MQSeries

IBM

Application Programming Guide

MQSeries

Application Programming Guide

IBM

```
┌─ Note! ──────────────────────────────────────────────────────────────────────┐
│                                                                                │
│  Before using this information and the product it supports, be sure to read the general information under Appendix J, "Notices" on │
│  page 545.                                                                     │
│                                                                                │
└────────────────────────────────────────────────────────────────────────────────┘
```

Ninth edition (February 1998)

This edition applies to the following products:

- MQSeries for AIX V5.0
- MQSeries for AS/400 V4R2
- 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 SINIX and DC/OSx V2.2
- MQSeries for SunOS V2.2
- MQSeries for Sun Solaris V5.0
- MQSeries for Tandem NonStop Kernel V2.2
- MQSeries for Windows V2.0
- MQSeries for Windows V2.1
- MQSeries for Windows NT V5.0

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.

Contents

Contents

Contents

Contents

Contents

Figures

1. Message queuing compared with traditional communication 7
2. Representation of a message 21
3. Group of logical messages 32
4. Segmented messages 33
5. How distribution lists work 112
6. Opening a distribution list in C 115
7. Opening a distribution list in COBOL 115
8. Logical order on a queue 126
9. Physical order on a queue 127
10. Skipping backout using MQGMO_MARK_SKIP_BACKOUT 147
11. Sample JCL used to invoke the CSQUCVX utility 167
12. Flow of application and trigger messages 199
13. Relationship of queues within triggering 201
14. Lotus Notes document and MQSeries message flow 240
15. Lotus Notes user input sample, form named Sample 241
16. Macro behind the SendData button 242
17. Example of a link database entry 244
18. Fragments of JCL to link-edit the object module in the batch environment 268
19. Fragments of JCL to link-edit the object module in the CICS environment 269
20. Fragments of JCL to link-edit the object module in the IMS environment 270
21. Dynamic linking using COBOL in the batch environment 271
22. Dynamic linking using COBOL in the CICS environment 272
23. Dynamic linking using COBOL in the IMS environment 272
24. Dynamic linking using assembler language in the batch environment . 273
25. Dynamic linking using assembler language in the CICS environment . 273
26. Dynamic linking using assembler language in the IMS environment . . 273
27. Dynamic linking using C language in the batch environment 273
28. Dynamic linking using C language in the CICS environment 274
29. Dynamic linking using C language in the IMS environment 274
30. Dynamic linking using PL/I in the batch environment 274
31. Dynamic linking using PL/I in the IMS environment 275
32. Running the reference message samples 334
33. Request and Inquire samples using triggering 343
34. Sample Client/Server (Echo) program flowchart 346
35. The database coordination samples 355
36. Example of ubbstxcx.cfg file for UNIX systems 367
37. Example of ubbstxcn.cfg file for Windows NT 369
38. Sample TUXEDO makefile for MQSeries for Windows NT 370
39. How TUXEDO samples work together 371
40. Example of a report from the Print Message sample application 397
41. Programs and panels for the TSO versions of the Mail Manager 405
42. Programs and panels for the CICS/ESA version of the Mail Manager . 406
43. Example of a panel showing a list of waiting messages 409
44. Example of a panel showing the contents of a message 409
45. Immediate Inquiry panel for the Credit Check sample application 411
46. Programs and queues for the Credit Check sample application (COBOL programs only) 413
47. Initial screen for Message Handler sample 425
48. Message list screen for Message Handler sample 425
49. Chosen message is displayed 426

Figures

Tables

About this book

The IBM MQSeries Level 2 products comprise:

- MQSeries for AIX
- MQSeries for AS/400
  (formerly known as MQSeries for OS/400)
- MQSeries for AT&T GIS UNIX** [1]
- MQSeries for Digital OpenVMS**
- MQSeries for HP-UX**
- MQSeries for MVS/ESA
- MQSeries for OS/2 Warp
- MQSeries for SINIX and DC/OSx**
- MQSeries for SunOS**
- MQSeries for Sun Solaris**
- MQSeries for Tandem NonStop Kernel
- MQSeries Three Tier
- MQSeries for Windows
- MQSeries for Windows NT

They are referred to in this book collectively as MQSeries. They provide application programming services that allow a new style of programming. This style enables you to write applications in which the constituent programs communicate with each other using *message queues*.

This book introduces the concepts of *messages* and *queues*, and shows you how to design and write applications that use the services that MQSeries provides. For information on the use of C++, see the *MQSeries Using C++* book.

For a full description of the MQSeries programming interface, see the MQSeries Application Programming Reference manual for your platform. The manuals are:

- *MQSeries Application Programming Reference*
- *MQSeries for AS/400 Application Programming Reference (RPG)*

For information on the MQSeries Level 1 products, see the books listed in "MQSeries Level 1 product publications" on page xxii.

In this book, references to "MQSeries on UNIX systems" include:

IBM MQSeries for AIX Version 5
IBM MQSeries for AT&T GIS UNIX Version 2.2
IBM MQSeries for HP-UX Version 5
IBM MQSeries for SINIX and DC/OSx Version 2.2
IBM MQSeries for SunOS Version 2.2
IBM MQSeries for Sun Solaris Version 5

The term "MQSeries Version 5 products" applies to the following MQSeries products:

IBM MQSeries for AIX Version 5
IBM MQSeries for HP-UX Version 5

---

[1] This platform has become NCR UNIX SVR4 MP-RAS, R3.0.

IBM MQSeries for OS/2 Warp Version 5
IBM MQSeries for Sun Solaris Version 5
IBM MQSeries for Windows NT Version 5

We ship sample programs with IBM MQSeries which are explained in Part 4, "Sample MQSeries programs" on page 309. You may find it useful to refer to these.

Who this book is for

This book is for the designers of applications that use message queuing techniques, and for the programmers who have to implement those designs.

What you need to know to understand this book

To write message queuing applications using MQSeries, you need to know how to write programs in at least one of the programming languages that MQSeries supports. These are:

- For MQSeries for MVS/ESA:

  – COBOL
  – Assembler language
  – C
  – PL/I

- For MQSeries for AS/400:

  – RPG
  – COBOL
  – C
  – C++

- For MQSeries for AT&T GIS UNIX, MQSeries for Digital OpenVMS, MQSeries for SINIX and DC/OSx, and MQSeries for SunOS:

  – COBOL
  – C

- For MQSeries for HP-UX and MQSeries for Sun Solaris:

  – COBOL
  – C
  – C++

- For MQSeries for AIX, MQSeries for OS/2 Warp, and MQSeries for Windows NT:

  – COBOL
  – C
  – C++
  – PL/I

- For MQSeries for Tandem NonStop Kernel:

  – COBOL
  – C
  – TAL

  • For MQSeries for Windows:

    – C
    – Visual Basic

If the applications you are writing will run within a CICS or IMS system, you must also be familiar with CICS or IMS, and their application programming interfaces.

To understand this book, you do not need to have written message queuing programs before.

How to use this book

This book contains guidance information to help you design an application, and procedural information to help you to write an application.

The book is divided into five parts:

Part 1, "Designing applications that use MQSeries" on page 1
Introduces the message queuing style of application design, describes MQSeries messages and queues, and shows how to design a message queuing application.

Part 2, "Writing an MQSeries application" on page 61
Describes how to use the IBM Message Queue Interface (MQI) to write the programs that comprise a message queuing application. The chapters guide you through the coding of each MQI call, showing you what information to supply as input and what returns to expect. These chapters first describe simple uses of the MQI calls, then go on to describe how to use all the features of each call.

Read Part 1, "Designing applications that use MQSeries" on page 1 to understand the concepts involved when designing MQSeries applications. The second part is self-contained: use an individual chapter when you are performing the task described in it.

Part 3, "Building an MQSeries application" on page 251
Explains how to build your MQSeries application on each platform.

Part 4, "Sample MQSeries programs" on page 309
Lists and explains how the sample programs work, for all platforms.

