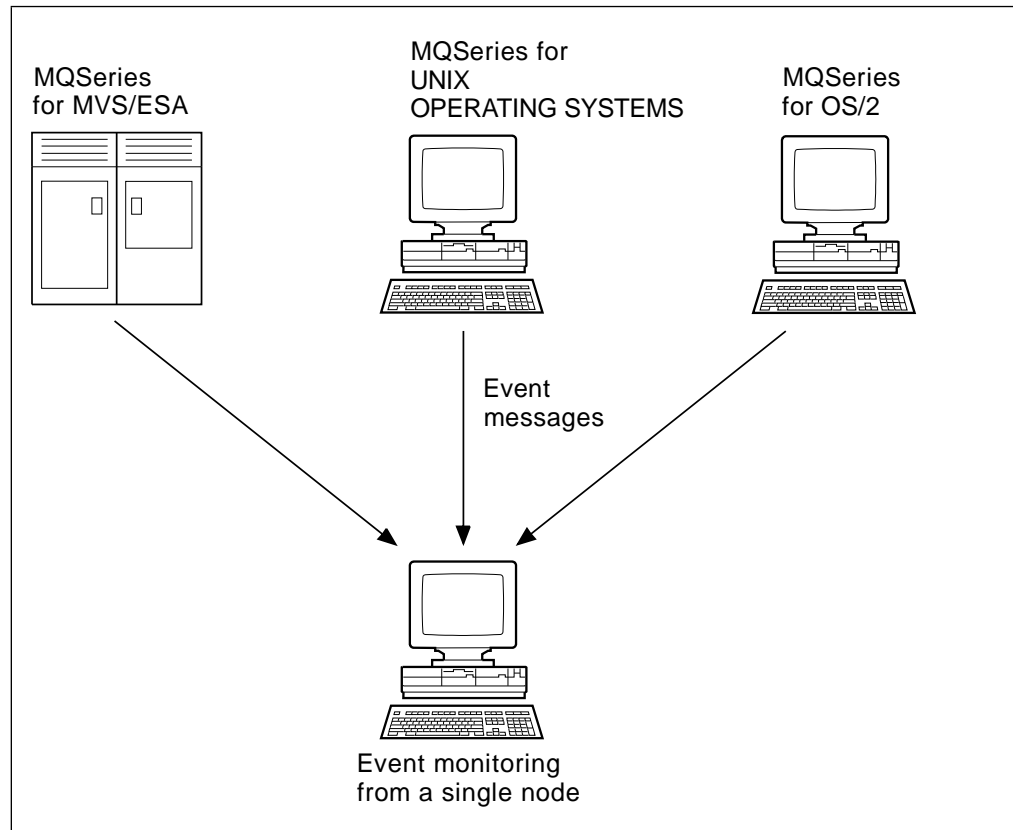


Figure 14. Monitoring queue managers across different platforms, on a single node

Types of event

MQSeries events are categorized as follows:

Queue manager events

These events are related to the definitions of resources within queue managers. For example, if an application attempts to update a resource but the associated user ID is not authorized to perform that operation, a queue manager event is generated.

Performance events

These events are notifications that a threshold condition has been reached by a resource. For example, a queue depth limit has been reached or, following an MQGET request, a queue has not been serviced within a predefined period of time.

Channel events

These events are reported by channels as a result of conditions detected during their operation. For example, a channel event is generated when a channel instance is stopped.

Trigger events

When we discuss triggering in this and other MQSeries books, we sometimes refer to a *trigger event*. This occurs when a queue manager detects that the conditions for a trigger event have been met. For example, a queue can be configured to generate a trigger event each time a message arrives. (The conditions for trigger events and instrumentation events are quite different.)

A trigger event causes a trigger message to be put on an initiation queue and, optionally, an application program is started.

Event notification through event queues

When an event occurs, the queue manager puts an event message on the appropriate event queue (if such a queue has been defined). The event message contains information about the event that you can retrieve by writing a suitable MQI application program that:

- Gets the message from the queue.
- Processes the message to extract the event data. For a description of event message formats, see the *MQSeries Programmable System Management* book.

Each category of event has its own event queue. All events in that category result in an event message being put onto the same queue.

This event queue...	Contains messages from...
SYSTEM.ADMIN.QMGR.EVENT	Queue manager events
SYSTEM.ADMIN.PERFM.EVENT	Performance events
SYSTEM.ADMIN.CHANNEL.EVENT	Channel events

You can define event queues as either local or remote queues. If you define all your event queues as remote queues on the same queue manager, you can centralize your monitoring activities.

Using triggered event queues

You can set up the event queues with triggers so that, when an event is generated, the event message being put onto the event queue starts a (user-written) monitoring application. This application can process the event messages and take appropriate action. For example, some events can require that an operator be informed, while others could start an application that performs some administration tasks automatically.

Enabling and disabling events

You enable and disable events by specifying the appropriate values for the queue manager, or queue attributes, or both, depending on the type of event. You do this using either of the following:

- MQSC commands. For more information, see the *MQSeries Command Reference* manual.
- PCF commands. For more information, see the *MQSeries Programmable System Management* manual.

Enabling an event depends on the category of the event:

- Queue manager events are enabled by setting attributes of the queue manager.
- Performance events as a whole must be enabled on the queue manager, or no performance events can occur. You enable the specific performance events by setting the appropriate queue attribute. You also have to identify the conditions, such as a queue depth high limit, that give rise to the event,
- Channel events occur automatically; they do not need to be enabled. If you do not want to monitor channel events, you can inhibit MQPUT requests to the channel event queue.

Event messages

Event messages contain information relating to the origin of an event, including the type of event, the name of the application that caused the event and, for performance events, a short statistics summary for the queue.

The format of event messages is similar to that of PCF response messages. The message data can be retrieved from them by user-written administration programs using the data structures described in the *MQSeries Programmable System Management* manual.

Chapter 10. Linking to Lotus Notes

The MQSeries Version 5 products provide a Lotus Notes server add-in task that gives Lotus Notes applications access to MQSeries messaging. This allows Lotus Notes users to communicate with other systems connected by MQSeries.

If you are installing or maintaining MQSeries in order to link to Lotus Notes, you should have Lotus Notes installed and have the documentation provided with Lotus Notes.

This chapter covers:

- “What is Lotus Notes?”
- “Linking applications”
- “Server or client?” on page 135
- “Disconnected requests” on page 137
- “Setting up your system” on page 137
- “Starting the server add-in task” on page 138
- “Verifying that Lotus Notes can link to MQSeries” on page 139
- “Stopping the link server task” on page 140
- “User notification” on page 140

What is Lotus Notes?

Lotus Notes is a networked application that users can use to share information. Lotus Notes has two main components; the server and the client. The Lotus Notes server provides services to Lotus Notes clients and to other servers. The services provided include storage and replication of shared databases and mail routing. Lotus Notes clients connect to a Lotus Notes server to use shared databases, and also to read and send mail.

The basic units of information in a Lotus Notes system are databases and the documents that they contain. A database can be used by one person, or shared among users who have common data requirements. Most databases in Lotus Notes reside on a Lotus Notes server.

Linking applications

MQSeries provides a Lotus Notes server add-in task that recognizes and interprets:

- Data from documents that Lotus Notes wants to send to MQSeries
- Messages from MQSeries sent in reply and used to update a Lotus Notes document

Figure 15 on page 134 illustrates the major components that are used to service the requests and responses.

Linking applications

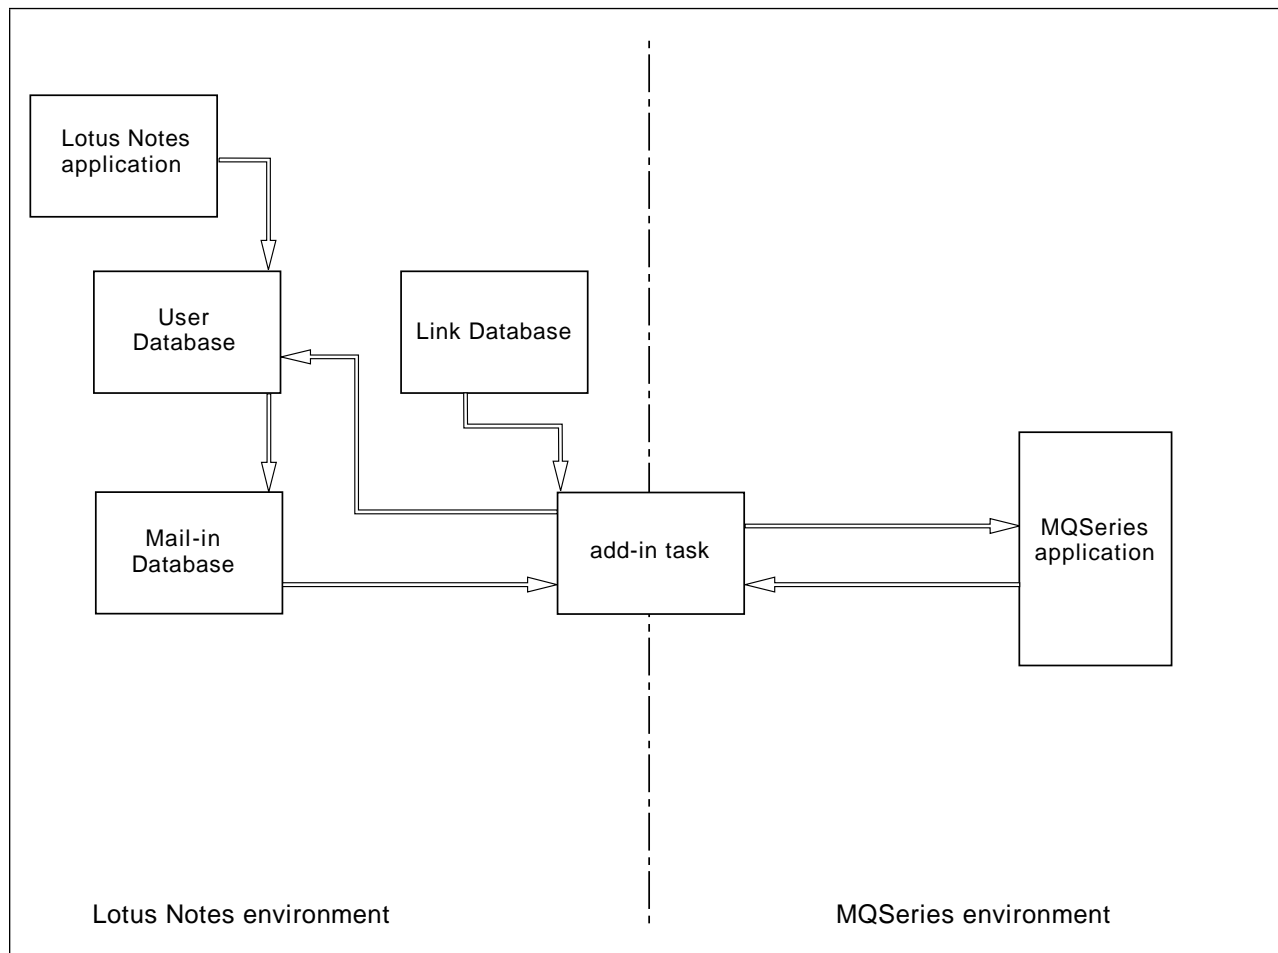


Figure 15. Lotus Notes and MQSeries

A Lotus Notes application consists of a database containing specially constructed documents. These documents contain formulas (or macros) that can be executed by the user. A formula that makes a link to MQSeries transfers parts of the document to a Mail-In database associated with the server add-in task.

The Lotus Notes server add-in task monitors the Mail-In database. Documents found there are used to construct MQSeries messages. The link database contains entries describing the relationship of a Lotus Notes document to an MQSeries message; that is, how a Lotus Notes document is mapped into an MQSeries message. The link database must be set up to define the mapping required for each type of document you want to use with MQSeries.

The name of the link database entry to use is specified in the mail message containing the request.

The Lotus Notes server add-in task can also monitor response queues based on any outstanding replies. The replies are interpreted and used to update the user document in the user database.

Refer to the *MQSeries Application Programming Guide* for information about writing the MQSeries applications and the link database entries.

Add-in task requirements

The add-in task requires:

- On the Lotus Notes Server:
 - A document in the link database for each document type that the task processes.
 - A Mail-In database from which the add-in task gets the information to send to an MQSeries-connected system.
- In the MQSeries queue manager:
 - A work queue (SYSTEM.NOTES.WORKQUEUE in Figure 16 on page 136). If you use a different name for this queue, you need to supply the name when you load the add-in task.
 - Any queues that the add-in task uses to get replies if they do not already exist. The names of these queues are contained in the link database documents in Lotus Notes.

Server or client?

The add-in task always runs on the Lotus Notes server machine. This machine will also be running either the MQSeries server or the MQSeries client code. Do not confuse the server and client relationship between the MQSeries components with that between a Lotus Notes server and its clients. The Lotus Notes server exists with an associated MQSeries that is either a server or an MQSeries client.

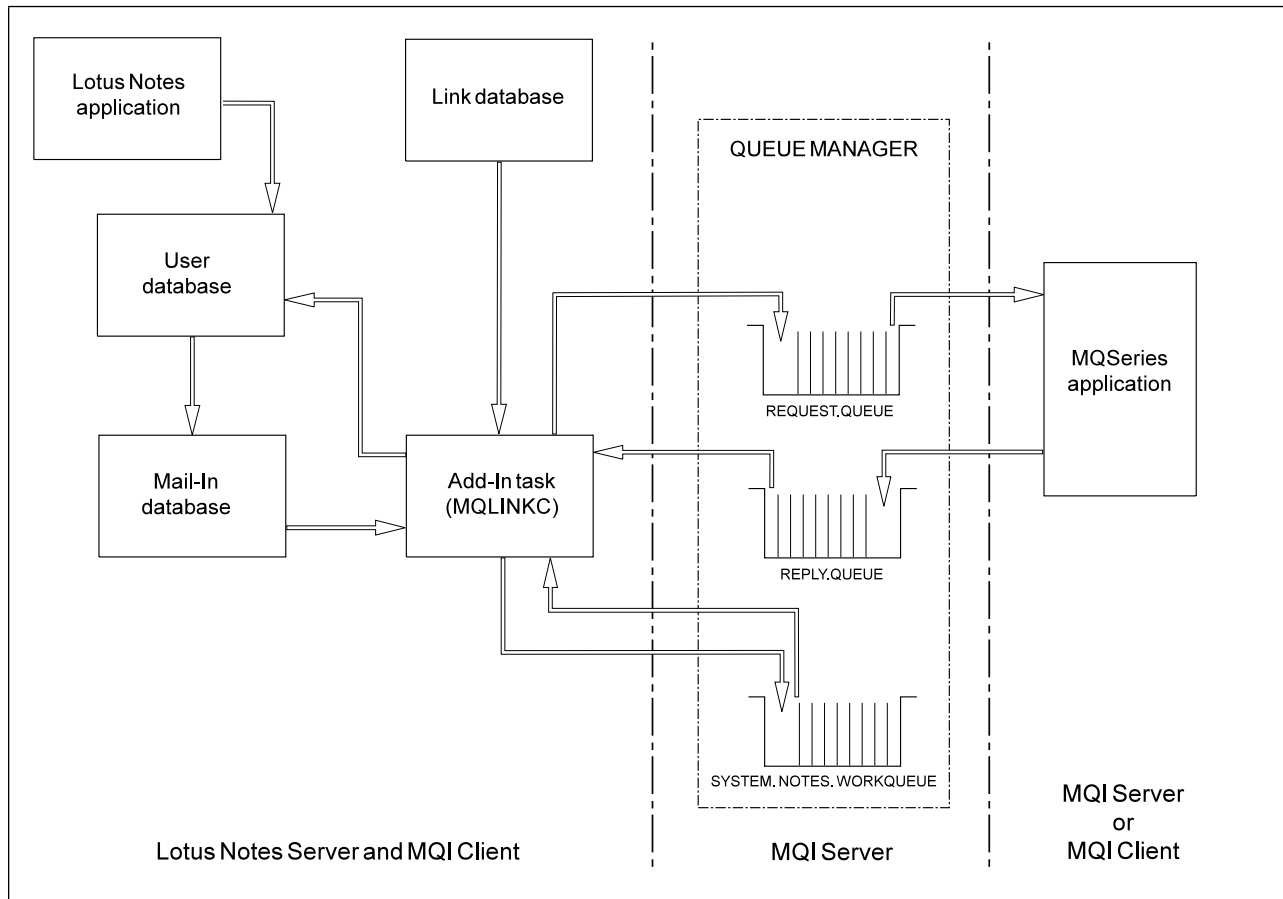


Figure 16. Configuration with the add-in task on the MQSeries client

Figure 16 shows a simple configuration with the add-in task running in an MQSeries client environment.

Notes:

1. The MQSeries application getting the messages from the REQUEST.QUEUE can be running from any MQSeries platform.
2. When the add-in task is started, it reads all the template documents in the link database and holds the information in memory. Therefore, any template you add after the add-in task has been started is not recognized until the add-in task has been stopped and restarted.

When the add-in task is running in a full MQSeries environment, the add-in task program is known as mqlink (in MQSeries for UNIX systems) or MQLINK (for MQSeries for OS/2 Warp and Windows NT).

When the add-in task is running in an MQSeries client environment, the add-in task program is known as mqlinkc (in MQSeries for UNIX systems) or MQLINKC (in MQSeries for OS/2 Warp and Windows NT).

Refer to the *MQSeries Intercommunication* book for information about channel requirements.

Disconnected requests

Lotus Notes provides two options for handling mail from a Lotus Notes client: server-based mail and workstation-based mail.

A Lotus Notes client can replicate a given user database that has been designed to generate MQSeries requests. Replication allows the client to disconnect from the server and work remotely. When the client is reconnected to the server, Lotus Notes can use the client-replicated database to update the server copy of the user database.

If a client has been working disconnected, it is important to ensure that, on reconnection, user databases are updated before any messages are passed to the add-in task for transmission by MQSeries. You can do this by ensuring that any user database is replicated before mail transfer takes place. A suggested method is to turn off the Transfer Outgoing Mail option on the Lotus Notes Tools Replicate window. Repeat the replication with the transfer enabled.

The Lotus Notes manuals document the uses and methods of replication, and also problem diagnosis. If you experience replication problems or need further information, please refer to those manuals.

Setting up your system

This section describes how to link MQSeries to a Lotus Notes server, and how to set up an MQSeries system to communicate with Lotus Notes. It is assumed that you have already installed both MQSeries and Lotus Notes V 4.0 (or a later level).

Lotus Notes setup

Enable MQSeries to communicate with Lotus Notes as follows:

1. Configure the Lotus Notes server.
2. Add the Lotus Notes server's Mail file to the Lotus Notes Workspace.
3. Set up Mail on the Lotus Notes Workspace.
4. Create a new Mail-In database.
5. Add the Mail-In database to "Address Book."
6. Add people entry to the address book with the same name as the server.

For further information about setting up Lotus Notes, see the appropriate Lotus Notes documentation.

Setting up the server add-in task

The steps you follow to prepare your Lotus Notes system to communicate with MQSeries depend on the MQSeries environment you are setting up and on your operating system.

In MQSeries for UNIX systems:

1. Copy `mqlink.nsf` from the MQSeries directory, `mqmtop/lib`, to your Lotus Notes directory.
2. Turn on write permission to the file using the following command:

```
chmod a+w mqlink.nsf
```

Starting the server

3. Copy amqsampl.nsf from the MQSeries sample directory, **mqmtop/samp**, to your Lotus Notes directory.
4. Turn on write permission to the file using the following command:

```
chmod a+w amqsampl.nsf
```
5. For a full MQSeries environment, copy mqlink from **mqmtop/bin** to your Lotus Notes directory as mqlink.
6. For an MQSeries client environment, copy mqlinkc from **mqmtop/bin** on the server to the MQSeries client Lotus Notes directory as mqlinkc.

In MQSeries for OS/2 Warp and Windows NT:

1. For a full MQSeries environment, copy MQLINK.EXE from C:\MQM\BIN to C:\NOTES as \$MQLINK.EXE.
2. For an MQSeries client environment, copy MQLINKC.EXE from C:\MQM\BIN on the server to the MQSeries client directory C:\NOTES as \$MQLINKC.EXE.
3. Copy MQLINK.NSF from the MQSeries directory, C:\MQM\TOOLS\LIB to the appropriate Lotus Notes directory (C:\NOTES\DATA on the server or client).
4. Copy AMQSAMPL.NSF from the MQSeries sample directory, C:\MQM\TOOLS\C\SAMPLES\BIN, to the appropriate Lotus Notes directory (C:\NOTES\DATA on the server or client).
5. If it is not already present, add C:\NOTES to your PATH.

Starting the server add-in task

You can start the server add-in task either by issuing a command or by updating the notes.ini file, in which case the server add-in task starts automatically. If you require the latter, the Lotus Notes notes.ini file must specify either `ServerTasks=mqlink` or `ServerTasks=mqlinkc` for the server or MQSeries client respectively (or add it to the list of tasks if `ServerTasks` is already defined). The program **must** exist in your Lotus Notes directory and have been copied there.

Note: In MQSeries for OS/2 Warp and Windows NT, if you are running the Lotus Notes server add-in task from an MQSeries client using MQLINKC, you must ensure that the Lotus Notes server shares the same environment variables as your MQSeries client. You do this by starting the Lotus Notes server from the window in which you set the MQSeries server variables, or by setting them in your CONFIG.SYS file. This ensures that Lotus Notes is aware of the channel, protocol, and port that your MQSeries client is set up to use by the SET MQSERVER command.

Note: In MQSeries for Windows NT, ensure that the directory containing the file nsdk.dll is in the path on the Lotus Notes server.

If you are starting the link server task manually, enter on the Lotus Notes server screen:

```
load mqlink [-t] [-q WorkQName] [-w WaitTime]
             [-d MailInDB] [-l LinkDB] [QMgrName]
```

Notes:

1. Start only one server add-in task on a Lotus Notes server.
2. In the command, optional parameters are enclosed in brackets. The flag indicator is a hyphen (-), and the flags are lowercase. In MQSeries for OS/2 Warp and Windows NT only, uppercase flags are valid, and the flag indicator can be a slash (/).
3. The flags can be in any order, but the queue manager name must be the last parameter.

The flags and parameters of the load command are:

Flag	Parameter
d	The name of the mail-in database. The default is mailindb.nsf.
l	The name of the link database. The default is mqlink.nsf.
q	The name of the work queue that is used by the link server task. The default is SYSTEM.NOTES.WORKQUEUE.
r	The error retry count. The default is 1.
s	The separator character as two hexadecimal characters. The default is FF.
t	If specified, trace link database data and field updates using MQSeries tracing. By default, tracing is set off.
w	The time in seconds to be allowed between scans of the mail-in database and response queue for work. The default is 30 seconds.
QMgrName	The queue manager name. This must be the last parameter.

Verifying that Lotus Notes can link to MQSeries

You can verify your configuration by running the amqslnk0 sample application, which is included in the supplied Lotus Notes database. It demonstrates a simple message exchange.

To run the sample, which is also described fully in its Help panels, you:

1. Create a queue manager and start it. The default queues are created automatically.
2. Create the additional queues used by the sample:
 - In MQSeries for UNIX systems:
runmqsc qmgrname < **mqm**top/samp/amqslnk0.tst
 - In MQSeries for OS/2 Warp and Windows NT:
runmqsc qmgrname < C:\MQM\TOOLS\MQSC\SAMPLES\AMQSLNK0.TST
3. Start the amqslnk0 program:

```
amqslnk0 [-q InputQueueName] [QMgrName]
```

- In MQSeries for UNIX systems, the program can be found in **mqm**top/samp/bin.
 - In MQSeries for OS/2 Warp and Windows NT, the program can be found in c:\mqm\tools\c\samples\bin.
4. Load mqlink (see "Starting the server add-in task" on page 138).
 5. Open the MQSeries sample database and compose a document.

Stopping the link server • User notification

6. Select the Send data button. Fields from the document are read and sent to the Mail-In database. The add-in task then reads the Mail-In database together with the link database.

MQSeries message data is constructed and placed on the MQSeries target queue specified in the link database document MQENTRY. The MQSeries sample program amqslnk0 reads the message and replies appropriately. The add-in task reads the reply queue and updates the original Lotus Notes document.

7. Refresh the screen to see the update that has been made to the document.

For more information about this sample, refer to the *MQSeries Application Programming Guide*.

Stopping the link server task

To stop or close down the link server add-in task, issue one of the following commands as appropriate from the Lotus Notes server.

```
tell mqlink quit  
  
tell mqlinkc quit
```

You should stop the add-in task before closing the queue manager.

User notification

If there is an error with the transfer of the request to the mail-in database, Lotus Notes reports the error. If the transfer to the mail-in database is successful, the add-in task attempts to process the request.

If there are no errors encountered when processing either the request or the reply, the target user database specified in the link database entry is updated. If MQSeries encounters an error, it attempts to inform the document's author via Lotus Notes electronic mail.

All add-in task activity is recorded in the Lotus Notes Server's Note log and in the MQSeries log.

Chapter 11. Transactional support

The *MQSeries Application Programming Guide* contains a complete introduction to the subject of this chapter. A brief introduction only is provided here.

An application program can group a set of updates into a *unit of work*. These updates are usually logically related and must all be successful for data integrity to be preserved. If one update succeeded while another failed then data integrity would be lost.

A unit of work **commits** when it completes successfully. At this point all updates made within that unit of work are made permanent or irreversible. If the unit of work fails then all updates are instead *backed out*. *Syncpoint coordination* is the process by which units of work are either committed or backed out with integrity.

A *local* unit of work is one in which the only resources updated are those of the MQSeries queue manager. Here syncpoint coordination is provided by the queue manager itself using a single-phase commit process.

A *global* unit of work is one in which resources belonging to other resource managers, such as XA-compliant databases, are also updated. Here, a two-phase commit procedure must be used and the unit of work may be coordinated by the queue manager itself, or externally by another XA-compliant transaction manager such as IBM CICS, Transarc Encina, or BEA Tuxedo.

In summary, queue manager resources can be updated as part of local or global units of work:

Local unit of work

Use local units of work when the only resources to be updated are those of the MQSeries queue manager. Updates are committed using the MQCMIT verb or backed out using MQBACK.

Global unit of work

Use global units of work when you also need to include updates to XA-compliant database managers. Here the coordination may be internal or external to the queue manager.

Queue manager coordination

Global units of work are started using the MQBEGIN verb and then committed using MQCMIT or backed out using MQBACK. A two-phase commit process is used whereby XA-compliant resource managers such as DB2 and Oracle are firstly all asked to prepare to commit. Only if all are prepared successfully will they then be asked to commit. If any resource manager signals that it cannot prepare to commit, each will be asked to back out instead.

External coordination

Here the coordination is performed by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or BEA Tuxedo. Units of work are started and committed under control of the transaction manager. The MQBEGIN, MQCMIT and MQBACK verbs are unavailable.

This chapter describes how to enable support for global units of work (support for local units of work does not need to be specifically enabled).

It contains these sections:

- “Database coordination”
- “DB2 configuration” on page 146
- “Oracle configuration” on page 152
- “Multiple database configurations” on page 157
- “Administration tasks” on page 159
- “External syncpoint coordination” on page 164
- “Using CICS” on page 166

Database coordination

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

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

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

The MQI verb, MQBEGIN, must be used to denote units of work that are also to involve database updates. The *MQSeries Application Programming Guide* describes sample programs that make MQSeries and database updates within the same unit of work.

The queue manager communicates with the database managers using the XA interface as described in *X/Open Distributed Transaction Processing: The XA Specification (ISBN 1 872630 24 3)*.

This means that the queue manager can communicate to database managers that also adhere to this standard. Such database managers are known as *XA-compliant* database managers.

Table 11 describes the XA-compliant database managers that are supported by the MQSeries Version 5 products.

Table 11. XA-compliant relational databases

MQSeries product	DB2	Oracle
MQSeries for AIX	√	√
MQSeries for HP-UX	√	√
MQSeries for OS/2 Warp	√	
MQSeries for Sun Solaris	√	√
MQSeries for Windows NT	√	

Restrictions

The following restrictions apply to the database coordination support:

- The ability to coordinate database updates within MQSeries units of work is **not** supported in an MQI client application.
- The MQI updates and database updates must be made on the same queue manager server machine.
- The database server may reside on a different machine from the queue manager server. In this case the database needs to be accessed via an XA-compliant client feature provided by the database manager itself.
- Although the queue manager itself is XA-compliant, it is not possible to configure another queue manager as a participant in global units of work. This is because only one connection at a time can be supported.

Database connections

An application that establishes a standard connection to the queue manager will be associated with a thread in a separate local queue manager agent process. When the application issues MQBEGIN then both it and the agent process will need to connect to the databases that are to be involved in the unit of work. The database connections are maintained while the application remains connected to the queue manager. This is an important consideration if the database only supports a limited number of users or connections.

One method of reducing the number of connections is for the application to use the MQCONNX call to request a fastpath binding. In this case the application and the local queue manager agent become the same process and consequently can share a single database connection. Before you do this, consult the *MQSeries Application Programming Guide* for a list of restrictions that apply to fastpath applications.

Configuring database managers

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

1. Create an XA switch load file for the database manager.

An *XA switch load* file is a dynamically loaded object that enables the queue manager and the database manager to communicate with each other.

2. Define the database manager in the queue manager's configuration file `qm.ini`.

Various items, including the name of the switch load file, need to be defined in `qm.ini`.

Creating switch load files

A sample makefile that can be used to build switch load files for the supported database managers is shipped with each of the MQSeries Version 5 products.

This makefile, together with all the associated files required to build the switch load files, is installed (for MQSeries for OS/2 Warp or Windows NT) in the `\mqm\tools\c\samples\xatm\` directory, or (for MQSeries for UNIX systems) in the `mqm\top\samp\xatm\` directory. Refer to your MQSeries installation documentation for more details about the installation procedure.

The sample source modules that are used to produce the switch load files all contain a single function called `MQStart`. When the switch load file is loaded, the queue manager calls this function and it returns the address of a structure called an XA switch. The switch load file is linked to a library provided by the database manager, which enables MQSeries to call that database manager.

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

- For DB2, `db2swit.c`
- For Oracle, `oraswit.c`

Defining database managers in the `qm.ini` file

When you have created a switch load file for your database manager, you must define it to your queue manager. This is done in the queue manager's `qm.ini` file in the `XAResourceManager` stanza.

Unless you are configuring MQSeries to coordinate updates to more than one database, you need to add only a single `XAResourceManager` stanza. More complicated configurations involving multiple databases, or different database managers, are discussed in "Multiple database configurations" on page 157.

The attributes of the `XAResourceManager` stanza are as follows:

Name

Identifies the database manager instance.

The name can be up to 31 characters in length, and is mandatory. It must be unique within the `qm.ini` file. It could simply be the name of the database manager, although to maintain its uniqueness in more complicated configurations it could, for example, also include the name of the database being updated.

The name that you choose should be meaningful because the queue manager uses it to refer to this database manager instance both in messages and in output when the **dspmqrn** command is used.

Once you have chosen a name, do not change this attribute. Information about changing the `qm.ini` file is provided in “Changing the `qm.ini` configuration file” on page 163.

SwitchFile

This is the fully-qualified name of the database manager’s XA switch load file. This is a mandatory attribute.

XAOpenString

This is a string of data that is passed to the database manager’s `xa_open` entry point. The format for this string depends on the particular database manager, but it should usually identify the name of the database that is to be updated.

This is an optional attribute; if it is omitted a blank string is assumed.

XACloseString

This is a string of data that is passed to the database manager’s `xa_close` entry point. The format for this string depends on the particular database manager.

This is an optional attribute; if it is omitted a blank string is assumed.

ThreadOfControl

This attribute applies only to the MQSeries for OS/2 and MQSeries for Windows NT products, where it is mandatory. The *ThreadOfControl* value can be `THREAD` or `PROCESS`, and the queue manager uses it for serialization purposes. It governs how the queue manager calls the database manager from within its own multithreaded processes.

If the database manager is thread safe the value for *ThreadOfControl* should be `THREAD`, and the queue manager can call the database manager from multiple threads at the same time.

If the database manager is not thread safe, the value for *ThreadOfControl* should be `PROCESS`, and the queue manager serializes all calls to the database manager so that only one call at a time is made from within a particular process.

“DB2 configuration” on page 146 and “Oracle configuration” on page 152 provide details of the specific tasks that need to be performed to configure MQSeries with each of the supported database managers.

DB2 configuration

The minimum supported level of DB2 is V2.1.1.

You need to perform the following tasks:

- Check the environment variable settings.
- Create the DB2 switch load file.
- Add an XAResourceManager stanza to the qm.ini file.
- Possibly change DB2 configuration parameters.

Checking the environment variable settings

You should ensure that your DB2 environment variables are set for queue manager processes as well as in your application processes. In particular, you must always set the following environment variable **before** you start the queue manager:

DB2INSTANCE Identifies the DB2 instance containing the DB2 databases that are being updated.

Creating the DB2 switch load file

The easiest method for creating the DB2 switch load file is to use the sample file xaswit.mak. The source code used to create the DB2 switch on most platforms is shown in Figure 17. The source for db2swit.c for Windows NT is different; it is shown in Figure 18.

```
#include <cmqc.h>
#include "xa.h"

extern struct xa_switch_t db2xa_switch;

struct xa_switch_t * MQENTRY MQStart(void)
{
    return(&db2xa_switch);
}
```

Figure 17. Source code for db2swit.c for platforms other than Windows NT

```
#include <cmqc.h>
#include "xa.h"

extern __declspec(dllimport) struct xa_switch_t db2xa_switch;

struct xa_switch_t * MQENTRY MQStart(void)
{
    return(&db2xa_switch);
}
```

Figure 18. Source code for db2swit.c on Windows NT (Microsoft Visual C++ specific)

The xa.h header file that is supplied with MQSeries is installed (for MQSeries for OS/2 Warp or MQSeries for Windows NT) in the \mqm\tools\c\samples\xatm directory, or (for MQSeries for AIX) in the **mqm**top/samp/xatm directory.

Creating db2swit.dll on OS/2

To create the DB2 switch load file on OS/2, db2swit.c must be compiled and linked against db2api.lib. The DEF file shown in Figure 19 is needed to produce the DLL:

```
LIBRARY DB2SWIT

CODE SHARED    LOADONCALL
DATA NONSHARED MULTIPLE

EXPORTS
  MQStart                @1
```

Figure 19. Source code for db2swit.def on OS/2

To create the DLL, first create a directory into which you want the switch file to be built. The switch file must be defined to MQSeries as a fully-qualified name so the DLL does not need to be built into a directory within LIBPATH.

Copy the following files from \mqm\tools\c\samples\xatm into your new directory:

- xa.h
- db2swit.c
- db2swit.def
- xaswit.mak

The following source code for xaswit.mak on OS/2 can be used to build the switch load file:

```
CFLAGS=/c /Ss /Gm /Ge- /Q /Sp1
LFLAGS=/NOFREE /noi /align:16 /exepack

.SUFFIXES: .c .obj

db2swit.dll: db2swit.obj db2swit.def {$(LIB)}db2api.lib

.obj.dll:
  $(CC) /B"$(LFLAGS)" /Fe $@ $**

.c.obj:
  $(CC) $(CFLAGS) $*.c
```

Figure 20. Makefile for DB2 switch on OS/2

Finally issue the following command to make db2swit.dll:

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

Creating db2swit.dll on Windows NT

To create the DB2 switch load file on Windows NT, db2swit.c must be compiled and linked against db2api.lib. The following DEF file is needed to produce the DLL:

```
LIBRARY DB2SWIT

EXPORTS
    MQStart
```

Figure 21. Source code for db2swit.def on Windows NT

To create the DLL, first create a directory into which you want the switch file to be built. The switch file must be defined to MQSeries as a fully-qualified name so the DLL does not need to be built into a directory with LIBPATH.

Copy the following files from \mqm\tools\c\samples\xatm into your new directory:

- xa.h
- db2swit.c
- db2swit.def
- xaswit.mak

The following source code forms part of xaswit.mak on Windows NT and can be used to build the switch load file:

```
!include <ntwin32.mak>

db2swit.lib db2swit.exp: *.obj *.def
                    $(implib) -machine:$(CPU) \
                    -def:*.def *.obj

db2swit.dll: *.obj *.def *.exp
            $(link) $(dllflags) \
            -base:0x1C000000 \
            *.exp *.obj \
            $(conlibsdl) db2api.lib

.c.obj:
    $(cc) $(cflags) $(cvarsdll) *.c
```

Figure 22. Makefile for DB2 switch on Windows NT

Finally issue the following command to make db2swit.dll using the Microsoft Visual C++ compiler:

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

Creating db2swit on UNIX systems

To create the DB2 switch load file on UNIX systems, db2swit.c must be compiled and linked against libdb2.

To build the switch load file, first create a directory into which db2swit will be built, then copy the following files from **mqmtp**/samp/xatm into this new directory:

- xa.h
- db2swit.c
- xaswit.mak

The following source forms part of xaswit.mak on AIX and is used to build the DB2 switch file:

```
DB2LIBS=-l db2
DB2LIBPATH=-L /usr/lpp/db2_02_01/lib

db2swit:
    $(CC) -e MQStart $(DB2LIBPATH) $(DB2LIBS) -o $@ db2swit.c
```

Figure 23. Makefile for DB2 switch on AIX

The following source forms part of xaswit.mak on Sun Solaris and is used to build the DB2 switch file:

```
DB2LIBS=-l db2
DB2LIBPATH=-L /opt/IBMDB2/V2.1/lib

db2swit:
    $(CC) -G -e MQStart $(DB2LIBPATH) $(DB2LIBS) -o $@ db2swit.c
```

Figure 24. Makefile for DB2 switch on Sun Solaris

The following source forms part of xaswit.mak on HP-UX and is used to build the DB2 switch file:

```
DB2LIBS=-l db2 -l c1
DB2LIBPATH=-L /opt/IBMDB2/V2.1/lib

db2swit: db2swit.c
    $(CC) -c -Ae +z db2swit.c

    ld -b -e MQStart $(DB2LIBPATH) $(DB2LIBS) -o db2swit db2swit.o
```

Figure 25. Makefile for DB2 switch on HP-UX

To build the DB2 switch using the sample makefile issue the following command:

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

Adding the XAResourceManager stanza to the qm.ini file

The next step is to modify the qm.ini configuration file of the queue manager to define DB2 as a participant in global units of work. You need to add an XAResourceManager stanza with the following attributes:

Name Choose a suitable name for this participant; you could include the name of the database being updated.

SwitchFile The fully-qualified name of the DB2 switch load file

XAOpenString

The XA open string for DB2 must be of the following format:

```
database_alias<,username,password>
```

where:

- database_alias is the name of the database, unless you have explicitly cataloged an alias name after the database was created in which case specify the alias instead.

The following two parameters are optional and are not normally needed when configuring DB2 with MQSeries as the transaction manager. They provide alternative authentication information to the database if it was set up with **authentication=server**.

- username specifies a valid operating system user ID.
- password is the password for the specified user ID.

See "Security considerations" on page 157 for more information about security.

XACloseString

DB2 does not require an XA close string.

ThreadOfControl

DB2 is thread-aware so specify *THREAD*.

In the sample XAResourceManager entries that follow the database to be updated is called MQBankDB, this name being specified as the *XAOpenString*.

```
XAResourceManager:
  Name=DB2 MQBankDB
  SwitchFile=c:\user\d11\db2swit.d11
  XAOpenString=MQBankDB
  ThreadOfControl=THREAD
```

Figure 26. Sample XAResourceManager entry for DB2 on OS/2 and Windows NT

In the following UNIX sample, it is assumed that the DB2 switch load file was copied to the /usr/bin directory after it had been created:

```
XAResourceManager:
  Name=DB2 MQBankDB
  SwitchFile=/usr/bin/db2swit
  XAOpenString=MQBankDB
```

Figure 27. Sample XAResourceManager entry for DB2 on UNIX platforms

Changing DB2 configuration parameters

Perform each of the following steps to each DB2 database that is being coordinated by the queue manager.

- *Database privileges*

The mqm user ID must be given connect authority to the DB2 database so that the queue manager can connect to DB2 from within its own processes.

For example, to give the mqm user ID connect authority to the MQBankDB database the following commands could be used:

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

See “Security considerations” on page 157 for more information about security.

- *tp_mon_name parameter*

For DB2 for OS/2 and DB2 for Windows NT only, the *TP_MON_NAME* configuration parameter must be updated to name the DLL that DB2 uses to call the queue manager for dynamic registration.

This can be achieved using the following command:

```
db2 update dbm cfg using TP_MON_NAME mqmax
```

This names MQMAX.DLL as the library that DB2 uses to call the queue manager. This must be present in a directory within LIBPATH.

- *maxappls parameter*

You may need to review your setting for the *maxappls* parameter, which limits the maximum number of applications that can be connected to a database.

Refer to “Database connections” on page 143.

Oracle configuration

You need to perform the following tasks:

- Check Oracle level and apply patches if you have not already done so.
- Check environment variable settings.
- Enable Oracle XA support.
- Create the Oracle switch load file.
- Add an XAResourceManager stanza to the qm.ini file.
- Possibly change the Oracle configuration parameters.

Checking Oracle level and applying patches

The minimum supported level of Oracle7 on AIX is 7.3.2.1. The minimum supported level of Oracle on HP-UX is 7.3.2.3. You also need to install Oracle patches 437448 and 441647. The minimum supported level of Oracle7 on Sun Solaris is 7.3.2.3.

Checking the environment variable settings

You should ensure that your Oracle environment variables are set for queue manager processes as well as in your application processes. In particular, the following environment variables should always be set **prior** to starting the queue manager:

ORACLE_HOME Is the Oracle home directory

ORACLE_SID Is the Oracle SID used during installation of Oracle

Enabling Oracle XA support

Before you can create the Oracle switch load file you need to ensure that Oracle XA support is enabled. In particular, an Oracle shared library must have been created; this happens during installation of the Oracle XA library. You may be prompted with:

```
Some TP Monitors require a shared version of the ORACLE7 libraries.  
Do you want to install a shared version of the libraries?
```

Make sure you answer Yes to this prompt. This creates a shared library called libclntsh in the \$ORACLE_HOME/lib directory. You should copy this to the /usr/lib directory.

Creating the Oracle switch load file

The simplest method for creating the Oracle switch load file is to use the sample file. The source code used to create the Oracle switch is as follows:

```
#include <cmqc.h>
#include "xa.h"

extern struct xa_switch_t xaosw;

struct xa_switch_t * MQENTRY MQStart(void)
{
    return(&xaosw);
}
```

Figure 28. Source code for oraswit.c

The xa.h header file that is included is shipped with MQSeries in the same directory as oraswit.c.

Creating oraswit on UNIX systems

To create the Oracle switch load file on UNIX systems, oraswit.c must be compiled and linked against libcIntsh.

To build the switch load file, first create the directory into which oraswit will be built, then copy the following files from **mqmtop/samp/xatm** into this directory:

- xa.h
- oraswit.c
- xaswit.mak

The following source forms part of xaswit.mak on AIX and is used to build the Oracle switch file:

```
ORALIBS=-l cIntsh -l m
ORALIBPATH=-L $(ORACLE_HOME)/lib -L /usr/lib

oraswit:
    xlc_r -e MQStart $(ORALIBPATH) $(ORALIBS) -o $@ oraswit.c
```

Figure 29. Makefile for Oracle switch on AIX

The following source forms part of xaswit.mak on Sun Solaris and is used to build the Oracle switch file:

```
ORALIBS=-l cIntsh -l m
ORALIBPATH=-L $(ORACLE_HOME)/lib -L /usr/lib

oraswit:
    $(CC) -G -e MQStart $(ORALIBPATH) $(ORALIBS) -o $@ oraswit.c
```

Figure 30. Makefile for Oracle switch on Sun Solaris

Oracle configuration

The following source forms part of xaswit.mak on HP-UX and is used to build the Oracle switch file:

```
ORALIBS=-l cIntsh -l m
ORALIBPATH=-L $(ORACLE_HOME)/lib -L /usr/lib

oraswit:
    $(CC) -c -Ae +z oraswit.c

    ld -b -e MQStart $(ORALIBPATH) $(ORALIBS) -o oraswit oraswit.o
```

Figure 31. Makefile for Oracle switch on HP-UX

To build the Oracle switch using the sample makefile, issue the following command:

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

Adding the XAResourceManager stanza to the qm.ini file

The next step is to modify the qm.ini configuration file of the queue manager to define Oracle as a participant in global units of work. You need to add an XAResourceManager stanza with the following attributes:

Name Choose a suitable name for this participant. You could include the name of the database being updated.

SwitchFile The fully-qualified name of the Oracle switch load file

XAOpenString

The XA open string for Oracle has the following format:

```
Oracle_XA+Acc=P//|P/userName/passWord
      +SesTm=sessionTimeLimit
      [+DB=dataBaseName]
      [+GPwd=P/groupPassWord]
      [+LogDir=logDir]
      [+MaxCur=maximumOpenCursors]
      [+SqlNet=connectString]
```

where:

Acc=

Is mandatory and is used to specify user access information. P// indicates that no explicit user or password information is provided and that the **ops\$login** form is to be used. P/userName/passWord indicates a valid ORACLE user ID and the corresponding password.

SesTm=

Is mandatory and is used to specify the maximum amount of time that a transaction can be inactive before the system automatically deletes it. The unit of time is in seconds.

DB=

Is used to specify the database name, where `dataBaseName` is the name Oracle precompilers use to identify the database. This field is required only when applications explicitly specify the database name (that is, use an `AT` clause in their SQL statements).

GPwd=

Is used to specify the server security password, where `P/groupPassWord` is the server security group password name. Server security groups provide an extra level of protection for different applications running against the same ORACLE instance. The default is an ORACLE-defined server security group.

LogDir=

Is used to specify the directory on a local machine where the Oracle XA library error and tracing information can be logged. If a value is not specified, the current directory is assumed.

MaxCur=

Is used to specify the number of cursors to be allocated when the database is opened. It serves the same purpose as the precompiler option, `maxopencursors`.

SqlNet=

Is used to specify the SQL*Net connect string that is used to log on to the system. The connect string can be either an SQL*Net V1 string, SQL*Net V2 string, or SQL*Net V2 alias. This field is required when you are setting up Oracle on a machine separate from the queue manager.

XACloseString

Oracle does not require an XA close string.

ThreadOfControl

You do not need to specify this parameter on UNIX platforms.

In Figure 32, the database to be updated is called `MQBankDB`. Note that it is recommended to add a `LogDir` to the XA open string so that all error and tracing information is logged to the same place. It is assumed that the Oracle switch load file was copied to the `/usr/bin` directory after it had been created.

```
XAResourceManager:
  Name=Oracle MQBankDB
  SwitchFile=/usr/bin/oraswit
  XAOpenString=Oracle_XA+Acc=P/scott/tiger+SesTm=35+LogDir=/tmp/ora.log+DB=MQBankDB
```

Figure 32. Sample XAResourceManager entry for Oracle on UNIX platforms

Changing the Oracle configuration parameters

The queue manager and user applications use the user ID specified in the XA open string when they connect to Oracle.

- *Database privileges*

The Oracle user ID specified in the open string must have the privileges to access the V\$XATRANS\$ table.

The necessary privilege can be given using the following command, where user is the user ID for which access is being given.

```
grant select on V$XATRANS$ to user
```

See “Security considerations” on page 157 for more information about security.

- *Additional database connections*

You may need to review your LICENSE_MAX_SESSIONS and PROCESSES settings to take into account the additional connections required by processes belonging to the queue manager. See “Database connections” on page 143 for details about the database connections that the queue manager needs for itself.

Multiple database configurations

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

If the databases are all managed by the same database manager, each stanza defines a separate database belonging to that database manager. Each stanza should specify the same `SwitchFile`, but the contents of the `XAOpenString` will be different because it specifies the name of the database being updated. For example, the stanzas shown in Figure 33 configure the queue manager with the DB2 databases `MQBankDB` and `MQFeeDB` on UNIX platforms.

```
XAResourceManager:
  Name=DB2 MQBankDB
  SwitchFile=/usr/bin/db2swit
  XAOpenString=MQBankDB

XAResourceManager:
  Name=DB2 MQFeeDB
  SwitchFile=/usr/bin/db2swit
  XAOpenString=MQFeeDB
```

Figure 33. Sample `XAResourceManager` entries for multiple DB2 databases

If the databases to be updated are managed by different database managers then once again an `XAResourceManager` stanza needs to be added for each. In this case, each stanza specifies a different `SwitchFile`. For example, if the `MQFeeDB` was managed by Oracle instead of DB2 then the following stanzas could be used:

```
XAResourceManager:
  Name=DB2 MQBankDB
  SwitchFile=/usr/bin/db2swit
  XAOpenString=MQBankDB

XAResourceManager:
  Name=Oracle MQFeeDB
  SwitchFile=/usr/bin/oraswit
  XAOpenString=Oracle_XA+Acc=P/scott/tiger+SesTm=35+LogDir=/tmp/ora.log+DB=MQFeeDB
```

Figure 34. Sample `XAResourceManager` entries for a DB2 and Oracle database

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

Security considerations

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

An application process denotes the start of a global unit of work using the `MQBEGIN` verb. The first `MQBEGIN` call that an application issues connects to each of the participating databases by calling them at their `xa_open` entry point. All of the database managers provide a mechanism for supplying a user ID and password in their `XAOpenString`.

Security considerations

If a user ID is specified in the XAOpenString then it is recommended that one with a minimal set of authorisations be chosen. Consult the documentation of the database manager to determine how the application can gain different privileges. In general this can usually be achieved using EXEC SQL CONNECT or EXEC SQL SET CONNECTION.

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

Administration tasks

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

- The queue manager can be started at any time without first starting each of the database managers.
- The queue manager does not need to be stopped and restarted if one of the database managers becomes unavailable.

This allows you to start and stop the queue manager independently from the database managers, and vice versa if the database manager supports it.

Whenever contact is lost between the queue manager and a database manager they need to resynchronize when both become available again.

Resynchronization is the process by which any in-doubt units of work involving that database are completed. In general this occurs automatically without the need for user intervention. The queue manager asks the database manager for a list of units of work in which it is in doubt. Next it instructs the database manager to either commit or rollback each of these in-doubt units of work.

When the queue manager stops, it needs to resynchronize with each database manager instance during restart. When an individual database manager becomes unavailable, only that database manager need be resynchronized the next time the queue manager notices that the database manager is available again.

The queue manager attempts to regain contact with an unavailable database manager automatically as new global units of work are started. Alternatively, the **rsvmqtrn** command can be used to resolve explicitly all in-doubt units of work.

In-doubt units of work

A database manager may be left with in-doubt units of work if contact with the queue manager is lost after the database manager has been instructed to PREPARE. Until the database manager receives the COMMIT or ROLLBACK outcome from the queue manager, it needs to retain the database locks associated with the updates.

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

If for some reason you cannot wait for the queue manager to resynchronize with the database automatically, you could use facilities provided by the database manager to commit or rollback the database updates manually. This is called making a *heuristic* decision and should be used only as a last resort because of the possibility of compromising data integrity; you may end up committing the database updates when all of the other participants rollback, or vice versa.

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

Using the `dspmqtrn` command

While a database manager is unavailable it is possible to use the `dspmqtrn` command to check the state of outstanding units of work involving that database.

When a database manager becomes unavailable, any in-flight units of work in which it was participating are rolled back. The database manager itself physically rolls back its in-flight updates when it next restarts.

The `dspmqtrn` command displays only those units of work in which one or more participants are in doubt, awaiting the COMMIT or ROLLBACK from the queue manager.

For each of these units of work the state of each of the participants is displayed. If the unit of work did not update the resources of a particular resource manager, it is not displayed.

With respect to an in-doubt unit of work, a resource manager is referred to as:

Prepared The resource manager is prepared to commit its updates.

Committed The resource manager has committed its updates.

Rolled-back
The resource manager has rolled back its updates.

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

Note that the queue manager does not remember the individual states of the participants when the queue manager restarts. If the queue manager is restarted, but is unable to contact a database manager, then the in-doubt units of work in which that database manager was participating are not resolved during restart. In this case, the database manager is reported as being in *prepared* state until such time as resynchronization has occurred.

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

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

Figure 35. Sample `dspmqtrn` output

“Using the `dspmqtrn` command” shows that there are three resource managers associated with the queue manager. The first is the resource manager 0, which is the queue manager itself. The other two resource manager instances are the `MQBankDB` and `MQFeeDB` DB2 databases.

The example shows only a single in-doubt unit of work. A message is issued for all three resource managers, which means that updates had been made to the queue manager and both DB2 databases within the unit of work.

The updates made to the queue manager, resource manager 0, have been *committed*, meaning that all message puts and message gets made under syncpoint have been actioned. The updates to the DB2 databases are in *prepared* state, which means that DB2 must have become unavailable before it was called to commit the updates to the `MQBankDB` and `MQFeeDB` databases.

The in-doubt unit of work has an external identifier called an `XID`. This is the identifier that DB2 associates with the updates.

Using the `rsvmqtrn` command

The output from the `dspmqtrn` command in “Using the `dspmqtrn` command” on page 160 showed a single in-doubt unit of work in which the commit decision had yet to be delivered to both DB2 databases.

In order to complete this unit of work, the queue manager and DB2 need to resynchronize when DB2 next becomes available. The queue manager uses the start of new units of work as an opportunity to attempt to regain contact with DB2. Alternatively, you can instruct the queue manager to resynchronize explicitly using the `rsvmqtrn` command. You should do this soon after DB2 has been restarted so that any database locks associated with the in-doubt unit of work are released as quickly as possible.

This is achieved using the `-a` option which tells the queue manager to resolve all in-doubt units of work. In the following example, DB2 had been restarted so the queue manager was able to resolve the in-doubt unit of work:

```
> rsvmqtrn -mMY_QMGR -a
```

```
Any in-doubt transactions have been resolved.
```

Mixed outcomes and errors

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

Units of work that complete with a mixed outcome have serious implications because shared resources are no longer in a consistent state.

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

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

- If a database manager rolled back instead of committing:
AMQ7606 A transaction has been committed but one or more resource managers have rolled back.
- If a database manager committed instead of rolling back:
AMQ7607 A transaction has been rolled back but one or more resource managers have committed.

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

Damage occurring due to software errors is less likely. Units of work affected in this way have their transaction number reported by message AMQ7112. The participants may be in an inconsistent state.

```
AMQ7107: Resource manager 0 is MQSeries.  
AMQ7107: Resource manager 1 is DB2 MQBankDB  
AMQ7107: Resource manager 2 is DB2 MQFeedB  
  
AMQ7112: Transaction number 0,1 has encountered an error.  
        XID: formatID 5067085, gtrid_length 12, bqual_length 4  
           gtrid [3291A5060000201374657374]  
           bqual [00000001]  
AMQ7105: Resource manager 0 has committed.  
AMQ7104: Resource manager 1 has prepared.  
AMQ7104: Resource manager 2 has rolled back.
```

Figure 36. Sample `dspmqrn` output for a transaction in error

The queue manager does not attempt to recover from such failures until the next queue manager restart. In Figure 36, this would mean that the updates to resource manager 1, the MQBankDB database, would be left in *prepared* state even if the `rsvmqtrn` was issued to resolve the unit of work.

Changing the `qm.ini` configuration file

After the queue manager has successfully started to coordinate global units of work you should be wary about making changes to any of the `XAResourceManager` stanzas in the `qm.ini` file. If you do need to change the `qm.ini` file you can do so at any time, but the changes do not take effect until after the queue manager has been restarted. For example, if you need to alter the XA open string passed to a database manager, you need to restart the queue manager for your change to take effect.

Note that if you remove an `XAResourceManager` stanza you are effectively removing the ability for the queue manager to contact that database manager.

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

Removing database manager instances

If you do need to remove a database or database manager from your configuration permanently, you should first ensure that the database is not in doubt. You should perform this check before you restart the queue manager. Most database managers provide commands for listing in-doubt transactions. If there are any in-doubt transactions, first allow the queue manager to resynchronize with the database manager before you remove its `XAResourceManager` stanza.

If you fail to observe this procedure the queue manager still remembers all in-doubt units of work involving that database. A warning message, AMQ7623, is issued every time the queue manager is restarted. If you are never going to configure this database with the queue manager again you can instruct it to forget about these in-doubt transactions using the `-r` option of the `rsvmqtrn` command.

There are times when you might need to remove an `XAResourceManager` stanza temporarily. This is best achieved by commenting out the stanza so that it can be easily reinstated at a later time. You may decide to take this action if you are suffering errors every time the queue manager contacts a particular database or database manager. Temporarily removing the `XAResourceManager` entry concerned allows the queue manager to start global units of work involving all of the other participants. An example of a commented out `XAResourceManager` stanza follows:

```
# This database has been temporarily removed
#XAResourceManager:
# Name=DB2 MQBankDB
# SwitchFile=/usr/bin/db2swit
# XAOpenString=MQBankDB
```

Figure 37. Commented out `XAResourceManager` stanza

External syncpoint coordination

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

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

1. An application informs the external syncpoint coordinator (for example, CICS) that it wants to start a transaction.
2. The syncpoint coordinator informs known resource managers, such as MQSeries, about the current transaction.
3. The application issues calls to resource managers that are associated with the current transaction. For example, the application could issue MQGET calls to MQSeries.
4. The application issues a commit or back-out request to the external syncpoint coordinator.
5. The syncpoint coordinator completes the transaction by issuing the appropriate calls to each resource manager, typically using two-phase commit protocols.

Table 12 lists the external syncpoint coordinators that can provide a two-phase commit process for transactions in which the MQSeries Version 5 products can participate. Minimum versions and releases are shown; later versions or releases, if any, may be used.

<i>Table 12. XA-compliant external syncpoint coordinators</i>	
MQSeries	External syncpoint coordinator
MQSeries for AIX	Transaction Server for AIX V4.0 BEA Tuxedo V5.1 or V6.1
MQSeries for HP-UX	CICS for HP9000, V2.1.1 BEA Tuxedo V5.1 or V6.1 HP Encina/9000, V1.2 Transarc Encina, V2.5
MQSeries for Sun Solaris	IBM CICS for the Solaris, V2.1.1 Transarc Encina Monitor V2.5 BEA Tuxedo V5.1 or V6.1
MQSeries for Windows NT	Transaction Server for Windows NT, V4.0 BEA TUXEDO V5.1 or V6.1

Note: For MQSeries for OS/2 Warp, and for MQSeries for Windows NT with CICS for Windows NT, a single-phase commit process only is supported. For more information, see “Using CICS” on page 166.

See the *MQSeries Application Programming Guide* for information about writing and building transactions to be coordinated by an external syncpoint coordinator.

The remainder of this chapter describes how to enable external units of work.

The MQSeries XA switch structure

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

MQSeries provides two versions of this structure:

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

In the MQSeries for UNIX systems, these structures are located in the following libraries:

```
libmqmxa.a      (nonthreaded)
libmqmxa_r.a    (threaded)
```

In MQSeries for Windows NT and MQSeries for OS/2 Warp the structures are located in the following libraries:

```
mqmxa.dll      (contains only the MQRMIXASwitch version)
mqmenc.dll     (for use with Encina for Windows NT)
mqmc4swi.dll   (for use with Transaction Server for Windows NT)
```

Some external syncpoint coordinators (not CICS) require that each resource manager participating in a unit of work supplies its name in the name field of the XA switch structure. The MQSeries resource manager name is "MQSeries XA RMI."

The way in which the MQSeries XA switch structure is linked to a specific syncpoint coordinator is defined by that coordinator. Information about linking the MQSeries XA switch structure with CICS is provided in "Using CICS" on page 166. For information about linking the MQSeries XA switch structure with other XA-compliant syncpoint coordinators, consult the documentation supplied with those products.

The following considerations apply to the use of MQSeries with all XA-compliant syncpoint coordinators:

- The `xa_info` structure passed on any `xa_open` call by the syncpoint coordinator includes the name of an MQSeries queue manager. The name takes the same form as the queue-manager name passed to the `MQCONN` call. If the name passed on the `xa_open` call is blank, the default queue manager is used.
- Only one queue manager at a time may participate in a transaction coordinated by an instance of an external syncpoint coordinator: the syncpoint coordinator is effectively connected to the queue manager, and is therefore subject to the rule that only one connection at a time is supported.
- All applications that include calls to an external syncpoint coordinator can connect only to the queue manager that is participating in the transaction managed by the external coordinator (because they are already effectively connected to that queue manager). However, such applications must issue an `MQCONN` call to obtain a connection handle, and should issue an `MQDISC` call before they exit.
- A queue manager whose resource updates are coordinated by an external syncpoint coordinator must be started before the external syncpoint coordinator

starts. Similarly, the syncpoint coordinator must be ended before the queue manager is ended.

- If you are using an external syncpoint coordinator that terminates abnormally, you should stop and restart your queue manager **before** restarting the syncpoint coordinator to ensure that any messaging operations uncommitted at the time of the failure are properly resolved.

Using CICS

The versions of CICS and Transaction Server that are XA-compliant (and use a two-phase commit process) are shown in Table 12 on page 164. The note following the table shows the versions that support only a single-phase commit process.

The CICS two-phase commit process

This process applies to those versions of MQSeries that support an XA-compliant external syncpoint coordinator as shown in Table 12 on page 164.

Requirements of the two-phase process

When you use the CICS two-phase commit process with MQSeries, note the following requirements:

- MQSeries and CICS must reside on the same physical machine.
- MQSeries does not support CICS on an MQSeries client.
- You must start the queue manager whose name is specified in the XAD resource definition stanza **before** you attempt to start CICS. Failure to do this will prevent you from starting CICS if you have added an XAD resource definition stanza for MQSeries to the CICS region.
- Only one MQSeries queue manager can be accessed at a time from a single CICS region.
- A CICS transaction must issue an MQCONN request before it can access MQSeries resources. The MQCONN call must specify the name of the MQSeries queue manager specified on the XAOpen entry of the XAD resource definition stanza for the CICS region. If this entry is blank, the MQCONN request must specify the default queue manager.
- A CICS transaction that accesses MQSeries resources must issue an MQDISC call from the transaction before returning to CICS. Failure to do this may mean that the CICS application server is still connected, leaving queues open.
- You must ensure that the CICS user ID (cics) is a member of the mqm group, so that the CICS code has the authority to call MQSeries.

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

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

Enabling the CICS two-phase commit process

To enable CICS to use a two-phase commit process to coordinate transactions that include MQI calls, you must add a CICS XAD resource definition stanza entry to the CICS region.

Here is an example of adding an XAD stanza entry for MQSeries for UNIX systems:

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

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

```
cicsadd -cxad -r<cics_region> \
  ResourceDescription="MQM XA Product Description" \
  SwitchLoadFile="<Drive>:\mqm\dll\mqmc4swi.dll" \
  XAOpen=<queue_manager_name>
```

For information about using the **cicsadd** command, see the *CICS on Open Systems Administration Reference* manual, SC33-1533 or the *Transaction Server for Windows NT Version 4: Administration Guide (CICS)*.

Calls to MQSeries on UNIX systems, and MQSeries for Windows NT can be included in a CICS transaction, and the MQSeries resources will be committed or rolled back as directed by CICS. This support is not available to client applications.

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

Enabling CICS user exits

Before you attempt to make use of a CICS user exit you should read the *Transaction Server for Windows NT Version 4: Administration Guide (CICS)*.

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

One of the user exits supplied with CICS is the “Task termination user exit (UE014015).” This exit can be invoked at normal and abnormal task termination (after any syncpoint has been taken).

MQSeries supplies a CICS task termination exit in source and executable form:

<i>Table 13. CICS task termination exits</i>		
MQSeries for...	Source	Executable
Windows NT	amqzscgn.c	mqmc1415.dll
UNIX systems	amqzscgx.c	amqzscg

If you are currently using this exit, you must add the MQSeries calls from the supplied exits to your current exits. Integrate the MQ calls into your existing exits at the appropriate place in the program logic. See the comments in the sample source file for help with this.

If you are not currently using this exit, you will need to add a CICS PD program definition stanza entry to the CICS region.

Here is an example of adding a PD stanza entry for UNIX:

```
cicsadd -cpd -r<cics_region> \
  PathName="mqmtop/lib/amqzscg" \
  UserExitNumber=15
```

Here is an example of adding a PD stanza entry for Windows NT:

```
cicsadd -cpd -r<cics_region> \
  PathName="<Drive>:\mqm\dll\mqmc4swi.dll" \
  UserExitNumber=15
```

The CICS single-phase commit process

The information in this section applies to CICS OS/2, V2.0 and to CICS for Windows NT, V2.0, which support a single-phase commit only.

Note the following:

- On a single physical machine, a CICS transaction can access any queue manager, subject to the restriction that any transaction can connect to only one queue manager at a time.
- CICS transactions distributed among multiple physical machines are not supported.
- For transactions running in a CICS environment, the queue manager changes its methods of authorization and determining context as follows:
 - For MQSeries for OS/2 Warp, the user ID remains os2.
 - In the message context, the application type is MQAT_CICS.
 - The application name in the context is copied from the CICS transaction name.
- To use CICS as an external syncpoint coordinator, you must install the MQSeries-supplied code for the appropriate CICS user exits.

Enabling CICS user exits

To enable the CICS single-phase commit process, you need to enable the CICS user exits 15 and 17 (see the information about user exits that customize the operator interface in the *CICS for Windows NT Version 2 Customization Guide* or the *CICS for OS/2 Version 2.0.1 Customization Guide*).

Sample exits, providing the minimum required function, are supplied in the forms shown in Table 14.

MQSeries for...	CICS for...	Sample source	Sample executable	Library linking service
OS/2	OS/2, V2.0.1	amqzsc52.c amqzsc72.c	FAAEXP15.DLL FAAEXP17.DLL	mqmcics.lib
OS/2	OS/2, V3	amqzsc53.c amqzsc73.c	FAAEX315.DLL FAAEX317.DLL	mqmcics3.lib
Windows NT	Windows NT, V2	amqzsc5n.c amqzsc7n.c	FAAEXP15.DLL FAAEXP17.DLL	mqmcics.lib

Using the sample exits

If you are not currently using these CICS exits, then copy the relevant DLLs into a directory from where they can be accessed by CICS at CICS runtime. This can be a directory referenced by your OS/2 LIBPATH setting or by your NT PATH setting.

If you are currently using these CICS exits, you must add the MQSeries calls from the supplied samples to your current exits. These MQSeries calls, which are valid only within the context of exits 15 or 17, enable support for CICS and disable the internal MQCMIT and MQBACK calls such that they will return MQRC_ENVIRONMENT_ERROR. Integrate the MQ calls (AMQ*) in your existing exits at the appropriate place in the program logic. See the comments in the sample source code for help with this.

Chapter 12. Recovery and restart

A messaging system ensures that messages entered into the system are delivered to their destination. This means that it must provide a method of tracking the messages in the system, and of recovering messages if the system fails for any reason.

MQSeries ensures that messages are not lost by maintaining records (logs) of the activities of the queue managers that handle the receipt, transmission, and delivery of messages. It uses these logs for three types of recovery:

1. *Restart recovery*, when you stop MQSeries in a planned way.
2. *Crash recovery*, when MQSeries is stopped by an unexpected failure.
3. *Media recovery*, to restore damaged objects.

In all cases, the recovery restores the queue manager to the state it was in when the queue manager stopped, except that any in-flight transactions are rolled back, removing from the queues any messages that were not committed at the time the queue manager stopped. Recovery restores all persistent messages; nonpersistent messages are lost during the process.

The remainder of this chapter introduces the concepts of recovery and restart in more detail, and tells you how to recover if problems occur. It covers the following topics:

- “Making sure that messages are not lost (logging)”
- “Checkpointing – ensuring complete recovery” on page 174
- “Managing logs” on page 194
- “Using the log for recovery” on page 197
- “Backup and restore” on page 200
- “Recovery scenarios” on page 201

Making sure that messages are not lost (logging)

MQSeries records all significant changes to the data controlled by the queue manager in a log. This includes the creation and deletion of objects (except channels), all persistent message updates, transaction states, changes to object attributes, and channel activities. Therefore, the log contains the information you need to recover all updates to message queues by:

- Keeping records of queue manager changes.
- Keeping records of queue updates for use by the restart process.
- Enabling you to restore data after a hardware or software failure.

What logs look like

An MQSeries log consists of two components:

1. One or more files of log data
2. A log control file

There are a number of log files that contain the data being recorded. You can define the number and size (as explained in Chapter 7, “Configuration files” on page 99), or take the system default of 3 files.

Logging

In MQSeries for UNIX systems, each of the three files defaults to 4 MB. In MQSeries for OS/2 Warp and Windows NT, each of the three files defaults to 1 MB.

When you create a queue manager, the number of log files you define is the number of *primary* log files allocated. If you do not specify a number, the default value is used.

In MQSeries for UNIX systems, if you have not changed the log path, log files are created in the directory:

```
/var/mqm/1og/QmName
```

In MQSeries for OS/2 Warp and Windows NT, if you have not changed the log path, log files are created in the directory:

```
C:\MQM\LOG\<QMgrName>
```

MQSeries starts with these primary log files, but, if the log starts to get full, allocates *secondary* log files. It does this dynamically, and removes them when the demand for log space reduces. By default, up to 2 secondary log files can be allocated. This default allocation can also be changed, as described in Chapter 7, “Configuration files” on page 99.

The log control file contains the information needed to monitor the use of log files, such as their size and location, the name of the next available file, and so on.

Note: You should ensure that the logs created when you start a queue manager are large enough to accommodate the size and volume of messages that your applications will handle. The default log numbers and sizes are likely to require modification to meet your requirements. How to change the default values is described in “Configuring the logs” on page 108.

Types of logging

In MQSeries, the number of files that are required for logging depends on the file size, the number of messages you have received, and the length of the messages. There are two ways of maintaining records of queue manager activities: circular logging and linear logging.

Circular logging

Use circular logging if all you want is restart recovery, using the log to roll back transactions that were in progress when the system stopped.

Circular logging keeps all restart data in a ring of log files. Logging fills the first file in the ring, then moves on to the next, and so on, until all the files are filled. It then goes back to the first file in the ring and starts again. This continues as long as the product is in use, and has the advantage that you never run out of log files.

The above is a simple explanation of circular logging. However, there is a complication. The log entries required to restart the queue manager without loss of data are kept until they are no longer required to ensure queue manager data recovery. The mechanism for releasing log files for reuse is described in “Checkpointing – ensuring complete recovery” on page 174. For now, you should know that MQSeries uses secondary log files to extend the log capacity as necessary.

Linear logging

Use linear logging if you want both restart recovery and media or forward recovery (recreating lost or damaged data by replaying the contents of the log).

Linear logging keeps the log data in a continuous sequence of files. Space is not reused, so you can always retrieve any record logged from the time that the queue manager was created.

As disk space is finite, you may have to think about some form of archiving. It is an administrative task to manage your disk space for the log, reusing or extending the existing space as necessary.

The number of log files used with linear logging can be very large, depending on your message flow and the age of your queue manager. However, there are a number of files that are said to be active. Active files contain the log entries required to restart the queue manager. The number of active log files is usually the same as the number of primary log files as defined in the configuration files. (See Chapter 7, “Configuration files” on page 99 for further details of how to define the number.)

The key event that controls whether a log file is termed active or not is a *checkpoint*. An MQSeries checkpoint is a group of log records containing information to allow a successful restart of the queue manager. Any information recorded previously is not required to restart the queue manager and can therefore be termed inactive. (See “Checkpointing – ensuring complete recovery” on page 174 for further information about checkpointing.)

You must decide when inactive log files are no longer required. You can archive them, or you can delete them if they are no longer of interest to your operation. Refer to “Managing logs” on page 194 for further information about the disposition of log files.

If a new checkpoint is recorded in the second, or later, primary log file, then the first file becomes inactive and a new primary file is formatted and added to the end of the primary pool, restoring the number of primary files available for logging. In this way the primary log file pool can be seen to be a current set of files in an ever extending list of log files. Again, it is an administrative task to manage the inactive files according to the requirements of your operation.

Although secondary log files are defined for linear logging, they are not used in normal operation. If a situation should arise when, probably due to long-lived transactions, it is not possible to free a file from the active pool because it may still be required for a restart, secondary files are formatted and added to the active log file pool.

If the number of secondary files available is used up, requests for most further operations requiring log activity will be refused with an MQRC_RESOURCE_PROBLEM being returned to the application.

Both types of logging can cope with unexpected loss of power assuming that there is no hardware failure.

Checkpointing – ensuring complete recovery

Persistent updates to message queues happen in two stages. First, the records representing the update are written to the log, then the queue file is updated. The log files can thus become more up-to-date than the queue files. To ensure that restart processing begins from a consistent point, MQSeries uses checkpoints. A checkpoint is a point in time when the record described in the log is the same as the record in the queue. The checkpoint itself consists of the series of log records needed to restart the queue manager; for example, the state of all transactions active at the time of the checkpoint.

Checkpoints are generated automatically by MQSeries. They are taken when the queue manager starts, at shutdown, when logging space is running low, and after every 1000 operations logged. As the queues handle further messages, the checkpoint record becomes inconsistent with the current state of the queues.

When MQSeries is restarted, it locates the latest checkpoint record in the log. This information is held in the checkpoint file that is updated at the end of every checkpoint. The checkpoint record represents the most recent point of consistency between the log and the data. The data from this checkpoint is used to rebuild the queues as they existed at the checkpoint time. When the queues are recreated, the log is then played forward to bring the queues back to the state they were in before system failure or close down.

MQSeries maintains internal pointers to the head and tail of the log. It moves the head pointer to the most recent checkpoint that is consistent with recovering message data.

Checkpoints are used to make recovery more efficient, and to control the reuse of primary and secondary log files.




Figure 38. Checkpointing. For simplicity, only the ends of the log files are shown.

In Figure 38, all records before the latest checkpoint, checkpoint 2, are no longer needed by MQSeries. The queues can be recovered from the checkpoint information and any later log entries. For circular logging, any freed files prior to the checkpoint can be reused. For a linear log, the freed log files no longer need to be accessed for normal operation and become inactive. In the example, the queue head pointer is moved to point at the latest checkpoint, checkpoint 2, which then becomes the new queue head, head 2. Log File 1 can now be reused.

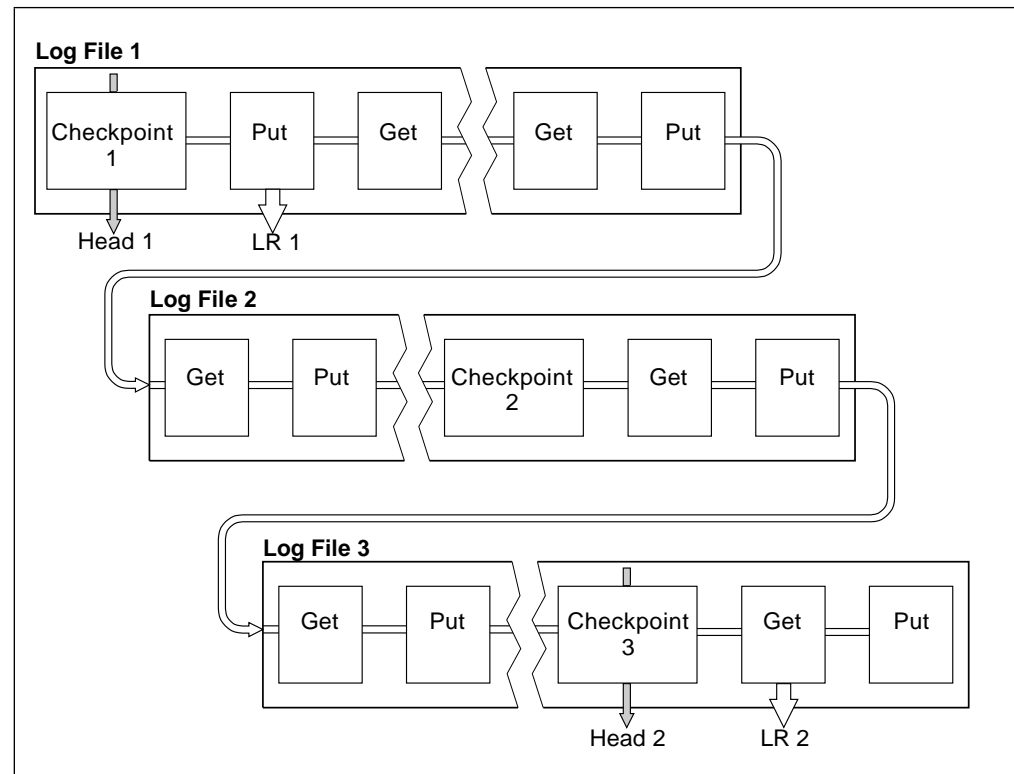


Figure 39. Checkpointing with a long-running transaction. For simplicity, only the ends of the log files are shown.

Figure 39 shows how a long-running transaction affects reuse of log files. In the example, a long-running transaction has caused an entry to the log, shown as LR 1, after the first checkpoint shown. The transaction does not complete, shown as LR 2, until after the third checkpoint. All the log information from LR 1 onwards is retained to allow recovery of that transaction, if necessary, until it has completed.

After the long-running transaction has completed, at LR 2, the head of the log is moved to checkpoint 3, the latest logged checkpoint. The files containing log records prior to checkpoint 3, Head 2, are no longer needed. If you are using circular logging, the space can be reused.

If the primary log files are completely filled before the long-running transaction completes, secondary log files are used to avoid the risk of a log full situation if possible.

When the log head is moved and you are using circular logging, the primary log files may become eligible for reuse and the logger, after filling the current file, reuses the first primary file available to it. If instead you are using linear logging, the log head is still moved down the active pool and the first file becomes inactive.

A new primary file is formatted and added to the bottom of the pool in readiness for future logging activities.

Dumping the contents of the log using the dmpmqlog command

The **dmpmqlog** command can be used to dump the contents of the queue manager log. By default all active log records are dumped, that is, the command starts dumping from the head of the log. Normally this is from the start of the last completed checkpoint.

The log can be dumped only when the queue manager is not running. Because the queue manager takes a checkpoint during shutdown, the active portion of the log usually contains a small number of log records. However, the **dmpmqlog** command can be instructed to dump more log records using one of the following options to change the start position of the dump:

- The simplest option is to start dumping from the *base* of the log. The base of the log is the first log record in the log file that contains the head of the log. The amount of additional data dumped in this case depends upon where the head of the log is positioned in the log file. If it is near to the start of the log file only a small amount of additional data is dumped. If the head is near to the end of the log file then significantly more data is dumped.
- Another option enables the start position of the dump to be specified as an individual log record. Each log record is identified by a unique *log sequence number (LSN)*. In the case of circular logging, this starting log record cannot be prior to the base of the log; this restriction does not apply to linear logs. Inactive log files may need to be reinstated before running the command. For this option a valid LSN must be specified as the start position. This must be taken from previous **dmpmqlog** output. For example, with linear logging you could specify the `nextlsn` from your last **dmpmqlog** output. The Next LSN appears in Log File Header and indicates the LSN of the next log record to be written. This can therefore be used as a start position to format all log records that have been written since the last time the log was dumped.
- The third option is for linear logs only. The dumper can be instructed to start formatting log records from any given log file extent. In this case the log dumper expects to find this log file, and each successive one, in the same directory as the active log files. This option does not apply to circular logs, because in this case the log dumper cannot access log records prior to the base of the log.

The output from the **dmpmqlog** command is the Log File Header and a series of formatted log records. The queue manager uses several log records to record changes to its data.

Some of the information that is formatted is of use only internally. The following list includes the most useful log records:

Log File Header

Each log has a single log file header, which is always the first thing formatted by the **dmpmqlog** command. It contains the following fields:

- | | |
|--------------------|--------------------------------------|
| <i>logactive</i> | The number of primary log extents. |
| <i>loginactive</i> | The number of secondary log extents. |

<i>logsize</i>	The number of 4 KB pages per extent.
<i>baselsn</i>	The first LSN in the log extent containing the head of the log.
<i>nextlsn</i>	The LSN of next log record to be written.
<i>headlsn</i>	The LSN of the log record at the head of the log.
<i>tailsn</i>	The LSN identifying the tail position of the log.
<i>hflag1</i>	Identifies whether log is CIRCULAR or LOG RETAIN (linear).
<i>HeadExtentID</i>	The log extent containing the head of the log.

Log Record Header

Each log record within the log has a fixed header containing the following information:

<i>LSN</i>	The log sequence number.
<i>LogRecdType</i>	The type of the log record.
<i>XTranid</i>	The transaction identifier associated with this log record (if any). A <i>TranType</i> of MQI indicates an MQ-only transaction. A <i>TranType</i> of XA is involved with other resource managers. Updates involved within the same unit of work have the same <i>XTranid</i> .
<i>QueueName</i>	The queue associated with this log record (if any).
<i>Qid</i>	The unique internal identifier for the queue.
<i>PrevLSN</i>	LSN of previous log record within the same transaction (if any).

Start Queue Manager

This logs that the queue manager has been started.

<i>StartDate</i>	The date that the queue manager was started.
<i>StartTime</i>	The time that the queue manager was started.

Stop Queue Manager

This logs that the queue manager has been stopped.

<i>StopDate</i>	The date that the queue manager was stopped.
<i>StopTime</i>	The time that the queue manager was stopped.
<i>ForceFlag</i>	The type of shutdown that was used.

Start Checkpoint

This denotes the start of a queue manager checkpoint.

End Checkpoint

This denotes the end of a queue manager checkpoint.