Part 5, "Appendixes" on page 429
Contains examples of how to use the MQI calls in each of the programming languages supported by MQSeries.

Appearance of text in this book

This book uses the following type style:

CompCode　　　　　Example of the name of a parameter of a call, or the attribute of an object

Terms used in this book

All new terms that the MQSeries library introduces are defined in the "Glossary of terms and abbreviations" on page 551. These terms are shown like *this* on their first occurrence in the text of this book.

In the body of this book, the following shortened names are used for these products:

CICS
> The CICS for MVS/ESA, CICS/MVS, CICS for AS/400, CICS for OS/2, CICS for AIX, CICS for HP 9000, CICS for Sun Solaris, and CICS for Windows NT products.

> CICS products were renamed to Transaction Server at certain versions: Transaction Server for OS/2, V4, Transaction Server for Windows NT, V4, and Transaction Server for AIX, V4.

IMS
> The IMS/ESA product.

MQSeries
> MQSeries for AIX, MQSeries for AT&T GIS UNIX, MQSeries for Digital OpenVMS, MQSeries for HP-UX, MQSeries for MVS/ESA, MQSeries for OS/2 Warp, MQSeries for AS/400, MQSeries for SINIX and DC/OSx, MQSeries for SunOS, MQSeries for Sun Solaris, and MQSeries for Windows NT.

MQSeries on UNIX systems
> MQSeries for AIX, MQSeries for AT&T GIS UNIX, MQSeries for HP-UX, MQSeries for SINIX and DC/OSx, MQSeries for SunOS, and MQSeries for Sun Solaris.

MVS
> The MVS/ESA System Product.

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 AS/400 V4R2
- 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 SINIX and DC/OSx V2.2
- MQSeries for SunOS V2.2
- MQSeries for Sun Solaris V5.0
- MQSeries for Tandem NonStop Kernel V2.2
- 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 xxii. 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 C++ is also supported by MQSeries for AS/400 V4R2.

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 AS/400

MQSeries for AS/400 Version 4 Release 2 Licensed Program Specifications,
GC33-1958

MQSeries for AS/400 Version 4 Release 2 Administration Guide, GC33-1956

*MQSeries for AS/400 Version 4 Release 2 Application Programming Reference
(RPG)*, SC33-1957

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 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 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 for Tandem NonStop Kernel

MQSeries for Tandem NonStop Kernel Version 2.2 System Management Guide, GC33-1893

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

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 product family Web site:

  `http://www.software.ibm.com/ts/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 web site ──────────────────────────────────────┐

  The MQSeries product family Web site is at:


     http://www.software.ibm.com/ts/mqseries/


  By following links from this Web site you can:
  • Obtain latest information about the MQSeries product family.
  • Access the MQSeries books in HTML format.
  • Download MQSeries SupportPacs.