<i>ChkPtLSN</i>	The LSN of the log record that started this checkpoint.
-----------------	---

Put Message

This logs a persistent message put to a queue. If the message was put under syncpoint, then the log record header contains a nonnull *XTranid*. The remainder of the record contains:

<i>SpcIndex</i>	An identifier for the message on the queue. It can be used to match the corresponding MQGET that was used to get this message from the queue. In this case a subsequent <i>Get Message</i> log record can be found containing the same <i>QueueName</i> and <i>SpcIndex</i> . At this point the <i>SpcIndex</i> identifier can be reused for a subsequent put message to that queue.
<i>Data</i>	Contained in the hex dump for this log record is various internal data followed by the Message Descriptor (eyecatcher MD) and the message data itself.

Put Part Persistent messages that are too large for a single log record are logged as a single *Put Message* record followed by multiple *Put Part* log records.

<i>Data</i>	Continues the message data where the previous log record left off.
-------------	--

Get Message

Only gets of persistent messages are logged. If the message was got under syncpoint then the log record header contains a nonnull *XTranid*. The remainder of the record contains:

<i>SpcIndex</i>	Identifies the message that was got from the queue. The most recent <i>Put Message</i> log record containing the same <i>QueueName</i> and <i>SpcIndex</i> identifies the message that was got.
<i>QPriority</i>	The priority of the message got from the queue.

Start Transaction

Indicates the start of a new transaction. A *TranType* of MQI indicates an MQ-only transaction. A *TranType* of XA indicates one that involves other resource managers. All updates made by this transaction will have the same *XTranid*.

Prepare Transaction

Indicates that the queue manager is prepared to commit the updates associated with the specified *XTranid*. This log record is written as part of a two-phase commit involving other resource managers.

Commit Transaction

Indicates that the queue manager has committed all updates made by a transaction.

Rollback Transaction

This log record denotes the queue manager's intention to roll back a transaction.

End Transaction

This log record denotes the end of a rolled-back transaction.

Transaction Table

This record is written during syncpoint. It records the state of each transaction that has made persistent updates. For each transaction the following information is recorded:

<i>XTranid</i>	Transaction identifier.
<i>FirstLSN</i>	LSN of first log record associated with transaction.
<i>LastLSN</i>	LSN of last log record associated with transaction.

Transaction Participants

This log record is written by the XA Transaction Manager component of the queue manager. It records the external resource managers that are participating in transactions. For each participant the following is recorded:

<i>RMName</i>	The name of the resource manager.
<i>RMIId</i>	Resource manager identifier. This is also logged in subsequent <i>Transaction Prepared</i> log records which record global transactions in which the resource manager is participating.
<i>SwitchFile</i>	The switch load file for this resource manager.
<i>XAOpenString</i>	The XA open string for this resource manager.
<i>XACloseString</i>	The XA open string for this resource manager.

Transaction Prepared

This log record is written by the XA Transaction Manager component of the queue manager. It indicates that the specified global transaction has been successfully prepared. Each of the participating resource managers will be instructed to commit. The *RMIId* of each prepared resource manager is recorded in the log record. If the queue manager itself is participating in the transaction a *Participant Entry* with an *RMIId* of zero will be present.

Transaction Forget

This log record is written by the XA Transaction Manager component of the queue manager. It follows the *Transaction Prepared* log record when the commit decision has been delivered to each participant.

Purge Queue

This logs the fact that all messages on a queue have been purged, for example, using the RUNMQSC CLEAR command.

Queue Attributes

This logs the initialization or change of the attributes of a queue

Create Object

Logs the creation of an MQSeries object

<i>ObjName</i>	The name of the object that was created.
<i>UserId</i>	The user ID performing the creation.

Delete Object

Logs the deletion of an MQSeries object

<i>ObjName</i>	The name of the object that was deleted.
----------------	--

Using dmpmqlog

Figure 40 shows example output from a **dmpmqlog** command. The dump, which started at the LSN of a specific log record, was produced using the following command:

```
dmpmqlog -mtestqm -s0:0:0:44162
```

```
AMQ7701: DMPMQLOG command is starting.
LOG FILE HEADER
*****

counter1 . . . . : 23           counter2 . . . . : 23
FormatVersion . . : 2           logtype . . . . : 10
logactive . . . . : 3           loginactive . . . : 2
logsize . . . . . : 1024        pages
baselsn . . . . . : <0:0:0:0>
nextlsn . . . . . : <0:0:0:60864>
lowtranlsn . . . . : <0:0:0:0>
minbufflsn . . . . : <0:0:0:58120>
headlsn . . . . . : <0:0:0:58120>
taillsn . . . . . : <0:0:0:60863>
logfilepath . . . : ""
hflag1 . . . . . : 1
                -> CONSISTENT
                -> CIRCULAR
HeadExtentID . . : 1           LastEID . . . . . : 846249092
LogId . . . . . : 846249061    LastCommit . . . : 0
FirstArchNum . . : 4294967295  LastArchNum . . . : 4294967295
nextArcFile . . . : 4294967295  firstRecFile . . : 4294967295
firstDlteFile . . : 4294967295  lastDeleteFile . : 4294967295
RecHeadFile . . . : 4294967295  FileCount . . . . : 3
frec_trunclsn . . : <0:0:0:0>
frec_readlsn . . . : <0:0:0:0>
frec_extnum . . . : 0           LastCid . . . . . : 0
onlineBkupEnd . . : 0           softmax . . . . . : 4194304

LOG RECORD - LSN <0:0:0:44162>
*****

HLG Header: lreclsize 212, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ALM Start Checkpoint (1025)
Eyecatcher . . . . : ALRH           Version . . . . . : 1
LogRecdLen . . . . : 192           LogRecdOwnr . . . : 1024   (ALM)
XTranid . . . . . : TranType: NULL
QueueName . . . . : NULL
Qid . . . . . . . : {NULL_QID}
ThisLSN . . . . . : <0:0:0:0>
PrevLSN . . . . . : <0:0:0:0>

No data for Start Checkpoint Record
```

Figure 40 (Part 1 of 13). Example dmpmqlog output

```

LOG RECORD - LSN <0:0:0:44374>
*****

HLG Header: lreclsize 220, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Transaction Table (773)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 200                      LogRecdOwnr . . : 768   (ATM)
XTranid . . . . : TranType: NULL
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 1
TranCount . . . : 0

LOG RECORD - LSN <0:0:0:44594>
*****

HLG Header: lreclsize 1836, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : Transaction Participants (1537)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 1816                    LogRecdOwnr . . : 1536  (T)
XTranid . . . . : TranType: NULL
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Id. . . . . : TLPH
Version . . . . : 1                        Flags . . . . . : 3
Count . . . . . : 2

Participant Entry 0
RMName . . . . : DB2 MQBankDB
RMId . . . . . : 1
SwitchFile . . : /Development/sbolam/build/devlib/tstxasw
XAOpenString . . :
XACloseString . . :

Participant Entry 1
RMName . . . . : DB2 MQBankDB
RMId . . . . . : 2
SwitchFile . . : /Development/sbolam/build/devlib/tstxasw
XAOpenString . . :
XACloseString . . :

```

Figure 40 (Part 2 of 13). Example dmpmqlog output

Using dmpmqlog

```
LOG RECORD - LSN <0:0:0:46448>
*****

HLG Header: lreclsize 236, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ALM End Checkpoint (1026)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 216                      LogRecdOwnr . . : 1024 (ALM)
XTranid . . . . : TranType: NULL
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

ChkPtLSN . . . . : <0:0:0:44162>
OldestLSN . . . . : <0:0:0:0>
MediaLSN . . . . : <0:0:0:0>

LOG RECORD - LSN <0:0:0:52262>
*****

HLG Header: lreclsize 220, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Start Transaction (769)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 200                      LogRecdOwnr . . : 768 (ATM)
XTranid . . . . : TranType: MQI    TranNum{High 0, Low 1}
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 1
SoftLogLimit . . : 10000
```

Figure 40 (Part 3 of 13). Example dmpmqlog output


```

LOG RECORD - LSN <0:0:0:52482>
*****

HLG Header: lreclsize 730, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : AQM Put Message (257)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 710                      LogRecdOwnr . . : 256      (AQM)
XTranid . . . . : TranType: MQI      TranNum{High 0, Low 1}
QueueName . . . : Queue1
Qid . . . . . : {Hash 196836031, Counter: 0}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:52262>

Version . . . . . : 3
SpIndex . . . . . : 1
PrevLink.Locn . . : 36                      PrevLink.Length : 8
PrevDataLink . . : {High 0, Low 2048}
Data.Locn . . . . : 2048                    Data.Length . . : 486
Data . . . . . :
00000: 41 51 52 48 00 00 00 04 FF FF FF FF FF FF FF FF  AQRH.....
00016: 00 00 00 00 00 00 00 00 00 00 00 01 00 01 01 C0  .....i
00032: 00 00 00 00 00 00 00 01 00 00 00 22 00 00 00 00  .....".
00048: 00 00 00 00 41 4D 51 20 74 65 73 74 71 6D 20 20  ....AMQ testqm
00064: 20 20 20 20 33 80 2D D2 00 00 10 13 00 00 00 00  .....
00080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00096: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00112: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 01  .....
00128: 00 00 00 00 00 00 00 22 00 00 00 00 00 00 00 00  .....".
00144: 00 00 00 00 00 00 00 C9 2C B5 C0 25 FF FF FF FF  ....., [i%....
00160: 4D 44 20 20 00 00 00 01 00 00 00 00 00 00 00 08  MD .....
00176: 00 00 00 00 00 00 01 11 00 00 03 33 20 20 20 20  .....3
00192: 20 20 20 20 00 00 00 00 00 00 00 01 20 20 20 20  .....
00208: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00224: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00240: 20 20 20 20 20 20 20 20 20 20 20 20 74 65 73 74  ..... test
00256: 71 6D 20 20 20 20 20 20 20 20 20 20 20 20 20 20  qm .....
00272: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00288: 20 20 20 20 20 20 20 20 20 20 20 20 73 62 6F 6C  ..... sbol
00304: 61 6D 20 20 20 20 20 20 04 37 34 38 30 00 00 00  am .....7480...
00320: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00336: 00 00 00 00 00 00 00 00 20 20 20 20 20 20 20 20  .....
00352: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00368: 20 20 20 20 20 20 20 20 00 00 00 06 75 74 7A 61  .....utza
00384: 70 69 20 20 20 20 20 20 20 20 20 20 20 20 20 20  pi .....
00400: 20 20 20 20 20 20 20 20 31 39 39 37 30 35 31 39  ..... 19970519
00416: 31 30 34 32 31 35 32 30 20 20 20 20 00 00 00 00  10421520 .....
00432: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00448: 50 65 72 73 69 73 74 65 6E 74 20 6D 65 73 73 61  Persistent messa
00464: 67 65 20 70 75 74 20 75 6E 64 65 72 20 73 79 6E  ge put under syn
00480: 63 70 6F 69 6E 74  cpoint

```

Figure 40 (Part 4 of 13). Example dmpmqlog output

Using dmpmqlog

```

LOG RECORD - LSN <0:0:0:53458>
*****

HLG Header: lreclsize 734, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : AQM Put Message (257)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 714                      LogRecdOwnr . . : 256    (AQM)
XTranid . . . . : TranType: NULL
QueueName . . . . : Queue2
Qid . . . . . : {Hash 184842943, Counter: 2}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 3
SpclIndex . . . : 1
PrevLink.Locn . : 36                      PrevLink.Length : 8
PrevDataLink . . : {High 0, Low 2048}
Data.Locn . . . : 2048                   Data.Length . . : 490
Data . . . . . :
00000: 41 51 52 48 00 00 00 04 FF FF FF FF FF FF FF FF   AQRH.....
00016: 00 00 00 00 00 00 00 00 00 00 00 01 00 01 01 C0   .....i
00032: 00 00 00 00 00 00 00 01 00 00 00 26 00 00 00 00   .....&....
00048: 00 00 00 00 41 4D 51 20 74 65 73 74 71 6D 20 20   ....AMQ testqm
00064: 20 20 20 20 33 80 2D D2 00 00 10 13 00 00 00 00   3ä-.....
00080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   .....
00096: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   .....
00112: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 01   .....
00128: 00 00 00 00 00 00 00 26 00 00 00 00 00 00 00 00   .....&.....
00144: 00 00 00 00 00 00 00 C9 2C B6 D8 DD FF FF FF FF   .....,”....
00160: 4D 44 20 20 00 00 00 01 00 00 00 00 00 00 00 08   MD .....
00176: 00 00 00 00 00 00 01 11 00 00 03 33 20 20 20 20   .....3
00192: 20 20 20 20 00 00 00 00 00 00 00 01 20 20 20 20   .....
00208: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20
00224: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20
00240: 20 20 20 20 20 20 20 20 20 20 20 20 74 65 73 74   test
00256: 71 6D 20 20 20 20 20 20 20 20 20 20 20 20 20 20   qm
00272: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20
00288: 20 20 20 20 20 20 20 20 20 20 20 20 73 62 6F 6C   sbol
00304: 61 6D 20 20 20 20 20 04 37 34 38 30 00 00 00 00   am .7480...
00320: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   .....
00336: 00 00 00 00 00 00 00 00 20 20 20 20 20 20 20 20   .....
00352: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20
00368: 20 20 20 20 20 20 20 20 00 00 00 06 75 74 7A 61   ....utza
00384: 70 69 20 20 20 20 20 20 20 20 20 20 20 20 20 20   pi
00400: 20 20 20 20 20 20 20 20 31 39 39 37 30 35 31 39   19970519
00416: 31 30 34 33 32 37 30 36 20 20 20 20 00 00 00 00   10432706 ....
00432: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   .....
00448: 50 65 72 73 69 73 74 65 6E 74 20 6D 65 73 73 61   Persistent messa
00464: 67 65 20 6E 6F 74 20 70 75 74 20 75 6E 64 65 72   ge not put under
00480: 20 73 79 6E 63 70 6F 69 6E 74   syncpoint

```

Figure 40 (Part 5 of 13). Example dmpmqlog output

```

LOG RECORD - LSN <0:0:0:54192>
*****

HLG Header: lrecsize 216, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Commit Transaction (774)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 196                      LogRecdOwnr . . : 768    (ATM)
XTranid . . . . : TranType: MQI    TranNum{High 0, Low 1}
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:52482>

Version . . . . : 1
LOG RECORD - LSN <0:0:0:54408>
*****

HLG Header: lrecsize 220, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Start Transaction (769)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 200                      LogRecdOwnr . . : 768    (ATM)
XTranid . . . . : TranType: MQI    TranNum{High 0, Low 3}
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 1
SoftLogLimit . . : 10000

LOG RECORD - LSN <0:0:0:54628>
*****

HLG Header: lrecsize 240, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : AQM Get Message (259)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 220                      LogRecdOwnr . . : 256    (AQM)
XTranid . . . . : TranType: MQI    TranNum{High 0, Low 3}
QueueName . . . : Queue1
Qid . . . . . : {Hash 196836031, Counter: 0}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:54408>

Version . . . . : 2
SpcIndex . . . . : 1                      QPriority . . . . : 0
PrevLink.Locn . . : 36                    PrevLink.Length : 8
PrevDataLink . . : {High 4294967295, Low 4294967295}

```

Figure 40 (Part 6 of 13). Example dmpmqlog output

Using dmpmqlog

```
LOG RECORD - LSN <0:0:0:54868>
*****

HLG Header: lreclsize 240, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : AQM Get Message (259)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 220                      LogRecdOwnr . . . : 256    (AQM)
XTranid . . . . : TranType: NULL
QueueName . . . . : Queue2
Qid . . . . . : {Hash 184842943, Counter: 2}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 2
Spcln . . . . : 1                          QPriority . . . . : 0
PrevLink.Locn . . : 36                      PrevLink.Length : 8
PrevDataLink . . : {High 4294967295, Low 4294967295}
LOG RECORD - LSN <0:0:0:55108>
*****

HLG Header: lreclsize 216, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Commit Transaction (774)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 196                      LogRecdOwnr . . . : 768    (ATM)
XTranid . . . . : TranType: MQI    TranNum{High 0, Low 3}
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:54628>

Version . . . . : 1

LOG RECORD - LSN <0:0:0:55324>
*****

HLG Header: lreclsize 220, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Start Transaction (769)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 200                      LogRecdOwnr . . . : 768    (ATM)
XTranid . . . . : TranType: XA
    XID: formatID 5067085, gtrid_length 14, bqual_length 4
        gtrid [3270BDB40000102374657374716D]
        bqual [00000001]
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 1
SoftLogLimit . . : 10000
```

Figure 40 (Part 7 of 13). Example dmpmqlog output

```

LOG RECORD - LSN <0:0:0:55544>
*****

HLG Header: lreclsize 738, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : AQM Put Message (257)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 718                      LogRecdOwnr . . : 256    (AQM)
XTranid . . . . : TranType: XA
  XID: formatID 5067085, gtrid_length 14, bqual_length 4
      gtrid [3270BDB40000102374657374716D]
      bqual [00000001]
QueueName . . . . : Queue2
Qid . . . . . : {Hash 184842943, Counter: 2}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:55324>

Version . . . . : 3
SpIndex . . . . : 1
PrevLink.Locn . . : 36                      PrevLink.Length : 8
PrevDataLink . . : {High 0, Low 2048}
Data.Locn . . . . : 2048                    Data.Length . . : 494
Data . . . . . :
00000: 41 51 52 48 00 00 00 04 FF FF FF FF FF FF FF FF  AQRH.....
00016: 00 00 00 00 00 00 00 00 00 00 00 01 00 01 01 C0  .....i
00032: 00 00 00 00 00 00 00 01 00 00 00 2A 00 00 00 00  .....*....
00048: 00 00 00 01 41 4D 51 20 74 65 73 74 71 6D 20 20  ....AMQ testqm
00064: 20 20 20 20 33 80 2D D2 00 00 10 13 00 00 00 00  .....
00080: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00096: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00112: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 01  .....
00128: 00 00 00 00 00 00 00 2A 00 00 00 00 00 00 00 00  .....*.....
00144: 00 00 00 00 00 00 00 C9 2C B8 3E E8 FF FF FF FF  .....,fl>....
00160: 4D 44 20 20 00 00 00 01 00 00 00 00 00 00 00 08  MD .....
00176: 00 00 00 00 00 00 01 11 00 00 03 33 20 20 20 20  .....3
00192: 20 20 20 20 00 00 00 00 00 00 00 01 20 20 20 20  .....
00208: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00224: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00240: 20 20 20 20 20 20 20 20 20 20 20 20 74 65 73 74  ..... test
00256: 71 6D 20 20 20 20 20 20 20 20 20 20 20 20 20 20  qm
00272: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00288: 20 20 20 20 20 20 20 20 20 20 20 20 73 62 6F 6C  ..... sbol
00304: 61 6D 20 20 20 20 20 04 37 34 38 30 00 00 00 00  am      .7480...
00320: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00336: 00 00 00 00 00 00 00 20 20 20 20 20 20 20 20 20  .....
00352: 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20 20  .....
00368: 20 20 20 20 20 20 20 20 00 00 00 06 75 74 7A 61  ..... utza
00384: 70 69 20 20 20 20 20 20 20 20 20 20 20 20 20 20  pi
00400: 20 20 20 20 20 20 20 20 31 39 39 37 30 35 31 39  ..... 19970519
00416: 31 30 34 34 35 38 37 32 20 20 20 20 00 00 00 00  10445872 ....
00432: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  .....
00448: 41 6E 6F 74 68 65 72 20 70 65 72 73 69 73 74 65  Another persiste
00464: 6E 74 20 6D 65 73 73 61 67 65 20 70 75 74 20 75  nt message put u
00480: 6E 64 65 72 20 73 79 6E 63 70 6F 69 6E 74      nder syncpoint

```

Figure 40 (Part 8 of 13). Example dmpmqlog output

Using dmpmqlog

```
LOG RECORD - LSN <0:0:0:56282>
*****

HLG Header: lreclsize 216, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Prepare Transaction (770)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 196                      LogRecdOwnr . . : 768   (ATM)
XTranid . . . . : TranType: XA
  XID: formatID 5067085, gtrid_length 14, bqual_length 4
      gtrid [3270BDB40000102374657374716D]
      bqual [00000001]
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:55544>

Version . . . . : 1

LOG RECORD - LSN <0:0:0:56498>
*****

HLG Header: lreclsize 708, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : Transaction Prepared (1538)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 688                      LogRecdOwnr . . : 1536 (T)
XTranid . . . . : TranType: XA
  XID: formatID 5067085, gtrid_length 14, bqual_length 4
      gtrid [3270BDB40000102374657374716D]
      bqual [00000001]
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Id. . . . . : TLPR
Version . . . . : 1                          Flags . . . . . : 1
Count . . . . . : 3

Participant Entry 0
RMId . . . . . : 0                          State . . . . . : 2

Participant Entry 1
RMId . . . . . : 1                          State . . . . . : 2

Participant Entry 2
RMId . . . . . : 2                          State . . . . . : 2
```

Figure 40 (Part 9 of 13). Example dmpmqlog output

```

LOG RECORD - LSN <0:0:0:57206>
*****

HLG Header: lrecsize 216, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Commit Transaction (774)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 196                      LogRecdOwnr . . . : 768   (ATM)
XTranid . . . . : TranType: XA
  XID: formatID 5067085, gtrid_length 14, bqual_length 4
      gtrid [3270BDB40000102374657374716D]
      bqual [00000001]
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:56282>

Version . . . . . : 1
LOG RECORD - LSN <0:0:0:57440>
*****

HLG Header: lrecsize 224, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : Transaction Forget (1539)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 204                      LogRecdOwnr . . . : 1536 (T)
XTranid . . . . : TranType: XA
  XID: formatID 5067085, gtrid_length 14, bqual_length 4
      gtrid [3270BDB40000102374657374716D]
      bqual [00000001]
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Id. . . . . : TLFG
Version . . . . : 1                          Flags . . . . . : 0

```

Figure 40 (Part 10 of 13). Example dmpmqlog output

Using dmpmqlog

```
LOG RECORD - LSN <0:0:0:58120>
*****

HLG Header: lreclsize 212, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ALM Start Checkpoint (1025)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 192                      LogRecdOwnr . . : 1024 (ALM)
XTranid . . . . : TranType: NULL
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

No data for Start Checkpoint Record

LOG RECORD - LSN <0:0:0:58332>
*****

HLG Header: lreclsize 220, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ATM Transaction Table (773)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 200                      LogRecdOwnr . . : 768 (ATM)
XTranid . . . . : TranType: NULL
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 1
TranCount . . . : 0
```

Figure 40 (Part 11 of 13). Example dmpmqlog output


```

LOG RECORD - LSN <0:0:0:58552>
*****

HLG Header: lreclsize 1836, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : Transaction Participants (1537)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 1816                     LogRecdOwnr . . : 1536 (T)
XTranid . . . . : TranType: NULL
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Id. . . . . : TLPH
Version . . . . : 1                          Flags . . . . . : 3
Count . . . . . : 2

Participant Entry 0
RMName . . . . : DB2 MQBankDB
RMId . . . . . : 1
SwitchFile . . . : /Development/sbolam/build/devlib/tstxasw
XAOpenString . . :
XACloseString . . :

Participant Entry 1
RMName . . . . : DB2 MQFeeDB
RMId . . . . . : 2
SwitchFile . . . : /Development/sbolam/build/devlib/tstxasw
XAOpenString . . :
XACloseString . . :

LOG RECORD - LSN <0:0:0:60388>
*****

HLG Header: lreclsize 236, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ALM End Checkpoint (1026)
Eyecatcher . . . : ALRH                      Version . . . . : 1
LogRecdLen . . . : 216                     LogRecdOwnr . . : 1024 (ALM)
XTranid . . . . : TranType: NULL
QueueName . . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

ChkPtLSN . . . . : <0:0:0:58120>
OldestLSN . . . . : <0:0:0:0>
MediaLSN . . . . : <0:0:0:0>

```

Figure 40 (Part 12 of 13). Example dmpmqlog output

Using dmpmqlog

```
LOG RECORD - LSN <0:0:0:60624>
*****

HLG Header: lreclsize 240, version 1, rmid 0, eyecatcher HLRH

LogRecdType . . . : ALM Stop Queue Manager (1028)
Eyecatcher . . . : ALRH                               Version . . . . : 1
LogRecdLen . . . : 220                               LogRecdOwnr . . : 1024 (ALM)
XTranid . . . . : TranType: NULL
QueueName . . . : NULL
Qid . . . . . : {NULL_QID}
ThisLSN . . . . : <0:0:0:0>
PrevLSN . . . . : <0:0:0:0>

Version . . . . : 1
StopDate . . . : 19970519                            StopTime . . . : 10490868
SessionNumber . : 0                                  ForceFlag . . . : Quiesce

AMQ7702: DMPMQLOG command has finished successfully.
```

Figure 40 (Part 13 of 13). Example dmpmqlog output

Notes:

1. The *headlsn* in the *Log File Header* has a value of <0:0:0:58120>. This is where the dump would have started had we not requested a different starting LSN.
2. The *nextlsn* is <0:0:0:60864> which will be the LSN of the first log record that the queue manager will write when it is next restarted.
3. The *HeadExtentID* is 1, indicating that the head of the log currently resides in log file S0000001.LOG.
4. The first log record formatted is a *Start Checkpoint* log record. The checkpoint spans a number of log records until the *End CheckPoint* record at <0:0:0:46448>.
5. One of the records logged during checkpoint is the *Transaction Participants* log record at <0:0:0:44594>. This details the resource managers that participate in global transactions coordinated by the queue manager.
6. The *Start Transaction* log record at <0:0:0:52262> denotes the start of a transaction. The *XTranid* shows a *TranType* of MQI, which indicates that it is a local transaction including MQSeries updates only.
7. The next log record is a *Put Message* log record that records the persistent MQPUT under the syncpoint that started the transaction. The MQPUT was made to the queue *Queue1* and the message data is logged as *Persistent message put* under syncpoint. This message has been allocated a *SpcIndex* of 1, which will be matched to the later MQGET of this message.
8. The next log record at LSN <0:0:0:53458> is also a *Put Message* record. This persistent message was put to a different queue, *Queue2*, but was not made under syncpoint since the *XTranid* is *NULL*. It too has a *SpcIndex* of 1, which is a unique identifier for this particular queue.
9. The next log record at LSN <0:0:0:54192> commits the message that was put under syncpoint.

10. In log records <0:0:0:54408> and <0:0:0:54628> a new transaction is started by an MQGET under syncpoint for queue *Queue1*. The *SpcIndex* in the *Get Message* log record is 1 indicating that this was the same message that was put to *Queue1* in <0:0:0:52262>.
11. The next log record gets the message that was put to *Queue2* by the other *Put Message* log record.
12. The MQGET under syncpoint has been committed as indicated by the *Commit Transaction* log record at <0:0:0:55108>.
13. Finally an MQBEGIN is used to start a global transaction in the *Start Transaction* log record at <0:0:0:55324>. The *XTranid* in this log record has a *TranType* of XA.
14. The following *Put Message* records a persistent message put to *Queue2*. This shares the same *XTranid* as the previous log record.
15. If a *Transaction Prepared* log record is written for this *Xtranid* then the transaction as a whole must be committed. The absence of such a log record can be taken as an indication that the transaction was rolled back. In this case a *Transaction Prepared* log record is found at <0:0:0:56498>. This records the queue manager itself as a participant with an *RMIId* of zero. There are two further participants, their *RMIIds* of 1 and 2 can be matched with the previous *Transaction Participants* log record.
16. During the commit phase the XA Transaction Manager component of the queue manager does not log individual responses from the participants. The log indicates only whether the queue manager updates were committed or not. The *Commit Transaction* log record at <0:0:0:57206> indicates that the message was indeed committed to *Queue2*.
17. The *Transaction Forget* log record at <0:0:0:57440> indicates that the commit decision was also delivered to the other two resource managers. Any failure of these resource managers to commit their updates will have been diagnosed in the queue manager's error logs.

Managing logs

Over time, some of the log records written become unnecessary for restarting the queue manager. If you are using circular logging, the queue manager reclaims freed space in the log files. This activity is transparent to the user and you do not usually see the amount of disk space used reduce because the space allocated is quickly reused.

Of the log records, only those written since the start of the last complete checkpoint, and those written by any active transactions, are needed to restart the queue manager. Thus, the log may fill if a checkpoint has not been taken for a long time, or if a long-running transaction wrote a log record a long time ago. The queue manager tries to take checkpoints sufficiently frequently to avoid the first problem.

When a long-running transaction fills the log, attempts to write log records fail and some MQI calls return `MQRC_RESOURCE_PROBLEM`. (Space is reserved to commit or rollback all in-flight transactions, so `MQCMIT` or `MQBACK` should not fail.)

The queue manager rolls back transactions that consume too much log space. An application whose transaction is rolled back in this way is unable to perform subsequent `MQPUT` or `MQGET` operations specifying syncpoint under the same transaction. An attempt to put or get a message under syncpoint in this state returns `MQRC_BACKED_OUT`. The application may then issue `MQCMIT`, which returns `MQRC_BACKED_OUT`, or `MQBACK` and start a new transaction. When the transaction consuming too much log space has been rolled back, its log space is released and the queue manager continues to operate normally.

If the log fills, message `AMQ7463` is issued. In addition, if the log fills because a long-running transaction has prevented the space being released, message `AMQ7465` is issued.

Finally, if records are being written to the log faster than the asynchronous housekeeping processes can handle them, message `AMQ7466` is issued. If you see this message, you should increase the number of log files or reduce the amount of data being processed by the queue manager.

What happens when a disk gets full

The queue manager logging component can cope with a full disk, and with full log files. If the disk containing the log fills, the queue manager issues message `AMQ6708` and an error record is taken.

The log files are created at their maximum size, rather than being extended as log records are written to them. This means that MQSeries can run out of disk space only when it is creating a new file. Therefore, it cannot run out of space when it is writing a record to the log. MQSeries always knows how much space is available in the existing log files, and manages the space within the files accordingly.

If you fill the drive containing the log files, you may be able to free some disk space. If you are using a linear log, there may be some inactive log files in the log directory, and you can copy these files to another drive or device. If you still run out of space, check that the configuration of the log in the queue manager configuration file is correct. You may be able to reduce the number of primary or

secondary log files so that the log does not outgrow the available space. Note that it is not possible to alter the size of the log files for an existing queue manager. The queue manager assumes that all log files are the same size.

Managing log files

If you are using circular logging, ensure that there is sufficient space to hold the log files. You do this when you configure your system (see “Log configuration stanzas” on page 109). The amount of disk space used by the log does not increase beyond the configured size, including space for secondary files to be created when required.

If you are using a linear log, the log files are added continually as data is logged, and the amount of disk space used increases with time. If the rate of data being logged is high, disk space is consumed rapidly by new log files.

Over time, the older log files for a linear log are no longer required to restart the queue manager or perform media recovery of any damaged objects. Periodically, the queue manager issues a pair of messages to indicate which of the log files is required:

- Message AMQ7467 gives the name of the oldest log file needed to restart the queue manager. This log file and all newer log files must be available during queue manager restart.
- Message AMQ7468 gives the name of the oldest log file needed to do media recovery.

Any log files older than these do not need to be online. You can copy them to an archive medium such as tape for disaster recovery, and remove them from the active log directory. Any log files needed for media recovery but not for restart can also be off-loaded to an archive.

If any log file that is needed cannot be found, operator message AMQ6767 is issued. Make the log file, and all subsequent log files, available to the queue manager and retry the operation.

Note: When performing media recovery, all the required log files must be available in the log file directory at the same time. Make sure that you take regular media images of any objects you may wish to recover to avoid running out of disk space to hold all the required log files.

Log file location

When choosing a location for your log files, remember that operation is severely impacted if MQSeries fails to format a new log because of lack of disk space. In MQSeries for OS/2 Warp, for example, put the log directory on a different drive from that used by the OS/2 swapper file: log files tend to be large, so could fill the disk and prevent expansion of the swapper file.

If you are using a circular log, ensure that there is sufficient space on the drive for at least the configured primary log files. You should also leave space for at least one secondary log file, which is needed if the log has to grow.

If you are using a linear log, you should allow considerably more space; the space consumed by the log increases continuously as data is logged.

Managing logs

Ideally, the log files should be placed on a separate disk drive from the queue manager data. This has benefits in terms of performance. It may also be possible to place the log files on multiple disk drives in a mirrored arrangement. This gives protection against failure of the drive containing the log. Without mirroring, you could be forced to go back to the last backup of your MQSeries system.

Using the log for recovery

There are several ways that your data can be damaged. MQSeries helps you recover from:

- A damaged data object
- A power loss in the system
- A communications failure
- A damaged log volume

This section looks at how the logs are used to recover from these problems.

Recovering from problems

MQSeries can recover from both communications failures and loss of power. In addition, it is sometimes possible to recover from other types of problem, such as inadvertent deletion of a file.

In the case of a communications failure, messages remain on queues until they are removed by a receiving application. If the message is being transmitted, it remains on the transmission queue until it can be successfully transmitted. To recover from a communications failure, it is normally sufficient simply to restart the channels using the link that failed.

If you lose power, when the queue manager is restarted MQSeries restores the queues to their committed state at the time of the failure. This ensures that no persistent messages are lost. Nonpersistent messages are discarded; they do not survive when MQSeries stops.

There are ways in which an MQSeries object can become unusable, for example due to inadvertent damage. You then have to recover either your complete system or some part of it. The action required depends on when the damage is detected, whether the log method selected supports media recovery, and which objects are damaged.

Media recovery

Media recovery is the re-creation of objects from information recorded in a linear log. For example, if an object file is inadvertently deleted, or becomes unusable for some other reason, media recovery can be used to recreate it. The information in the log required for media recovery of an object is called a *media image*. Media images can be recorded manually, using the **rctdmqimg** command, or automatically in some circumstances.

A media image is a sequence of log records containing an image of an object from which the object itself can be recreated.

The first log record required to recreate an object is known as its *media recovery record*; it is the start of the latest media image for the object. The media recovery record of each object is one of the pieces of information recorded during a checkpoint.

When an object is recreated from its media image, it is also necessary to replay any log records describing updates performed on the object since the last image was taken.

Consider, for example, a local queue that has an image of the queue object taken before a persistent message is put onto the queue. In order to recreate the latest image of the object, it is necessary to replay the log entries recording the putting of the message to the queue, as well as replaying the image itself.

When an object is created, the log records written contain enough information to completely recreate the object. These records make up the object's first media image. Subsequently, media images are recorded automatically by the queue manager at the following times:

- Images of all process objects and non-local queues are taken at each shutdown.
- Local queue images are taken when the queue becomes empty.

Media images can also be recorded manually using the **rcdmqimg** command, described in “rcdmqimg (Record media image)” on page 265. Issuing this command causes a media image of the MQSeries object to be written. Once this has been done, only the logs that hold the media image, and all the logs created after this time, are needed to recreate damaged objects. The benefit of doing this depends on such factors as the amount of free storage available, and the speed at which log files are created.

Recovering media images

MQSeries automatically recovers some objects from their media image if it finds that they are corrupt or damaged. In particular, this applies to objects found to be damaged during the normal queue manager startup. If any transaction was incomplete at the time of the last shutdown of the queue manager, any queue affected is also recovered automatically in order to complete the startup operation.

You must recover other objects manually, using the **rcrmqobj** command. This command replays the records in the log to recreate the MQSeries object. The object is recreated from its latest image found in the log, together with all applicable log events between the time the image was saved and the time the recreate command is issued. Should an MQSeries object become damaged, the only valid actions that can be performed are either to delete it or to recreate it by this method. Note, however, that nonpersistent messages cannot be recovered in this way.

See “rcrmqobj (Recreate object)” on page 267 for further details of the **rcrmqobj** command.

It is important to remember that you must have the log file containing the media recovery record, and all subsequent log files, available in the log file directory when attempting media recovery of an object. If a required file cannot be found, operator message AMQ6767 is issued and the media recovery operation fails. If you do not take regular media images of the objects that you may wish to recreate, you can get into the situation where you have insufficient disk space to hold all the log files required to recreate an object.

Recovering damaged objects during startup

If the queue manager discovers a damaged object during startup, the action it takes depends on the type of object and whether the queue manager is configured to support media recovery.

If the queue manager object is damaged, the queue manager cannot start unless it can recover the object. If the queue manager is configured with a linear log, and thus supports media recovery, MQSeries automatically tries to recreate the MQSeries object from its media images. If the log method selected does not support media recovery, you can either restore a backup of the queue manager or delete the queue manager.

If any transactions were active when the queue manager stopped, the local queues containing the persistent, uncommitted messages put or got inside these transactions are also needed to start the queue manager successfully. If any of these local queues is found to be damaged, and the queue manager supports media recovery, it automatically attempts to recreate them from their media images. If any of the queues cannot be recovered, MQSeries cannot start.

If any damaged local queues containing uncommitted messages are discovered during startup processing on a queue manager that does not support media recovery, the queues are marked as damaged objects and the uncommitted messages on them are ignored. This is because it is not possible to perform media recovery of damaged objects on such a queue manager and the only action left is to delete them. Message AMQ7472 is issued to report any damage.

Recovering damaged objects at other times

Media recovery of objects is automatic only during startup. At other times, when object damage is detected, operator message AMQ7472 is issued and most operations using the object fail. If the queue manager object is damaged at any time after the queue manager has started, the queue manager performs a preemptive shutdown. When an object has been damaged you may delete it or, if the queue manager is using a linear log, attempt to recover it from its media image using the **rcrmqobj** command (see “rcrmqobj (Recreate object)” on page 267 for further details).

Backup and restore

Periodically, you may want to take a backup of your queue manager data to provide protection against possible corruption due to hardware failures. However, because message data is often short-lived, you may choose not to take backups.

Backing up MQSeries

To take a backup of a queue manager's data, you must:

1. Ensure that the queue manager is not running.

If your queue manager is running, stop it with the **endmqm** command.

Note: If you try to take a backup of a running queue manager, the backup may not be consistent due to updates in progress when the files were copied.

2. Locate the directories under which the queue manager places its data and its log files.

You can use the information in the configuration files to determine these directories. For more information about this, see Chapter 7, "Configuration files" on page 99.

Note: You may have some difficulty in understanding the names that appear in the directory. This is because the names are transformed to ensure that they are compatible with the platform on which you are using MQSeries. For more information about name transformations, see "Understanding MQSeries file names" on page 30.

3. Take copies of all the queue manager's data and log file directories, including all subdirectories.

Make sure that you do not miss any of the files, especially the log control file and the configuration files. Some of the directories may be empty, but they will all be required if you restore the backup at a later date, so it is advisable to save them too.

4. Ensure that you preserve the ownerships of the files. For MQSeries for UNIX systems, you can do this with the **tar** command.

Restoring MQSeries

To restore a backup of a queue manager's data, you must:

1. Ensure that the queue manager is not running.
2. Locate the directories under which the queue manager places its data and its log files. This information is held in the configuration file.
3. Clear out the directories into which you are going to place the backed up data.
4. Copy the backed up queue manager data and log files into the correct places.

Check the resulting directory structure to ensure that you have all of the required directories.

See Appendix B, "Directory structure (UNIX systems)" on page 303 for more information about MQSeries directories and subdirectories.

Make sure that you have a log control file as well as the log files. Also check that the MQSeries and queue manager configuration files are consistent so that MQSeries can look in the correct places for the restored data.

If the data was backed up and restored correctly, the queue manager will now start.

Note: Even though the queue manager data and log files are held in different directories, you should back up and restore the directories at the same time. If the queue manager data and log files have different ages, the queue manager is not in a valid state and will probably not start. If it does start, your data will almost certainly be corrupt.

Recovery scenarios

This section looks at a number of possible problems and indicates how to recover from them.

Disk drive failures

You may suffer problems with a disk drive containing either the queue manager data, the log, or both. Problems can include data loss or corruption. The three cases differ only in the part of the data that survives, if any.

In *all* cases you must first check the directory structure for any damage and, if necessary, repair such damage. If you lose queue manager data, there is a danger that the queue manager directory structure has been damaged. If so, you must recreate the directory tree manually before you try to restart the queue manager. Having checked for structural damage, there are a number of alternative things you can do, depending on the type of logging that you use.

- **Where there is major damage to the directory structure or any damage to the log**, remove all the old files back to the QMgrName level, including the configuration files, the log, and the queue manager directory, restore the last backup, and try to restart the queue manager.
- **For linear logging with media recovery**, ensure the directory structure is intact and try to restart the queue manager. If the queue manager does not restart, restore a backup. If the queue manager restarts, check whether any other objects have been damaged using MQSC commands, such as DISPLAY QUEUE. Recover those you find, using the **rcrmqobj** command. For example:

```
rcrmqobj -m QMgrName -t all *
```

where QMgrName is the queue manager being recovered. -t all * indicates that all objects of any type (except channels) are to be recovered. If only one or two objects have been reported as damaged, you may want to specify those objects by name and type here.

- **For linear logging with media recovery and with an undamaged log**, you may be able to restore a backup of the queue manager data leaving the existing log files and log control file unchanged. Starting the queue manager applies the changes from the log to bring the queue manager back to its state when the failure occurred.

Recovery scenarios

This method relies on two facts. Firstly, it is vital that the checkpoint file be restored as part of the queue manager data. This file contains the information determining how much of the data in the log must be applied to give a consistent queue manager.

Secondly, you must have the oldest log file that was required to start the queue manager at the time of the backup, and all subsequent log files, available in the log file directory.

If this is not possible, you must restore a backup of both the queue manager data and the log, both of which were taken at the same time.

- **For circular logging, or linear logging without media recovery**, you must restore the queue manager from the latest backup that you have. Once you have restored the backup, restart the queue manager and check as above for damaged objects. However, because you do not have media recovery, you must find other ways of recreating the damaged objects.

Damaged queue manager object

If the queue manager object has been reported as damaged during normal operation, the queue manager performs a preemptive shutdown. There are two ways of recovering in these circumstances depending on the type of logging you use:

- **For linear logging only**, manually delete the file containing the damaged object and restart the queue manager. (You can use the **dspmqls** command to determine the real, file-system name of the damaged object.) Media recovery of the damaged object is automatic.
- **For circular or linear logging**, restore the last backup of the queue manager data and log and restart the queue manager.

Damaged single object

If a single object is reported as damaged during normal operation, there are two ways of recovering, depending on the type of logging you use:

- **For linear logging**, recreate the object from its media image.
- **For circular logging**, restore the last backup of the queue manager data and log and restart the queue manager.

Automatic media recovery failure

If a local queue required for queue manager startup with a linear log is damaged, and the automatic media recovery fails, restore the last backup of the queue manager data and log and restart the queue manager.

Chapter 13. Problem determination

This chapter suggests reasons for some of the problems you may have using MQSeries. You usually start with a symptom, or set of symptoms, and trace them back to their cause.

Problem determination is not problem solving. However, the process of problem determination often enables you to solve a problem. For example, if you find that the cause of the problem is an error in an application program, you can solve the problem by correcting the error.

The process of problem determination is that you start with the symptoms and trace them back to their cause.

Not all problems can be solved immediately, for example, performance problems caused by the limitations of your hardware. Also, if you think that the cause of the problem is in the MQSeries code, contact your IBM Support Center. This chapter contains these sections:

- “Preliminary checks”
- “Common programming errors” on page 206
- “What to do next” on page 207
- “Application design considerations” on page 211
- “Incorrect output” on page 212
- “Error logs” on page 215
- “Dead-letter queues” on page 219
- “Configuration files and problem determination” on page 219
- “Tracing” on page 219
- “First-failure support technology (FFST)” on page 226
- “Problem determination with clients” on page 230

Preliminary checks

Before you start problem determination in detail, it is worth considering the facts to see if there is an obvious cause of the problem, or a likely area in which to start your investigation. This approach to debugging can often save a lot of work by highlighting a simple error, or by narrowing down the range of possibilities.

The cause of your problem could be in:

- MQSeries
- The network
- The application

The sections that follow raise some fundamental questions that you need to consider. As you work through the questions, make a note of anything that might be relevant to the problem. Even if your observations do not suggest a cause immediately, they could be useful later if you have to carry out a systematic problem determination exercise.

Has MQSeries run successfully before?

If MQSeries has not run successfully before, it is likely that you have not yet set it up correctly. See the *MQSeries Quick Beginnings* book for your MQSeries product to check that you installed the product correctly, and ensure that the Installation Verification Test (IVT) has been run. See also the *MQSeries Intercommunication* book for information about post-installation configuration of MQSeries.

Are there any error messages?

MQSeries uses error logs to capture messages concerning the operation of MQSeries itself, any queue managers that you start, and error data coming from the channels that are in use. Check the error logs to see if any messages have been recorded that are associated with your problem.

See “Error logs” on page 215 for information about the contents of the error logs, and their locations.

Are there any return codes explaining the problem?

If your application gets a return code indicating that a Message Queue Interface (MQI) call has failed, refer to the *MQSeries Application Programming Reference* manual for a description of that return code.

Can you reproduce the problem?

If you can reproduce the problem, consider the conditions under which it is reproduced:

- Is it caused by a command or an equivalent administration request?
Does the operation work if it is entered by another method? If the command works if it is entered on the command line, but not otherwise, check that the command server has not stopped, and that the queue definition of the SYSTEM.ADMIN.COMMAND.QUEUE has not been changed.
- Is it caused by a program? Does it fail on all MQSeries systems and all queue managers, or only on some?
- Can you identify any application that always seems to be running in the system when the problem occurs? If so, examine the application to see if it is in error.

Have any changes been made since the last successful run?

When you are considering changes that might recently have been made, think about the MQSeries system, and also about the other programs it interfaces with, the hardware, and any new applications. Consider also the possibility that a new application that you are not aware of might have been run on the system.

- Have you changed, added, or deleted any queue definitions?
- Have you changed or added any channel definitions? Changes may have been made to either MQSeries channel definitions or any underlying communications definitions required by your application.
- Do your applications deal with return codes that they might get as a result of any changes you have made?
- Have you changed any component of the operating system that could affect the operation of MQSeries? For example, have you modified the Windows NT Registry hive?

Has the application run successfully before?

If the problem appears to involve one particular application, consider whether the application has run successfully before.

Before you answer **Yes** to this question, consider the following:

- Have any changes been made to the application since it last ran successfully?

If so, it is likely that the error lies somewhere in the new or modified part of the application. Take a look at the changes and see if you can find an obvious reason for the problem. Is it possible to retry using a back level of the application?

- Have all the functions of the application been fully exercised before?

Could it be that the problem occurred when part of the application that had never been invoked before was used for the first time? If so, it is likely that the error lies in that part of the application. Try to find out what the application was doing when it failed, and check the source code in that part of the program for errors.

If a program has been run successfully on many previous occasions, check the current queue status, and the files that were being processed when the error occurred. It is possible that they contain some unusual data value that causes a rarely used path in the program to be invoked.

- Does the application check all return codes?

Has your MQSeries system been changed, perhaps in a minor way, such that your application does not check the return codes it receives as a result of the change. For example, does your application assume that the queues it accesses can be shared? If a queue has been redefined as exclusive, can your application deal with return codes indicating that it can no longer access that queue?

- Does the application run on other MQSeries systems?

Could it be that there is something different about the way that this MQSeries system is set up which is causing the problem? For example, have the queues been defined with the same message length or priority?

If the application has not run successfully before

If your application has not yet run successfully, you need to examine it carefully to see if you can find any errors.

Before you look at the code, and depending upon which programming language the code is written in, examine the output from the translator, or the compiler and linkage editor, if applicable, to see if any errors have been reported.

If your application fails to translate, compile, or link-edit into the load library, it will also fail to run if you attempt to invoke it. See the *MQSeries Application Programming Guide* for information about building your application.

If the documentation shows that each of these steps was accomplished without error, you should consider the coding logic of the application. Do the symptoms of the problem indicate the function that is failing and, therefore, the piece of code in error? See “Common programming errors” on page 206 for some examples of common errors that cause problems with MQSeries applications.

Common programming errors

The errors in the following list illustrate the most common causes of problems encountered while running MQSeries programs. You should consider the possibility that the problem with your MQSeries system could be caused by one or more of these errors:

- Assuming that queues can be shared, when they are in fact exclusive.
- Passing incorrect parameters in an MQI call.
- Passing insufficient parameters in an MQI call. This may mean that MQI cannot set up completion and reason codes for your application to process.
- Failing to check return codes from MQI requests.
- Passing variables with incorrect lengths specified.
- Passing parameters in the wrong order.
- Failing to initialize *MsgId* and *CorrelId* correctly.
- Failing to initialize *Encoding* and *CodedCharSetId* following `MQRC_TRUNCATED_MSG_ACCEPTED`.

Problems with commands

You should be careful when including special characters, for example, back slash (\) and double quote (") characters, in descriptive text for some commands. If you use either of these characters in descriptive text, precede them with a \, that is, enter \\ or \" if you want \ or " in your text.

Does the problem affect specific parts of the network?

You might be able to identify specific parts of the network that are affected by the problem (remote queues, for example). If the link to a remote message queue manager is not working, the messages cannot flow to a remote queue.

Check that the connection between the two systems is available, and that the intercommunication component of MQSeries has been started.

Check that messages are reaching the transmission queue, and check the local queue definition of the transmission queue and any remote queues.

Have you made any network-related changes, or changed any MQSeries definitions, that might account for the problem?

Does the problem occur at specific times of the day?

If the problem occurs at specific times of day, it could be that it is dependent on system loading. Typically, peak system loading is at mid-morning and mid-afternoon, so these are the times when load-dependent problems are most likely to occur. (If your MQSeries network extends across more than one time zone, peak system loading might seem to occur at some other time of day.)

Is the problem intermittent?

An intermittent problem could be caused by failing to take into account the fact that processes can run independently of each other. For example, a program may issue an MQGET call, without specifying a wait option, before an earlier process has completed. An intermittent problem may also be seen if your application tries to get a message from a queue while the call that put the message is in-doubt (that is, before it has been committed or backed out).

Have you applied any service updates?

If a service update has been applied to MQSeries, check that the update action completed successfully and that no error message was produced.

- Did the update have any special instructions?
- Was any test run to verify that the update had been applied correctly and completely?
- Does the problem still exist if MQSeries is restored to the previous service level?
- If the installation was successful, check with the IBM Support Center for any PTF error.
- If a PTF has been applied to any other program, consider the effect it might have on the way MQSeries interfaces with it.

What to do next

Perhaps the preliminary checks have enabled you to find the cause of the problem. If so, you should now be able to resolve it, possibly with the help of other books in the MQSeries library (see “MQSeries publications” on page xi) and in the libraries of other licensed programs.

If you have not yet found the cause, you must start to look at the problem in greater detail.

The purpose of this section is to help you identify the cause of your problem if the preliminary checks have not enabled you to find it.

When you have established that no changes have been made to your system, and that there are no problems with your application programs, choose the option that best describes the symptoms of your problem.

- “Have you obtained incorrect output?” on page 208
- “Have you failed to receive a response from a PCF command?” on page 208
- “Does the problem affect only remote queues?” on page 210
- “Is your application or system running slowly?” on page 210

If none of these symptoms describe your problem, consider whether it might have been caused by another component of your system.

Have you obtained incorrect output?

In this book, “incorrect output” refers to your application:

- Not receiving a message that it was expecting.
- Receiving a message containing unexpected or corrupted information.
- Receiving a message that it was not expecting, for example, one that was destined for a different application.

In all cases, check that any queue or queue manager aliases that your applications are using are correctly specified and accommodate any changes that have been made to your network.

If an MQSeries error message is generated, all of which are prefixed with the letters “AMQ,” you should look in the error log. See “Error logs” on page 215 for further information.

Have you failed to receive a response from a PCF command?

If you have issued a command but you have not received a response, consider the following questions:

- Is the command server running?

Work with the **dspmqcsv** command to check the status of the command server.

- If the response to this command indicates that the command server is not running, use the **strmqcsv** command to start it.
- If the response to the command indicates that the SYSTEM.ADMIN.COMMAND.QUEUE is not enabled for MQGET requests, enable the queue for MQGET requests.

- Has a reply been sent to the dead-letter queue?

The dead-letter queue header structure contains a reason or feedback code describing the problem. See the *MQSeries Application Programming Reference* manual for information about the dead-letter queue header structure (MQDLH).

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

- Has a message been sent to the error log?

See “Error logs” on page 215 for further information.

- Are the queues enabled for put and get operations?

- Is the *WaitInterval* long enough?

If your MQGET call has timed out, a completion code of MQCC_FAILED and a reason code of MQRC_NO_MSG_AVAILABLE are returned. (See the *MQSeries Application Programming Reference* manual for information about the *WaitInterval* field, and completion and reason codes from MQGET.)

- If you are using your own application program to put commands onto the SYSTEM.ADMIN.COMMAND.QUEUE, do you need to take a syncpoint?

Unless you have specifically excluded your request message from syncpoint, you need to take a syncpoint before attempting to receive reply messages.

- Are the MAXDEPTH and MAXMSGL attributes of your queues set sufficiently high?
- Are you using the *CorrelId* and *MsgId* fields correctly?

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

Try stopping the command server and then restarting it, responding to any error messages that are produced.

If the system still does not respond, the problem could be with either a queue manager or the whole of the MQSeries system. First try stopping individual queue managers to try and isolate a failing queue manager. If this does not reveal the problem, try stopping and restarting MQSeries, responding to any messages that are produced in the error log.

If the problem still occurs after restart, contact your IBM Support Center for help.

Are some of your queues failing?

If you suspect that the problem occurs with only a subset of queues, check the local queues that you think are having problems:

1. Display the information about each queue. You can use the MQSC command DISPLAY QUEUE to display the information.
2. Use the data displayed to do the following checks:
 - If CURDEPTH is at MAXDEPTH, this indicates that the queue is not being processed. Check that all applications are running normally.
 - If CURDEPTH is not at MAXDEPTH, check the following queue attributes to ensure that they are correct:
 - If triggering is being used:
 - Is the trigger monitor running?
 - Is the trigger depth too great? That is, does it generate a trigger event often enough?
 - Is the process name correct?
 - Is the process available and operational?
 - Can the queue be shared? If not, another application could already have it open for input.
 - Is the queue enabled appropriately for GET and PUT?
 - If there are no application processes getting messages from the queue, determine why this is so. It could be because the applications need to be started, a connection has been disrupted, or the MQOPEN call has failed for some reason.

Check the queue attributes IPPROCS and OPPROCS. These attributes indicate whether the queue has been opened for input and output. If a value is zero, it indicates that no operations of that type can occur. Note that the values may have changed and that the queue was open but is now closed.

What next

You need to check the status at the time you expect to put or get a message.

If you are unable to solve the problem, contact your IBM Support Center for help.

Does the problem affect only remote queues?

If the problem affects only remote queues, check the following:

- Check that required channels have been started and are triggerable, and that any required initiators are running.
- Check that the programs that should be putting messages to the remote queues have not reported problems.
- If you use triggering to start the distributed queuing process, check that the transmission queue has triggering set on. Also, check that the trigger monitor is running.
- Check the error logs for messages indicating channel errors or problems.
- If necessary, start the channel manually. See the *MQSeries Intercommunication* book for information about how to do this.

See the *MQSeries Intercommunication* book for information about how to define channels.

Is your application or system running slowly?

If your application is running slowly, this could indicate that it is in a loop, or waiting for a resource that is not available.

This could also be caused by a performance problem. Perhaps it is because your system is operating near the limits of its capacity. This type of problem is probably worst at peak system load times, typically at mid-morning and mid-afternoon. (If your network extends across more than one time zone, peak system load might seem to occur at some other time.)

A performance problem may be caused by a limitation of your hardware.

If you find that performance degradation is not dependent on system loading, but happens sometimes when the system is lightly loaded, a poorly designed application program is probably to blame. This could manifest itself as a problem that only occurs when certain queues are accessed.

The following symptoms might indicate that MQSeries is running slowly:

- Your system is slow to respond to MQSeries commands.
- Repeated displays of the queue depth indicate that the queue is being processed slowly for an application with which you would expect a large amount of queue activity.

If the performance of your system is still degraded after reviewing the above possible causes, the problem may lie with MQSeries itself. If you suspect this, you need to contact your IBM Support Center for assistance.

Application design considerations

There are a number of ways in which poor program design can affect performance. These can be difficult to detect because the program can appear to perform well, while impacting the performance of other tasks. Several problems specific to programs making MQSeries calls are discussed in the following sections.

For more information about application design, see the *MQSeries Application Programming Guide*.

Effect of message length

The amount of data in a message can affect the performance of the application that processes the message. To achieve the best performance from your application, you should send only the essential data in a message; for example, in a request to debit a bank account, the only information that may need to be passed from the client to the server application is the account number and the amount of the debit.

Effect of message persistence

Persistent messages are logged. Logging messages reduces the performance of your application, so you should use persistent messages for essential data only. If the data in a message can be discarded if the queue manager stops or fails, use a nonpersistent message.

Searching for a particular message

The MQGET call usually retrieves the first message from a queue. If you use the message and correlation identifiers (*MsgId* and *CorrelId*) in the message descriptor to specify a particular message, the queue manager has to search the queue until it finds that message. Using the MQGET call in this way affects the performance of your application.

Queues that contain messages of different lengths

If the messages on a queue are of different lengths, to determine the size of a message, your application could use the MQGET call with the *BufferLength* field set to zero so that, even though the call fails, it returns the size of the message data. The application could then repeat the call, specifying the identifier of the message it measured in its first call and a buffer of the correct size. However, if there are other applications serving the same queue, you might find that the performance of your application is reduced because its second MQGET call spends time searching for a message that another application has retrieved in the time between your two calls.

If your application cannot use messages of a fixed length, another solution to this problem is to use the MQINQ call to find the maximum size of messages that the queue can accept, then use this value in your MQGET call. The maximum size of messages for a queue is stored in the *MaxMsgLength* attribute of the queue. This method could use large amounts of storage, however, because the value of this queue attribute could be as high as 100 MB, the maximum allowed by MQSeries. Note also that if you do not set the *MaxMsgLength* attribute explicitly, it defaults to 4 MB, which may be very inefficient.

Incorrect output

Frequency of syncpoints

Programs that issue numerous MQPUT calls within syncpoint, without committing them, can cause performance problems. Affected queues can fill up with messages that are currently inaccessible, while other tasks might be waiting to get these messages. This has implications in terms of storage, and in terms of threads tied up with tasks that are attempting to get messages.

Use of the MQPUT1 call

Use the MQPUT1 call only if you have a single message to put on a queue. If you want to put more than one message, use the MQOPEN call, followed by a series of MQPUT calls and a single MQCLOSE call.

Number of threads in use

For MQSeries for OS/2 Warp and MQSeries for Windows NT, an application may require a large number of threads. Each queue manager process is allocated a maximum allowable number of threads.

If some applications are troublesome, it could be due to their design using too many threads. Consider whether the application takes into account this possibility and that it takes actions either to stop or to report this type of occurrence.

The maximum number of threads that OS/2 allows is 4095. However, the default is 64. The default can be changed with the THREADS=xxxx parameter in CONFIG.SYS. MQSeries makes available up to 63 threads to its processes.

Incorrect output

The term “incorrect output” can be interpreted in many different ways. For the purpose of problem determination within this book, the meaning is explained in “Have you obtained incorrect output?” on page 208.

Two types of incorrect output are discussed in this section:

- Messages that do not appear when you are expecting them
- Messages that contain the wrong information, or information that has been corrupted

Additional problems that you might find if your application includes the use of distributed queues are also discussed.

Messages that do not appear on the queue

If messages do not appear when you are expecting them, check for the following:

- Has the message been put on the queue successfully?
 - Has the queue been defined correctly. For example, is MAXMSGL sufficiently large?
 - Is the queue enabled for putting?
 - Is the queue already full? This could mean that an application was unable to put the required message on the queue.
 - Has another application got exclusive access to the queue?

- Are you able to get any messages from the queue?
 - Do you need to take a syncpoint?

If messages are being put or retrieved within syncpoint, they are not available to other tasks until the unit of recovery has been committed.
 - Is your wait interval long enough?

You can set the wait interval as an option for the MQGET call. You should ensure that you are waiting long enough for a response.
 - Are you waiting for a specific message that is identified by a message or correlation identifier (*MsgId* or *CorrelId*)?

Check that you are waiting for a message with the correct *MsgId* or *CorrelId*. A successful MQGET call sets both these values to that of the message retrieved, so you may need to reset these values in order to get another message successfully.

Also, check whether you can get other messages from the queue.
 - Can other applications get messages from the queue?
 - Was the message you are expecting defined as persistent?

If not, and MQSeries has been restarted, the message has been lost.
 - Has another application got exclusive access to the queue?

If you are unable to find anything wrong with the queue, and MQSeries is running, make the following checks on the process that you expected to put the message on to the queue:

- Did the application get started?

If it should have been triggered, check that the correct trigger options were specified.
- Did the application stop?
- Is a trigger monitor running?
- Was the trigger process defined correctly?
- Did the application complete correctly?

Look for evidence of an abnormal end in the job log.
- Did the application commit its changes, or were they backed out?

If multiple transactions are serving the queue, they can conflict with one another. For example, suppose one transaction issues an MQGET call with a buffer length of zero to find out the length of the message, and then issues a specific MQGET call specifying the *MsgId* of that message. However, in the meantime, another transaction issues a successful MQGET call for that message, so the first application receives a reason code of MQRC_NO_MSG_AVAILABLE. Applications that are expected to run in a multi-server environment must be designed to cope with this situation.

Consider that the message could have been received, but that your application failed to process it in some way. For example, did an error in the expected format of the message cause your program to reject it? If this is the case, refer to “Messages that contain unexpected or corrupted information” on page 214.

Messages that contain unexpected or corrupted information

If the information contained in the message is not what your application was expecting, or has been corrupted in some way, consider the following points:

- Has your application, or the application that put the message onto the queue, changed?

Ensure that all changes are simultaneously reflected on all systems that need to be aware of the change.

For example, the format of the message data may have been changed, in which case, both applications must be recompiled to pick up the changes. If one application has not been recompiled, the data will appear corrupt to the other.

- Is an application sending messages to the wrong queue?

Check that the messages your application is receiving are not really intended for an application servicing a different queue. If necessary, change your security definitions to prevent unauthorized applications from putting messages on to the wrong queues.

If your application has used an alias queue, check that the alias points to the correct queue.

- Has the trigger information been specified correctly for this queue?

Check that your application should have been started; or should a different application have been started?

If these checks do not enable you to solve the problem, you should check your application logic, both for the program sending the message, and for the program receiving it.

Problems with incorrect output when using distributed queues

If your application uses distributed queues, you should also consider the following points:

- Has MQSeries been correctly installed on both the sending and receiving systems, and correctly configured for distributed queuing?
- Are the links available between the two systems?

Check that both systems are available, and connected to MQSeries. Check that the connection between the two systems is active.

You can use an MQSeries PING command against either the queue manager (PING QMGR) or the channel (PING CHANNEL) to verify that the link is operable.

- Is triggering set on in the sending system?
- Is the message you are waiting for a reply message from a remote system?

Check that triggering is activated in the remote system.

- Is the queue already full?

This could mean that an application was unable to put the required message onto the queue. If this is so, check if the message has been put onto the dead-letter queue.

The dead-letter queue header contains a reason or feedback code explaining why the message could not be put onto the target queue. See the *MQSeries Application Programming Reference* manual for information about the dead-letter queue header structure.

- Is there a mismatch between the sending and receiving queue managers?
For example, the message length could be longer than the receiving queue manager can handle.
- Are the channel definitions of the sending and receiving channels compatible?
For example, a mismatch in sequence number wrap stops the distributed queuing component. See *MQSeries Intercommunication* for more information about distributed queuing.
- Is data conversion involved? If the data formats between the sending and receiving applications differ, data conversion is necessary. Automatic conversion occurs when the MQGET is issued if the format is recognized as one of the built-in formats.

If the data set is not recognized for conversion, the data conversion exit is taken to allow you to perform the translation with your own routines.

Refer to the *MQSeries Application Programming Guide* for further details of data conversion.

Error logs

MQSeries uses a number of error logs to capture messages concerning the operation of MQSeries itself, any queue managers that you start, and error data coming from the channels that are in use.

The location of the error logs depends on whether the queue manager name is known and whether the error is associated with a client.

In MQSeries for UNIX systems:

- If the queue manager name is known and the queue manager is available, error logs are located in:
`/var/mqm/qmgrs/qmname/errors`
- If the queue manager is not available, error logs are located in:
`/var/mqm/qmgrs/@SYSTEM/errors`
- If an error has occurred with a client application, error logs are located on the client's root drive in
`/var/mqm/errors`

Error logs

In MQSeries for OS/2 Warp and Windows NT, and assuming that MQSeries has been installed on the C drive in the MQM directory:

- If the queue manager name is known and the queue manager is available, error logs are located in:
c:\mqm\qmgrs\qmname\errors
- If the queue manager is not available, error logs are located in:
c:\mqm\qmgrs\@SYSTEM\errors
- If an error has occurred with a client application, error logs are located on the client's root drive in:
c:\mqm\errors

In MQSeries for Windows NT only, an indication of the error is also added to the Application Log, which can be examined with the Event Viewer application provided with Windows NT.

You can also examine the Registry to help resolve any errors. The Registry Editor supplied with Windows NT allows you to filter errors that are placed in the Event Log by placing the code in the following Registry entry:

HKEY_LOCAL_MACHINE\SOFTWARE\IBM\MQSeries\CurrentVersion\IgnoredErrorCodes

For example, to ignore error 5000, add AMQ5000 to the list.

Log files

At installation time an @SYSTEM errors subdirectory is created in the qmgrs file path. The errors subdirectory can contain up to three error log files named:

- AMQERR01.LOG
- AMQERR02.LOG
- AMQERR03.LOG

After you have created a queue manager, three error log files are created when they are needed by the queue manager. These files have the same names as the @SYSTEM ones, that is AMQERR01, AMQERR02, and AMQERR03, and each has a capacity of 256 KB. The files are placed in the errors subdirectory of each queue manager that you create.

As error messages are generated, they are placed in AMQERR01. When AMQERR01 gets bigger than 256 KB it is copied to AMQERR02. Before the copy, AMQERR02 is copied to AMQERR03.LOG. The previous contents, if any, of AMQERR03 are discarded.

The latest error messages are thus always placed in AMQERR01, the other files being used to maintain a history of error messages.

All messages relating to channels are also placed in the appropriate queue manager's errors files unless the name of their queue manager is unknown or the queue manager is unavailable. When the queue manager name is unavailable or its name cannot be determined, channel-related messages are placed in the @SYSTEM errors subdirectory.

To examine the contents of any error log file, use your usual system editor.

Early errors

There are a number of special cases where the above error logs have not yet been established and an error occurs. MQSeries attempts to record any such errors in an error log. The location of the log depends on how much of a queue manager has been established.

If, due to a corrupt configuration file for example, no location information can be determined, errors are logged to an errors directory that is created at installation time on the root directory (`/var/mqm` or `C:\MQM`).

If the MQSeries configuration file is readable, and the `DefaultPrefix` attribute of the `AllQueueManagers` stanza is readable, errors are logged in the errors subdirectory of the directory identified by the `DefaultPrefix` attribute. For example, if the `DefaultPrefix` is `C:\MQM`, errors are logged in `C:\MQM\ERRORS`.

For further information about configuration files, see Chapter 7, “Configuration files” on page 99.

Operator messages

Operator messages identify normal errors, typically caused directly by users doing things like using parameters that are not valid on a command. Operator messages are national language (NLS) enabled, with message catalogs installed in standard locations.

These messages are written to the associated window, if any. In addition, some operator messages are written to the `AMQERR01.LOG` file in the queue manager directory, and others to the `@SYSTEM` directory copy of the error log.

An example error log

Figure 41 on page 218 shows a typical extract from an MQSeries error log.

```
...
08/01/97 11:41:56 AMQ8003: MQSeries queue manager started.
EXPLANATION: MQSeries queue manager Janet started.
ACTION: None.
-----
08/01/97 11:56:52 AMQ9002: Channel program started.
EXPLANATION: Channel program 'JANET' started.
ACTION: None.
-----
08/01/97 11:57:26 AMQ9208: Error on receive from host 'camelot
(9.20.12.34)'.
EXPLANATION: An error occurred receiving data from 'camelot
(9.20.12.34)' over TCP/IP. This may be due to a communications failure.
ACTION: Record the TCP/IP return code 232 (X'E8') and tell the
systems administrator.
-----
08/01/97 11:57:27 AMQ9999: Channel program ended abnormally.
EXPLANATION: Channel program 'JANET' ended abnormally.
ACTION: Look at previous error messages for channel program
'JANET' in the error files to determine the cause of the failure.
-----
08/01/97 14:28:57 AMQ8004: MQSeries queue manager ended.
EXPLANATION: MQSeries queue manager Janet ended.
ACTION: None.
-----
08/02/97 15:02:49 AMQ9002: Channel program started.
EXPLANATION: Channel program 'JANET' started.
ACTION: None.
-----
08/02/97 15:02:51 AMQ9001: Channel program ended normally.
EXPLANATION: Channel program 'JANET' ended normally.
ACTION: None.
08/02/97 15:09:27 AMQ7030: Request to quiesce the queue manager
accepted. The queue manager will stop when there is no further
work for it to perform.
EXPLANATION: You have requested that the queue manager end when
there is no more work for it. In the meantime, it will refuse
new applications that attempt to start, although it allows those
already running to complete their work.
ACTION: None.
-----
08/02/97 15:09:32 AMQ8004: MQSeries queue manager ended.
EXPLANATION: MQSeries queue manager Janet ended.
ACTION: None.
...

```

Figure 41. Extract from an MQSeries error log

The MQSeries log-dump utility

For a description of the **dmpmqlog** command, see “dmpmqlog (Dump log)” on page 246.

Dead-letter queues

Messages that cannot be delivered for some reason are placed on the dead-letter queue. You can check whether the queue contains any messages by issuing an MQSC DISPLAY QUEUE command. If the queue contains messages, you can use the provided browse sample application (amqsbcbg) to browse messages on the queue using the MQGET call. The sample application steps through all the messages on a named queue for a named queue manager, displaying both the message descriptor and the message context fields for all the messages on the named queue. See “Browsing queues” on page 53 for more information about running this sample and about the kind of output it produces.

You must decide how to dispose of any messages found on the dead-letter queue, depending on the reasons for the messages being put on the queue.

Problems may occur if you do not associate a dead-letter queue with each queue manager. For more information about dead-letter queues, see Chapter 8, “The MQSeries dead-letter queue handler” on page 115.

Configuration files and problem determination

Configuration file errors typically prevent queue managers from being found, and result in “queue manager unavailable” type errors. Ensure that the configuration files exist, and that the MQSeries configuration file references the correct queue manager and log directories.

Tracing

This section describes how to produce a trace for each of the MQSeries Version 5 products.

Tracing MQSeries for AIX

MQSeries for AIX uses the standard AIX system trace. Tracing is a two-step process:

1. Gathering the data
2. Formatting the results

MQSeries uses two trace hook identifiers:

- | | |
|---------------|---|
| X'30D' | This event is recorded by MQSeries on entry to or exit from a subroutine. |
| X'30E' | This event is recorded by MQSeries to trace data such as that being sent or received across a communications network. |

Trace provides detailed execution tracing to help you to analyze problems. IBM service support personnel may ask for a problem to be recreated with trace enabled. The files produced by trace can be **very** large so it is important to qualify a trace, where possible. For example, you can optionally qualify a trace by time and by component.

There are two ways to run trace:

1. Interactively.

The following sequence of commands runs an interactive trace on the program myprog and ends the trace.

```
trace -j30D,30E -o trace.file  
->!myprog  
->q
```

2. Asynchronously.

The following sequence of commands runs an asynchronous trace on the program myprog and ends the trace.

```
trace -a -j30D,30E -o trace.file  
myprog  
trcstop
```

You can format the trace file with the command:

```
trcrpt -t mqmtop/lib/amqtrc.fmt trace.file > report.file
```

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

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

Selective component tracing

You should set the environment variable MQS_TRACE_OPTIONS only if you have been instructed to do so by your service personnel.

The environment variable MQS_TRACE_OPTIONS can be used to activate the high detail and parameter tracing functions individually. Because it enables tracing to be active without these functions, you can use it to reduce the overhead on execution speed when you are trying to reproduce a problem with tracing switched on. Table 15 on page 221 defines the trace behavior under the various settings of MQS_TRACE_OPTIONS.

<i>Table 15. MQS_TRACE_OPTIONS settings</i>	
MQS_TRACE_OPTIONS Value	What will be traced
Unset (default)	Default trace (all except high detail)
0	No MQSeries trace
262148	Entry, exit and parameter trace
786436	Entry, exit, parameter, and high detail trace
3407871	Default trace without parameter trace
3670015	Default trace, including parameter trace
4194303	All tracing, including high detail trace

Notes:

1. Typically MQS_TRACE_OPTIONS must be set in the process that starts the queue manager, and before the queue manager is started, or it is not recognized.
2. MQS_TRACE_OPTIONS must be set before tracing starts. If it is set after tracing starts it is not recognized.

An example of MQSeries for AIX trace data

The following example is an extract of an AIX trace:

ID	ELAPSED_SEC	DELTA_MSEC	APPL	SYSCALL	KERNEL	INTERRUPT
...						
30D	1.189295104	0.000000	MQS FNC	Exit...	17726.1	xllListenSel
30D	1.189341184	0.046080	MQS CEI	Entry...	17726.1	xllSpinLockR
30D	1.189364992	0.023808	MQS FNC	Exit...	17726.1	xllSpinLockR
30D	1.189380096	0.015104	MQS CEI	Entry...	17726.1	xllSpinLockR
30D	1.189394816	0.014720	MQS FNC	Exit...	17726.1	xllSpinLockR
30D	1.189408512	0.013696	MQS CEI	Entry...	17726.1	xllSpinLockR
30D	1.189427328	0.018816	MQS FNC	Exit....	17726.1	xllSpinLockR
30D	1.189444480	0.017152	MQS CEI	Entry...	17726.1	xcsFreeQuick
30D	1.189461120	0.016640	MQS CEI	Entry....	17726.1	xllSpinLock
30D	1.189480320	0.019200	MQS FNC	Exit....	17726.1	xllSpinLock
30D	1.189592192	0.111872	MQS FNC	Entry....	17726.1	xstFreeCell
30D	1.189608448	0.016256	MQS FNC	Exit....	17726.1	xstFreeCell
30D	1.189658496	0.050048	MQS CEI	Entry....	17726.1	xllSpinLock
30D	1.189672832	0.014336	MQS FNC	Exit....	17726.1	xllSpinLock
30D	1.189691520	0.018688	MQS CEI	Exit...	17726.1	xcsFreeQuick
30D	1.189704064	0.012544	MQS CEI	Entry...	17726.1	xllSpinLockR
30D	1.189717504	0.013440	MQS FNC	Exit...	17726.1	xllSpinLockR
30D	1.189729536	0.012032	MQS FNC	Exit!..	17726.1	xllWaitSocket
30D	1.189744512	0.014976	MQS FNC	Exit!.	17726.1	xcsWaitEventSe
30D	1.189765376	0.020864	MQS CEI	Exit!	17726.1	zcpReceiveOnLin
30D	1.189792128	0.026752	MQS FNC	Entry	17726.1	zapInquireStatu
30D	1.189814400	0.022272	MQS FNC	Entry.	17726.1	xcsRequestMute
30D	1.189832064	0.017664	MQS FNC	Entry..	17726.1	xllSemGetVal
30D	1.189898240	0.066176	MQS FNC	Exit...	17726.1	xllSemGetVal
30E	1.204718976	14.820736	xcsRequestMutexSem	phmtx:30000f3c	Tim	
...						

Figure 42. Sample AIX trace

Notes:

1. In this example the data is truncated. In a real trace, the complete function names and return codes are present.
2. The return codes are given as values, not literals.

Tracing MQSeries for HP-UX and MQSeries for Sun Solaris

In MQSeries for HP-UX and Sun Solaris, you enable or modify tracing using the **strmqtrc** control command, which is described in “strmqtrc (Start MQSeries trace)” on page 295. To stop tracing, you use the **endmqtrc** control command, which is described in “endmqtrc (End MQSeries trace)” on page 263. You can display formatted trace output using the **dspmqtrc** control command, which is described in “dspmqtrc (Display MQSeries formatted trace output)” on page 255.

Trace files

All trace files are created in the directory `/var/mqm/trace`.

Note: It is possible to accommodate production of large trace files by mounting a temporary file system over this directory.

Trace-file names have the following format:

`AMQppppp.TRC`

where `ppppp` is the process identifier (PID) of the process producing the trace.

Notes:

1. The process identifier can contain fewer, or more, digits than shown in the example.
2. There is one trace file for each process running as part of the entity being traced.

Example trace data

Figure 43 on page 223 shows an extract from an MQSeries for HP-UX trace:


```

...
ID      ELAPSED_MICROSEC DELTA_MICROSEC  APPL      SYSCALL  KERNEL  INTERRUPT
30d     0                0                MQS FNC  Exit..... 18855.1  xcsChec
30d     292            292             MQS CEI  Entry..... 18855.1  xcsHSHM
30d     363            71              MQS CEI  Exit..... 18855.1  xcsHSHM
30d     420            57              MQS CEI  Entry..... 18855.1  xcsHSHM
30d     482            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d     539            57              MQS CEI  Entry..... 18855.1  xcsHSHM
30d     602            63              MQS CEI  Exit..... 18855.1  xcsHSHM
30d     659            57              MQS CEI  Entry..... 18855.1  xcsHSHM
30d     721            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d     779            58              MQS CEI  Entry..... 18855.1  xcsHSHM
30d     841            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d     899            58              MQS CEI  Entry..... 18855.1  xcsHSHM
30d     961            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d    1018            57              MQS CEI  Entry..... 18855.1  xcsHSHM
30d    1080            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d    1138            58              MQS CEI  Entry..... 18855.1  xcsHSHM
30d    1200            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d    1257            57              MQS CEI  Entry..... 18855.1  xcsHSHM
30d    1319            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d    1377            58              MQS CEI  Entry..... 18855.1  xcsHSHM
30d    1439            62              MQS CEI  Exit..... 18855.1  xcsHSHME
30d    1498            59              MQS FNC  Entry..... 18855.1  xcsAlloc
30d    1554            56              MQS CEI  Entry..... 18855.1  xcsHSHM
30d    1616            62              MQS CEI  Exit..... 18855.1  xcsHSHM
30d    1674            58              MQS FNC  Entry..... 18855.1  xllSpin
30d    1733            59              MQS FNC  Exit..... 18855.1  xllSpin
30e    1825            92              MQS      Signals Blocked with mask:
30e    1967            142             MQS      Data from xcsAllocateQuickCell Le
          FFFFFFFBFF FFFFFFFF FFFFFFFF FFFFFFFF
          FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
...

```

Figure 43. Sample HP-UX trace

Figure 44 shows an extract from an MQSeries for Sun Solaris trace:

ID	ELAPSED_MICROSEC	DELTA_MICROSEC	APPL	SYSCALL	KERNEL	INTERRUPT
...						
30d	0	0	MQS FNC	Exit..	5814.1	xcsCheckProcess
30d	247	247	MQS CEI	Entry..	5814.1	xcsHSHMEMBtoPTR
30d	301	54	MQS CEI	Exit..	5814.1	xcsHSHMEMBtoPTR
30d	343	42	MQS CEI	Entry..	5814.1	xcsHSHMEMBtoPTR
30d	387	44	MQS CEI	Exit..	5814.1	xcsHSHMEMBtoPTR
30d	428	41	MQS CEI	Entry..	5814.1	xcsHSHMEMBtoPTR
30d	472	44	MQS CEI	Exit..	5814.1	xcsHSHMEMBtoPTR
30d	514	42	MQS FNC	Entry..	5814.1	xcsAllocateQuic
30d	554	40	MQS CEI	Entry..	5814.1	xcsHSHMEMBtoPT
30d	598	44	MQS CEI	Exit...	5814.1	xcsHSHMEMBtoPT
30d	639	41	MQS FNC	Entry..	5814.1	xllSpinLockReq
30d	684	45	MQS FNC	Exit...	5814.1	xllSpinLockReq
30e	764	80	MQS	Signals Blocked with mask:		
30e	882	118	MQS	Data from xcsAllocateQuickCell Le		
				FFFFFFFF	00000FFF	00000000 00000000
30d	956	74	MQS CEI	Entry..	5814.1	xcsHSHMEMBtoPT
30d	1000	44	MQS CEI	Exit...	5814.1	xcsHSHMEMBtoPT
30d	1040	40	MQS FNC	Entry..	5814.1	xstAllocateCel
30d	1082	42	MQS FNC	Exit...	5814.1	xstAllocateCel
30e	1125	43	MQS	Signals Unblocked with mask:		
30e	1222	97	MQS	Data from xcsAllocateQuickCell Le		
				00024007	00000000	00000000 00000000
30d	1373	151	MQS FNC	Entry..	5814.1	xllSpinLockRel
...						

Figure 44. Sample MQSeries for Sun Solaris trace

Notes:

1. In these examples, the data is truncated. In a real trace, the complete function names and return codes are present.
2. The return codes are given as values, not literals.

Tracing MQSeries for OS/2 Warp and MQSeries for Windows NT

In MQSeries for OS/2 Warp and Windows NT, you enable or modify tracing using the **strmqtrc** control command, which is described in “strmqtrc (Start MQSeries trace)” on page 295. To stop tracing, you use the **endmqtrc** control command, which is described in “endmqtrc (End MQSeries trace)” on page 263.

Trace files

During the installation process, you can choose the drive on which trace files are to be located. However, the trace files are always placed in the directory `\<mqmwork\errors`, where `<mqmwork>` is the directory selected when MQSeries was installed to hold MQSeries data files.

Trace-file names have the following format:

```
AMQppppp.TRC
```

where *ppppp* is the process identifier (PID) of the process producing the trace.

Notes:

1. The process identifier can contain fewer, or more, digits than shown in the example.
2. There is one trace file for each process running as part of the entity being traced.