└──────────────────────────────────────────────────────────┘
```

Related publications

This section describes the documentation available for some related products and issues mentioned in this book.

CICS

For information about those aspects of CICS for MVS/ESA that this book refers to, see the following books:

CICS for MVS/ESA V4.1 Application Programming Reference, SC33-1170
CICS for MVS/ESA V4.1 Customization Guide, SC33-1165
CICS for MVS/ESA V4.1 CICS-Supplied Transactions, SC33-1168
CICS for MVS/ESA V4.1 System Definition Guide, SC33-1164

For information about those aspects of CICS/MVS that this book refers to, see the following book:

CICS/MVS Application Programmer's Reference, SC33-0512

For information about those aspects of CICS for MVS/ESA Version 4.1 that this book refers to, see the following books:

CICS/ESA V3.3 Application Programming Reference, SC33-0676
CICS/ESA V3.3 Customization Guide, SC33-0665
CICS/ESA V3.3 CICS-Supplied Transactions, SC33-0669
CICS/ESA V3.3 System Definition Guide, SC33-0664

For information about CICS programs that this book refers to, see the following books:

CICS on Open Systems Application Programming Guide, SC33-1568-00
CICS for OS/2 V2.0.1 Application Programming, SC33-0883
Transaction Server for OS/2 Warp, V4 Application Programming, SC33-1585
CICS for OS/400 V3 Application Programming, SC33-1386
CICS for Windows NT V2.0 Application Programming, SC33-1425
Transaction Server for Windows NT V4 Application Programming Guide, SC33-1888

IMS

For information about those aspects of IMS that this book refers to, see the following books:

IMS/ESA Version 4 Application Programming: DL/I Calls, SC26-3062
IMS/ESA Version 4 Application Programming: Design Guide, SC26-3066
IMS/ESA Version 5 Application Programming: Database Manager, SC26-8015
IMS/ESA Version 5 Application Programming: Design Guide, SC26-8016
IMS/ESA Version 5 Application Programming: Transaction Manager, SC26-8017

MVS/ESA

For information about those aspects of MVS that this book refers to, see the following book:

MVS/ESA Application Development Guide: Assembler Language Programs, GC28-1644

Design

For information on how to design panel-driven application interfaces, see the following book:

Systems Application Architecture, Common User Access: Basic Interface Design Guide, SC26-4583

C

For information about C programming, see the following books:

Guide to Tools for Programming in C, U6296-J-Z145-2-7600
SNI Programmer, U6401-J-Z145-3-7600

C++

For information about C++ programming, see the following books:

C Set ++ for AIX: User's Guide, SC09-1968
VisualAge C++ for OS/2 User's Guide, S25H-6961
VisualAge C++ for OS/2 Programming Guide, S25H-6958
VisualAge for C++ for Windows User's Guide, S33H-5031
VisualAge for C++ for Windows Programming Guide, S33H 5032
VisualAge for C++ for AS/400 : C++ User's Guide, SC09-2416

COBOL

For information about COBOL programming that this book refers to, see the following books:

COBOL V3.2 SINIX pocket guide, U21709-J-Z145-2-7600
IBM COBOL Set for AIX Programming Guide, SC26-8423
IBM COBOL for MVS and VM, IBM VisualAge for COBOL for OS/2, IBM COBOL Set for AIX Language Reference, SC26-4769

Related publications

Summary of Changes

Changes for this edition are marked with the | character in the left-hand margin.

| Changes for this edition

| Changes for edition number SC33-0807-08 include:

- | • A new release of:
  - | – MQSeries for AS/400 V4R2

- | • Inclusion of a new MQSeries product:
  - | – MQSeries for Tandem NonStop Kernel V2.2

Changes for the eighth edition

Changes for edition number SC33-0807-07 include:

- • New releases of the following 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
  - – MQSeries for OS/400 V3R7
  - – MQSeries for MVS/ESA V1.2

- • Inclusion of new MQSeries products:
  - – MQSeries for Digital OpenVMS V2.2
  - – MQSeries for Windows V2.0
  - – MQSeries for Windows V2.1

- • Support of programming languages:
  - – C++
  - – PL/I on MQSeries for AIX, OS/2 Warp, and Windows NT

- • New function:
  - – Distribution lists
  - – Handling large messages
  - – External resource managers
  - – Report options for application acknowledgement
  - – Indexed queues (MQSeries for MVS/ESA only)

- • Enhancement of existing function:
  - – Change in triggering rules
  - – Lotus Notes link with AIX
  - – Support for Encina by MQSeries on Windows NT
  - – Support for signaling on MQSeries for Windows V2.1

Changes for the seventh edition

Changes for edition number SC33-0807-06 include:

- Inclusion of new platforms, and new releases of existing platforms:

  - MQSeries for HP-UX Version 2.2.1
  - MQSeries for MVS/ESA Version 1.1.4
  - MQSeries for OS/400 Version 3.2
  - MQSeries for SINIX and DC/OSx Version 2.2
  - MQSeries for SunOS Version 2.2
  - MQSeries for Sun Solaris Version 2.2
  - MQSeries for Windows NT Version 2.3

- Data-conversion exits (previously described in the *MQSeries Distributed Queuing Guide*), and support for data conversion by MQSeries for MVS/ESA.

- New sample programs in Chapter 31, "Sample programs for MQSeries for MVS/ESA" on page 375.

- Information about writing MQSeries-IMS bridge applications.

Part 1. Designing applications that use MQSeries

Chapter 1. Introduction to message queuing

The MQSeries products enable programs to communicate with one another across a network of unlike components – processors, operating systems, subsystems and communication protocols – using a consistent application programming interface.

Applications designed and written using this interface are known as *message queuing* applications, as they use the *Messaging and Queuing* style:

Messaging   Because programs communicate by sending each other data in messages rather than calling each other directly.

Queuing   Because the messages are placed on queues in storage, allowing programs to run independently of each other, at different speeds and times, in different locations, and without having a logical connection between them.

What is a message?

In message queuing, a *message* is simply a collection of data sent by one program and intended for another program.

MQSeries defines four types of message:

Datagram   A simple message for which no reply is expected
Request   A message for which a reply is expected
Reply   A reply to a request message
Report   A message that describes an event such as the occurrence of an error

See "Types of message" on page 23 for more information about these messages.

Message descriptor

An MQSeries message consists of control information and application data. The control information is defined in a *message descriptor* structure (MQMD) and contains such things as:

- The type of the message
- An identifier for the message
- The priority for delivery of the message

The structure and content of the application data is determined by the participating programs, not by MQSeries.

Message channel agent

A message channel agent moves messages from one queue manager to another. References are made to them in this book when dealing with report messages and you will need to consider them when designing your application. See the *MQSeries Intercommunication* book for more information.

What is a message queue?

A *message queue* is a named destination to which messages can be sent. Messages accumulate on queues until they are retrieved by programs that service those queues.

Queues reside in, and are managed by, a queue manager (see "What is a queue manager?"). The physical nature of a queue depends on the operating system on which the queue manager is running. A queue can either be a volatile buffer area in the memory of a computer, or a data set on a permanent storage device (such as a disk). The physical management of queues is the responsibility of the queue manager and is not made apparent to the participating application programs.

Programs access queues only through the external services of the queue manager. They can open a queue, put messages on it, get messages from it, and close the queue. They can also set, and inquire about, the attributes of queues.

What is a queue manager?

A *queue manager* is a system program that provides queuing services to applications. It provides an application programming interface so that programs can put messages on, and get messages from, queues. A queue manager provides additional functions so that administrators can create new queues, alter the properties of existing queues, and control the operation of the queue manager.

For MQSeries message queuing services to be available on a system, there must be a queue manager running:

- On MVS, OS/2, Windows NT, Digital OpenVMS, and UNIX systems, you can have more than one queue manager running on a single system (for example, to separate a test system from a "live" system). To an application, each queue manager is identified by a *connection handle* (*Hconn*).

- On the AS/400 and Windows platforms, you can have only one queue manager running on a single system. *Hconn* is still used, but only to give compatibility with other MQSeries platforms.

Many different applications can make use of the queue manager's services at the same time and these applications can be entirely unrelated. For a program to use the services of a queue manager, it must establish a connection to that queue manager.

For applications to be able to send messages to applications that are connected to other queue managers, the queue managers must be able to communicate among themselves. MQSeries implements a *store-and-forward* protocol to ensure the safe delivery of messages between such applications.

What is an MQSeries client?

An *MQSeries client* is an independently installable component of an MQSeries product. It allows you to run MQSeries applications, by means of a communications protocol, to interact with one or more MQI servers on other platforms and to connect to their queue managers.

For full details on how to install the MQSeries client component and use the environment, see the *MQSeries Clients* book.

What is message queuing?