An example of MQSeries for Windows NT trace data

Figure 45 shows an extract from an MQSeries for Windows NT trace:

```

MQSeries Trace - Version 050000
!! - BuildDate Apr 27 1997

...
00A2FE ----> (0153) zcpCreatePacket
00A2FF -----> (0153) xcsAllocateMemBlock
00A300 -----> (0153) xstAllocateMemBlock
00A301 -----> (0153) xstAllocBlockInSharedMemSet
00A302 -----> (0153) xstAllocBlockInAnExtentOnList
00A303 -----> (0153) xstSerialiseExtentList
00A304 -----> (0153) xllSpinLockRequest
00A305 -----> (0153) xllAccessHandle
00A306 <----- (0153) xllAccessHandle (rc=OK)
00A307 <----- (0153) xllSpinLockRequest (rc=OK)
00A308 <----- (0153) xstSerialiseExtentList (rc=OK)
00A309 -----> (0153) xstAllocBlockInExtent
00A30A -----> (0153) xstSerialiseExtent
00A30B -----> (0153) xllSpinLockRequest
00A30C -----> (0153) xllAccessHandle
00A30D <----- (0153) xllAccessHandle (rc=OK)
00A30E <----- (0153) xllSpinLockRequest (rc=OK)
00A30F <----- (0153) xstSerialiseExtent (rc=OK)
00A310 -----> (0153) xstInitialiseBlock
00A311 <----- (0153) xstInitialiseBlock (rc=OK)

...

```

Figure 45. Sample MQSeries for Windows NT trace

Notes:

1. In this example the data is truncated. In a real trace, the complete function names and return codes are present.
2. The return codes are given as values, not literals.

First-failure support technology (FFST)

This section describes the role of first-failure support technology (FFST) in each of the MQSeries Version 5 products.

FFST: MQSeries for UNIX systems

For MQSeries for UNIX systems, FFST information is recorded in a file in the `/var/mqm/errors` directory.

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

The files are named `AMQnnnnn.mm.FDC`, where:

`nnnnn` Is the ID of the process reporting the error
`mm` Is a sequence number, normally 0

When a process creates an FFST record, it also sends a record to syslog. The record contains the name of the FFST file to assist in automatic problem tracking.

The syslog entry is made at the “user.error” level. See the operating-system documentation about `syslog.conf` for information about configuring this.

Some typical FFST data is shown in Figure 46.

```

MQSeries First Failure Symptom Report
=====

Date/Time      :- Friday July 14 14:06:52 BST 1995
Host Name     :- unknown
PIDS          :- 5697-249
LVLS          :- 220
Product Long Name :- MQSeries for Sun Solaris
Vendor        :- IBM
Probe Id      :- XC130003
Application Name :- MQM
Component     :- xehExcepti
Build Date    :- Jul 14 1995
Userid        :- 00000231 (mqm)
Process       :- 00015967
Major Errorcode :- xecSTOP
Minor Errorcode :- OK
Probe Type    :- HALT6109
Probe Severity :- 1
Probe Description :- AMQ6125: An internal MQSeries error has occurred.
Arith1        :- 11 b

MQM Function Stack
xllTidyUpSems
xcsFFST

MQM Trace History
...

```

Figure 46. Sample MQSeries for Sun Solaris First Failure Symptom Report

The Function Stack and Trace History are used by IBM to assist in problem determination. In most cases there is little that the system administrator can do when an FFST report is generated, apart from raising problems through the IBM Support Centers.

However, there are some problems that the system administrator might be able to solve. If the FFST shows “out of resource” or “out of space on device” descriptions when calling one of the IPC functions (for example, **semop** or **shmget**), it is likely that the relevant kernel parameter limit has been exceeded.

If the FFST report shows a problem with **setitimer**, it is likely that a change to the kernel timer parameters is needed.

To resolve these problems, increase the IPC limits, rebuild the kernel, and restart the machine. See the *MQSeries Quick Beginnings* book for your MQSeries product for further information.

FFST: MQSeries for OS/2 Warp and Windows NT

In MQSeries for OS/2 Warp and Windows NT, FFST information is recorded in a file in the `c:\mqm\errors` directory.

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

FFST files are named `AMQnnnnn.mm.FDC`, where:

<code>nnnnn</code>	Is the ID of the process reporting the error
<code>mm</code>	Is a sequence number, normally 0

When a process creates an FFST record it also sends a record to the Event Log. The record contains the name of the FFST file to assist in automatic problem tracking. The Event log entry is made at the “application” level.

A typical FFST log is shown in Figure 47 on page 228.

```

MQSeries First Failure Symptom Report
=====

Date/Time      :- Friday July 14 14:06:52 BST 1995
Host Name     :- unknown
PIDS          :- 5697175
LVLS          :- 220
Product Long Name :- MQSeries for WINDOWSNT
Vendor        :- IBM
Probe Id      :- XC130003
Application Name :- MQM
Component     :- xehExcepti
Build Date    :- Jul 14 1995
Userid       :- 00000231 (mqm)
Process      :- 00015967
Major Errorcode :- xecSTOP
Minor Errorcode :- OK
Probe Type    :- HALT6109
Probe Severity :- 1
Probe Description :- AMQ6125: An internal MQSeries error has occurred.
Arith1       :- 11 b

MQM Function Stack
x11TidyUpSems
xcsFFST

MQM Trace History
...

```

Figure 47. Sample MQSeries for Windows NT First Failure Symptom Report

The Function Stack and Trace History are used by IBM to assist in problem determination. In most cases there is little that the system administrator can do when an FFST record is generated, apart from raising problems through the IBM Support Center.

FFST: MQSeries for OS/2 Warp

In addition to the FFST files described in “FFST: MQSeries for OS/2 Warp and Windows NT” on page 227, MQSeries for OS/2 Warp also produces dumps in FFST/2 format.

FFST/2 (First Failure Support Technology for OS/2) is an IBM licensed program that improves availability for IBM software applications by providing:

- Immediate event notification
- First failure data capture for software events
- Automated event tracking and management

FFST/2 is a corequisite to MQSeries for OS/2 Warp installation. At installation time, MQSeries checks whether FFST/2 exists on your system and, if so, at what level. Refer to *MQSeries for OS/2 Warp V5.0 Quick Beginnings* for information about how FFST/2 is installed with MQSeries for OS/2 Warp.

Ensure that duplicate customized dumps are suppressed for FFST/2. This avoids the situation where you could receive multiple dumps when a software probe associated with FFST/2 is triggered more than once.

For information about using FFST/2, and how to suppress the duplicate dumps, refer to the *FFST/2 Administration Guide, S96F-8593*.

FFST/2 operation

There are a number of software probes included in the MQSeries code. These probes are triggered by specific events, usually error conditions. When a probe is triggered, the FFST/2 program can create the following output:

- A customized dump
- A symptom record
- A generic alert

You can use the diagnostic output to help you identify, track, and analyze an event. The diagnostic outputs are generated by default. You can override the defaults by making changes to the probe control table (PCT). For information about how to do this, refer to the *FFST/2 Administration Guide*.

MQSeries for OS/2 Warp entries in the PCT include the following fields:

program_id	6539B42
application	MQM
probe_id	A probe identifier that is set by MQSeries
options	The action that you want the FFST/2 program to take for the specified range of software probes

Using a symptom record

When an event triggers a software probe, the FFST/2 program creates a symptom record. When created, the symptom record can contain:

- The date and time the software probe was triggered
- Hardware and software vital product data
- Any error code information provided by the software probe
- The name of the dump file, if the software probe creates a customized dump file
- A unique problem identifier for each event
- The message number and the first 32 bytes of the message text, if the software probe specifies that a message is issued
- The primary symptom string, which uniquely identifies the event
- The secondary symptom string, if the software probe has one
- A brief description of the software probe, if one is available

The FFST/2 program places all symptom records in the OS/2 system error log. Parts of the symptom record are also included in the customized dump file. The maximum size of an entry in the OS/2 system error log is 4 KB. If the symptom record exceeds this limit, it is truncated.

To display the symptom records, display the system error log. This is accessed by opening the FFST/2 icon that is installed on your desktop and opening (double clicking on) the system error log. Note that after installation the FFST/2 icon could have been moved to another folder.

Using a customized dump

When the FFST/2 program creates a system dump for MQSeries following the triggering of a software probe, the dump file is named OS2SYSxx.DMP, where xx ranges from 00 to the number set when you configured FFST/2.

Whenever FFST/2 generates a dump, xx is incremented by 01. After xx reaches the maximum number of dumps, the FFST/2 program resets the number to 00 and begins overwriting existing dumps with new dumps.

If FFST/2 wraps, it adds a second batch of dump information to an already used dump file. When this dump file is viewed, the dump shown is of the oldest dump in the file. In order to view separate dumps within the dump file, you need to use the "identification" field. This will give a list of the problem identifiers stored in that dump file. You can then select the one you want to display.

For further information about FFST/2 records and how to use them, refer to the *FFST/2 Administration Guide*.

Problem determination with clients

An MQI client application receives MQRC_* reason codes in the same way as non-client MQI applications. However, there are additional reason codes for error conditions associated with clients. For example:

- Remote machine not responding
- Communications line error
- Invalid machine address

The most common time for errors to occur is when an application issues an MQCONN and receives the response MQRC_Q_MQR_NOT_AVAILABLE. An error message, written to the client log file, explains the cause of the error. Messages may also be logged at the server depending on the nature of the failure.

Terminating clients

Even though a client has terminated, it is still possible for the process at the server to be holding its queues open. Normally, this will only be for a short time until the communications layer detects that the partner has gone.

Error messages with clients

When an error occurs with a client system, error messages are put into the error files associated with the server, if possible. If an error cannot be placed there, the client code attempts to place the error message in an error log in the root directory of the client machine.

OS/2 and UNIX-systems clients

Error messages for OS/2 and UNIX-systems clients are placed in the error logs in the same way as they are for the respective MQSeries server systems. Typically these files appear in `/var/mqm/errors`.

DOS and Windows clients

The location of the log file AMQERR01.LOG is set by the MQDATA environment variable. The default location, if not overridden by MQDATA, is the C drive. Working in the DOS environment involves the environment variable MQDATA.

This is the default library used by the client code to store trace and error information; it also holds the directory name in which the qm.ini file is stored. (needed for NetBIOS setup). If not specified, it defaults to the C drive.

The names of the default files held in this library are:

AMQERR01.LOG For error messages.

AMQERR01.FDC For First Failure Data Capture messages.

For more information about clients, see the *MQSeries Clients* book.

Client problem determination

Part 2. Reference

Chapter 14. MQSeries control commands	235
Names of MQSeries objects	235
How to read syntax diagrams	236
Example syntax diagram	237
Syntax help	237
crtmqcvx (Data conversion)	238
crtmqm (Create queue manager)	240
dltmqm (Delete queue manager)	244
dmpmqlog (Dump log)	246
dspmqaut (Display authority)	248
dspmqcsv (Display command server)	252
dspmqfls (Display MQSeries files)	253
dspmqtrc (Display MQSeries formatted trace output)	255
dspmqtrn (Display MQSeries transactions)	256
endmqcsv (End command server)	258
endmqlsr (End listener)	260
endmqm (End queue manager)	261
endmqtrc (End MQSeries trace)	263
rcdmqimg (Record media image)	265
rcrmqobj (Recreate object)	267
rsvmqtrn (Resolve MQSeries transactions)	269
runmqchi (Run channel initiator)	271
runmqchl (Run channel)	272
runmqdlq (Run dead-letter queue handler)	273
runmqlsr (Run listener)	275
runmqsc (Run MQSeries commands)	277
runmqtrc (Start client trigger monitor)	280
runmqtrm (Start trigger monitor)	281
scmmqm (Add the queue manager to the Windows NT Service Control Manager)	282
setmqaut (Set/reset authority)	284
setmqprd (Enroll production license)	290
setmqtry (Start trial period)	291
strmqcsv (Start command server)	292
strmqm (Start queue manager)	293
strmqtrc (Start MQSeries trace)	295

Chapter 14. MQSeries control commands

This chapter contains reference material for the control commands supported by the MQSeries Version 5 products. Please note the following environment-specific information:

MQSeries for AIX

All commands in this chapter can be issued from an AIX shell. These commands are case-sensitive.

MQSeries for HP-UX

All commands in this chapter can be issued from an HP-UX shell. These commands are case sensitive.

MQSeries for OS/2 Warp

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

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

MQSeries for Sun Solaris

All commands in this chapter can be issued from a Solaris shell. These commands are case sensitive.

MQSeries for Windows NT

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

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

Names of MQSeries objects

In general, the names of MQSeries objects can have up to 48 characters. This rule applies to all the following objects:

- Queue managers
- Queues
- Process definitions

The maximum length of channel names is 20 characters.

The characters that can be used for all MQSeries names are:

- Uppercase A–Z
- Lowercase a–z
- Numerics 0–9
- Period (.)
- Underscore (_)
- Forward slash (/) (see note 1)
- Percent sign (%) (see note 1)

Notes:

1. Forward slash and percent are special characters. If you use either of these characters in a name, the name must be enclosed in double quotation marks whenever it is used.
2. Leading or embedded blanks are not allowed.
3. National language characters are not allowed.
4. Names may be enclosed in double quotation marks, but this is essential only if special characters are included in the name.

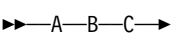
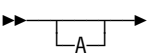
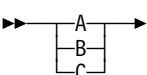
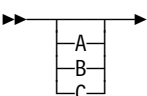
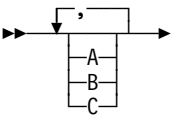
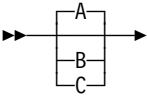
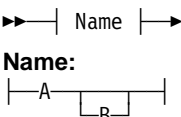
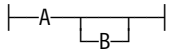
How to read syntax diagrams

This chapter contains syntax diagrams (sometimes referred to as “railroad” diagrams).

Each syntax diagram begins with a double right arrow and ends with a right and left arrow pair. Lines beginning with a single right arrow are continuation lines. You read a syntax diagram from left to right and from top to bottom, following the direction of the arrows.

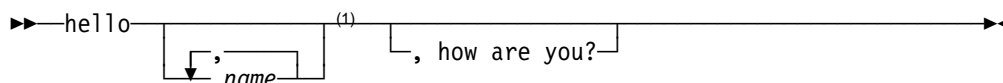
Other conventions used in syntax diagrams are:

Table 16. How to read syntax diagrams

Convention	Meaning
	You must specify values A, B, and C. Required values are shown on the main line of a syntax diagram.
	You may specify value A. Optional values are shown below the main line of a syntax diagram.
	Values A, B, and C are alternatives, one of which you must specify.
	Values A, B, and C are alternatives, one of which you may specify.
	You may specify one or more of the values A, B, and C. Any required separator for multiple or repeated values (in this example, the comma (,)) is shown on the arrow.
	Values A, B, and C are alternatives, one of which you may specify. If you specify none of the values shown, the default A (the value shown above the main line) is used.
 Name: 	The syntax fragment Name is shown separately from the main syntax diagram.

Example syntax diagram

Here is an example syntax diagram that describes the **hello** command:



Note:

¹ You can code up to three names.

According to the syntax diagram, these are all valid versions of the **hello** command:

```

hello
hello name
hello name, name
hello name, name, name
hello, how are you?
hello name, how are you?
hello name, name, how are you?
hello name, name, name, how are you?
  
```

Note that the space before the *name* value is significant, and that if you do not code *name* at all, you must still code the comma before how are you?.

Syntax help

You can obtain help for the syntax of any of the commands in this chapter by entering the command followed by a question mark. MQSeries responds by listing the syntax required for the selected command. The syntax shows all the parameters and variables associated with the command. Different forms of parentheses are used to indicate whether a parameter is required. For example:

```
CmdName [-x OptParam ] ( -c | -b ) argument
```

where:

CmdName Is the command name for which help has been requested.

[-x OptParam] Square brackets enclose one or more optional parameters. Where square brackets enclose multiple parameters, you can select no more than one of them.

(-c | -b) Brackets enclose multiple values, one of which you must select. In this example, you must select either flag c or flag b.

argument A mandatory argument.

Examples

1. Result of entering endmqm ?

```
endmqm [-z] [-c | -i | -p] QMgrName
```

2. Result of entering rcdmqimg ?

```
rcdmqimg [-z] [-m QMgrName] -t ObjType [GenericObjName]
```

crtmqcvx (Data conversion)

Purpose

Use the **crtmqcvx** command to create a fragment of code that performs data conversion on data type structures. The command generates a C function that can be used in an exit to convert C structures.

The command reads an input file containing structures to be converted, and writes an output file containing code fragments to convert those structures.

For information about using this command, see the *MQSeries Application Programming Guide*.

Syntax

```
▶▶—crtmqcvx—SourceFile—TargetFile————▶▶
```

Required parameters

SourceFile

Specifies the input file containing the C structures to be converted.

TargetFile

Specifies the output file containing the code fragments generated to convert the structures.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

Examples

The following example shows the results of using the data conversion command against a source C structure. The command issued is:

```
crtmqcvx source.tmp target.c
```

The input file, `source.tmp` looks like this:

```
/* This is a test C structure which can be converted by the */
/* crtmqcvx utility                                         */

struct my_structure
{
    int    code;
    MQLONG value;
};
```


The output file, `target.c`, produced by the command is shown below. You can use these code fragments in your applications to convert data structures. However, if you do so, be aware that the fragment uses macros supplied in the MQSeries header file `amqsvmha.h`.

```
MQLONG Convertmy_structure(
    PMQBYTE *in_cursor,
    PMQBYTE *out_cursor,
    PMQBYTE in_lastbyte,
    PMQBYTE out_lastbyte,
    MQHCONN hConn,
    MQLONG opts,
    MQLONG MsgEncoding,
    MQLONG ReqEncoding,
    MQLONG MsgCCSID,
    MQLONG ReqCCSID,
    MQLONG CompCode,
    MQLONG Reason)
{
    MQLONG ReturnCode = MQRC_NONE;

    ConvertLong(1); /* code */

    AlignLong();
    ConvertLong(1); /* value */

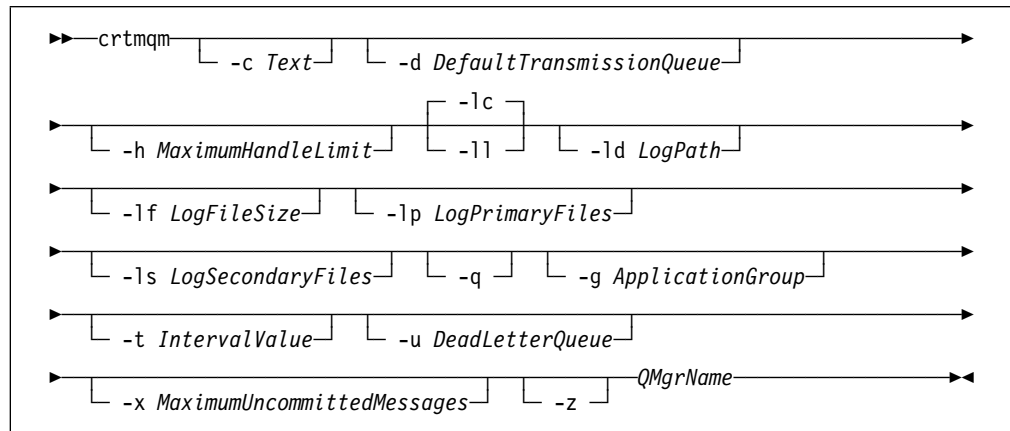
Fail:
    return(ReturnCode);
}
```

crtmqm (Create queue manager)

Purpose

Use the **crtmqm** command to create a local queue manager and define the default and system objects. The objects created by **crtmqm** are listed in Appendix A, “System and default objects” on page 301. When a queue manager has been created, use the **strmqm** command to start it.

Syntax



Required parameters

QMgrName

Specifies the name of the queue manager to be created. The name can contain up to 48 characters. This must be the last item in the command.

Optional parameters

-c *Text*

Specifies some descriptive text for this queue manager. You can use up to 64 characters; the default is all blanks.

If special characters are required, the description must be enclosed in double quotes. The maximum number of characters is reduced if the system is using a double-byte character set (DBCS).

-d *DefaultTransmissionQueue*

Specifies the name of the local transmission queue that remote messages are placed on if a transmission queue is not explicitly defined for their destination. There is no default.

-h *MaximumHandleLimit*

Specifies the maximum number of handles that any one application can have open at the same time.

Specify a value in the range 1 through 999 999 999. The default value is 256.

The next six parameter descriptions relate to logging, which is described in “Using the log for recovery” on page 197.

-lc Circular logging is to be used. This is the default logging method.

-ll Linear logging is to be used.

-ld *LogPath*

Specifies the directory to be used to hold log files.

In MQSeries for UNIX systems, the default is `/var/mqm/log`.

User ID `mqm` and group `mqm` must have full authorities to the log files. If you change the locations of these files, you must give these authorities yourself. This occurs automatically if the log files are in their default locations.

In MQSeries for OS/2 Warp and Windows NT, the default is `C:\MQMLOG` (assuming that C is your data drive).

-lf *LogFileSize*

Specifies the size of the log files in units of 4 KB.

In MQSeries for UNIX systems, the minimum value is 64, and the maximum is 16384. The default value is 1024, giving a default log size of 4 MB.

In MQSeries for OS/2 Warp and Windows NT, the minimum value is 32, and the maximum is 4095. The default value is 256, giving a default log size of 1 MB.

-lp *LogPrimaryFiles*

Specifies the number of primary log files to be allocated. The default value is 3, the minimum is 2, and the maximum is 62.

-ls *LogSecondaryFiles*

Specifies the number of secondary log files to be allocated. The default value is 2, the minimum is 1, and the maximum is 61.

Note: The total number of log files is restricted to 63, regardless of the number requested.

-q Specifies that this queue manager is to be made the default queue manager. The new queue manager replaces any existing default queue manager.

If you accidentally use this flag and wish to revert to an existing queue manager as the default queue manager, you can edit the *DefaultQueueManager* stanza in the MQSeries configuration file. See Chapter 7, "Configuration files" on page 99 for information about configuration files.

-g *ApplicationGroup*

Specifies the name of the group whose members are allowed to:

- Run MQI applications
- Update all IPCC resources
- Change the contents of some queue manager directories

This option applies only to MQSeries for AIX, Sun Solaris, and HP-UX.

The default value is `-g all`, which allows unrestricted access.

The `-g ApplicationGroup` value is recorded in the queue manager configuration file, `qm.ini`.

-t *IntervalValue*

Specifies the trigger time interval in milliseconds for all queues controlled by this queue manager. This value specifies the time after the receipt of a trigger-generating message when triggering is suspended. That is, if the arrival of a message on a queue causes a trigger message to be put on the initiation queue, any message arriving on the same queue within the specified interval does not generate another trigger message.

You can use the trigger time interval to ensure that your application is allowed sufficient time to deal with a trigger condition before it is alerted to deal with another on the same queue. You may wish to see all trigger events that happen; if so, set a low or zero value in this field.

Specify a value in the range 0 through 999 999 999. The default is 999 999 999 milliseconds, a time of more than 11 days. Allowing the default to be used effectively means that triggering is disabled after the first trigger message. However, triggering can be reenabled by an application servicing the queue using an alter queue command to reset the trigger attribute.

-u *DeadLetterQueue*

Specifies the name of the local queue that is to be used as the dead-letter (undelivered-message) queue. Messages are put on this queue if they cannot be routed to their correct destination.

The default if the attribute is omitted is no dead-letter queue.

-x *MaximumUncommittedMessages*

Specifies the maximum number of uncommitted messages under any one syncpoint. That is, the sum of:

- The number of messages that can be retrieved from queues
- The number of messages that can be put on queues
- Any trigger messages generated within this unit of work

This limit does not apply to messages that are retrieved or put outside a syncpoint.

Specify a value in the range 1 through 10 000. The default value is 1000 uncommitted messages.

-z Suppresses error messages.

This flag is normally used within MQSeries to suppress unwanted error messages. As use of this flag could result in loss of information, it is recommended that you do not use it when entering commands on a command line.

Return codes

0	Queue manager created
8	Queue manager already exists
49	Queue manager stopping
69	Storage not available
70	Queue space not available
71	Unexpected error
72	Queue manager name error
100	Log location invalid
111	Queue manager created. However, there was a problem processing the default queue manager definition in the product configuration file. The default queue manager specification may be incorrect.
115	Invalid log size

Examples

1. This command creates a default queue manager named `Paint.queue.manager`, which is given a description of `Paint shop`, and creates the system and default objects. It also specifies that linear logging is to be used:

```
crtmqm -c "Paint shop" -ll -q Paint.queue.manager
```

2. This command creates a default queue manager named `Paint.queue.manager`, creates the system and default objects, and requests two primary and three secondary log files:

```
crtmqm -c "Paint shop" -ll -lp 2 -ls 3 -q Paint.queue.manager
```

3. This command creates a queue manager called `travel`, creates the system and default objects, sets the trigger interval to 5000 milliseconds (or 5 seconds), and specifies `SYSTEM.DEAD.LETTER.QUEUE` as its dead-letter queue.

```
crtmqm -t 5000 -u SYSTEM.DEAD.LETTER.QUEUE travel
```

Related commands

<code>strmqm</code>	Start queue manager
<code>endmqm</code>	End queue manager
<code>dlmqm</code>	Delete queue manager

dltmqm (Delete queue manager)

Purpose

Use the **dltmqm** command to delete a specified queue manager. All objects associated with this queue manager are also deleted. Before you can delete a queue manager you must end it using the **endmqm** command.

In MQSeries for Windows NT, if you attempt to delete a queue manager when queue manager files are open, an error may occur. In this case, close the files and reissue the command.

Syntax

```
▶▶ dltmqm [-z] QMgrName ▶▶
```

Required parameters

QMGrName

Specifies the name of the queue manager to be deleted.

Optional parameters

-z Suppresses error messages.

Return codes

0	Queue manager deleted
3	Queue manager being created
5	Queue manager running
16	Queue manager does not exist
49	Queue manager stopping
69	Storage not available
71	Unexpected error
72	Queue manager name error
100	Log location invalid
112	Queue manager deleted. However, there was a problem processing the default queue manager definition in the product configuration file. The default queue manager specification may be incorrect.

Examples

1. The following command deletes the queue manager saturn.queue.manager.

```
dltmqm saturn.queue.manager
```

2. The following command deletes the queue manager `travel` and also suppresses any messages caused by the command.

```
dltmqm -z travel
```

Related commands

<code>crtmqm</code>	Create queue manager
<code>strmqm</code>	Start queue manager
<code>endmqm</code>	End queue manager


dmpmqlog (Dump log)

Purpose

Use the **dmpmqlog** command to dump a formatted version of the MQSeries system log.

The log to be dumped must have been created on the same type of operating system as that being used to issue the command.

Syntax



Optional parameters

Dump start point

Use one of the following parameters to specify the log sequence number (LSN) at which the dump should start. If no start point is specified, dumping starts by default from the LSN of the first record in the active portion of the log.

-b Specifies that dumping should start from the base LSN. The base LSN identifies the start of the log extent that contains the start of the active portion of the log.

-s *StartLSN*
Specifies that dumping should start from the specified LSN. The LSN is specified in the format `nnnn:nnnn:nnnn:nnnn`. If you are using a circular log, the LSN value must be equal to or greater than the base LSN value of the log.

-n *ExtentNumber*
Specifies that dumping should start from the specified extent number. The extent number must be in the range 0–9 999 999.

This parameter is valid only for queue managers whose *LogType* (as recorded in the configuration file, `qm.ini`) is LINEAR.

-e *EndLSN*
Specifies that dumping should end at the specified LSN. The LSN is specified in the format `nnnn:nnnn:nnnn:nnnn`.

-f *LogFilePath*
Is the absolute, rather than the relative, directory path name to the log files. The specified directory must contain the log header file (`amqh1ct1.1fh`) and a subdirectory called `active`. The `active` subdirectory must contain the log files. By default, log files are assumed to be in the directories specified in the `mqs.ini` and `qm.ini` files. If this option is used then queue names, associated with queue identifiers, will only be shown in the dump if a queue manager name is specified explicitly for the `-m` option and that queue manager has the

object catalog file in its directory path. On a system that supports long file names this file is named `qmqmobjcat` and, in order to map the queue identifiers to queue names, it must be the file used when the log files were created. As an example, for a queue manager named `qm1`, the object catalog file is located in the directory `..\qmgrs\qm1\qmanager\`. To achieve this mapping, it may be necessary to create a temporary queue manager, for example named `tmpq`, replace its object catalog with the one associated with the specific log files, and then start `dmpmqlog`, specifying `-m tmpq` and `-f` with the absolute directory path name to the log files.

-m QMgrName

Is the name of the queue manager. If this parameter is omitted, the name of the default queue manager is used.

The queue manager you specify, or default to, must not be running when the **dmpmqlog** command is issued. Similarly, the queue manager must not be started while **dmpmqlog** is running.

dspmqaout (Display authority)

Purpose

Use the **dspmqaout** command to display the current authorizations to a specified object.

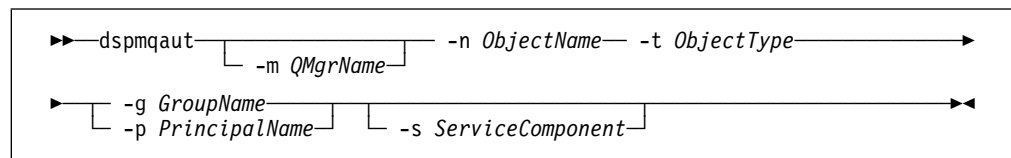
If a user ID is a member of more than one group in MQSeries for UNIX systems, display the combined authorizations of all of the groups.

Only one group or principal may be specified.

You can use this command meaningfully in MQSeries for OS/2 Warp only if an authorization service component has been installed for the current queue manager. For MQSeries for OS/2 Warp, this does not occur by default, and no such component is supplied with the product. Furthermore, when you issue this command, the results always indicate that any group or principal has all authorizations.

For more information about authorization service components, see the *MQSeries Programmable System Management* book.

Syntax



Required parameters

-n *ObjectName*

Specifies the name of a queue manager, queue, or process definition on which the inquiry is to be made.

This is a required parameter, **unless** the inquiry relates to the queue manager itself, in which case it must not be included.

-t *ObjectType*

Specifies the type of object on which the inquiry is to be made. Possible values are:

- queue** or **q** A queue or queues matching the object type parameter
- qmgr** A queue manager object
- process** or **prcs** A process

Optional parameters

-m *QMgrName*

Specifies the name of the queue manager on which the inquiry is to be made.

-g *GroupName*

Specifies the name of the user group on which the inquiry is to be made. You can specify only one name, which must be the name of an existing user group.

-p *PrincipalName*

Specifies the name of a user whose authorizations to the specified object are to be displayed.

-s *ServiceComponent*

This parameter applies only if you are using installable authorization services, otherwise it is ignored.

If installable authorization services are supported, this parameter specifies the name of the authorization service to which the authorizations apply. This parameter is optional; if it is not specified, the authorization update is made to the first installable component for the service.

Returned parameters

This command returns an authorization list, which can contain none, one, or more authorization values. Each authorization value returned means that any user ID in the specified group has the authority to perform the operation defined by that value.

Table 17 shows the authorities that can be given to the different object types.

Authority	Queue	Process	Qmgr
all	√	√	√
alladm	√	√	√
allmqi	√	√	√
altusr			√
browse	√		
chg	√	√	√
chgaut	√	√	√
clr	√		
connect			√
crt	√	√	√
dlt	√	√	√
dsp	√	√	√
get	√		
inq	√	√	√
passall	√		
passid	√		
put	√		
set	√	√	√
setall	√		√
setid	√		√

The following list defines the authorizations associated with each value:

all	Use all operations relevant to the object.
alladm	Perform all administration operations relevant to the object.
allmqi	Use all MQI calls relevant to the object.
altusr	Specify an alternate user ID on an MQI call.
browse	Retrieve a message from a queue by issuing an MQGET call with the BROWSE option.
chg	Change the attributes of the specified object, using the appropriate command set.
chgaut	Specify authorizations for other groups of users on the object, using the setmqaut command.
clr	Clear a queue (PCF command Clear queue only).
connect	Connect the application to the specified queue manager by issuing an MQCONN call.
crt	Create objects of the specified type, using the appropriate command set.
dlt	Delete the specified object, using the appropriate command set.
dsp	Display the attributes of the specified object, using the appropriate command set.
get	Retrieve a message from a queue by issuing an MQGET call.
inq	Make an inquiry on a specific queue by issuing an MQINQ call.
passall	Pass all context.
passid	Pass the identity context.
put	Put a message on a specific queue by issuing an MQPUT call.
set	Set attributes on a queue from the MQI by issuing an MQSET call.
setall	Set all context on a queue.
setid	Set the identity context on a queue.

The authorizations for administration operations, where supported, apply to these command sets:

- Control commands
- MQSC commands
- PCF commands

Return codes

0	Successful operation
36	Invalid arguments supplied
40	Queue manager not available
49	Queue manager stopping
69	Storage not available
71	Unexpected error
72	Queue manager name error
133	Unknown object name
145	Unexpected object name
146	Object name missing
147	Object type missing
148	Invalid object type
149	Entity name missing

Examples

The following example shows a command to display the authorizations on queue manager saturn.queue.manager associated with user group staff:

```
dspmqaout -m saturn.queue.manager -t qmgr -g staff
```

The results from this command are:

```
Entity staff has the following authorizations for object :  
  get  
  browse  
  put  
  inq  
  set  
  connect  
  altusr  
  passid  
  passall  
  setid
```

Related commands

setmqaut

Set or reset authority

dspmqcsv (Display command server)

Purpose

Use the **dspmqcsv** command to display the status of the command server for the specified queue manager.

The status can be one of the following:

- Starting
- Running
- Running with SYSTEM.ADMIN.COMMAND.QUEUE not enabled for gets
- Ending
- Stopped

Syntax

```
▶▶—dspmqcsv—QMgrName—————▶▶
```

Required parameters

QMgrName

Specifies the name of the local queue manager for which the command server status is being requested.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

Examples

The following command displays the status of the command server associated with `venus.q.mgr`:

```
dspmqcsv venus.q.mgr
```

Related commands

<code>strmqcsv</code>	Start a command server
<code>endmqcsv</code>	End a command server

dspmqfls (Display MQSeries files)

Purpose

Use the **dspmqfls** command to display the real file system name for all MQSeries objects that match a specified criterion. You can use this command to identify the files associated with a particular MQSeries object. This is useful for backing up specific objects. See “Understanding MQSeries file names” on page 30 for further information about name transformation.

Syntax

```

  ──▶ dspmqfls [ -m QMgrName ] [ -t ObjType ] GenericObjName ──▶

```

Required parameters

GenericObjName

Specifies the name of the MQSeries object. The name is a string with no flag and is a required parameter. If the name is omitted an error is returned.

This parameter supports a wild card character * at the end of the string.

Optional parameters

-m *QMgrName*

Specifies the name of the queue manager for which files are to be examined. If omitted, the command operates on the default queue manager.

-t *ObjType*

Specifies the MQSeries object type. The following list shows the valid object types. The abbreviated name is shown first followed by the full name.

* or all	All object types; this is the default
q or queue	A queue or queues matching the object name parameter
ql or qlocal	A local queue
qa or qalias	An alias queue
qr or qremote	A remote queue
qm or qmodel	A model queue
qmgr	A queue manager object
prcs or process	A process

Notes:

1. The **dspmqfls** command displays the directory containing the queue, **not** the name of the queue itself.
2. In MQSeries for UNIX systems, you need to prevent the shell from interpreting the meaning of special characters, for example, '*'. To accomplish this, use 'quoting'.

dspmqfls

There are a number of ways of 'quoting', depending on your shell. For example, single quotation marks, double quotation marks, or a backslash are used by some shells.

Return codes

0	Command completed normally
10	Command completed but not entirely as expected
20	An error occurred during processing

Examples

1. The following command displays the details of all objects with names beginning SYSTEM.ADMIN that are defined on the default queue manager.

```
dspmqfls SYSTEM.ADMIN*
```

2. The following command displays file details for all processes with names beginning PROC defined on queue manager RADIUS.

```
dspmqfls -m RADIUS -t prcs PROC*
```


dspmqtrc (Display MQSeries formatted trace output)

Special note


The **dspmqtrc** command is not supported by these MQSeries products:

- MQSeries for AIX
- MQSeries for OS/2 Warp
- MQSeries for Windows NT

Purpose

Use the **dspmqtrc** command to display MQSeries formatted trace output.

Syntax



Required parameters

InputFileName

Specifies the name of the file containing the unformatted trace. For example `/var/mqm/trace/AMQ12345.TRC`.

Optional parameters

-t *FormatTemplate*

Specifies the name of the template file containing details of how to display the trace. The default value is `mqmtop/lib/amqtrc.fmt`.

-h Omit header information from the report.

-o *output_filename*

The name of the file into which to write formatted data.

Related commands

<code>endmqtrc</code>	End MQSeries trace
<code>strmqtrc</code>	Start MQSeries trace

dspmqrn (Display MQSeries transactions)

Purpose

Use the **dspmqrn** command to display details of in-doubt transactions. Such transactions can be either internally or externally coordinated.

For each in-doubt transaction, a transaction number (a human-readable transaction identifier), the transaction state, and the transaction ID are displayed. (Transaction IDs can be up to 128 characters long, hence the need for a transaction number.)

Syntax

```

▶▶ dspmqrn [ -e ] [ -i ] [ -m QMgrName ] ▶▶

```

Optional parameters

- e** Requests details of externally coordinated, in-doubt transactions. Such transactions are those for which MQSeries has been asked to prepare to commit, but has not yet been informed of the transaction outcome.
- i** Requests details of internally coordinated, in-doubt transactions. Such transactions are those for which each resource manager has been asked to prepare to commit, but MQSeries has yet to inform the resource managers of the transaction outcome.

Information about the (deduced) state of the transaction in each of its participating resource managers is displayed. This information can help you assess the effects of failure in a particular resource manager.

- m** *QMgrName*
Specifies the name of the queue manager whose transactions are to be displayed. If no name is specified, the default queue manager's transactions are displayed.

Note: If you specify neither **-e** nor **-i**, details of both internally and externally coordinated in-doubt transactions are displayed.

Return codes

0	Successful operation
36	Invalid arguments supplied
40	Queue manager not available
49	Queue manager stopping
69	Storage not available
71	Unexpected error
72	Queue manager name error
102	No transactions found

Related commands

rsvmqtrn


Resolve MQSeries transaction

endmqcsv (End command server)

Purpose

Use the **endmqcsv** command to stop the command server on the specified queue manager.

Syntax



Required parameters

QMgrName

Specifies the name of the queue manager for which the command server is to be ended.

Optional parameters

- c Specifies that the command server is to be stopped in a controlled manner. The command server is allowed to complete the processing of any command message that it has already started. No new message is read from the command queue.
This is the default.
- i Specifies that the command server is to be stopped immediately. Actions associated with a command message currently being processed may not be completed.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

Examples

1. The following command stops the command server on queue manager saturn.queue.manager:

```
endmqcsv -c saturn.queue.manager
```

The command server can complete processing any command it has already started before it stops. Any new commands received remain unprocessed in the command queue until the command server is restarted.

2. The following command stops the command server on queue manager `pluto` immediately:

```
endmqcsv -i pluto
```

Related commands

`strmqcsv`
`dspmqcsv`

Start a command server
Display the status of a command server

endmq1sr (End listener)

Special note

The **endmq1sr** command is not supported by these MQSeries products:


- MQSeries for AIX
- MQSeries for HP-UX
- MQSeries for Sun Solaris

Purpose

The **endmq1sr** command ends all listener process for the specified queue manager.

The queue manager must be stopped before the **endmq1sr** command is issued.

Syntax



Optional parameters

-m *QMgrName*

Specifies the name of the queue manager. If no name is specified, the processing will be done for the default queue manager.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

endmqm (End queue manager)

Purpose


Use the **endmqm** command to end (stop) a specified local queue manager. This command stops a queue manager in one of three modes:

- Normal or quiesced shutdown
- Immediate shutdown
- Preemptive shutdown

The attributes of the queue manager and the objects associated with it are not affected. You can restart the queue manager using the **strmqm** (Start queue manager) command.

To delete a queue manager, you must stop it and then use the **dltmqm** (Delete queue manager) command.

Syntax



Required parameters

QMGrName

Specifies the name of the message queue manager to be stopped.

Optional parameters

- c Controlled (or quiesced) shutdown. The queue manager stops but only after all applications have disconnected. Any MQI calls currently being processed are completed. This is the default.
- i Immediate shutdown. The queue manager stops after it has completed all the MQI calls currently being processed. Any MQI requests issued after the command has been issued fail. Any incomplete units of work are rolled back when the queue manager is next started.

- p Preemptive shutdown.

Use this type of shutdown only in exceptional circumstances. For example, when a queue manager does not stop as a result of a normal **endmqm** command.

The queue manager stops without waiting for applications to disconnect or for MQI calls to complete. This can give unpredictable results for MQSeries applications. All processes in the queue manager that fail to stop are terminated 30 seconds after the command is issued.

- z Suppresses error messages on the command.

endmqm

Return codes

0	Queue manager ended
3	Queue manager being created
16	Queue manager does not exist
40	Queue manager not available
49	Queue manager stopping
69	Storage not available
71	Unexpected error
72	Queue manager name error

Examples

The following examples show commands that end (stop) the specified queue managers.

1. This command ends the queue manager named `mercury.queue.manager` in a controlled way. All applications currently connected are allowed to disconnect.

```
endmqm mercury.queue.manager
```

2. This command ends the queue manager named `saturn.queue.manager` immediately. All current MQI calls complete, but no new ones are allowed.

```
endmqm -i saturn.queue.manager
```

Related commands

<code>crtmqm</code>	Create a queue manager
<code>strmqm</code>	Start a queue manager
<code>dltmqm</code>	Delete a queue manager

endmqtrc (End MQSeries trace)

Special note

The **endmqtrc** command is not supported by MQSeries for AIX.

Purpose

Use the **endmqtrc** command to end tracing for the specified entity or all entities.

Syntax

The syntax of this command in MQSeries for HP-UX and Sun Solaris is as follows:

```

▶▶—endmqtrc— [ -m QMgrName ] [ -e ] [ -a ] ▶▶

```

The syntax of this command in MQSeries for OS/2 Warp and Windows NT is as follows:

```

▶▶—endmqtrc— ▶▶

```

Optional parameters

The following parameters can be specified in MQSeries for HP-UX and Sun Solaris only:

-m *QMgrName*

Is the name of the queue manager for which tracing is to be ended.

A maximum of one -m flag and associated queue manager name can be supplied on the command.

A queue manager name and -m flag can be specified on the same command as the -e flag.

-e If this flag is specified, early tracing is ended.

-a If this flag is specified all tracing is ended.

This flag **must** be specified alone.

Return codes

AMQ5611 This message is issued if arguments that are not valid are supplied to the command.

Examples

This command ends tracing of data for a queue manager called QM1.

```
endmqtrc -m QM1
```

endmqtrc

Related commands

dspmqtrc
strmqtrc

Display formatted trace output
Start MQSeries trace

rcdmqimg (Record media image)

Purpose

Use the **rcdmqimg** command to write an image of an MQSeries object, or group of objects, to the log for use in media recovery. Use the associated command **rcrmqobj** to recreate the object from the image.

This command is used with an active queue manager. Further activity on the queue manager is logged so that, although the image becomes out of date, the log records reflect any changes to the object.

Syntax

```

  ►► rcdmqimg [ -m QMgrName ] [ -z ] -t ObjectType GenericObjName ►►

```

Required parameters

GenericObjName

Specifies the name of the object that is to be recorded. This parameter can have a trailing asterisk to indicate that any objects with names matching the portion of the name prior to the asterisk are to be recorded.

This parameter is required **unless** you are recording a queue manager object or the channel synchronization file. If you specify an object name for the channel synchronization file, it is ignored.

-t *ObjectType*

Specifies the types of object whose images are to be recorded. Valid object types are:

prcs or process	Processes
q or queue	All types of queue
ql or qlocal	Local queues
qa or qalias	Alias queues
qr or qremote	Remote queues
qm or qmodel	Model queues
qmgr	Queue manager object
syncfile	Channel synchronization file
* or all	All of the above

Note: When using MQSeries for UNIX systems, you need to prevent the shell from interpreting the meaning of special characters, for example, '*'. To accomplish this, use 'quoting'. There are a number of ways of 'quoting' depending on your shell. For example, single quotation marks, double quotation marks, or a backslash are used by some shells.

Optional parameters

- m** *QMgrName*
Specifies the name of the queue manager for which images are to be recorded. If omitted, the command operates on the default queue manager.
- z** Suppresses error messages.

Return codes

0	Successful operation
36	Invalid arguments supplied
40	Queue manager not available
49	Queue manager stopping
68	Media recovery is not supported
69	Storage not available
71	Unexpected error
72	Queue manager name error
119	User not authorized
128	No objects processed
131	Resource problem
132	Object damaged
135	Temporary object cannot be recorded

Examples

The following command records an image of the queue manager object saturn.queue.manager in the log.

```
rctmqimg -t qmgr -m saturn.queue.manager
```

Related commands

rctmqobj Recreate a queue manager object

rcrmqobj (Recreate object)

Purpose

Use the **rcrmqobj** command to recreate an object, or group of objects, from their images contained in the log. Use the associated command, **rcdmqimg**, to record the object images to the log.

This command must be used on a running queue manager. All activity on the queue manager after the image was recorded is logged. To recreate an object you must replay the log to recreate events that occurred after the object image was captured.

Syntax

```

  ►► rcrmqobj [ -m QMGrName ] [ -z ] -t ObjectType GenericObjName ►►

```

Required parameters

GenericObjName

Specifies the name of the object that is to be recreated. This parameter can have a trailing asterisk to indicate that any objects with names matching the portion of the name prior to the asterisk are to be recreated.

This parameter is required **unless** the object type is the channel synchronization file; if an object name is supplied for this object type, it is ignored.

-t ObjectType

Specifies the types of object to be recreated. Valid object types are:

prcs or process	Processes
q or queue	All types of queue
ql or qlocal	Local queues
qa or qalias	Alias queues
qr or qremote	Remote queues
qm or qmodel	Model queues
syncfile	The channel synchronization file
* or all	All the above

Note: When using MQSeries for UNIX systems, you need to prevent the shell from interpreting the meaning of special characters, for example, '*'. To accomplish this, use 'quoting'. There are a number of ways of 'quoting' depending on your shell. For example, single quotation marks, double quotation marks, or a backslash are used by some shells.

Optional parameters

- m** *QMgrName*
Specifies the name of the queue manager for which objects are to be recreated. If omitted, the command operates on the default queue manager.
- z** Suppresses error messages.

Return codes

0	Successful operation
36	Invalid arguments supplied
40	Queue manager not available
49	Queue manager stopping
66	Media image not available
68	Media recovery is not supported
69	Storage not available
71	Unexpected error
72	Queue manager name error
119	User not authorized
128	No objects processed
135	Temporary object cannot be recovered
136	Object in use

Examples

1. The following command recreates all local queues for the default queue manager:

```
rcrmqobj -t ql *
```

2. The following command recreates all remote queues associated with queue manager store:

```
rcrmqobj -m store -t qr
```

Related commands

- rcdmqimg Record an MQSeries object in the log


rsvmqtrn (Resolve MQSeries transactions)

Purpose

Use the **rsvmqtrn** command to commit or backout internally or externally coordinated in-doubt transactions.

Use this command only when you are certain that transactions cannot be resolved by the normal protocols. Issuing this command may result in the loss of transactional integrity between resource managers for a distributed transaction.

Syntax



Required parameters

-m *QMgrName*

Specifies the name of the queue manager. This parameter is mandatory.

Optional parameters

- a** Specifies that the queue manager should attempt to resolve all internally coordinated, in-doubt transactions (that is, all global units of work).
- b** Specifies that the named transaction is to be backed out. This flag is valid for externally coordinated transactions (that is, for external units of work) only.
- c** Specifies that the named transaction is to be committed. This flag is valid for externally coordinated transactions (that is, for external units of work) only.
- r** *RMIId*
Identifies the resource manager to which the commit or backout decision applies. This flag is valid for internally coordinated transactions only, and for resource managers that are no longer configured in the queue manager's qm.ini file. The outcome delivered will be consistent with the decision reached by MQSeries for the transaction.

Transaction

Is the transaction number of the transaction being committed or backed out. To discover the relevant transaction number, use the **dspmqtrn** command. This parameter is required with the **-b**, **-c**, and **-r** *RMIId* parameters.

rsvmqtrn

Return codes

0	Successful operation
32	Transactions could not be resolved
34	Resource manager not recognized
35	Resource manager not permanently unavailable
36	Invalid arguments supplied
40	Queue manager not available
49	Queue manager stopping
69	Storage not available
71	Unexpected error
72	Queue manager name error
85	Transactions not known

Related commands

dspmqtrn	Display list of prepared transactions
----------	---------------------------------------

runmqchi (Run channel initiator)

Purpose

Use the **runmqchi** command to run a channel initiator process. For more information about the use of this command, refer to the *MQSeries Intercommunication* book.

Syntax

```

▶▶—runmqchi—┌─q InitiationQName─┐ ┌─m QMgrName─┐▶▶

```

Optional parameters

-q *InitiationQName*

Specifies the name of the initiation queue to be processed by this channel initiator. If not specified, SYSTEM.CHANNEL.INITQ is used.

-m *QMgrName*

Specifies the name of the queue manager on which the initiation queue exists. If the name is omitted, the default queue manager is used.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

If errors occur that result in return codes of either 10 or 20, you should review the queue manager error log that the channel is associated with for the error messages. You should also review the @SYSTEM error log, as problems that occur before the channel is associated with the queue manager are recorded there. For more information about error logs, see “Error logs” on page 215.

runmqchl (Run channel)

Purpose

Use the **runmqchl** command to run either a Sender (SDR) or a Requester (RQSTR) channel.

The channel runs synchronously. To stop the channel, issue the MQSC command STOP CHANNEL.

Syntax

```
▶▶—runmqchl— -c ChannelName [ -m QMgrName ]▶▶
```

Required parameters

-c *ChannelName*

Specifies the name of the channel to run.

Optional parameters

-m *QMgrName*

Specifies the name of the queue manager with which this channel is associated. If no name is specified, the default queue manager is used.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing


If return codes 10 or 20 are generated, review the error log of the associated queue manager for the error messages. You should also review the @SYSTEM error log because problems that occur before the channel is associated with the queue manager are recorded there.

runmqdlq (Run dead-letter queue handler)

Purpose

Use the **runmqdlq** command to start the dead-letter queue (DLQ) handler, a utility that you can run to monitor and handle messages on a dead-letter queue.

Syntax



Description

The dead-letter queue handler can be used to perform various actions on selected messages by specifying a set of rules that can both select a message and define the action to be performed on that message.

The **runmqdlq** command takes its input from `stdin`. When the command is processed, the results and a summary are put into a report that is sent to `stdout`.

By taking `stdin` from the keyboard, you can enter **runmqdlq** rules interactively.

By redirecting the input from a file, you can apply a rules table to the specified queue. The rules table must contain at least one rule.

If the DLQ handler is used without redirecting `stdin` from a file (the rules table), the DLQ handler:

- Reads its input from the keyboard. In MQSeries for AIX, HP-UX, and Sun Solaris, it does not start to process the named queue until it receives an `end_of_file` (Ctrl+D) character. In MQSeries for OS/2 Warp and Windows NT, it does not start to process the named queue until exactly the following key sequence is pressed: Ctrl+Z, Enter, Ctrl+Z, Enter.

For more information about rules tables and how to construct them, see “The DLQ handler rules table” on page 116.

Note: For MQSeries for OS/2 Warp, the final rule must be terminated by an end-of-line character.

Optional parameters

The MQSC rules for comment lines and for joining lines also apply to the DLQ handler input parameters.

QName Specifies the name of the queue to be processed.

If no name is specified the dead-letter queue defined for the local queue manager is used. If one or more blanks (' ') are used, the dead-letter queue of the local queue manager is explicitly assigned.

A DLQ handler can be used to select particular messages on a dead-letter queue for special processing. For example, you could redirect the messages to different dead-letter queues. Subsequent processing with

runmqdlq

another instance of the DLQ handler might then process the messages, according to a different rules table.

QMgrName

The name of the queue manager that owns the queue to be processed.

If no name is specified, the default queue manager for the installation is used. If one or more blanks (' ') are used, the default queue manager for this installation is explicitly assigned.

runmqtsr (Run listener)

Special note

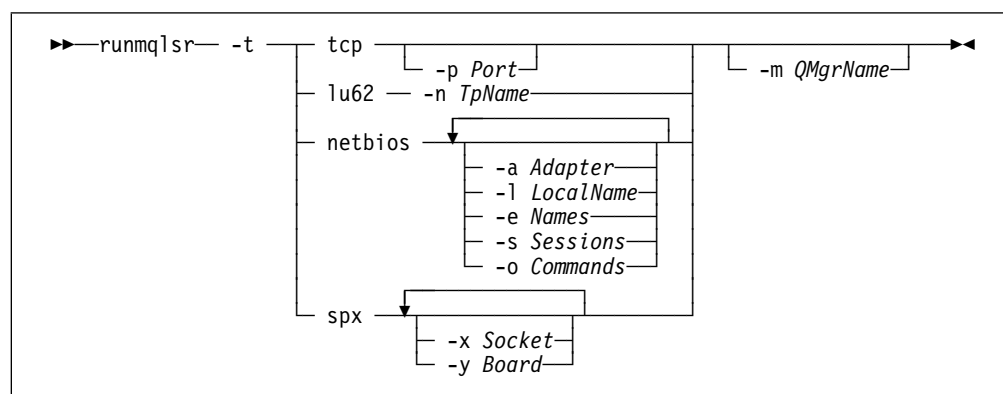
The **runmqtsr** command is not supported by these MQSeries products:

- MQSeries for AIX
- MQSeries for HP-UX
- MQSeries for Sun Solaris

Purpose

The **runmqtsr** (Run listener) command runs a listener process.

Syntax



Required parameters

- t** Specifies the transmission protocol to be used:
- | | |
|---------|--|
| tcp | Transmission Control Protocol / Internet Protocol (TCP/IP) |
| lu62 | SNA LU 6.2 |
| netbios | NetBIOS |
| spx | SPX |

Optional parameters

- p *Port*** Port number for TCP/IP. This flag is valid only for TCP/IP. If a value is not specified, the value is taken from the queue manager configuration file, or from defaults in the program. The default value is 1414.
- n *TpName*** LU 6.2 transaction program name. This flag is valid only for the LU 6.2 transmission protocol. If a value is not specified, the value is taken from the queue manager configuration file. If a value is not given, the command fails.
- a *Adapter*** Specifies the adapter number on which NetBIOS listens. The default value is 0, that is, the listener uses adapter 0.

runmq1sr

- l *LocalName*
Specifies the NetBIOS local name that the listener uses. The default is specified in the queue manager configuration file.
- e *Names*
Specifies the number of names that the listener can use. The default value is specified in the queue manager configuration file.
- s *Sessions*
Specifies the number of sessions that the listener can use. The default value is specified in the queue manager configuration file.
- o *Commands*
Specifies the number of commands that the listener can use. The default value is specified in the queue manager configuration file.
- x *Socket*
Specifies the SPX socket on which SPX listens. The default value is E135.
- y *Board*
Specifies the adapter number for SPX.

This parameter applies to MQSeries for OS/2 Warp only.
- m *QMgrName*
Specifies the name of the queue manager. If no name is specified, the command operates on the default queue manager.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

Examples

The following command runs a listener on the default queue manager using the NetBIOS protocol. Five names, five commands, and five sessions are specified for this listener, indicating the maximum number of each that this listener can use. These resources must be within the limits set in the queue manager configuration file.


```
runmq1sr -t netbios -e 5 -s 5 -o 5
```

runmqsc (Run MQSeries commands)

Purpose

Use the **runmqsc** command to issue MQSC commands to a queue manager. MQSC commands enable you to perform administration tasks, for example defining, altering, or deleting a local queue object. MQSC commands and their syntax are described in the *MQSeries Command Reference*.

Syntax



Description

You can invoke the **runmqsc** command in three modes:

- Verify mode** MQSC commands are verified but not actually run. An output report is generated indicating the success or failure of each command. This mode is available only on a local queue manager.
- Direct mode** MQSC commands are sent directly to a local queue manager.
- Indirect mode** MQSC commands are run on a remote queue manager. These commands are put on the command queue on a remote queue manager and are run in the order in which they were queued. Reports from the commands are returned to the local queue manager.

Indirect mode operation is performed through the default queue manager.

The **runmqsc** command takes its input from `stdin`. When the commands are processed, the results and a summary are put into a report that is sent to `stdout`.

By taking `stdin` from the keyboard, you can enter MQSC commands interactively.

By redirecting the input from a file you can run a sequence of frequently-used commands contained in the file. You can also redirect the output report to a file.

Note: To run this command in MQSeries for UNIX systems, your user ID must belong to user group `mqm`.

Optional parameters

- e** Prevents source text for the MQSC commands from being copied into a report. This is useful when you enter commands interactively.
- v** Specifies verification mode; this verifies the specified commands without performing the actions. This mode is only available locally. The **-w** and **-x** flags are ignored if they are specified at the same time.

runmqsc

-w *WaitTime*

Specifies indirect mode, that is, the MQSC commands are to be run on another queue manager. You must have the required channel and transmission queues set up for this. See “Preparing channels and transmission queues for remote administration” on page 65 for more information.

WaitTime Specifies the time, in seconds, that **runmqsc** waits for replies. Any replies received after this are discarded, however, the MQSC commands are still run. Specify a time between 1 and 999 999 seconds.

Each command is sent as an Escape PCF to the command queue (SYSTEM.ADMIN.COMMAND.QUEUE) of the target queue manager.

The replies are received on queue SYSTEM.MQSC.REPLY.QUEUE and the outcome is added to the report. This can be defined as either a local queue or a model queue.

Indirect mode operation is performed through the default queue manager.

This flag is ignored if the -v flag is specified.

-x Specifies that the target queue manager is running under MVS/ESA. This flag applies only in indirect mode. The -w flag must also be specified. In indirect mode, the MQSC commands are written in a form suitable for the MQSeries for MVS/ESA command queue.

QMgrName

Specifies the name of the target queue manager on which the MQSC commands are to be run. If omitted, the MQSC commands run on the default queue manager.

Return codes

00	MQSC command file processed successfully
10	MQSC command file processed with errors—report contains reasons for failing commands
20	Error—MQSC command file not run

Examples

1. Enter this command at the command prompt:

```
runmqsc
```

Now you can enter MQSC commands directly at the command prompt. No queue manager name is specified, therefore the MQSC commands are processed on the default queue manager.

2. Use one of these commands, as appropriate in your environment, to specify that MQSC commands are to be verified only:


```
runmqsc -v BANK < /u/users/commfile.in  
runmqsc -v BANK < c:\users\commfile.in
```

This command verifies the MQSC commands in file `commfile.in`. The queue manager name is `BANK`. The output is displayed in the current window.

3. These commands run the MQSC command file `mqscfile.in` against the default queue manager.

```
runmqsc < /var/mqm/mqsc/mqscfile.in > /var/mqm/mqsc/mqscfile.out  
runmqsc < c:\mqm\mqsc\mqscfile.in > c:\mqm\mqsc\mqscfile.out
```

In this example, the output is directed to file `mqscfile.out`.

runmqtmc (Start client trigger monitor)


Special note

The **runmqtmc** command is available on OS/2 and AIX clients only.

Purpose

Use the **runmqtmc** command to invoke a trigger monitor for a client. For further information about using trigger monitors, refer to the *MQSeries Application Programming Guide*.

Syntax



Optional parameters

-m *QMgrName*

Specifies the name of the queue manager on which the client trigger monitor operates. If the name is omitted, the client trigger monitor operates on the default queue manager.

-q *InitiationQName*

Specifies the name of the initiation queue to be processed. If the name is omitted, SYSTEM.DEFAULT.INITIATION.QUEUE is used.

Return codes

0	Not used. The client trigger monitor is designed to run continuously and therefore not to end. The value is reserved.
10	Client trigger monitor interrupted by an error.
20	Error—client trigger monitor not run.

runmqtrm (Start trigger monitor)

Purpose

Use the **runmqtrm** command to invoke a trigger monitor. For further information about using trigger monitors, refer to the *MQSeries Application Programming Guide*.

Syntax

```

▶▶—runmqtrm [ -m QMgrName ] [ -q InitiationQName ]

```

Optional parameters

-m *QMgrName*

Specifies the name of the queue manager on which the trigger monitor operates. If the name is omitted, the trigger monitor operates on the default queue manager.

-q *InitiationQName*

Specifies the name of the initiation queue to be processed. If the name is omitted, SYSTEM.DEFAULT.INITIATION.QUEUE is used.

Return codes

- | | |
|-----------|--|
| 0 | Not used. The trigger monitor is designed to run continuously and therefore not to end. Hence a value of 0 would not be seen. The value is reserved. |
| 10 | Trigger monitor interrupted by an error. |
| 20 | Error—trigger monitor not run. |

scmmqm (Add the queue manager to the Windows NT Service Control Manager)

Special note

The **scmmqm** command is supported by MQSeries for Windows NT only.

Purpose

Use the **scmmqm** command to add a queue manager to, or delete a queue manager from, the list of those that will be started automatically, by the IBM MQSeries service, when the system starts.

Syntax

```

  ▶▶ scmmqm [-z] [-a | -d] [-s Commandfile] QMgrName ▶▶

```

Required parameters

- a Adds a queue manager to the list of those that start automatically.
- d Removes a queue manager from the list of those that start automatically.
- s *Commandfile*
Locates the command file that is executed when a queue manager is automatically started.
- QMgrName* Name of the queue manager concerned.

Optional parameters

- z Suppresses error messages.

Return codes

- 0 Windows NT successful operation (Service Control Manager updated)
 - 36 Invalid arguments supplied
 - 71 Unexpected error
- Note:** If this happens when you are attempting to remove a queue manager from the automatic start-up list, follow the steps in Appendix E, "Stopping and removing queue managers manually" on page 315 to remove it manually.
- 72 Queue manager name error

Examples

The following examples show commands that update the Windows NT Service Control Manager:

1. Add a queue manager to the automatic start-up list.

```
scmmqm -a -s c:\mqm\saturnstartup.cmd saturn.queue.manager
```

2. Modify the command file, newstartup.cmd that is executed when a queue manager is automatically started.

```
scmmqm -a -s c:\mqm\newstartup.cmd saturn.queue.manager
```

3. Remove a queue manager from the automatic start-up list.

```
scmmqm -d saturn.queue.manager
```

Note: The command files can be stored in any location under any name. Each queue manager needs its own command file.

For information about starting a queue manager automatically, see “Starting a queue manager automatically” on page 33.

setmqaut (Set/reset authority)

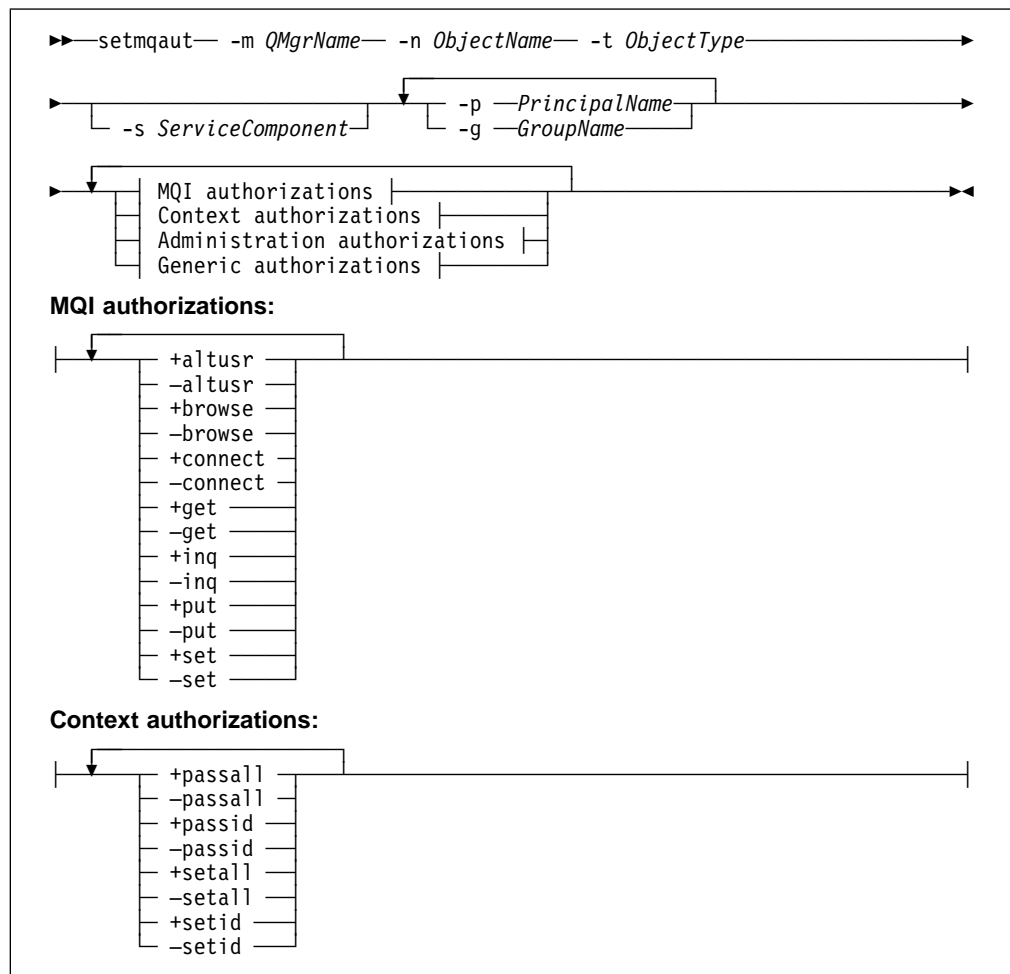
Purpose

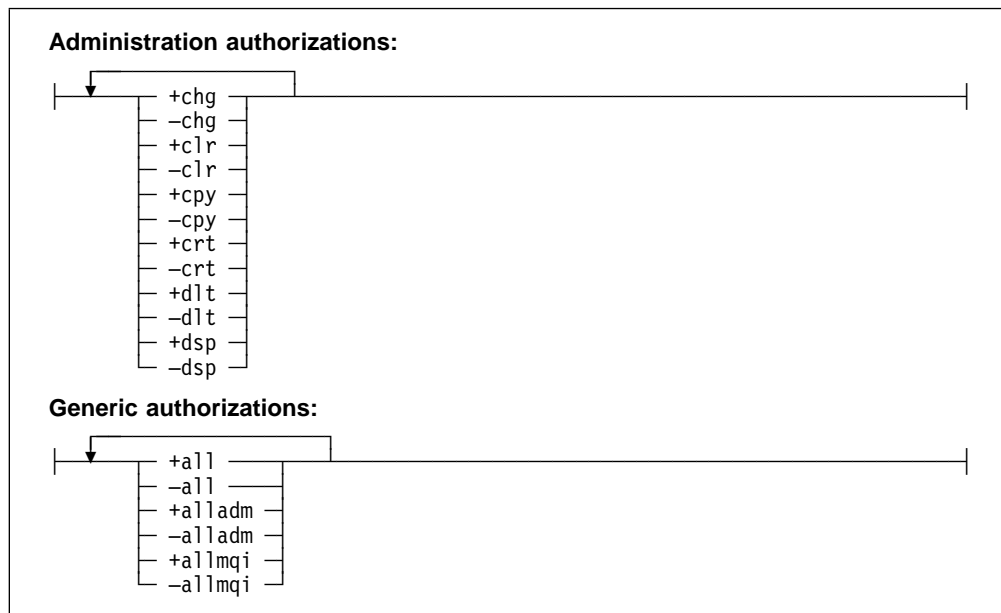
Use the **setmqaut** command to change the authorizations to an object or to a class of objects. Authorizations can be granted to, or revoked from, any number of principals or groups.

In MQSeries for OS/2 Warp only, you can use this command to specify authorizations only if an authorization service component has been installed for the current queue manager. By default, no such component is supplied with MQSeries for OS/2 Warp.

For more information about authorization service components, see the *MQSeries Programmable System Management* book.

Syntax





Description

You can use this command both to *set* an authorization, that is, give a user group or principal permission to perform an operation, and to *reset* an authorization, that is, remove the permission to perform an operation. You must specify the user groups and principals to which the authorizations apply and also the queue manager, object type, and object name of the object. You can specify any number of groups and principals in a single command.

Note: In MQSeries for UNIX systems, if you specify a set of authorizations for a principal, the same authorizations are given to all principals in the same primary group.

The authorizations that can be given are categorized as follows:

- Authorizations for issuing MQI calls
- Authorizations for MQI context
- Authorizations for issuing commands for administration tasks
- Generic authorizations

Each authorization to be changed is specified in an authorization list as part of the command. Each item in the list is a string prefixed by '+' or '-'. For example, if you include +put in the authorization list, you are giving authority to issue MQPUT calls against a queue. Alternatively, if you include -put in the authorization list, you are removing the authorization to issue MQPUT calls.

Authorizations can be specified in any order provided that they do not clash. For example, specifying allmqi with set causes a clash.

You can specify as many groups or authorizations as you require in a single command.

If a user ID is a member of more than one group, the authorizations that apply are the union of the authorizations of each group to which that user ID belongs.

Required parameters

-n *ObjectName*

Specifies the name of the object for which the authorizations are to be changed. You must not use a generic name.

This parameter is optional if you are changing the authorizations of your default queue manager.

-t *ObjectType*

Specifies the type of object for which the authorizations are to be changed.

Possible values are:

- **q** or **queue**
- **prcs** or **process**
- **qmgr**

-m *QMgrName*

Specifies the name of the queue manager of the object for which the authorizations are to be changed. The name can contain up to 48 characters.

This parameter is optional if you are changing the authorizations of your default queue manager.

Optional parameters

-p *PrincipalName*

Specifies the name of the principal for which the authorizations are to be changed.

You must have at least one principal or one group.

-g *GroupName*

Specifies the name of the user group whose authorizations are to be changed. You can specify more than one group name, but each name must be prefixed by the -g flag.

-s *ServiceComponent*

This parameter applies only if you are using installable authorization services, otherwise it is ignored.

If installable authorization services are supported, this parameter specifies the name of the authorization service to which the authorizations apply.

This parameter is optional; if it is not specified, the authorization update is made to the first installable component for the service.

Authorizations

Specifies the authorizations to be given or removed. Each item in the list is prefixed by a '+' indicating that authority is to be given, or a '-', indicating that authorization is to be removed. For example, to give authority to issue an MQPUT call from the MQI, specify +put in the list. To remove authority to issue an MQPUT call, specify -put.

Table 18 on page 287 shows the authorities that can be given to the different object types.

Authority	Queue	Process	Qmgr
all	√	√	√
alladm	√	√	√
allmqi	√	√	√
altusr			√
browse	√		
chg	√	√	√
clr	√		
connect			√
crt	√	√	√
cpy	√	√	√
dlt	√	√	√
dsp	√	√	√
put	√		
inq	√	√	√
get	√		
passall	√		
passid	√		
set	√	√	√
setall	√		√
setid	√		√

Authorizations for MQI calls:

altusr	Allows another user's authority to be used for MQOPEN and MQPUT1 calls. See the <i>MQSeries Application Programming Guide</i> for more information about alternate user IDs.
browse	Retrieve a message from a queue by issuing an MQGET call with the BROWSE option.
connect	Connect the application to the specified queue manager by issuing an MQCONN call.
get	Retrieve a message from a queue by issuing an MQGET call.
inq	Make an inquiry on a specific queue by issuing an MQINQ call.
put	Put a message on a specific queue by issuing an MQPUT call.
set	Set attributes on a queue from the MQI by issuing an MQSET call.

Note: If you open a queue for multiple options, you have to be authorized for each of them.

Authorizations for context:

passall	Pass all context on the specified queue. All the context fields are copied from the original request.
passid	Pass identity context on the specified queue. The identity context is the same as that of the request.
setall	Set all context on the specified queue. This is used by special system utilities.
setid	Set identity context on the specified queue. This is used by special system utilities.

Authorizations for commands:

chg	Change the attributes of the specified object.
clr	Clear the specified queue (PCF Clear queue command only).
cpy	Copy the attributes of the specified object (PCF Copy commands only).
cr	Create objects of the specified type.
dlt	Delete the specified object.
dsp	Display the attributes of the specified object.

Authorizations for generic operations:

all	Use all operations applicable to the object.
alladm	Perform all administration operations applicable to the object.
allmqi	Use all MQI calls applicable to the object.

Return codes

0	Successful operation
36	Invalid arguments supplied
40	Queue manager not available
49	Queue manager stopping
69	Storage not available
71	Unexpected error
72	Queue manager name error
133	Unknown object name
145	Unexpected object name
146	Object name missing
147	Object type missing
148	Invalid object type
149	Entity name missing
150	Authorization specification missing
151	Invalid authorization specification

Examples

1. This example shows a command that specifies that the object on which authorizations are being given is the queue orange.queue on queue manager saturn.queue.manager.

```
setmqaut -m saturn.queue.manager -n orange.queue -t queue -g tango +inq +alladm
```

The authorizations are being given to user group tango and the associated authorization list specifies that user group tango:

- Can issue MQINQ calls
 - Has authority to perform all administration operations on that object
2. In this example, the authorization list specifies that user group foxy:
 - Cannot issue any calls from the MQI to the specified queue
 - Has authority to perform all administration operations on the specified queue

```
setmqaut -m saturn.queue.manager -n orange.queue -t queue -g foxy -allmqi +alladm
```

3. In this example, the authorization list specifies that user group waltz has authority to create and delete queue manager saturn.queue.manager.

```
setmqaut -m saturn.queue.manager -t qmgr -g waltz +crt +dlt
```

Related commands

dspmqaut

Display authority

setmqprd (Enroll production license)

Purpose

Use the **setmqprd** command to enroll a production license. For further information about enrolling production licenses, refer to the *MQSeries Quick Beginnings* book for your operating system.

Syntax

```
▶▶—setmqprd—┌LicenseFilename┐▶▶
```

Optional parameters

LicenseFilename

Specifies the fully-qualified name of the production license certificate file (usually amqpcert.lic in MQSeries for OS/2 Warp, Windows NT, and AIX, and amqppswd.lic in MQSeries for HP-UX and Sun Solaris). If the name is omitted, the file name defaults to the name and path of the certificate file as installed from the purchased CD.

Return codes

0	Production license installed <i>or</i> Production license previously installed
36	Invalid arguments
71	License expired <i>or</i> License not found or invalid <i>or</i> Unexpected error

setmqtry (Start trial period)

Purpose

Use the **setmqtry** command to start an MQSeries trial period. For further information about trial periods, refer to the *MQSeries Quick Beginnings* book for your operating system.

Syntax

```
▶▶—setmqtry—————▶▶
```

Return codes

0	Trial period started
36	Invalid arguments
71	License expired <i>or</i> Agreement rejected <i>or</i> In production mode <i>or</i> Unexpected error

strmqcsv (Start command server)

Purpose

Use the **strmqcsv** command to start the command server for the specified queue manager. This enables MQSeries to process commands sent to the command queue.

Syntax

```
▶▶—strmqcsv—QMgrName————▶▶
```

Required parameters

QMgrName

Specifies the name of the queue manager for which the command server is to be started.

Return codes

0	Command completed normally
10	Command completed with unexpected results
20	An error occurred during processing

Examples

The following command starts a command server for queue manager earth:

```
strmqcsv earth
```

Related commands

endmqcsv	End a command server
dsprmqcsv	Display the status of a command server

strmqm

Related commands

crtmqm
dlmqm
endmqm

Create a queue manager
Delete a queue manager
End a queue manager

strmqtrc (Start MQSeries trace)

Special note


The **strmqtrc** command is not supported by MQSeries for AIX.

Purpose


Use the **strmqtrc** command to enable tracing. This command can be run regardless of whether tracing is enabled. If tracing is already enabled, the trace options in effect are modified to those specified on the latest invocation of the command.

Syntax

The syntax of this command in MQSeries for HP-UX and Sun Solaris is as follows:



The syntax of this command in MQSeries for OS/2 Warp and Windows NT is as follows:



Description

Different levels of trace detail can be requested. For each flow tracetype value you specify, including -t all, specify either -t parms or -t detail to obtain the appropriate level of trace detail. If you do not specify either -t parms or -t detail for any particular trace type, only a default-detail trace is generated for that trace type.

For examples of trace data generated by this command see “Tracing” on page 219.

Optional parameters

-m *QMgrName*

Is the name of the queue manager to be traced.

A queue manager name and the -m flag can be specified on the same command as the -e flag. If more than one trace specification applies to a given entity being traced, the actual trace includes all of the specified options.

It is an error to omit the -m flag and queue manager name, unless the -e flag is specified.

This parameter is not valid in MQSeries for OS/2 Warp and Windows NT.

- e If this flag is specified, early tracing is requested. Consequently, it is possible to trace the creation or startup of a queue manager. Any process, belonging to any component of any queue manager, traces its early processing if this flag is specified. The default, if this flag is not specified, is not to perform early tracing.

This parameter is not valid in MQSeries for OS/2 Warp and Windows NT.

-t *TraceType*

Identifies the points to be traced, and specifies the amount of trace detail to be recorded. If this flag is omitted, all trace points are enabled, and a default-detail trace is generated.

Alternatively, one or more of the options in the following list can be supplied.

If multiple trace types are supplied, each **must** have its own -t flag. Any number of -t flags can be specified, provided that each has a valid trace type associated with it.

It is not an error to specify the same trace type on multiple -t flags.

all Output data for every trace point in the system. This is also the default if the -t flag is not specified. The all parameter activates tracing at default detail level.

api Output data for trace points associated with the MQI and major queue manager components.

commentary

Output data for trace points associated with comments in the MQSeries components.

comms

Output data for trace points associated with data flowing over communications networks.

csdata

Output data for trace points associated with internal data buffers in common services.

csflows

Output data for trace points associated with processing flow in common services.

detail

Activates tracing at high-detail level for flow processing trace points.

lqmdata

Output data for trace points associated with internal data buffers in the local queue manager.

lqmflows

Output data for trace points associated with processing flow in the local queue manager.

otherdata

Output data for trace points associated with internal data buffers in other components.

otherflows

Output data for trace points associated with processing flow in other components.

parms

Activates tracing at default-detail level for flow processing trace points.

remotedata

Output data for trace points associated with internal data buffers in the communications component.

remoteflows

Output data for trace points associated with processing flow in the communications component.

servicedata

Output data for trace points associated with internal data buffers in the service component.

serviceflows

Output data for trace points associated with processing flow in the service component.

versiondata

Output data for trace points associated with the version of MQSeries running.

Return codes

- AMQ7024** This message is issued if arguments that are not valid are supplied to the command.
- AMQ8304** The maximum number of nine concurrent traces is already running.

Examples

This command enables tracing of processing flow from common services and the local queue manager for a queue manager called QM1 in MQSeries for UNIX systems. Trace data is generated at the default level of detail.

```
strmqtrc -m QM1 -t csflows -t lqmfows -t parms
```

This command enables high-detail tracing of the processing flow for all components in MQSeries for OS/2 Warp or Windows NT:

```
strmqtrc -t all -t detail
```

Related commands

dspmqtrc	Display formatted trace output
endmqtrc	End MQSeries trace

strmqtrc

Part 3. Appendixes

Appendix A. System and default objects	301
Appendix B. Directory structure (UNIX systems)	303
Queue manager log directory structure	305
Appendix C. Directory structure (OS/2)	307
Queue manager log directory structure	309
Appendix D. Directory structure (Windows NT)	311
Queue manager log directory structure	313
Appendix E. Stopping and removing queue managers manually	315
Stopping a queue manager manually	315
Stopping queue managers in MQSeries for UNIX systems	315
Stopping queue managers in MQSeries for Windows NT	316
Stopping queue managers in MQSeries for OS/2 Warp	316
Removing queue managers manually	316
Removing queue managers in MQSeries for UNIX systems	316
Removing queue managers in MQSeries for Windows NT	317
Removing queue managers in MQSeries for OS/2 Warp	318
Appendix F. User identifier service	319
Appendix G. Notices	321
Trademarks	322

Appendix A. System and default objects

When you create a queue manager using the **crtmqm** control command, the system objects and the default objects are created automatically.

- The system objects are those MQSeries objects required for the operation of a queue manager or channel.
- The default objects define all of the attributes of an object. When you create an object, such as a local queue, any attributes that you do not specify explicitly are inherited from the default object.

Table 19 lists the system and default objects created by **crtmqm**.

<i>Table 19. The system and default objects</i>	
Object name	Description
SYSTEM.DEFAULT.ALIAS.QUEUE	Default alias queue.
SYSTEM.DEFAULT.LOCAL.QUEUE	Default local queue.
SYSTEM.DEFAULT.MODEL.QUEUE	Default model queue.
SYSTEM.DEFAULT.REMOTE.QUEUE	Default remote queue.
SYSTEM.DEAD.LETTER.QUEUE	Dead-letter (undelivered-message) queue.
SYSTEM.DEFAULT.PROCESS	Default process definition.
SYSTEM.DEF.SENDER	Default sender channel.
SYSTEM.DEF.SERVER	Default server channel.
SYSTEM.DEF.RECEIVER	Default receiver channel.
SYSTEM.DEF.REQUESTER	Default requester channel.
SYSTEM.DEF.SVRCONN	Default server-connection channel.
SYSTEM.DEF.CLNTCONN	Default client-connection channel.
SYSTEM.AUTO.RECEIVER	Dynamic receiver channel.
SYSTEM.AUTO.SVRCONN	Dynamic server-connection channel.
SYSTEM.CHANNEL.INITQ	Channel initiation queue.
SYSTEM.DEFAULT.INITIATION.QUEUE	Default initiation queue.
SYSTEM.CICS.INITIATION.QUEUE	Default CICS initiation queue.
SYSTEM.ADMIN.COMMAND.QUEUE	Administration command queue. Used for remote MQSC commands and PCF commands.
SYSTEM.MQSC.REPLY.QUEUE	MQSC reply-to queue. This is a model queue that creates a temporary dynamic queue for replies to remote MQSC commands.
SYSTEM.ADMIN.QMGR.EVENT	Event queue for queue manager events.
SYSTEM.ADMIN.PERFM.EVENT	Event queue for performance events.
SYSTEM.ADMIN.CHANNEL.EVENT	Event queue for channel events.