Message queuing has been used in data processing for many years. It is most commonly used today in electronic mail. Without queuing, sending an electronic message over long distances requires every node on the route to be available for forwarding messages, and the addressees to be logged on and conscious of the fact that you are trying to send them a message. In a queuing system, messages are stored at intermediate nodes until the system is ready to forward them. At their final destination they are stored in an electronic mailbox until the addressee is ready to read them.

Even so, many complex business transactions are processed today without queuing. In a large network, the system might be maintaining many thousands of connections in a ready-to-use state. If one part of the system suffers a problem, many parts of the system become unusable.

You can think of message queuing as being electronic mail for programs. In a message queuing environment, each program from the set that makes up an application suite is designed to perform a well-defined, self-contained function in response to a specific request. To communicate with another program, a program must put a message on a predefined queue. The other program retrieves the message from the queue, and processes the requests and information contained in the message. So message queuing is a style of program-to-program communication.

Queuing is the mechanism by which messages are held until an application is ready to process them. Queuing allows you to:

- Communicate between programs (which may each be running in different environments) without having to write the communication code.

- Select the order in which a program processes messages.

- Balance loads on a system by arranging for more than one program to service a queue when the number of messages exceeds a threshold.

- Increase the availability of your applications by arranging for an alternative system to service the queues if your primary system is unavailable.

Main features of message queuing

The main features of applications that use message queuing techniques are:

- There are no direct connections between programs.
- Communication between programs can be time independent.
- Work can be carried out by small, self-contained programs.
- Communication can be driven by events.
- Applications can assign a priority to a message.
- Security.
- Syncpoint support.
- Recovery support.

No direct connections between programs

Message queuing is a technique for indirect program-to-program communication. It can be used within any application where programs communicate with each other. Communication occurs by one program putting messages on a queue (owned by a queue manager) and another program getting the messages from the queue.

Programs can get messages that were put on a queue by other programs. The other programs can be connected to the same queue manager as the receiving program, or to another queue manager. This other queue manager might be on another system, a different computer system, or even within a different business or enterprise.

There are no physical connections between programs that communicate using message queues. A program sends messages to a queue owned by a queue manager, and another program retrieves messages from the queue (see).
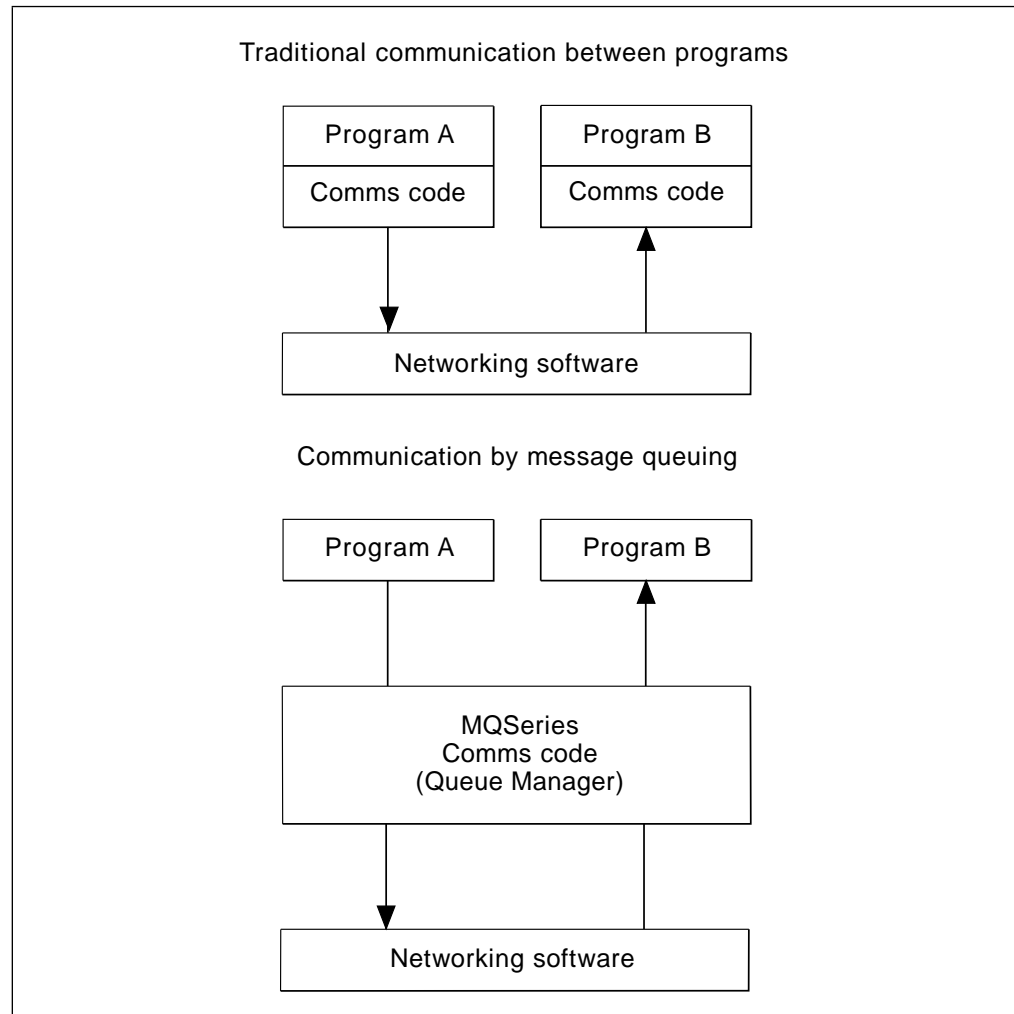
Figure 1. Message queuing compared with traditional communication

As with electronic mail, the individual messages that may be part of a transaction, travel through a network on a store-and-forward basis. If a link between nodes fails, the message is kept until the link is restored, or the operator or program redirects the message.

The mechanism by which a message moves from queue to queue is hidden from the programs. Therefore the programs are simpler.

Time-independent communication
Programs requesting others to do work do not have to wait for the reply to a request. They can do other work, and process the reply either when it arrives or at a later time. When writing a messaging application, you need not know (or be concerned) when a program sends a message, or when the target is able to receive the message. The message is not lost; it is retained by the queue manager until the target is ready to process it. The message stays on the queue until it is removed by a program.

Small programs

Message queuing allows you to exploit the advantages of using small, self-contained programs. Instead of a single, large program performing all the parts of a job sequentially, you can spread the job over several smaller, independent programs. The requesting program sends messages to each of the separate programs, asking them to perform their function; when each program is complete, the results are sent back as one or more messages.

Event-driven processing

Programs can be controlled according to the state of queues. For example, you can arrange for a program to start as soon as a message arrives on a queue, or you can specify that the program does not start until there are, for example, 10 messages above a certain priority on the queue, or 10 messages of any priority on the queue.

Message priority

A program can assign a priority to a message when it puts the message on a queue. This determines the position in the queue at which the new message is added.

Programs can get messages from a queue either in the order in which the messages appear in the queue, or by getting a specific message. (A program may want to get a specific message if it is looking for the reply to a request it sent earlier.)

Security

Authorization checks are carried out on each resource, using the tables that are set up and maintained by the MQSeries administrator.

- RACF or other external security managers may be used within MQSeries for MVS/ESA.

- Native OS/400 is used within MQSeries for AS/400.

- There is no authorization checking within MQSeries for OS/2 Warp; however, an interface is provided to enable you to build and install your own.

- Within MQSeries on UNIX systems, Digital OpenVMS, and Windows NT, a security manager, the Object Authority Manager (OAM), is provided as an installable service. By default, the OAM is active.

Syncpoint support

The synchronization of the start and end of units of work is fully supported as an option on each MQGET/MQPUT, allowing the results of the call to be committed or rolled back. Syncpoint support operates either internally or externally to MQSeries depending on the form of syncpoint coordination selected for the application.

Recovery support

In order for recovery to be possible, all persistent MQSeries updates are logged. Hence, in the event that recovery is necessary, all persistent messages will be restored, all in-flight transactions will be rolled back and any syncpoint commit and backouts will be handled in the normal way of the syncpoint manager in control. For more information on persistent messages, see "Message persistence" on page 33.

MQSeries clients and servers

A server application will not have to be changed to be able to support additional MQSeries clients on new platforms.

Similarly, the MQSeries client will, without change, be able to function with additional types of server.

Lotus Notes linking with MQSeries

MQSeries Version 5 products provide a Lotus Notes add-in task that gives a Lotus Notes application access to an MQSeries application.

Lotus Notes uses the MQSeries add-in task to send a document (data) to the connected MQSeries system. MQSeries processes the data it receives, returning data to Lotus Notes when requested in the original Lotus Notes document.

Note: Any MQSeries product can process the data (message).

Installation of the components involved is explained in the:

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

See the guide for your platform.

Configuration and management are explained in the *MQSeries System Administration* book.

For an explanation of how MQSeries and Lotus Notes work together, covering aspects to consider when designing and writing applications, see Chapter 16, "Lotus Notes" on page 239.

Benefits of message queuing to the application designer and developer

Some of the benefits of message queuing are:

- You can design applications using small programs that you can share between many applications.

- You can quickly build new applications by reusing these building blocks.

- Applications written to use message queuing techniques are not affected by changes in the way queue managers work.

- You do not need to use any communication protocols. The queue manager deals with all aspects of communication for you.

- Programs that receive messages need not be running at the time messages are sent to them. The messages are retained on queues.

Designers can reduce the cost of their applications because development is faster, fewer developers are needed, and demands on programming skill are lower than those for applications that do not use message queuing.

What can you do with MQSeries products?

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

MQSeries for MVS/ESA

With MQSeries for MVS/ESA you can write applications that:

- Use message queuing within CICS or IMS.

- Send messages between batch, CICS, and IMS applications, selecting the most appropriate environment for each function.

- Send messages to applications that run on other MQSeries platforms.

- Process several messages together as a single unit of work that can be committed or backed out.

- Send messages to and interact with IMS applications by means of the IMS bridge.

You can write programs in COBOL, assembler, PL/I, and C.

Each environment within MVS has its own characteristics, advantages, and disadvantages. The advantage of MQSeries for MVS/ESA is that applications are not tied to any one environment, but can be distributed to take advantage of the benefits of each environment. For example, you can develop end-user interfaces using TSO or CICS, you can run processing-intensive modules in MVS batch, and you can run database applications in IMS or CICS. In all cases, the various parts of the application can communicate using messages and queues.

Designers of MQSeries applications must be aware of the differences and limitations imposed by these environments. For example:

- MQSeries provides facilities that allow intercommunication between queue managers (this is known as *distributed queuing*).

- Methods of committing and backing out changes differ between the batch and CICS environments.

- MQSeries for MVS/ESA provides support in the IMS environment for online message processing programs (MPPs), interactive fast path programs (IFPs), and batch message processing programs (BMPs). If you are writing batch DL/I programs, follow the guidance given in this book for MVS batch programs.

- Although multiple instances of MQSeries for MVS/ESA can exist on a single MVS system, a CICS region can connect to only one queue manager at a time. However, more than one CICS region can be connected to the same queue manager. In the IMS and MVS batch environments, programs can connect to more than one queue manager.

The differences between the supported environments, and their limitations, are discussed further in Chapter 15, "Applications on MQSeries for MVS/ESA" on page 219.

MQSeries for non-MVS/ESA platforms

With MQSeries for non-MVS/ESA platforms you can write applications that:

- Send messages to other applications running under the same operating systems. The applications can be on either the same or another system.

- Send messages to applications that run on other MQSeries platforms.

- Use message queuing from within CICS for OS/2, CICS for AS/400, CICS for AIX, CICS for HP 9000, CICS for Siemens Nixdorf SINIX, CICS for Solaris, and CICS for Windows NT, DOS, and Windows 3.1 applications.

- Use message queuing from within Encina for AIX, HP-UX, SINIX, Sun Solaris, and Windows NT.

- Use message queuing from within Tuxedo for AIX, AT&T, HP-UX, SINIX and DC/OSx, Sun Solaris, Sun OS, and Windows NT.

- Process several messages together as a single unit of work that can be committed or backed out.

- Run from a full MQSeries environment, or run from an MQSeries client environment on the following platforms:

  - Digital OpenVMS
  - DOS
  - OS/2
  - UNIX systems
  - VM/ESA
  - Windows NT
  - Windows 3.1
  - Windows 95

You can write programs in COBOL and C. In addition:

- For MQSeries for AS/400, you can write applications in RPG.

- For MQSeries for Windows, you can write applications in Visual Basic.

- For MQSeries for AS/400, OS/2 Warp, Windows NT, AIX, HP-UX, and Sun Solaris, you can write applications in C++.

- For MQSeries for AIX, MQSeries for OS/2 Warp, and MQSeries for Windows NT, you can write applications in PL/I.

- For MQSeries for OS/2 Warp, Windows NT, AIX, HP-UX, and Sun Solaris, you can write applications in Java. Documentation for the MQSeries Java Bindings and MQSeries Java Client is supplied in HTML format with the MQSeries Version 5 products products. The MQSeries product READ.ME file tells you where to point your browser.

- For MQSeries for Tandem NSK you can write applications in TAL.

Uses of MQSeries

Chapter 2. Application design overview

This chapter discusses how to design MQSeries applications. It introduces:

- The decisions you should take during the design process
- The features that MQSeries provides for you to work with
- The programming techniques that you should consider
- The Message Queue Interface (MQI)

These subjects are discussed in greater detail in the remaining chapters of this book.

Planning

When you have decided how your applications are able to take advantage of the platforms and environments available to you, you need to decide how to use the features offered by MQSeries. Some of the key aspects are:

What types of queue should you use?

Do you want to create a queue each time you need one, or do you want to use queues that have already been set up? Do you want to delete a queue when you have finished using it, or is it going to be used again? Do you want to use alias queues for application independence? To see what types of queues are supported, refer to "Queues" on page 40.

What types of message should you use?

You may want to use datagrams for simple messages, but request messages (for which you expect replies) for other situations. You may want to assign different priorities to some of your messages.

How can you control your MQSeries programs?

You may want to start some programs automatically or make programs wait until a particular message arrives on a queue, (using the MQSeries *triggering* feature, see Chapter 14, "Starting MQSeries applications using triggers" on page 197). Alternatively, you may want to start up another instance of an application when the messages on a queue are not getting processed fast enough (using the MQSeries *instrumentation events* feature as described in the *MQSeries Programmable System Management* book).

Will your application run on an MQSeries client?

The full MQI is supported in the client environment and this enables almost any MQSeries application to be relinked to run on an MQSeries client. Link the application on the MQSeries client to the MQIC library, rather than to the MQI library. The exceptions are:

- An application that needs syncpoint coordination with other resource managers.

- Get(signal) on MVS/ESA is not supported.

Note: An application running on an MQSeries client may connect to more than one queue manager concurrently, or use a queue manager name with an asterisk (*) on an MQCONN or MQCONNX call. The application will have to be changed if you want to link to the queue

manager libraries instead of the client libraries, as this function will not be available.

See the *MQSeries Clients* book for more information.

How can you secure your data and maintain its integrity?

You can use the context information that is passed with a message to test that the message has been sent from an acceptable source. You can use the syncpointing facilities provided by MQSeries or your operating system to ensure that your data remains consistent with other resources (see Chapter 13, "Committing and backing out units of work" on page 185 for further details). You can use the *persistence* feature of MQSeries messages to assure the delivery of important messages.

How should you handle exceptions and errors?

You need to consider how to process messages that cannot be delivered, and how to resolve error situations that are reported to you by the queue manager. For some reports, you must set report options on MQPUT.

The remainder of this chapter introduces the features and techniques that MQSeries provides to help you answer questions like these.

Using MQSeries objects

The MQI uses the following types of object:

- Queue managers
- Queues
- Namelists (MQSeries for MVS/ESA only)
- Process definitions
- Channels
- Storage classes (MVS/ESA only)

These objects are discussed in Chapter 4, "MQSeries objects" on page 39.

Each object is identified by an *object descriptor* (MQOD), which you use when you write MQSeries programs. However, with the exception of dynamic queues, these objects must be defined to the queue manager before you can work with them.

You define objects using the PCF commands, the MQSC facility, or the MQSeries for MVS/ESA operations and control panels. The commands are described in the *Programmable System Management* guide, the *MQSeries Command Reference* manual, and the *MQSeries for MVS/ESA System Management Guide*. You can also display or alter the attributes of objects, or delete the objects.

Alternatively, for sequences of MQSeries for MVS/ESA commands you use regularly, you can write administration programs that create messages containing commands and that put these messages on the system-command input queue. The queue manager processes the messages on this queue in the same way that it processes commands entered from the command line or from the operations and control panels. This technique is described in the *MQSeries for MVS/ESA System Management Guide*, and demonstrated in the Mail Manager sample application delivered with MQSeries for MVS/ESA. For a description of this sample, see Chapter 31, "Sample programs for MQSeries for MVS/ESA" on page 375.

For sequences of MQSeries for AS/400 commands you use regularly, you can write CL programs that contain the commands, as demonstrated in the CL sample program (AMQSDEF4) delivered with MQSeries for AS/400. For a description of this sample, see the *MQSeries for AS/400 Administration Guide*.

For sequences of MQSeries commands on OS/2, Windows NT, and UNIX systems, you can use the MQSC facility to run a series of commands held in a file. For information on how to use this facility, see the *MQSeries Command Reference* manual.

Designing your messages

You create a message when you use an MQI call to put the message on a queue. As input to the call, you supply some control information in a *message descriptor* (MQMD) and the data that you want to send to another program. But at the design stage, you need to consider the following questions, because they affect the way you create your messages:

What type of message should I use?
Are you designing a simple application in which you can send a message, then take no further action? Or are you asking for a reply to a question? If you are asking a question, you may include in the message descriptor the name of the queue on which you want to receive the reply.

Do you want your request and reply messages to be synchronous? This implies that you set a timeout period for the reply to answer your request, and if you do not receive the reply within that period, it is treated as an error.

Or would you prefer to work asynchronously, so that your processes do not have to depend upon the occurrence of specific events, such as common timing signals?

Another consideration is whether you have all your messages inside a unit of work.

Should I assign different priorities to some of the messages I create?
You can assign a priority value to each message, and define the queue so that it maintains its messages in order of their priority. If you do this, when another program retrieves a message from the queue, it always gets the message with the highest priority. If the queue does not maintain its messages in priority order, a program that retrieves messages from the queue will retrieve them in the order in which they were added to the queue.