Default objects

Appendix B. Directory structure (UNIX systems)

Figure 48 shows the general layout of the data and log directories associated with a specific queue manager. The directories shown apply to the default installation. If you change this, the locations of the files and directories will be modified accordingly. For information about the location of the product files, see the *MQSeries Installation* booklet for your MQSeries product.

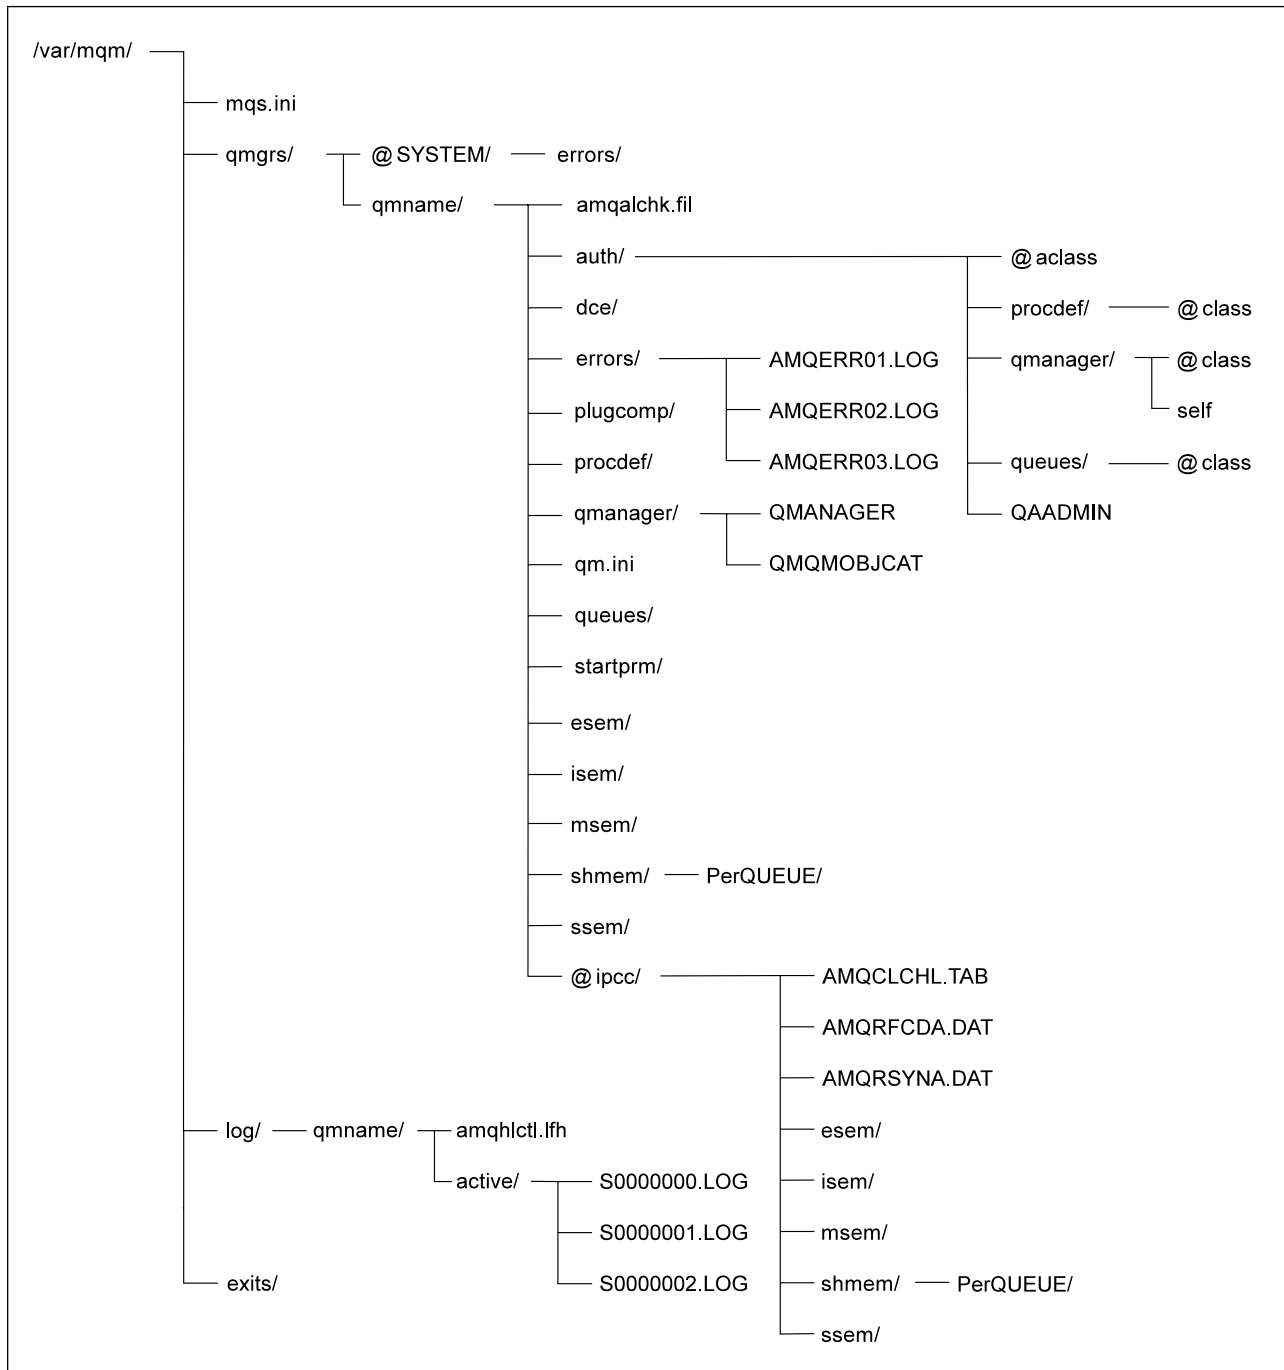


Figure 48. Default directory structure (UNIX systems) after a queue manager has been started

Directory structure (UNIX)

In Figure 48, the layout is representative of MQSeries after a queue manager has been in use for some time. The actual structure that you have depends on which operations have occurred on the queue manager.

By default, the following directories and files located in the directory `/var/mqm/qmgrs/qmname/`.

amqalchk.fil	Checkpoint file containing information about last checkpoint.
auth/	This directory contains subdirectories and files associated with authority. <ul style="list-style-type: none">@aclass This file contains the authority stanzas for all classes.procdef/ This directory contains a file for each process definition. Each file contains the authority stanzas for the associated process definition.<ul style="list-style-type: none">@class This file contains the authority stanzas for the process definition class.qmanager/<ul style="list-style-type: none">@class This file contains the authority stanzas for the queue manager class.self This file contains the authority stanzas for the queue manager object.queues/ This directory contains a file for each queue. Each file contains the authority stanzas for the associated queue.<ul style="list-style-type: none">@class This file contains the authority stanzas for the queue class.
	QAADMIN File used internally for controlling authorizations.
dce/	Empty directory reserved for use by DCE support.
errors/	The operator message files, from newest to oldest: <ul style="list-style-type: none">AMQERR01.LOGAMQERR02.LOGAMQERR03.LOG
plugcomp/	Empty directory reserved for use by installable services.
procdef/	Each MQSeries process definition is associated with a file in this directory. The file name matches the process definition name—subject to certain restrictions; see “Understanding MQSeries file names” on page 30.
qmanager/	<ul style="list-style-type: none">QMANAGER The queue manager object.QMQMBOBJCAT The object catalogue containing the list of all MQSeries objects—used internally.
qm.ini	Queue manager configuration file.

queues/	Each queue has a directory in here containing a single file called 'q'. The file name matches the queue name—subject to certain restrictions; see “Understanding MQSeries file names” on page 30.	
startprm/	Directory containing temporary files used internally.	
esem/	Directories containing files used internally.	
isem/		
msem/		
shmem/		
PerQUEUE/		Directory containing files used internally.
ssem/	Directory containing files used internally.	
@ipcc/		
	AMQCLCHL.TAB	Client channel table file.
	AMQRFDA.DAT	Channel table file.
	AMQRSYNA.DAT	Channel synchronization file.
	esem/	Directories containing files used internally.
	msem/	
	ssem/	
	isem/	
	shmem/	
		PerQUEUE/ Directory containing files used internally.

Queue manager log directory structure

By default, the following directories and files are found in `/var/mqm/log/qmname/`.

The following subdirectories and files exist after you have installed MQSeries, created and started a queue manager, and have been using that queue manager for some time.

amqhctl.lfh	Log control file.
active/	This directory contains the log files numbered S0000000.LOG, S0000001.LOG, S0000002.LOG, and so on.

Directory structure (UNIX)

Appendix C. Directory structure (OS/2)

The following directories and files are found under the root C:\MQM\QMGRS\QMNAME\.
If you have installed MQSeries for OS/2 Warp under different directories, the root is modified appropriately.




Figure 49. Default file tree (OS/2) after a queue manager has been started. If you are using a FAT system, the name QMQMOBJCAT will be transformed.

Directory structure (OS/2)

Figure 49 shows the general layout of the data and log directories. The layout is representative of MQSeries after a queue manager has been in use for some time. However, the actual structure that you have depends on the operations that have occurred on the queue manager. A brief description of the files follows.

Notes:

1. The directory and file names are all shown in uppercase. The case depends on the file system you are using (FAT or HPFS).
2. The queue manager names may have been transformed. See "Understanding MQSeries file names" on page 30 for more information about name transformation.

AMQALCHK.FIL

Checkpoint file containing information about last checkpoint.

AUTH\

PROCDEF\ Empty directory reserved for authority parameters.

QMANAGER\

Empty directory reserved for authority parameters.

QUEUES\ Empty directory reserved for authority parameters.

DCE\

Empty directory reserved for use by DCE support.

ERRORS\

The operator message files from newest to oldest.

AMQERR01.LOG

AMQERR02.LOG

AMQERR03.LOG

PLUGCOMP\

Empty directory reserved for use by installable services.

PROCDEF\

Each MQSeries process definition has a file in here.

Where possible, the file name matches the associated process definition name but some characters have to be altered.

There may be a directory called @MANGLED here containing process definitions with transformed or mangled names.

QMANAGER\

QMANAGER The queue manager object.

QMOMOBJCAT

The object catalogue containing the list of all MQSeries objects, used internally.

Note: If you are using a FAT system, this name will be transformed and a subdirectory created containing the file with its name transformed.

QAADMIN File used internally for controlling authorizations.

QM.INI

Queue manager configuration file

QUEUES Each queue has a directory here containing a single file called Q.
Where possible, the directory name matches the associated queue name but some characters have to be altered.

There may be a directory called @MANGLED here containing queues with transformed or mangled names.

STARTPRM Directory containing temporary files used internally.

@IPCC

AMQCLCHL.TAB

File containing the client channel table

AMQRFEDA.DAT

File containing the channel table

AMQRSYNA.DAT

Channel synchronization file

Queue manager log directory structure

The following directories and files are found under C:\MQM\LOG\QMNAME\. If you have installed the product under different directories or specified different log paths in the configuration file, the root will be modified appropriately.

The following subdirectories and files will exist after you have installed MQSeries, created and started a queue manager, and have been using that queue manager for some time.

AMQHLCTL.LFH

Log control file.

ACTIVE This directory contains the log files numbered S0000000.LOG, S0000001.LOG, S0000002.LOG, and so on.

Directory structure (OS/2)

Appendix D. Directory structure (Windows NT)

Figure 50 shows some of the directories and files found under the root `c:\mqm\`. If you have installed MQSeries for Windows NT under different directories, the root is modified appropriately.

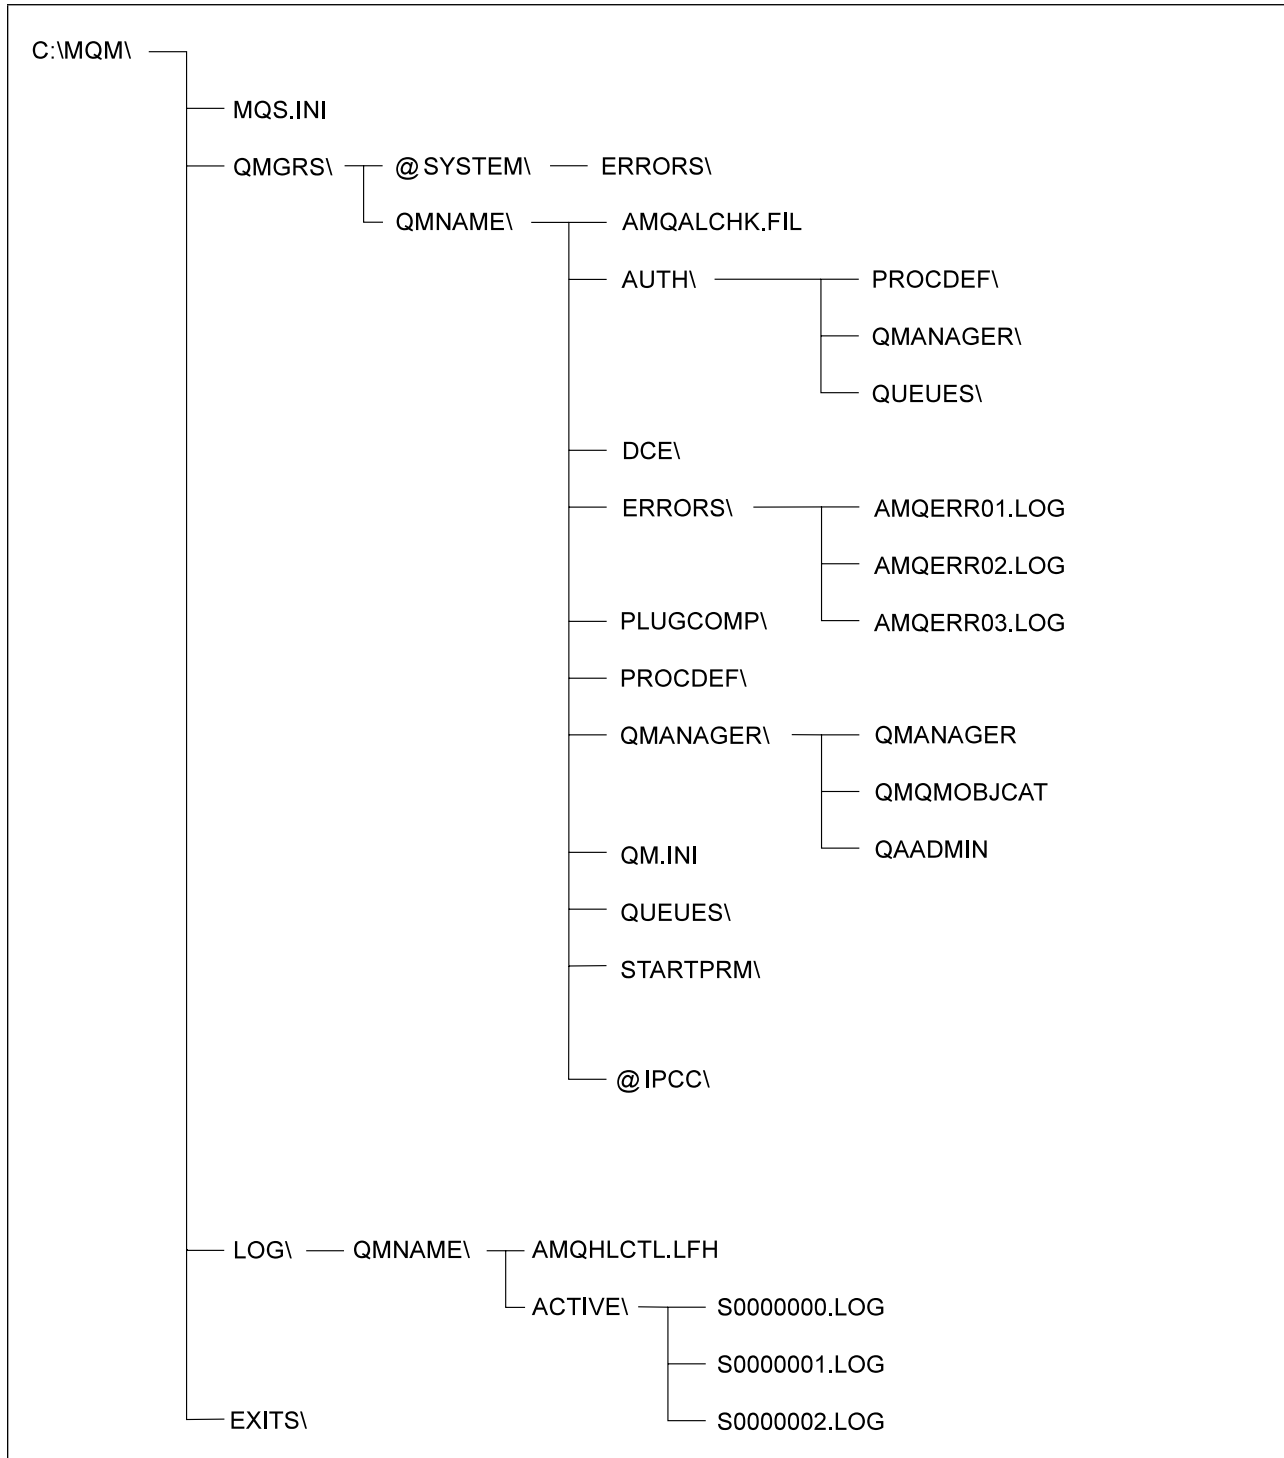


Figure 50. Default file tree (Windows NT) after a queue manager has been started. If you are using a FAT system, the name `QMCMOBYCAT` is transformed.

Directory structure (Windows NT)

Figure 50 shows the general layout of the data and log directories. The layout is representative of MQSeries after a queue manager has been in use for some time. However, the actual structure that you have depends on the operations that have occurred on the queue manager. A brief description of the files follows.

Notes:

1. The directory and file names are all shown in upper case. The case depends on the file system you are using (NTFS, HPFS or FAT).
2. The queue manager names may have been transformed. See "Understanding MQSeries file names" on page 30 for more information about name transformation.

AMQALCHK.FIL

Checkpoint file containing information about last checkpoint.

AUTH\

PROCDEF\ Empty directory reserved for authority parameters.

QMANAGER\

Empty directory reserved for authority parameters.

QUEUES\ Empty directory reserved for authority parameters.

DCE\

Empty directory reserved for use by DCE support.

ERRORS\

The operator message files, from newest to oldest:

AMQERR01.LOG

AMQERR02.LOG

AMQERR03.LOG

PLUGCOMP\

Empty directory reserved for use by installable services.

PROCDEF\

Each MQSeries process definition has a file in here.

Where possible, the file name matches the associated process definition name but some characters have to be altered.

There may be a directory called @MANGLED here containing process definitions with transformed or mangled names.

QMANAGER\

QMANAGER

The queue manager object.

QMQM OBJCAT

The object catalogue containing the list of all MQSeries objects, used internally.

Note: If you are using a FAT system, this name will be transformed and a subdirectory created containing the file with its name transformed.

QAADMIN File used internally for controlling authorizations.

QM.INI

Queue manager configuration file.

QUEUES Each queue has a directory here containing a single file called Q.
Where possible, the directory name matches the associated queue name but some characters have to be altered.
There may be a directory called @MANGLED here containing queues with transformed or mangled names.

STARTPRM Directory containing temporary files used internally.

@IPCC

AMQCLCHL.TAB
File containing the client channel table.

AMQRFEDA.DAT
File containing the channel table.

AMQRSYNA.DAT
Channel synchronization file.

Queue manager log directory structure

The following directories and files are found under C:\MQM\LOG\QMNAME\. If you have installed the product under different directories or specified different log paths in the configuration file, the root will be modified appropriately.

The following subdirectories and files will exist after you have installed MQSeries, created and started a queue manager, and have been using that queue manager for some time.

AMQHLCTL.LFH
Log control file.

ACTIVE This directory contains the log files numbered S0000000.LOG, S0000001.LOG, S0000002.LOG, and so on.

Directory structure (Windows NT)

Appendix E. Stopping and removing queue managers manually

If the standard methods for stopping and removing queue managers fail, you can resort to the more drastic methods described here.

Stopping a queue manager manually

The standard way of stopping queue managers, using the **endmqm** command, should work even in the event of failures within the queue manager. In exceptional circumstances, if this method of stopping a queue manager fails, use one of the procedures described here to stop it manually.

Stopping queue managers in MQSeries for UNIX systems

To stop a queue manager running under MQSeries for UNIX systems:

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

```
ps -ef | grep QMNAME
```

2. End any queue manager processes that are still running. Use the **kill** command, specifying the process IDs discovered using the **ps** command.

Note: Processes that fail to stop can be ended using **kill -9**.

End the processes in the following order:

amqhasmx	logger
amqharmx	log formatter (LINEAR logs only)
amqzllp0	checkpoint processor
amqzlaa0	queue manager agents
amqzma0	processing controller

Note: Manual stopping of the queue manager may result in FFSTs being taken, and the production of FDC files in `/var/mqm/errors`. This should not be regarded as a defect in the queue manager.

The queue manager should restart normally, even after having been stopped using this method.

Stopping queue managers in MQSeries for Windows NT

To stop a queue manager running under MQSeries for Windows NT:

1. List the names (IDs) of the processes currently running using the Windows NT Process Viewer (PView)
2. Stop the processes using PView in the following order (if they are running):

AMQHASMN.EXE	The logger
AMQHARMN.EXE	Log formatter (LINEAR logs only)
AMQZLLP0.EXE	Checkpoint process
AMQZLAA0.EXE	LQM agents
AMQZTRCN.EXE	Trace
AMQZXMA0.EXE	Execution controller
AMQXSSVN.EXE	Shared memory servers

3. Stop the queue manager service using the Windows NT Control Panel.
4. If you have tried all methods and the queue manager has not stopped, reboot your system.

Stopping queue managers in MQSeries for OS/2 Warp

To stop a queue manager running under MQSeries for OS/2 Warp:

1. If you have access to an appropriate utility (equivalent to the UNIX **ps** command), list the names (IDs) of the processes currently running.
2. If you have access to a utility that stops processes (equivalent to the UNIX **kill** command), stop them in the following order:

AMQHASM2.EXE	The logger
AMQHARM2.EXE	Log formatter (LINEAR logs only)
AMQZLLP0.EXE	Checkpoint process
AMQZLAA0.EXE	LQM agents
AMQZXMA0.EXE	Execution controller
AMQXSSV2.EXE	Shared memory servers

Note: If you do not have access to a suitable utility, and you have tried all other methods, you must reboot your system.

Removing queue managers manually

If you want to delete the queue manager after stopping it manually, use the **dltmqm** command as normal. If, for some reason, this command fails to delete the queue manager, the manual processes described here can be used.

Removing queue managers in MQSeries for UNIX systems

You should note that manual removal of a queue manager is potentially very disruptive, particularly if multiple queue managers are being used on a single system. This is because complete removal of a queue manager requires deletion of files, shared memory and semaphores. As it is impossible to identify which shared memory and semaphores belong to a particular queue manager, it is necessary to stop all running queue managers.

If you need to delete a queue manager manually, use the following procedure:

1. Stop all queue managers running on the machine from which you need to remove the queue manager.
2. Locate the queue manager directory from the configuration file `/var/mqm/mqs.ini` and look for the `QueueManager` stanza naming the queue manager to be deleted.

Its `Prefix` and `Directory` attributes identify the queue manager directory. For a `Prefix` attribute of `<Prefix>` and a `Directory` attribute of `<Directory>`, the full path to the queue manager directory is

`<Prefix>/qmgrs/<Directory>`

3. Locate the queue manager log directory from the `qm.ini` configuration file in the queue manager directory. The `LogPath` attribute of the `Log` stanza identifies this directory.
4. Delete the queue manager directory, all subdirectories and files.
5. Delete the queue manager log directory, all subdirectories and files.
6. Remove the queue manager's `QueueManager` stanza from the `/var/mqm/mqs.ini` configuration file.
7. If the queue manager being deleted is also the default queue manager, remove the `DefaultQueueManager` stanza from the `/var/mqm/mqs.ini` configuration file.
8. Either remove all shared memory and semaphores owned by the `mqm` user ID and `mqm` group, or restart the machine. Shared resources can be identified using the `ipcs` command, and can be removed with the `ipcrm` command.

Removing queue managers in MQSeries for Windows NT

If you encounter problems with the `dltmqm` command in MQSeries for Windows NT, use the following procedure to delete a queue manager:

1. Locate the queue manager directory from the `MQS.INI` configuration file. By default, this location is:

`C:\MQM\QMGRS\<QMgrName>`

where `<QMgrName>` (or its transformed equivalent) is the name of the queue manager to be deleted.

2. Delete this directory, all subdirectories and files.
3. Locate the associated log directory from the `mqs.ini` file.
4. Delete the directory, all subdirectories and files.
5. Remove its `QueueManager` stanza from `MQS.INI`.
6. Remove the `DefaultQueueManager` stanza, if the queue manager being deleted is the default queue manager.
7. If appropriate, remove the queue manager from the automatic start-up list by using the `scmmqm` command or by following the steps in "Removing queue managers from the automatic start-up list" on page 318.

Removing queue managers from the automatic start-up list

If the **scmmqm** command fails to remove the queue manager from the list of those that start automatically when the system starts, use the following procedure:

1. Type REGEDT32 from the command prompt.
2. Select the **HKEY_LOCAL_MACHINE** window.
3. Navigate the tree structure to find the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\MQSeries\CurrentVersion
4. Select the **autostart** value to start the editor.
5. Remove the entry referring to the queue manager that you want to remove.
6. Close the registry editor.

Removing queue managers in MQSeries for OS/2 Warp

If you encounter problems with the **dltmqm** command in MQSeries for OS/2 Warp, use the following procedure to delete a queue manager:

1. Locate the queue manager directory from the `mqs.ini` configuration file. By default, this location is:
C:\MQM\QMGRS\`<QMgrName>`
where `<QMgrName>` (or its transformed equivalent) is the name of the queue manager to be deleted.
2. Delete this directory, all subdirectories and files.
3. Locate the associated log directory from the `mqs.ini` file.
4. Delete the directory, all subdirectories and files.
5. Remove its QueueManager stanza from `mqs.ini`.
6. Remove the DefaultQueueManager stanza, if the queue manager being deleted is the default queue manager.

Appendix F. User identifier service

This information applies to MQSeries for OS/2 Warp only.

By default, the user ID associated with applications running under OS/2 is `os2`. The queue manager inserts this user ID into the context fields of any messages that are sent by the application.

The user identifier service enables a user-defined user ID to be substituted for the default. When the user identifier service is active, this user ID is accessed by the local queue manager when an application issues an MQCONN request. Thereafter, the queue manager inserts the new user ID in the context field of any message sent by the application.

The mechanism for defining the new user ID must be user-written. For example, in the sample `AMQSZFC0`, which is supplied with MQSeries for OS/2 Warp, the user ID is defined in an OS/2 environment variable. Once you have added the appropriate stanza to the queue manager configuration file, you simply type this in at the keyboard. To get the queue manager to recognize the new user ID, you must restart the queue manager. The user identifier service is described in detail in the *MQSeries Programmable System Management* book.

Appendix G. Notices

The following paragraph does not apply to any 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 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.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. 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 of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

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 Laboratory Counsel, MP151, IBM United Kingdom Laboratories, 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.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, New York 10594, U.S.A.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States, or other countries, or both:

400	MQ
AIX	MQSeries
AIX/6000	MQSeries Three Tier
AS/400	MVS
BookManager	MVS/ESA
CICS	NetView
CICS/6000	Operating System/2
DB2	OS/2
FFST	OS/400
FFST/2	RS/6000
First Failure Support Technology	VSE/ESA
IBM	

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

C-bus is a trademark of Corollary, Inc.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Java and HotJava are trademarks of Sun Microsystems, Inc.

Other company, product, and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

Part 4. Glossary and index

Glossary of terms and abbreviations

This glossary defines MQSeries terms and abbreviations used in this book. If you do not find the term you are looking for, see the Index or the *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.

This glossary includes terms and definitions from the *American National Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42 Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.

A

add-in task. A function provided by MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT that coordinates the passing of data between a Lotus Notes application and an MQSeries application.

administrator commands. MQSeries commands used to manage MQSeries objects, such as queues, processes, and namelists.

alert. A message sent to a management services focal point in a network to identify a problem or an impending problem.

alias queue object. An MQSeries object, the name of which is an alias for a base queue defined to the local queue manager. When an application or a queue manager uses an alias queue, the alias name is resolved and the requested operation is performed on the associated base queue.

alternate user security. A security feature in which the authority of one user ID can be used by another user ID; for example, to open an MQSeries object.

APAR. Authorized program analysis report.

application log. In Windows NT, a log that records significant application events.

application queue. A queue used by an application.

asynchronous messaging. A method of communication between programs in which programs place messages on message queues. With asynchronous messaging, the sending program proceeds with its own processing without waiting for a reply to its message. Contrast with *synchronous messaging*.

attribute. One of a set of properties that defines the characteristics of an MQSeries object.

authorization checks. Security checks that are performed when a user tries to open an MQSeries object.

authorization file. In MQSeries on UNIX systems, a file that provides security definitions for an object, a class of objects, or all classes of objects.

authorization service. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a service that provides authority checking of commands and MQI calls for the user identifier associated with the command or call.

authorized program analysis report (APAR). A report of a problem caused by a suspected defect in a current, unaltered release of a program.

B

backout. An operation that reverses all the changes made during the current unit of recovery or unit of work. After the operation is complete, a new unit of recovery or unit of work begins. Contrast with *commit*.

browse. In message queuing, to use the MQGET call to copy a message without removing it from the queue. See also *get*.

browse cursor. In message queuing, an indicator used when browsing a queue to identify the message that is next in sequence.

C

call back. In MQSeries, a requester message channel initiates a transfer from a sender channel by first calling the sender, then closing down and awaiting a call back.

CCF. Channel control function.

CCSID. Coded character set identifier.

CDF. Channel definition file.

channel. See *message channel*.

channel control function (CCF). In MQSeries, a program to move messages from a transmission queue to a communication link, and from a communication link to a local queue, together with an operator panel interface to allow the setup and control of channels.

channel definition file (CDF) • dead-letter queue (DLQ)

channel definition file (CDF). In MQSeries, a file containing communication channel definitions that associate transmission queues with communication links.

channel event. An event indicating that a channel instance has become available or unavailable. Channel events are generated on the queue managers at both ends of the channel.

checkpoint. (1) A time when significant information is written on the log. Contrast with *syncpoint*. (2) In MQSeries on UNIX systems, the point in time when a data record described in the log is the same as the data record in the queue. Checkpoints are generated automatically and are used during the system restart process.

CICS transaction. In CICS, a unit of application processing, usually comprising one or more units of work.

circular logging. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, the process of keeping all restart data in a ring of log files. Logging fills the first file in the ring and then moves on to the next, until all the files are full. At this point, logging goes back to the first file in the ring and starts again, if the space has been freed or is no longer needed. Circular logging is used during restart recovery, using the log to roll back transactions that were in progress when the system stopped. Contrast with *linear logging*.

client. A run-time component that provides access to queuing services on a server for local user applications. The queues used by the applications reside on the server. See also *MQSeries client*.

client application. An application, running on a workstation and linked to a client, that gives the application access to queuing services on a server.

client connection channel type. The type of MQI channel definition associated with an MQSeries client. See also *server connection channel type*.

coded character set identifier (CCSID). The name of a coded set of characters and their code point assignments.

command. In MQSeries, an instruction that can be carried out by the queue manager.

command processor. The MQSeries component that processes commands.

command server. The MQSeries component that reads commands from the system-command input queue, verifies them, and passes valid commands to the command processor.

commit. An operation that applies all the changes made during the current unit of recovery or unit of work. After the operation is complete, a new unit of recovery or unit of work begins. Contrast with *backout*.

completion code. A return code indicating how an MQI call has ended.

configuration file. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a file that contains configuration information related to, for example, logs, communications, or installable services. Synonymous with *.ini file*. See also *stanza*.

connect. To provide a queue manager connection handle, which an application uses on subsequent MQI calls. The connection is made either by the MQCONN call, or automatically by the MQOPEN call.

connection handle. The identifier or token by which a program accesses the queue manager to which it is connected.

context. Information about the origin of a message.

context security. In MQSeries, a method of allowing security to be handled such that messages are obliged to carry details of their origins in the message descriptor.

control command. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a command that can be entered interactively from the operating system command line. Such a command requires only that the MQSeries product be installed; it does not require a special utility or program to run it.

controlled shutdown. See *quiesced shutdown*.

D

data conversion interface (DCI). The MQSeries interface to which customer- or vendor-written programs that convert application data between different machine encodings and CCSIDs must conform. A part of the MQSeries Framework.

datagram. The simplest message that MQSeries supports. This type of message does not require a reply.

DCE. Distributed Computing Environment.

DCI. Data conversion interface.

dead-letter queue (DLQ). A queue to which a queue manager or application sends messages that it cannot deliver to their correct destination.

dead-letter queue handler. An MQSeries-supplied utility that monitors a dead-letter queue (DLQ) and processes messages on the queue in accordance with a user-written rules table.

default object. A definition of an object (for example, a queue) with all attributes defined. If a user defines an object but does not specify all possible attributes for that object, the queue manager uses default attributes in place of any that were not specified.

distributed application. In message queuing, a set of application programs that can each be connected to a different queue manager, but that collectively constitute a single application.

Distributed Computing Environment (DCE). Middleware that provides some basic services, making the development of distributed applications easier. DCE is defined by the Open Software Foundation (OSF).

distributed queue management (DQM). In message queuing, the setup and control of message channels to queue managers on other systems.

DLQ. Dead-letter queue.

DQM. Distributed queue management.

dynamic queue. A local queue created when a program opens a model queue object. See also *permanent dynamic queue* and *temporary dynamic queue*.

E

event. See *channel event*, *instrumentation event*, *performance event*, and *queue manager event*.

event data. In an event message, the part of the message data that contains information about the event (such as the queue manager name, and the application that gave rise to the event). See also *event header*.

event header. In an event message, the part of the message data that identifies the event type of the reason code for the event.

event log. See *application log*.

event message. Contains information (such as the category of event, the name of the application that caused the event, and queue manager statistics) relating to the origin of an instrumentation event in a network of MQSeries systems.

event queue. The queue onto which the queue manager puts an event message after it detects an event. Each category of event (queue manager,

performance, or channel event) has its own event queue.

Event Viewer. A tool provided by Windows NT to examine and manage log files.

F

FFST. First Failure Support Technology.

FIFO. First-in-first-out.

First Failure Support Technology (FFST). Used by MQSeries on UNIX systems, MQSeries for OS/2 Warp, MQSeries for Windows NT, and MQSeries for OS/400 to detect and report software problems.

first-in-first-out (FIFO). A queuing technique in which the next item to be retrieved is the item that has been in the queue for the longest time. (A)

Framework. In MQSeries, a collection of programming interfaces that allow customers or vendors to write programs that extend or replace certain functions provided in MQSeries products. The interfaces are:

- MQSeries data conversion interface (DCI)
- MQSeries message channel interface (MCI)
- MQSeries name service interface (NSI)
- MQSeries security enabling interface (SEI)
- MQSeries trigger monitor interface (TMI)

G

get. In message queuing, to use the MQGET call to remove a message from a queue. See also *browse*.

H

handle. See *connection handle* and *object handle*.

I

immediate shutdown. In MQSeries, a shutdown of a queue manager that does not wait for applications to disconnect. Current MQI calls are allowed to complete, but new MQI calls fail after an immediate shutdown has been requested. Contrast with *quiesced shutdown* and *preemptive shutdown*.

.ini file. See *configuration file*.

initiation queue. A local queue on which the queue manager puts trigger messages.

input/output parameter. A parameter of an MQI call in which you supply information when you make the

input parameter • message channel agent (MCA)

call, and in which the queue manager changes the information when the call completes or fails.

input parameter. A parameter of an MQI call in which you supply information when you make the call.

installable services. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, additional functionality provided as independent components. The installation of each component is optional: in-house or third-party components can be used instead. See also *authorization service*, *name service*, and *user identifier service*.

instrumentation event. A facility that can be used to monitor the operation of queue managers in a network of MQSeries systems. MQSeries provides instrumentation events for monitoring queue manager resource definitions, performance conditions, and channel conditions. Instrumentation events can be used by a user-written reporting mechanism in an administration application that displays the events to a system operator. They also allow applications acting as agents for other administration networks to monitor reports and create the appropriate alerts.

L

linear logging. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, the process of keeping restart data in a sequence of files. New files are added to the sequence as necessary. The space in which the data is written is not reused until the queue manager is restarted. Contrast with *circular logging*.

listener. In MQSeries distributed queuing, a program that monitors for incoming network connections.

local definition. An MQSeries object belonging to a local queue manager.

local definition of a remote queue. An MQSeries object belonging to a local queue manager. This object defines the attributes of a queue that is owned by another queue manager. In addition, it is used for queue-manager aliasing and reply-to-queue aliasing.

locale. On UNIX systems, a subset of a user's environment that defines conventions for a specific culture (such as time, numeric, or monetary formatting and character classification, collation, or conversion). The queue manager CCSID is derived from the locale of the user ID that created the queue manager.

local queue. A queue that belongs to the local queue manager. A local queue can contain a list of messages waiting to be processed. Contrast with *remote queue*.

local queue manager. The queue manager to which a program is connected and that provides message queuing services to the program. Queue managers to which a program is not connected are called *remote queue managers*, even if they are running on the same system as the program.

log. In MQSeries, a file recording the work done by queue managers while they receive, transmit, and deliver messages.

log control file. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, the file containing information needed to monitor the use of log files (for example, their size and location, and the name of the next available file).

log file. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a file in which all significant changes to the data controlled by a queue manager are recorded. If the primary log files become full, MQSeries allocates secondary log files.

logical unit of work (LUW). See *unit of work*.

M

mail-in database. A Lotus Notes database for sole use by the add-in task. It holds the request from a Lotus Notes application before the request is passed to the MQSeries application.

MCA. Message channel agent.

MCI. Message channel interface.

media image. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, the sequence of log records that contain an image of an object. The object can be recreated from this image.

message. (1) In message queuing applications, a communication sent between programs. See also *persistent message* and *nonpersistent message*. (2) In system programming, information intended for the terminal operator or system administrator.

message channel. In distributed message queuing, a mechanism for moving messages from one queue manager to another. A message channel comprises two message channel agents (a sender and a receiver) and a communication link. Contrast with *MQI channel*.

message channel agent (MCA). A program that transmits prepared messages from a transmission queue to a communication link, or from a communication link to a destination queue.

message channel interface (MCI). The MQSeries interface to which customer- or vendor-written programs that transmit messages between an MQSeries queue manager and another messaging system must conform. A part of the MQSeries Framework.

message descriptor. Control information describing the message format and presentation that is carried as part of an MQSeries message. The format of the message descriptor is defined by the MQMD structure.

message priority. In MQSeries, an attribute of a message that can affect the order in which messages on a queue are retrieved, and whether a trigger event is generated.

message queue. Synonym for *queue*.

message queue interface (MQI). The programming interface provided by the MQSeries queue managers. This programming interface allows application programs to access message queuing services.

message queuing. A programming technique in which each program within an application communicates with the other programs by putting messages on queues.

message sequence numbering. A programming technique in which messages are given unique numbers during transmission over a communication link. This enables the receiving process to check whether all messages are received, to place them in a queue in the original order, and to discard duplicate messages.

messaging. See *synchronous messaging* and *asynchronous messaging*.

model queue object. A set of queue attributes that act as a template when a program creates a dynamic queue.

MQI. Message queue interface.

MQI channel. Connects an MQSeries client to a queue manager on a server system, and transfers only MQI calls and responses in a bidirectional manner. Contrast with *message channel*.

MQSC. MQSeries commands.

MQSeries. A family of IBM licensed programs that provides message queuing services.

MQSeries client. Part of an MQSeries product that can be installed on a system without installing the full queue manager. The MQSeries client accepts MQI calls from applications and communicates with a queue manager on a server system.

MQSeries commands (MQSC). Human readable commands, uniform across all platforms, that are used

to manipulate MQSeries objects. Contrast with *programmable command format (PCF)*.

N

name service. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, the facility that determines which queue manager owns a specified queue.

name service interface (NSI). The MQSeries interface to which customer- or vendor-written programs that resolve queue-name ownership must conform. A part of the MQSeries Framework.

name transformation. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, an internal process that changes a queue manager name so that it is unique and valid for the system being used. Externally, the queue manager name remains unchanged.

New Technology File System (NTFS). A Windows NT recoverable file system that provides security for files.

nonpersistent message. A message that does not survive a restart of the queue manager. Contrast with *persistent message*.

NSI. Name service interface.

NTFS. New Technology File System.

null character. The character that is represented by X'00'.

O

OAM. Object authority manager.

object. In MQSeries, an object is a queue manager, a queue, a process definition, a channel, a namelist (MVS/ESA only), or a storage class (MVS/ESA only).

object authority manager (OAM). In MQSeries on UNIX systems and MQSeries for Windows NT, the default authorization service for command and object management. The OAM can be replaced by, or run in combination with, a customer-supplied security service.

object descriptor. A data structure that identifies a particular MQSeries object. Included in the descriptor are the name of the object and the object type.

object handle. The identifier or token by which a program accesses the MQSeries object with which it is working.

output parameter • quiescing

output parameter. A parameter of an MQI call in which the queue manager returns information when the call completes or fails.

P

PCF. Programmable command format.

PCF command. See *programmable command format*.

pending event. An unscheduled event that occurs as a result of a connect request from a CICS adapter.

percolation. In error recovery, the passing along a preestablished path of control from a recovery routine to a higher-level recovery routine.

performance event. A category of event indicating that a limit condition has occurred.

performance trace. An MQSeries trace option where the trace data is to be used for performance analysis and tuning.

permanent dynamic queue. A dynamic queue that is deleted when it is closed only if deletion is explicitly requested. Permanent dynamic queues are recovered if the queue manager fails, so they can contain persistent messages. Contrast with *temporary dynamic queue*.

persistent message. A message that survives a restart of the queue manager. Contrast with *nonpersistent message*.

ping. In distributed queuing, a diagnostic aid that uses the exchange of a test message to confirm that a message channel or a TCP/IP connection is functioning.

platform. In MQSeries, the operating system under which a queue manager is running.

preemptive shutdown. In MQSeries, a shutdown of a queue manager that does not wait for connected applications to disconnect, nor for current MQI calls to complete. Contrast with *immediate shutdown* and *quiesced shutdown*.

principal. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a term used for a user identifier. Used by the object authority manager for checking authorizations to system resources.

process definition object. An MQSeries object that contains the definition of an MQSeries application. For example, a queue manager uses the definition when it works with trigger messages.

programmable command format (PCF). A type of MQSeries message used by:

- User administration applications, to put PCF commands onto the system command input queue of a specified queue manager
- User administration applications, to get the results of a PCF command from a specified queue manager
- A queue manager, as a notification that an event has occurred

Contrast with *MQSC*.

program temporary fix (PTF). A solution or by-pass of a problem diagnosed by IBM field engineering as the result of a defect in a current, unaltered release of a program.

PTF. Program temporary fix.

Q

queue. An MQSeries object. Message queuing applications can put messages on, and get messages from, a queue. A queue is owned and maintained by a queue manager. Local queues can contain a list of messages waiting to be processed. Queues of other types cannot contain messages—they point to other queues, or can be used as models for dynamic queues.

queue manager. (1) A system program that provides queuing services to applications. It provides an application programming interface so that programs can access messages on the queues that the queue manager owns. See also *local queue manager* and *remote queue manager*. (2) An MQSeries object that defines the attributes of a particular queue manager.

queue manager event. An event that indicates:

- An error condition has occurred in relation to the resources used by a queue manager. For example, a queue is unavailable.
- A significant change has occurred in the queue manager. For example, a queue manager has stopped or started.

queuing. See *message queuing*.

quiesced shutdown. (1) In MQSeries, a shutdown of a queue manager that allows all connected applications to disconnect. Contrast with *immediate shutdown* and *preemptive shutdown*. (2) A type of shutdown of the CICS adapter where the adapter disconnects from MQSeries, but only after all the currently active tasks have been completed. Contrast with *forced shutdown*.

quiescing. In MQSeries, the state of a queue manager prior to it being stopped. In this state,

programs are allowed to finish processing, but no new programs are allowed to start.

R

RBA. Relative byte address.

reason code. A return code that describes the reason for the failure or partial success of an MQI call.

receiver channel. In message queuing, a channel that responds to a sender channel, takes messages from a communication link, and puts them on a local queue.

Registry. In Windows NT, a secure database that provides a single source for system and application configuration data.

Registry Editor. In Windows NT, the program item that allows the user to edit the Registry.

Registry Hive. In Windows NT, the structure of the data stored in the Registry.

remote queue. A queue belonging to a remote queue manager. Programs can put messages on remote queues, but they cannot get messages from remote queues. Contrast with *local queue*.

remote queue manager. To a program, a queue manager that is not the one to which the program is connected.

remote queue object. See *local definition of a remote queue*.

remote queuing. In message queuing, the provision of services to enable applications to put messages on queues belonging to other queue managers.

reply message. A type of message used for replies to request messages.

reply-to queue. The name of a queue to which the program that issued an MQPUT call wants a reply message or report message sent.

report message. A type of message that gives information about another message. A report message can indicate that a message has been delivered, has arrived at its destination, has expired, or could not be processed for some reason.

requester channel. In message queuing, a channel that may be started remotely by a sender channel. The requester channel accepts messages from the sender channel over a communication link and puts the messages on the local queue designated in the message. See also *server channel*.

request message. A type of message used to request a reply from another program.

resolution path. The set of queues that are opened when an application specifies an alias or a remote queue on input to an MQOPEN call.

resource manager. An application, program, or transaction that manages and controls access to shared resources such as memory buffers and data sets. MQSeries, CICS, and IMS are resource managers.

responder. In distributed queuing, a program that replies to network connection requests from another system.

resynch. In MQSeries, an option to direct a channel to start up and resolve any in-doubt status messages, but without restarting message transfer.

return codes. The collective name for completion codes and reason codes.

rollback. Synonym for *back out*.

rules table. A control file containing one or more rules that the dead-letter queue handler applies to messages on the DLQ.

S

security enabling interface (SEI). The MQSeries interface to which customer- or vendor-written programs that check authorization, supply a user identifier, or perform authentication must conform. A part of the MQSeries Framework.

SEI. Security enabling interface.

sender channel. In message queuing, a channel that initiates transfers, removes messages from a transmission queue, and moves them over a communication link to a receiver or requester channel.

sequential delivery. In MQSeries, a method of transmitting messages with a sequence number so that the receiving channel can reestablish the message sequence when storing the messages. This is required where messages must be delivered only once, and in the correct order.

sequential number wrap value. In MQSeries, a method of ensuring that both ends of a communication link reset their current message sequence numbers at the same time. Transmitting messages with a sequence number ensures that the receiving channel can reestablish the message sequence when storing the messages.

server • trigger message

server. (1) In MQSeries, a queue manager that provides queue services to client applications running on a remote workstation. (2) The program that responds to requests for information in the particular two-program, information-flow model of client/server. See also *client*.

server channel. In message queuing, a channel that responds to a requester channel, removes messages from a transmission queue, and moves them over a communication link to the requester channel.

server connection channel type. The type of MQI channel definition associated with the server that runs a queue manager. See also *client connection channel type*.

service interval. A time interval, against which the elapsed time between a put or a get and a subsequent get is compared by the queue manager in deciding whether the conditions for a service interval event have been met. The service interval for a queue is specified by a queue attribute.

service interval event. An event related to the service interval.

shutdown. See *immediate shutdown*, *preemptive shutdown*, and *quiesced shutdown*.

single-phase backout. A method in which an action in progress must not be allowed to finish, and all changes that are part of that action must be undone.

single-phase commit. A method in which a program can commit updates to a queue without coordinating those updates with updates the program has made to resources controlled by another resource manager. Contrast with *two-phase commit*.

stanza. A group of lines in a configuration file that assigns a value to a parameter modifying the behavior of a queue manager, client, or channel. In MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a configuration (.ini) file may contain a number of stanzas.

store and forward. The temporary storing of packets, messages, or frames in a data network before they are retransmitted toward their destination.

symptom string. Diagnostic information displayed in a structured format designed for searching the IBM software support database.

synchronous messaging. A method of communication between programs in which programs place messages on message queues. With synchronous messaging, the sending program waits for

a reply to its message before resuming its own processing. Contrast with *asynchronous messaging*.

syncpoint. An intermediate or end point during processing of a transaction at which the transaction's protected resources are consistent. At a syncpoint, changes to the resources can safely be committed, or they can be backed out to the previous syncpoint.

system.command.input queue. A local queue on which application programs can put MQSeries commands. The commands are retrieved from the queue by the command server, which validates them and passes them to the command processor to be run.

system control commands. Commands used to manipulate platform-specific entities such as buffer pools, storage classes, and page sets.

T

temporary dynamic queue. A dynamic queue that is deleted when it is closed. Temporary dynamic queues are not recovered if the queue manager fails, so they can contain nonpersistent messages only. Contrast with *permanent dynamic queue*.

thread. In MQSeries, the lowest level of parallel execution available on an operating system platform.

time-independent messaging. See *asynchronous messaging*.

TMI. Trigger monitor interface. See also *global trace* and *performance trace*.

tranid. See *transaction identifier*.

transaction. See *unit of work* and *CICS transaction*.

transaction identifier. In CICS, a name that is specified when the transaction is defined, and that is used to invoke the transaction.

transmission program. See *message channel agent*.

transmission queue. A local queue on which prepared messages destined for a remote queue manager are temporarily stored.

trigger event. An event (such as a message arriving on a queue) that causes a queue manager to create a trigger message on an initiation queue.

triggering. In MQSeries, a facility allowing a queue manager to start an application automatically when predetermined conditions on a queue are satisfied.

trigger message. A message containing information about the program that a trigger monitor is to start.

trigger monitor. A continuously-running application serving one or more initiation queues. When a trigger message arrives on an initiation queue, the trigger monitor retrieves the message. It uses the information in the trigger message to start a process that serves the queue on which a trigger event occurred.

trigger monitor interface (TMI). The MQSeries interface to which customer- or vendor-written trigger monitor programs must conform. A part of the MQSeries Framework.

two-phase commit. A protocol for the coordination of changes to recoverable resources when more than one resource manager is used by a single transaction. Contrast with *single-phase commit*.

U

UIS. User identifier service.

undelivered-message queue. See *dead-letter queue*.

undo/redo record. A log record used in recovery. The redo part of the record describes a change to be made to an MQSeries object. The undo part describes how to back out the change if the work is not committed.

unit of recovery. A recoverable sequence of operations within a single resource manager. Contrast with *unit of work*.

unit of work. A recoverable sequence of operations performed by an application between two points of consistency. A unit of work begins when a transaction starts or after a user-requested syncpoint. It ends either at a user-requested syncpoint or at the end of a transaction. Contrast with *unit of recovery*.

user identifier service (UIS). In MQSeries for OS/2 Warp, the facility that allows MQI applications to associate a user ID, other than the default user ID, with MQSeries messages.

utility. In MQSeries, a supplied set of programs that provide the system operator or system administrator with facilities in addition to those provided by the MQSeries commands. Some utilities invoke more than one function.

X

X/Open XA. The X/Open Distributed Transaction Processing XA interface. A proposed standard for distributed transaction communication. The standard specifies a bidirectional interface between resource managers that provide access to shared resources within transactions, and between a transaction service that monitors and resolves transactions.

Index

A

ACTION keyword, rules table 120
 action keywords, rules table 120
 add queue manager to the Windows NT Service Control Manager 282
 administration
 authorizations 91
 command sets 19
 control commands 19
 MQSeries commands (MQSC) 20
 programmable command format commands (PCF) 21
 local 39
 remote 64
 channels 65
 objects 63
 transmission queues 65
 alias queues 55
 authorizations to 85
 description 11
 aliases
 queue manager 74
 reply-to queues 74
 alter queue manager attributes 44
 alternate-user authority 85
 amqsdlq, the sample DLQ handler 116
 application
 data 8
 design considerations 211
 MQI administration support 39
 programming errors, examples of 206
 time-independent 7
 APPLIDAT keyword, rules table 118
 APPLNAME keyword, rules table 119
 APPLTYPE keyword, rules table 119
 attributes
 ALL attribute 51
 altering 44
 changing 52
 default 51
 displaying queue manager 42
 MQSC and PCFs compared 22
 queue manager
 altering 44
 displaying 42
 queues 11
 authority
 alternate-user 85
 commands 84
 context 86
 installable services 84

authority (*continued*)
 set/reset command 284
 authorization
 administration 91
 dspmqaut command 84
 lists 83
 MQI 89
 setmqaut command 84
 user groups 81
 authorization files
 all class 97
 authorization to 97
 class 97
 contents 96
 directories 94
 managing 97
 paths 94
 understanding 94
 auto-definition of channels 68

B

bibliography xi
 binding
 fastpath 61
 standard 61
 BookManager xv
 browsing queues 53

C

case-sensitive control commands 20
 CCSID 75
 ccsid.tbl 75
 changing queue attributes 52
 channel
 auto-definition 68
 command security requirements 87
 commands 87
 configuration 104
 defining 66
 description 14, 63
 escape command authorizations 92
 events 129
 remote administration 65
 run command 272
 run initiator command 271
 security 87
 starting 67
 character set, specifying 75
 CICS
 enabling user exits 167

Index

- CICS (*continued*)
 - user exits 167
 - user ID 85
 - with MQSeries 166
- circular logging 172
- clearing a local queue 53
- client based mail 137
- clients 15
 - error messages on DOS and Windows 231
 - problem determination 230
 - trigger monitor start command 280
- coded character set
 - specifying 75
- command errors 206
- command files 45
- command queue 13
- command server
 - display command 252
 - displaying status 37
 - end command 258
 - remote administration 37
 - start command 292
 - starting 37
 - stopping 38
- command set
 - administration 19
 - comparison 22
- commands
 - add a queue manager to Windows NT Service Control Manager 33
 - add queue manager to the Windows NT Service Control Manager 282
 - comparison of sets 22
 - control 19
 - create queue manager (crtmqm) 240
 - data conversion (crtmqcvx) 238
 - delete queue manager (dlmqm) 244
 - display authority (dspmqaut) 248
 - display command server (dspmqcsv) 252
 - display MQSeries files (dspmqfls) 253
 - display MQSeries formatted trace (dspmqtrc) 255
 - display MQSeries transactions (dspmqtrn) 256
 - dump log (dmpmqlog) 246
 - end command server (endmqcsv) 258
 - end listener (endmqslr) 260
 - end MQSeries trace (endmqtrc) 263
 - end queue manager (endmqm) 261
 - enroll production license (setmqprd) 290
 - help with syntax 237
- MQSC
 - ALTER QLOCAL 52
 - ALTER QREMOTE 73
 - DEFINE CHANNEL 66
 - DEFINE QALIAS 55
 - DEFINE QLOCAL 51
 - DEFINE QLOCAL LIKE 52
 - DEFINE QLOCAL REPLACE 52
- commands (*continued*)
 - MQSC (*continued*)
 - DEFINE QMODEL 58
 - DEFINE QREMOTE 71
 - DELETE QLOCAL 53
 - DELETE QREMOTE 73
 - DISPLAY QREMOTE 73
 - MQSC command files 46
 - input 45
 - output reports 46
 - MQSeries (MQSC)
 - using 21
 - verifying 47
 - MQSeries commands (MQSC) 20
 - programmable command format (PCF) 21
 - record media image (rcdmqimg) 265
 - recreate object (rcrmqobj) 267
 - resolve MQSeries transactions (rsvmqtrn) 269
 - run channel (runmqchl) 272
 - run channel initiator (runmqchi) 271
 - run dead-letter queue handler 273
 - run DLQ handler (runmqdlq) 115
 - run listener (runmqslr) 275
 - run MQSeries commands (runmqsc) 277
 - runmqsc 41
 - security commands
 - dspmqaut 84
 - setmqaut 82
 - set/reset authority (setmqaut) 83, 284
 - start client trigger monitor (runmqtrmc) 280
 - start command server (strmqcsv) 292
 - start MQSeries trace (strmqtrc) 295
 - start queue manager (strmqm) 293
 - start trial period (setmqtry) 291
 - start trigger monitor (runmqtrm) 281
- configuration
 - size and location of log 108
- configuration files
 - LogPrimaryFile value 110
- MQSeries (mqseries.ini)
 - contents 99
 - LogBufferPages 111
 - LogDefaultPath 111
 - LogDefaults 109
 - LogFilePages 110
 - logging parameters 112
 - LogPath 112
 - LogSecondaryFiles 110
 - overview 99
 - path 47
- overview 99
- queue manager (qm.ini)
 - contents 103
 - disabling the object authority manager 82
 - Log stanza 109
 - LogBufferPages 111
 - LogDefaultPath 111

- configuration files (*continued*)
 - queue manager (qm.ini) (*continued*)
 - LogFilePages 110
 - logging parameters 112
 - LogPath 112
 - LogSecondaryFiles 110
 - stanzas 103
 - configuration, Lotus Notes 137
 - contents of
 - mqs.ini 99
 - qm.ini 103
 - context authority 86
 - control commands 19
 - case-sensitive 20
 - runmqsc 41
 - controlled shutdown 34
 - CorrelId, performance considerations when using 211
 - creating
 - crtmqm command 240
 - process definitions 60
 - queue manager 27, 32
 - crtmqcvx command
 - examples 238
 - parameters 238
 - return codes 238
 - crtmqm command 240
 - examples 243
 - parameters 240
 - related commands 243
 - return codes 243
 - current queue depth 51
- D**
- daemon, inetd 85
 - data conversion 75
 - crtmqcvx command 238
 - DCE 16
 - dead-letter header, MQDLH 115
 - dead-letter queue
 - description 12
 - handler 273
 - specifying 29
 - debugging
 - command syntax errors 206
 - common command errors 206
 - common programming errors 206
 - preliminary checks 203
 - secondary checks 207—210
 - default
 - attributes of objects 51
 - data conversion 75
 - objects 14, 301
 - overriding the configuration file 108
 - queue manager 28
 - accidental change 36
 - accidental deletion 241
 - default (*continued*)
 - queue manager (*continued*)
 - changing 36, 44
 - commands processed 41
 - transmission queue 29, 73
 - user group for authority 81
 - DEFINE QUEUE command, REPLACE attribute
 - defining queues 11
 - deleting 53
 - dltmqm command 244
 - local queue 53
 - queue manager 36
 - DESTQ keyword, rules table 119
 - DESTQM keyword, rules table 119
 - directories 85
 - authorization 94
 - queue manager 85
 - directory structure 303, 307, 311
 - disabling events 130
 - disabling the object authority manager 82
 - disconnected requests 137
 - display
 - authority command 248
 - command server command 252
 - MQSeries files command 253
 - MQSeries formatted trace output command 255
 - MQSeries transactions command 256
 - process definitions 61
 - queue manager attributes 42
 - status of command server 37
 - display authority command
 - See dspmqaut command
 - distributed queuing
 - dead-letter queue 12
 - incorrect output 214
 - undelivered-message queue 12
 - DLQ handler
 - invoking 115
 - rules table 116
 - sample, amqsdlq 116
 - dltmqm command 244
 - examples 244
 - parameters 244
 - related commands 245
 - return codes 244
 - dmpmqlog command 246
 - parameters 246
 - DOS clients error messages 231
 - dspmqaut command 248
 - examples 251
 - parameters 248
 - related commands 251
 - return codes 250
 - using 82, 84
 - dspmqcsv command 252
 - examples 252

Index

- dspmqcsv command (*continued*)
 - parameters 252
 - related commands 252
 - return codes 252
- dspmqls command 253
 - examples 254
 - parameters 253
 - return codes 254
- dspmqrtrc command 255
 - parameters 255
 - related commands 255
- dspmqrtrn command 256
 - parameters 256
 - related commands 257
 - return codes 256
- dynamic definition of channels 68
- dynamic queues 9
 - authorizations to 85

E

- enabling
 - events 130
 - security 82
- end listener 260
- end MQSeries trace 263
- ending
 - queue manager 34
- ending interactive MQSC commands 42
- endmqcsv command 258
 - examples 258
 - parameters 258
 - related commands 259
 - return codes 258
- endmqslr command 260
 - parameters 260
 - return codes 260
- endmqm command 34, 261
 - examples 262
 - parameters 261
 - related commands 262
 - return codes 262
- endmqtrc command 263
 - examples 263
 - parameters 263
 - related commands 264
 - return codes 263
- environment variable
 - MQS_TRACE_OPTIONS 220
- environment variable, disabling security 82
- environment variables
 - MQDATA 231
 - MQIBindType 104
 - MQSNOAUT 82
 - MQSPREFIX 108

- error log 215
 - error occurring before established 217
 - example 217
- error messages 42
- escape PCFs 22
- event queue 13
- event-driven processing 7
- events
 - channel 129
 - instrumentation
 - description 127
 - enabling and disabling 130
 - message 131
 - types of 129
 - what they are 127
 - why use them 128
 - queues 130
 - trigger 130
 - types of 129
- examples
 - crtmqcvx command 238
 - crtmqm command 243
 - dltmqm command 244
 - dspmqaout command 251
 - dspmqcsv command 252
 - dspmqls command 254
 - endmqcsv command 258
 - endmqm command 262
 - endmqtrc command 263
 - error log 217
 - programming errors 206
 - rcdmqimg command 266
 - rcrmqobj command 268
 - runmqslr command 276
 - runmqsc command 278
 - scmmqm command 283
 - setmqaut command 289
 - strmqcsv command 292
 - strmqm command 293
 - strmqtrc command 297
 - trace data (AIX) 221

F

- fastpath binding 61
- feedback from MQSC commands 42
- FEEDBACK keyword, rules table 119
- FFST
 - customized dump 230
 - in problem determination 228
 - symptom records 229
 - using 229
- FFST, examining 226, 227
- FFST/2
 - further information 228

file sizes, for logs 112

files

authorization

all class 97

authorizations to 97

class 97

contents 96

managing 97

paths 94

understanding 94

configuration

in problem determination 219

overview 99

log control 172

MQSeries configuration 99

queue manager configuration 103

understanding names 30

First Failure Support Technology

see FFST

FORMAT keyword, rules table 119

format of logs 171

FWDQ keyword, rules table 120

FWDQM keyword, rules table 120

G

glossary 325

group set authorizations 81

group sets, for authority 80

H

HEADER keyword, rules table 121

help for syntax 237

HTML (Hypertext Markup Language) xv

Hypertext Markup Language (HTML) xv

I

incorrect output 212

indirect mode, of runmqsc 69

inetd daemon 85

Information Presentation Facility (IPF) xvi

initiation queue

defining 60

description 12

input, standard 41

INPUTQ keyword, rules table 117

INPUTQM keyword, rules table 117

installable component

authority manager (OAM) 79

installable services

disabling object authority manager 82

disabling 82

object authority manager 79

user identifier 319

installation directory, mqmtop x

instrumentation event

description 127

enabling 130

messages 131

types of 129

why use them 128

interactive MQSC

ending 42

feedback from 42

using 41

IPF (Information Presentation Facility) xvi

issuing MQSeries commands 40

L

LIKE attribute 51

linear logging 173

linking to Lotus Notes 133

local administration 39

local queues

clearing 53

command 13

copying definitions 51

dead-letter 12

defining one 49

deleting 53

description 11

initiation 12

transmission 12

undelivered-message 12

log

configuration 108

directory structure 309, 313

error 215

error, example of 217

file

control 172

path 111

reuse 174

size 110

file size 112

format 171

managing 194

number of buffers 111

overheads 112

parameters 30

primary files 110

queue manager 171

secondary files 110

type of 111

using for recovery 197

logging

checkpoints 174

circular 172

linear 173

Index

- logging (*continued*)
 - media recovery 198
 - types of 172
- Lotus Notes
 - configuration 137
 - mapping 133
 - setup 137
 - with MQSeries 133
- M**
- managing access 80
- managing log files 195
- managing objects for triggering 59
- mapping Lotus Notes 133
- maximum line length for MQSC commands 45
- media images
 - description 197
 - record 198
 - record command 265
 - recovering 198
- message
 - containing unexpected information 214
 - description 8
 - descriptor 8
 - errors on DOS and Windows clients 231
 - for instrumentation events 131
 - lengths of 8
 - not appearing on queues 212
 - operator 217
 - performance considerations
 - lengths of 211
 - persistent 211
 - queuing 7
 - retrieval algorithms 9
 - searching for particular 211
 - undelivered 219
 - variable length 211
- message length, decreasing 52
- message queue interface (MQI) 7
- message queuing 7
- message-driven processing 7
- model queues
 - defining 58
 - description 12
 - working with 57
- monitoring
 - queue managers 128
- MQDATA environment variable 231
- MQDLH, dead-letter header 115
- MQI
 - authorizations 88, 89
 - description 7
 - local administration support 39
 - queue manager calls 11
- MQIBindType environment variable 104
- mqlink 136
- mqlinkc 136
- mqm
 - user group 77, 78
 - user ID 77, 78, 85
- mqmtop, the installation directory x
- MQOPEN authorizations 89
- MQPUT and MQPUT1, performance considerations 212
- MQPUT authorizations 89
- MQS_TRACE_OPTIONS environment variable 220
- mqsc.ini
 - See configuration files
- mqsc.ini, path to 47
- MQSC
 - attributes 22
 - command files
 - input 45
 - output reports 46
 - running 47
 - commands 20
 - ending interactive input 42
 - escape PCFs 22
 - how to issue commands 40
 - issuing remotely 69
 - maximum line length 45
 - problems
 - local 47
 - remote 70
 - redirecting input and output 44
 - security requirements on channels 87
 - timed out command responses 69
 - using commands 44
 - verifying commands 47
- MQSC commands
 - ALTER QLOCAL 52
 - ALTER QREMOTE 73
 - DEFINE CHANNEL 66
 - DEFINE QALIAS 55
 - DEFINE QLOCAL 51
 - DEFINE QLOCAL LIKE 52
 - DEFINE QLOCAL REPLACE 52
 - DEFINE QMODEL 58
 - DEFINE QREMOTE 71
 - DELETE QLOCAL 53
 - DELETE QREMOTE 73
 - DISPLAY QREMOTE 73
 - issuing interactively 41
 - maximum line length 45
 - using 21
- MQSeries
 - super user, mqm 77, 78
- MQSeries configuration file
 - See configuration files

MQSeries publications xi
 MQSNOAUT environment variable 82
 MQSPREFIX environment variable 108
 MQZAO constants and authority 89
 Msgld, performance considerations when using 211
 MSGTYPE keyword, rules table 119
 MVS/ESA queue manager 69

N

name service 16
 names
 allowed for objects 235
 objects 9
 naming conventions, national language support 235
 national language support
 naming conventions 235
 nobody, default user group 81
 Notes configuration 137
 Notes setup 137
 notification of events 130

O

OAM 79
 object authority manager 79
 default user group 81
 disabling 82
 dspmqaut command 84
 how it works 80
 principals 80
 sensitive operations 85
 setmqaut command 82, 83
 objects
 access to 77
 default
 attributes 51
 for triggering 59
 media image 197
 names of 9
 naming conventions 235
 process definition 13
 queue 11
 queue manager
 MQI calls 11
 recovery 198
 recreate command 267
 remote administration 63
 system default 14
 types of 9
 operating system variable, disabling security 82
 operator commands, no response from 208
 operator messages 217
 OS/2 user identifier service 319
 output, standard 41

overheads, for logs 112
 overrides
 in configuration files 108

P

parameters
 runmqslr command 275
 scmmqm command 282
 pattern-matching keywords, rules table 118
 PCF
 See programmable command format (PCF)
 performance considerations
 trace 219, 224
 performance considerations when using trace 222
 performance events 129
 permanent queues 9
 PERSIST keyword, rules table 119
 PostScript format xv
 predefined queues 9
 preemptive queue manager shutdown 35
 primary group authorizations 81
 primary group, for authority 80
 primary log files 110
 principals
 belonging to more than one group 81
 managing access to 80
 problem determination 203
 clients 230
 command errors 206
 configuration files 219
 FFST 228
 FFST/2 operation 229
 further checks 207—210
 incorrect output
 messages containing unexpected
 information 214
 messages not appearing on queues 212
 with distributed queuing 214
 no response from commands 208
 programming errors 206
 things to check first 203—207
 trace 219, 222, 224
 problems
 recovering from 197
 running MQSC commands 47
 using MQSC locally 47
 using MQSC remotely 70
 process definitions
 creating 60
 description 13
 displaying 61
 processing, event-driven 7
 production license, enrolling 290
 programmable command format
 See programmable command format (PCF)

Index

- programmable command format (PCF)
 - administration with 21
 - attributes 22
 - authorizations 88
 - commands 21
 - description 21
 - escape PCFs 22
 - security requirements 87
- programming errors, examples of 206
- protected resources 81
- publications
 - MQSeries xi
 - related xvi
- PUTAUT keyword, rules table 121

Q

- qm.ini
 - See configuration files
- queue depth
 - current 51
 - determining 51
- queue manager
 - alias
 - remote queue 74
 - authorization directories 94
 - authorizations 85
 - CICS 166
 - circular logging
 - restart recovery 173
 - command server 37
 - configuration file 103
 - configuration files
 - LogPath 112
 - specifying 30
 - creating 27, 32
 - crtmqm command 240
 - default 28
 - accidental change 36
 - accidental deletion 241
 - changing 36
 - deleting 36
 - dltmqm command 244
 - description 10
 - directories 85
 - dmpmqlog command 246
 - endmqm command 261
 - events 129
 - immediate shutdown 34
 - linear logging 173
 - local administration 39
 - logs 171
 - monitoring 128
 - MVS/ESA 69
 - name transformation 31
 - numbers of 28
- queue manager (*continued*)
 - object authority manager
 - description 79
 - disabling 82
 - objects
 - MQI calls 11
 - preemptive shutdown 34, 35
 - recording media images 198
 - remote administration 63
 - removing
 - manually 316
 - restart 35
 - shutdown
 - controlled 34
 - immediate 34
 - preemptive 34
 - quiesced 34
 - specifying on runmqsc 44
 - starting 33
 - automatically 33
 - stopping 34
 - manually 315
 - unique name 28
- queue manager configuration file
 - See configuration files
- queues
 - alias 11
 - aliases, working with 55
 - application
 - defining for triggering 59
 - attributes 11
 - changing 52
 - authorizations to 85
 - browsing 53
 - command 13
 - dead-letter 12
 - specifying 29
 - defining 11
 - description 8
 - distributed, incorrect output from 214
 - dynamic 9
 - event 13
 - event notification 130
 - for MQSeries applications 39
 - initiation
 - defining 60
 - trigger messages 12
 - local 11
 - clearing 53
 - copying 51
 - defining 49
 - deleting 53
 - model 12
 - defining 58
 - working with 57
 - objects
 - alias 11

- queues (*continued*)
 - objects (*continued*)
 - local 11
 - model 12
 - remote 11
 - predefined 9
 - remote 11
 - creating 71
 - queue manager alias 74
 - working with 74
 - reply-to 13, 74
 - temporary 9
 - transmission 12
 - creating 73
 - default 29, 73
 - defining 66
 - remote administration 65
 - undelivered-message 12
 - specifying 29
 - working with 49
 - quiesced shutdown 34
- R**
- railroad diagrams, how to read 236
 - rcdmqimg command 265
 - examples 266
 - parameters 265
 - related commands 266
 - return codes 266
 - rcrmqobj command 267, 268
 - examples 268
 - parameters 267
 - related commands 268
 - return codes 268
 - REASON keyword, rules table 119
 - receiver channel, automatic definition of 68
 - recovering
 - media images 198
 - recovery
 - scenarios 201
 - damaged queue manager object 202
 - damaged single object 202
 - disk drive failures 201
 - redirecting input and output, on MQSC commands 44
 - related publications xvi
 - remote
 - administration 64
 - of objects 63
 - issuing of MQSC commands 69
 - queue definition
 - creating 71
 - queue object
 - working with 74
 - queues
 - as queue manager aliases 74
 - as reply-to queue aliases 74
 - remote (*continued*)
 - queuing
 - description 63
 - recommendations 70
 - security considerations 87
 - remote administration
 - command server 37
 - initial problems 70
 - remote queues
 - authorizations to 85
 - description 11
 - removing
 - queue manager
 - manually 316
 - REPLACE attribute, DEFINE commands 45
 - replication 137
 - reply-to queue 13
 - reply-to queue aliases 74
 - REPLYQ keyword, rules table 119
 - REPLYQM keyword, rules table 119
 - resources
 - protected 81
 - why protect 79
 - restart queue manager 35
 - restart recovery
 - with circular logging 173
 - restrictions 77
 - access to MQM objects 77
 - for UNIX systems x
 - object names 235
 - retrieval algorithms for messages 9
 - RETRY keyword, rules table 121
 - RETRYINT keyword, rules table 117
 - return codes 204
 - crtmqcvx command 238
 - crtmqm command 243
 - dltmqm command 244
 - dspmqaout command 250
 - dspmqcsv command 252
 - dspmqls command 254
 - dspmqrn command 256
 - endmqcsv command 258
 - endmqslr command 260
 - endmqm command 262
 - endmqtrc command 263
 - rcdmqimg command 266
 - rcrmqobj command 268
 - rsvmqtrn command 270
 - runmqchi command 271
 - runmqchl command 272
 - runmqslr command 276
 - runmqsc command 278
 - runmqtrmc command 280
 - runmqtrm command 281
 - scmmqm command 282
 - setmqaut command 288

Index

- return codes (*continued*)
 - setmqprd command 290
 - setmqtry command 291
 - strmqcsv command 292
 - strmqm command 293
 - strmqtrc command 297
 - rsvmqtrn command 269
 - parameters 269
 - related commands 270
 - return codes 270
 - rules table, DLQ handler 116
 - See also* DLQ handler
 - control data entry 117
 - INPUTQ keyword 117
 - INPUTQM keyword 117
 - RETRYINT keyword 117
 - WAIT keyword 117
 - example 125
 - patterns and actions (rules) 118
 - ACTION keyword 120
 - APPLIDAT keyword 118
 - APPLNAME keyword 119
 - APPLTYPE keyword 119
 - DESTQ keyword 119
 - DESTQM keyword 119
 - FEEDBACK keyword 119
 - FORMAT keyword 119
 - FWDQ keyword 120
 - FWDQM keyword 120
 - HEADER keyword 121
 - MSGTYPE keyword 119
 - PERSIST keyword 119
 - PUTAUT keyword 121
 - REASON keyword 119
 - REPLYQ keyword 119
 - REPLYQM keyword 119
 - RETRY keyword 121
 - USERID keyword 120
 - processing of 123
 - syntax 121
 - run listener (runmqtsr command) 275
 - runmqchi command 271
 - parameters 271
 - return codes 271
 - runmqchl command 272
 - parameters 272
 - return codes 272
 - runmqdlq command 115
 - runmqtsr command 275
 - example 276
 - parameters 275
 - return codes 276
 - runmqsc
 - ending 42
 - feedback 42
 - indirect mode 69
 - runmqsc (*continued*)
 - issuing MQSC commands 40
 - problems 47
 - specifying a queue manager 44
 - using 44
 - using interactively 41
 - verifying 47
 - runmqsc command 277
 - examples 278
 - parameters 277
 - redirecting input and output 44
 - return codes 278
 - runmqtsmc command 280
 - parameters 280
 - return codes 280
 - runmqtsm command 281
 - parameters 281
 - return codes 281
- ## S
- samples
 - trace data (HP-UX) 222
 - trace data (MQSeries) 225
 - trace data (Sun Solaris) 224
 - scmmqm command 33, 282
 - parameters 282
 - return codes 282
 - secondary log files 110
 - security 77
 - enabling 82
 - remote 87
 - using the commands 82, 84
 - server based mail 137
 - server-connection channel, automatic definition of 68
 - set/reset authority command
 - See* setmqaut command
 - setmqaut command 284
 - examples 289
 - installable services 84
 - parameters 286
 - related commands 289
 - return codes 288
 - using 82, 83
 - setmqprd command 290
 - parameters 290
 - return codes 290
 - setmqtry command 291
 - return codes 291
 - setup, Lotus Notes 137
 - shell commands for MQSeries 19
 - shutdown queue manager
 - controlled 34
 - immediate 34
 - preemptive 34
 - quiesced 34

- softcopy books xv
 - specifying coded character set 75
 - standard binding 61
 - stanzas
 - mqs.ini 99
 - qm.ini 103
 - start MQSeries trace command 295
 - start queue manager command 293
 - starting
 - a queue manager 33
 - a queue manager automatically 33
 - channels 67
 - command server 37
 - stdin, on runmqsc 44
 - stdout, on runmqsc 44
 - stopping
 - command server 38
 - queue manager
 - manually 315
 - using endmqm command 34
 - strmqcsv command 292
 - examples 292
 - parameters 292
 - related commands 292
 - return codes 292
 - strmqm command 293
 - examples 293
 - parameters 293
 - related commands 294
 - return codes 293
 - strmqtrc command 295
 - examples 297
 - parameters 295
 - related commands 297
 - return codes 297
 - super user (MQSeries)
 - mqm 77, 78
 - symptom records, FFST 229
 - syncpoint, performance considerations 212
 - syntax
 - help 237
 - syntax diagrams, how to read 236
 - syntax error, in MQSC commands 42
 - system
 - restrictions for UNIX x
 - system default objects 14
 - system dump, FFST 230
 - system objects 301
 - system restrictions x
 - system setup 137
- T**
- temporary queues 9
 - terminology used in this book 325
 - time-independent applications 7
 - timed out responses from MQSC commands 69
 - trace
 - data sample (AIX) 221
 - data sample (HP-UX) 222
 - data sample (Sun Solaris) 224
 - data sample (Windows NT) 225
 - performance considerations 219, 222, 224
 - using 224
 - when using trace 224
 - tracing
 - OS/2 trace 224
 - Windows NT trace 224
 - transactions
 - display MQSeries command 256
 - resolve MQSeries command 269
 - transmission queue 73
 - creating 73
 - default 29, 73
 - defining 66
 - description 12
 - remote administration 65
 - trial period 291
 - trigger
 - event queues 130
 - events
 - compared with instrumentation events 130
 - messages on initiation queue 12
 - monitor
 - description 12
 - start command 281
 - triggering
 - application queue
 - defining 59
 - managing objects for 59
 - trusted application 61
 - types of event 129
 - types of objects 9
- U**
- unauthorized access, protecting from 79
 - undelivered message queue
 - See dead-letter queue
 - UNIX
 - system restrictions x
 - updating coded character sets 75
 - user exits
 - using CICS 169
 - user group
 - default for authority 81
 - default, nobody 81
 - for authorization 81
 - mqm 77, 78
 - user ID
 - authority 77, 78

Index

user ID (*continued*)
 authorization 85
 belonging to group nobody 81
 for authorization 85
 logged-in user 85
user identifier service
 description 319
user-defined message formats 76
USERID keyword, rules table 120
users
 groups
 principals 80

V

verifying MQSC commands 47

W

WAIT keyword, rules table 117
Windows clients error messages 231
Windows Help xvi

X

XA switch load files
 MQSeries 165

Sending your comments to IBM

MQSeries

System Administration

SC33-1873-00

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 request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should 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, use the Readers' Comment Form
- By fax:
 - From outside the U.K., after your international access code use 44 1962 870229
 - From within the U.K., use 01962 870229
- Electronically, use the appropriate network ID:
 - IBM Mail Exchange: GBIBM2Q9 at IBMMAIL
 - IBMLink: WINVMD(IDRCF)
 - Internet: idrcf@winvmd.vnet.ibm.com

Whichever you use, ensure that you include:

- The publication number and title
- The page number or topic to which your comment applies
- Your name and address/telephone number/fax number/network ID.

Readers' Comments

MQSeries

System Administration

SC33-1873-00

Use this form to tell us what you think about this manual. If you have found errors in it, or if you want to express your opinion about it (such as organization, subject matter, appearance) or make suggestions for improvement, this is the form to use.

To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer. This form is provided for comments about the information in this manual and the way it is presented.

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.

Be sure to print your name and address below if you would like a reply.

Name

Address

Company or Organization

Telephone

Email



You can send your comments POST FREE on this form from any one of these countries:

Australia	Finland	Iceland	Netherlands	Singapore	United States
Belgium	France	Israel	New Zealand	Spain	of America
Bermuda	Germany	Italy	Norway	Sweden	
Cyprus	Greece	Luxembourg	Portugal	Switzerland	
Denmark	Hong Kong	Monaco	Republic of Ireland	United Arab Emirates	


If your country is not listed here, your local IBM representative will be pleased to forward your comments to us. Or you can pay the postage and send the form direct to IBM (this includes mailing in the U.K.).

1 Cut along this line

2 Fold along this line

By air mail
Par avion

IBRS/CCR NUMBER: PHQ - D/1348/SO



REPONSE PAYEE
GRANDE-BRETAGNE

IBM United Kingdom Laboratories
Information Development Department (MP095)
Hursley Park,
WINCHESTER, Hants
SO21 2ZZ United Kingdom

3 Fold along this line

From: Name _____
Company or Organization _____
Address _____

EMAIL _____
Telephone _____

1 Cut along this line

4 Fasten here with adhesive tape



Printed in U.S.A.

SC33-1873-00