Programs can also select a message using the identifier that the queue manager assigned when the message was put on the queue. Alternatively, you can generate your own identifiers for each of your messages.

Will my messages be discarded when the queue manager restarts?
The queue manager preserves all persistent messages, recovering them when necessary from the MQSeries log files, when it is restarted. Nonpersistent messages and temporary dynamic queues are not preserved. Any messages that you do not want discarded must be defined as persistent at the time they are created. When writing an application for MQSeries for OS/2 Warp, MQSeries for Windows NT, or MQSeries on UNIX systems, make sure that you know how your system has been set up in respect of log file allocation to reduce the risk of designing an application that will run to the log file limits.

Do I want to give information about myself to the recipient of my messages?
Normally, the queue manager sets the user ID, but suitably authorized applications can also set this field, so that you can include your own user ID and other information that the receiving program can use for accounting or security purposes.

MQSeries techniques

For a simple MQSeries application, you need to decide which MQSeries objects to use in your application, and which types of message you want to use. For a more advanced application, you may want to use some of the techniques introduced in the following sections.

Waiting for messages

A program that is serving a queue can await messages by:

- Making periodic calls on the queue to see whether a message has arrived (*polling*)

- Waiting until either a message arrives, or a specified time interval expires (see "Waiting for messages" on page 142)

- setting a signal so that the program is informed when a message arrives (MQSeries for MVS/ESA and MQSeries for Windows V2.1 only). For information about this, see "Signaling" on page 143.

Correlating replies

In MQSeries applications, when a program receives a message that asks it to do some work, the program usually sends one or more reply messages to the requester. To help the requester to associate these replies with its original request, an application can set a *correlation identifier* field in the descriptor of each message. Programs should copy the message identifier of the request message into the correlation identifier field of their reply messages.

Setting and using context information

Context information is used for associating messages with the user who generated them, and for identifying the application that generated the message. Such information is useful for security, accounting, auditing, and problem determination.

When you create a message, you can specify an option that requests that the queue manager associates default context information with your message.

For more information on using and setting context information, see "Message context" on page 35.

Starting MQSeries programs automatically

MQSeries *triggering* enables a program to be started automatically when messages arrive on a queue. You can set trigger conditions on a queue so that a program is started to process that queue:

- Every time a message arrives on the queue
- When the first message arrives on the queue
- When the number of messages on the queue reaches a predefined number

For more information on triggering, see Chapter 14, "Starting MQSeries applications using triggers" on page 197.

Note: This is just one way of doing this. For example, you can start a program automatically on a timer using non-MQSeries facilities.

Generating MQSeries reports

You can request the following reports within an application:

- Exception reports
- Expiry reports
- Confirm-on-arrival (COA) reports
- Confirm-on-delivery (COD) reports
- Positive action notification (PAN) reports
- Negative action notification (NAN) reports

These are described in "Report messages" on page 24.

Application programming

MQSeries supports IBM's Message Queue Interface (MQI). The MQI includes a set of calls with which you can send and receive messages, and manipulate MQSeries objects.

Call interface

The MQI calls allow you to:

- Connect programs to, and disconnect programs from, a queue manager

- Open and close objects (such as queues, queue managers, namelists, and processes)

- Put messages on queues

- Receive messages from a queue, or browse them (leaving them on the queue)

- Inquire about the attributes (or properties) of MQSeries objects, and set some of the attributes of queues

- Commit and back out changes made within a unit of work, in environments where there is no natural syncpoint support, for example, OS/2 and UNIX systems

- Coordinate queue manager updates and updates made by other resource managers

The MQI provides *structures* (groups of fields) with which you supply input to, and get output from, the calls. It also provides a large set of named constants to help you supply options in the parameters of the calls. The definitions of the calls, structures, and named constants are supplied in data definition files for each of the supported programming languages. Also, default values are set within the MQI calls.

Design for performance - hints and tips

Here are a few ideas to help you design efficient applications:

- Design your application so that processing goes on in parallel with a user's thinking time:

  - Display a panel and allow the user to start typing while the application is still initializing.

  - Don't be afraid to get the data you need in parallel from different servers.

- Keep connections and queues open if you are going to reuse them instead of repeatedly opening and closing, connecting and disconnecting.

  Note: However, a server application which is putting only one message should use MQPUT1.

- Keep your messages within a unit of work, so that they can be committed or backed out simultaneously.

- Use the nonpersistent option for messages that do not need to be recoverable.

Programming platforms

MQSeries for MVS/ESA

MQSeries for MVS/ESA operates under MVS/ESA Version 4.3 and subsequent compatible releases. You can run MQSeries for MVS/ESA programs in the CICS for MVS/ESA, IMS/ESA, and MVS environments. You can write these programs in COBOL, assembler, PL/I, and C.

UNIX systems

MQSeries for AIX operates under AIX/6000 Version 4.1.3 and subsequent compatible releases. You can run MQSeries for AIX programs from within CICS for AIX, Transaction Server for AIX, Encina for AIX, and Tuxedo for AIX. Applications using threads are supported by MQSeries for AIX.

MQSeries for AT&T GIS UNIX operates under AT&T GIS UNIX Version 3 [2] and subsequent compatible releases. You can run MQSeries for AT&T GIS UNIX programs from within Tuxedo for AT&T.

MQSeries for Digital OpenVMS operates under VMS Version 6.2 and subsequent compatible releases.

MQSeries for HP-UX operates under HP-UX Version 10.01 and subsequent compatible releases. You can run MQSeries for HP-UX programs from within CICS, Encina, and Tuxedo for HP-UX. Applications using threads are supported by MQSeries for HP-UX V5.0.

MQSeries for SINIX and DC/OSx operates under SINIX and DC/OSx Version 2.1 and subsequent compatible releases. You can run MQSeries for SINIX and DC/OSx programs from within CICS for Siemens Nixdorf SINIX, and Tuxedo for SINIX and DC/OSx. You can also run MQSeries for SINIX programs from withing Encina for SINIX.

MQSeries for Sun Solaris operates under Sun Solaris Version 2.5 and subsequent compatible releases. You can run MQSeries for Sun Solaris programs from within CICS, Transaction Server, Encina, and Tuxedo for Sun

---

[2] This platform has become NCR UNIX SVR4 MP-RAS, R3.0.

Solaris. Applications using threads are supported by MQSeries for Sun Solaris V5.0.

MQSeries for SunOS operates under SunOS Version 4.1.3 and subsequent compatible releases. You can run MQSeries for SunOS programs from within Tuxedo for Sun Solaris.

You can write MQSeries applications on UNIX systems in COBOL and C. Also, you can write MQSeries applications on AIX, HP-UX, and Sun Solaris in C++, and applications on AIX in PL/I.

MQSeries for AS/400

MQSeries for AS/400 operates under OS/400 Version 4 Release 2 and subsequent compatible releases. You can run MQSeries for AS/400 programs in the CICS for AS/400 environment. You can write MQSeries for AS/400 programs in RPG, COBOL, C, and C++

MQSeries for OS/2 Warp

MQSeries for OS/2 Warp operates under OS/2 Warp Version 3.0 and subsequent compatible releases. You can run MQSeries for OS/2 Warp programs in the CICS and Transaction Server environment. You can write MQSeries for OS/2 Warp programs in COBOL, C, C++, and PL/I.

MQSeries for Tandem NonStop Kernel

MQSeries for Tandem NonStop Kernel V2.2 operates under Tandem NSK operating system version D3x, D4x, or G02 with TMF and PATHWAY, and subsequent compatible releases. You can write MQSeries for Tandem NSK programs in COBOL, C and TAL.

MQSeries for Windows

MQSeries for Windows V2.0 operates under Windows Version 3.1 and Windows 95. MQSeries for Windows V2.1 operates under Windows 95 and Windows NT V4. You can write MQSeries for Windows programs in C and Visual Basic.

MQSeries for Windows NT

MQSeries for Windows NT operates under Intel and Windows NT Version 3.5 and subsequent compatible releases. You can run MQSeries for Windows NT programs from within CICS, Transaction Server, Encina, and Tuxedo for Windows NT. You can write MQSeries for Windows NT programs in COBOL, C, C++, and PL/I.

Applications for more than one platform

Will your application run on more than one platform? Do you have a strategy to move to a different platform from the one you use today? If the answer to either of these questions is "yes," you need to make sure that you code your programs for platform independence.

If you are using C, make sure that you code in ANSI standard C. Do not use platform specific function just because it is faster or more efficient if there is an equivalent standard C library function. The exception is when efficiency in the code is paramount, when you should code for both situations using #ifdef. For example:

```
#ifdef _OS2
     OS/2 specific code
#else
     generic code
#endif
```

When the time comes to move the code to another platform, you can now search the source for #ifdef with the platform specific identifiers, in this example _OS2, and add/change code as necessary.

It is worth considering keeping portable code in separate source files from the platform specific code, and using a simple naming convention to split the categories.

Testing MQSeries applications

The application development environment for MQSeries programs is no different from that for any other application, so you can use the same development tools as well as the MQSeries trace facilities. This is most noticeable on OS/2 and UNIX systems where there is a wide selection.

When testing CICS applications on MVS/ESA, you can use the CICS Execution Diagnostic Facility (CEDF). CEDF traps the entry and exit of every MQI call as well as calls to all CICS services. Also, in the CICS for MVS/ESA environment, you can write an API-crossing exit program to provide diagnostic information before and after every MQI call. For information on how to do this, see Chapter 15, "Applications on MQSeries for MVS/ESA" on page 219.

When testing AS/400 applications, you can use the Extended Program Model Debugger. To start this, use the STRDBG command.

Chapter 3. MQSeries messages

An MQSeries message consists of two parts:

- Message descriptor
- Application data

Figure 2 represents a message and shows how it is logically divided into message data and application data.

Figure 2. Representation of a message

The application data carried in an MQSeries message is not changed by a queue manager unless data conversion is carried out on it. Also, MQSeries does not put any restrictions on the content of this data. The length of the data in each message cannot exceed the value of the *MaxMsgLength* attribute of both the queue and queue manager. In each MQSeries product, the *MaxMsgLength* defaults to 4 MB (where MB equals 1048576 bytes). However, you should make your messages slightly shorter than the value of the *MaxMsgLength* attribute in some circumstances (see "The data in your message" on page 107 for more information).

You create a message when you use the MQPUT or MQPUT1 MQI call. As input to these calls, you supply the control information (such as the priority of the message, and the name of a reply queue) and your data. These calls put the message on a queue. See the *MQSeries Application Programming Reference* for more information on these calls.

Message descriptor

You can access message control information using the MQMD structure, which defines the *message descriptor*. Here is a brief description of that information, listed in the order of the fields in the MQMD structure. For a full description of the MQMD structure, see the *MQSeries Application Programming Reference* manual.

Note for RPG programmers

The names of MQI calls, call parameters, data types, fields of structures, and constants are shown in this book with long names. When writing RPG programs, use the equivalent names given in Appendix B, "MQI names in RPG" on page 437.

Strucid
  Identifies the structure as a Message Descriptor structure. Always specify the value of MQMD_STRUC_ID.

Version

Describes the version number of the structure. MQMD_VERSION_1 is the default. If you wish to use segmentation or grouping facilities (see on page 136), you must specify MQMD_VERSION_2. MQMD_CURRENT_VERSION sets your application to use the most recent level.

Report

The type of report required for the message (see "Report messages" on page 24). The report options also specify the action that an application processing a request message, or creating a reply message, should take. See "Request messages" on page 24 and "Reply messages" on page 24.

MsgType

See "Types of message" on page 23.

Expiry

This may be either set to unlimited (using MQEI_UNLIMITED) or a period of time expressed in tenths of a second, set at the time the message is put. If it is still on the queue after this time has elapsed, it will be removed from the queue.

Feedback

Used only for report messages. This is a code that describes how another message was handled. See "Report messages" on page 24.

Encoding

Describes the physical representation of the message data. You will normally use the default value, MQENC_NATIVE. For more information, see "Format of message data" on page 29 and the description of the *Encoding* field in the *MQSeries Application Programming Reference* manual.

CodedCharSetId

Describes the appearance of characters in the message data. See "Format of control information and message data" on page 29.

Format

An user-defined name of the message data format. See "Format of control information and message data" on page 29.

Priority

Use the priority to influence the order in which messages are retrieved from the queue. See "Message priorities" on page 31.

Persistence

Whether or not the message survives restarts of the queue manager. See "Message persistence" on page 33.

MsgId

An identifier by which the message can be located. See "Selecting messages from queues" on page 34.

CorrelId

An application-specified field that you can use, for example, to relate a response message to a request message. See "Selecting messages from queues" on page 34.

BackoutCount

This is the number of times a message has been taken from the queue under syncpoint and then backed out. See "Messages that are backed out" on page 34.

ReplyToQ

   The name of the queue to which reply and report messages should be sent. See "Reply-to queue and queue manager" on page 35.

ReplyToQMgr

   The name of the queue manager owning the reply-to queue. See "Reply-to queue and queue manager" on page 35.

Message context

   Information about the origin of the message held in eight fields within the MQMD:

   - *UserIdentifier*
   - *AccountingToken*
   - *ApplIdentityData*
   - *PutApplType*
   - *PutApplName*
   - *PutDate*
   - *PutTime*
   - *ApplOriginData*

   See "Message context" on page 35 for a description of how to use these fields.

Additional information for grouping and segmenting messages (see "Message groups" on page 32) is provided in Version 2 of the Message Descriptor (or the MQMDE). This is the same as the Version 1 Message Descriptor as described previously, but has additional fields as described in the *MQSeries Application Programming Reference* manual.

Types of message

There are four types of message defined by MQSeries:

- Datagram
- Request
- Reply
- Report

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

Each type of message is identified by an MQMT_$\star$ value. You can also define your own types of message. For the range of values you can use, see the description of the *MsgType* field in the *MQSeries Application Programming Reference* manual.

Datagrams

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

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

Request messages

You should use a *request message* when you want a reply from the application that receives the message.

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

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

- You can give your application the responsibility of ensuring that it puts information into the reply message that relates to the request message.

- You can use the report field in the message descriptor of your request message to specify the content of the *MsgId* and *CorrelId* fields of the reply message:

  - You can request that either the *MsgId* or the *CorrelId* of the original message is to be copied into the *CorrelId* field of the reply message (the default action is to copy *MsgId*)

  - You can request that either a new *MsgId* is generated for the reply message, or that the *MsgId* of the original message is to be copied into the *MsgId* field of the reply message (the default action is to generate a new message identifier)

Reply messages

You should use a *reply message* when you reply to another message.

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

Report messages

Report messages inform applications about events such as the occurrence of an error when processing a message. They can be generated by:

- a queue manager,

- a message channel agent (for example, if they cannot deliver the message),

or

- an application (for example, if it cannot use the data in the message).

Note that report messages can be generated at any time, and they may arrive on a queue when your application is not expecting them.

Types of report message

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

- An *exception report message*. This is sent in response to a message that had the exceptions flag set. It is generated by the message channel agent (MCA) or the application.

- An *expiry report message*. This indicates that a message has been discarded before being retrieved by an application due to reaching its expiry threshold. This type of report is generated by the queue manager.

- A *confirmation of arrival (COA) report message*. This indicates that the message has reached its target queue. It is generated by the queue manager.

- A *confirmation of delivery (COD) report message*. This indicates that the message has been retrieved by a receiving application. It is generated by the queue manager.

- A *positive action notification (PAN) report message*. This indicates that a request has been successfully serviced (that is, the action requested in the message has been performed successfully). This type of report is generated by the application.

- A *negative action notification (NAN) report message*. This indicates that a request has **not** been successfully serviced (that is, the action requested in the message has **not** been performed successfully). This type of report is generated by the application.

Note: Each type of report message is about this message, containing the original message, the first 100 bytes of data in the original message, or no data from the original message.

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

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

Report message options

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

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

When an exception report message is generated, it inherits the persistence of the original message.

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

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

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

- You can request that either the *MsgId* or the *CorrelId* of the original message is to be copied into the *CorrelId* field of the report message (the default action is to copy the message identifier)

- You can request that either a new *MsgId* is generated for the report message, or that the *MsgId* of the original message is to be copied into the *MsgId* field of the report message (the default action is to generate a new message identifier)

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

For more information on how to use report messages, see the description of the *Report* field in the *MQSeries Application Programming Reference* manual.

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

For more information on feedback and reason codes, see the description of the *Feedback* field in the *MQSeries Application Programming Reference* manual. manual.

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

Reports and segmented messages

These are supported on MQSeries Version 5 products and MQSeries for AS/400 V4R2 only.

If a message is segmented (see "Message segmentation" on page 136 for a description of this) and you ask for reports to be generated, you may receive more reports than you would have done had the message not been segmented.

MQSeries-generated reports

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

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

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

Application-generated reports

If your application generates reports, you should always copy the MQSeries headers that are present at the start of the original message data to the report message data. Then add none, 100 bytes, or all of the original message data (or whatever other amount you would normally include) to the report message data.

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

- MQMDE
- MQDLH
- MQXQH
- MQIIH
- MQH*

MQH* means any name starting with the characters MQH.

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

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

Retrieval of reports

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

This technique can be used even if there are several different report types present on the queue (for example, both COA and COD), because an MQGET with MQGMO_COMPLETE_MSG reassembles report messages only if they have the

same *Feedback* code. Note, however, that you cannot normally use the technique for exception reports, since in general these have different *Feedback* codes.

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

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

One approach is as follows:

- Ask for COD reports and exception reports.

- After a specific time, check whether a complete set of COD reports has been received using MQGMO_COMPLETE_MSG. If so, your application knows that the entire message has been processed.

- If not, and exception reports relating to this message are present, the problem should be handled just as for unsegmented messages, though provision must also be made for 'orphan' segments to be cleaned up at some point.

- If there are segments for which there are no reports of any kind, the original segments (or the reports) may be waiting for a channel to be reconnected, or the network might be overloaded at some point. If no exception reports at all have been received (or if you think that the ones you have may be temporary only), you may decide to let your application wait a little longer.

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

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

Back-level queue managers

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

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

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

Format of control information and message data

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

Format of message control information

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

Format of message data

You can specify any of the following:

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

To do this, use these fields:

Format

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

When the queue manager creates a message, in some circumstances it uses the *Format* field to identify the format of that message. For example, when a queue manager cannot deliver a message, it puts the message on a dead-letter (undelivered-message) queue. It adds a header (containing more control information) to the message, and changes the *Format* field to show this.

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

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

CodedCharSetId

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

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

Encoding
>   This describes the format of numeric message data that contains binary integers, packed-decimal integers, and floating point numbers. It is usually encoded according to the particular machine on which the queue manager is running.
>
>   When you put a message on a queue, you should normally specify the constant MQENC_NATIVE in the *Encoding* field. This means that the encoding of your message data is the same as that of the machine on which your application is running.
>
>   When you get a message from a queue, you should compare the value of the *Encoding* field in the message descriptor with the value of the constant MQENC_NATIVE on your machine. If the two values differ, you may need to convert any numeric data in the message or use a data-conversion message exit if one is available.

Application data conversion

Application data may need to be converted to the character set and the encoding required by another application where different platforms are concerned. It may be converted at the sending queue manager, or at the receiving queue manager. If the library of built-in formats does not meet your needs, you must define your own. The type of conversion depends on the message format which is specified in the format field of the message descriptor, MQMD.

Conversion at the sending queue manager

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

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

Built-in formats: These include:

- Messages that are all characters (using the format name MQFMT_STRING)

- MQSeries defined messages, for example Programmable Command Formats

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

Note: Messages with MQFMT_NONE specified are not converted.

The queue manager built-in formats all have names beginning with MQFMT. They are listed and described in the *MQSeries Application Programming Reference* manual, under the *Format* field of the Message descriptor (MQMD).

Application-defined formats: For user-defined formats, application data conversion must be performed by a data-conversion exit program (for more information, see Chapter 11, "Data-conversion exits" on page 157). In a client-server environment, the exit is loaded at the server and conversion takes place there.

Conversion at the receiving queue manager

Application message data may be converted by the receiving queue manager for the built-in formats and user-defined formats. The conversion is performed during the processing of an MQGET call if the MQGMO_CONVERT option is specified. For details see the *MQSeries Application Programming Reference* manual.

Coded character sets

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

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

For general data conversion, if the underlying operating system supports DBCS code pages then MQSeries is able to use it.

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

You need to consider application data conversion, format names and user exits when writing applications that span multiple platforms. For details of the MQGET call, the Convert characters call, the MQGMO_CONVERT option, and the built-in formats, see the *MQSeries Application Programming Reference* manual. See Chapter 11, "Data-conversion exits" on page 157 for information about invoking and writing data-conversion exits.

Message priorities

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

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

The *DefPriority* attribute of a queue sets the default priority value for messages being put on that queue. This value is set when the queue is created, but it can be changed afterwards. Alias queues, and local definitions of remote queues, may have different default priorities from the base queues to which they resolve. If there is more than one queue definition in the resolution path (see "Name resolution" on page 96), the default priority is taken from the value (at the time of the put operation) of the *DefPriority* attribute of the queue specified in the open command.

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

Message groups

| *These are supported on MQSeries Version 5 products and MQSeries for AS/400*
| *V4R2 only .*

Messages can occur within groups. This allows ordering of messages (see "Logical and physical ordering" on page 125), and segmentation of large messages (see "Message segmentation" on page 136) within the same group.

The hierarchy within a group is as follows:

Group

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

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

Figure 3 shows a group of logical messages:

```
                         Group
       ┌───────────────────┼───────────────────┐
       │                   │                   │
   LOGMSG1             LOGMSG2             LOGMSG3
```

Figure 3. Group of logical messages

Logical message

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

Logical messages within a group can be used to:

- Ensure ordering (if this is not guaranteed under the circumstances in which the message is transmitted).

- Allow applications to group together similar messages (for example, those that must all be processed by the same server instance).

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

Segment

Segments are used to handle messages that are too large for either the putting or getting application or the queue manager (including intervening queue managers through which the message passes). For more information about this, see "Message segmentation" on page 136.

A segment of a message is identified by the *GroupId*, *MsgSeqNumber*, and *Offset* fields. The *Offset* field starts at zero for the first segment within a message.

Each segment consists of one physical message that may or may not belong to a group (4 shows an example of messages within a group). A segment is logically part of a single message, so only the *MsgId*, *Offset*, and *SegmentFlag* fields in the MQMD should differ between separate segments of the same message.

Figure 4 shows a group of logical messages, some of which are segmented:

```
                              Group
                                |
           +--------------------+--------------------+
           |                    |                    |
       LOGMSG1              LOGMSG2              LOGMSG3
           |                                         |
    +------+------+                   +--------------+--------------+
    |             |                   |              |              |
  SEG1          SEG2               SEG1            SEG2           SEG3
```

Figure 4. Segmented messages

For a description of logical and physical messages, see "Logical and physical ordering" on page 125. For further information about segmenting messages, see "Message segmentation" on page 136.

Message persistence

When you create a message, if you initialize the message descriptor (MQMD) using the defaults, the persistence for the message will be taken from the *DefPersistence* attribute of the queue specified in the MQOPEN command. Alternatively, you may set the persistence of the message using the *Persistence* field of the MQMD structure to define the message as persistent or not persistent.

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

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

- A persistent message, outside the current unit of work, is written to disk on every put and get operation. See Chapter 13, "Committing and backing out units of work" on page 185.

- In MQSeries on UNIX systems, MQSeries for Digital OpenVMS, MQSeries for MVS/ESA, MQSeries for OS/2 Warp, and MQSeries for Windows NT, a persistent message within the current unit of work is logged only when the unit of work is committed (and the unit of work could contain many queue operations).

Nonpersistent messages can be used for fast messaging if retrieved outside syncpoint. See the *MQSeries Application Programming Reference* book and the *MQSeries Intercommunication* book for further information about fast messages.

Selecting messages from queues

To get a particular message from a queue, you need to use the *MsgId* and *CorrelId* fields of the message descriptor. If you specify Version 2 of the MQMD, the *GroupId* can also be used. (See "Getting a particular message" on page 133.)

The message identifier is usually generated by the queue manager when the message is put on a queue. The queue manager tries to ensure that message identifiers are unique. However, an MQSeries application can specify a particular value for the message identifier.

You can use the correlation identifier in any way you like. However, an intended use of this field is for applications to copy the message identifier of a request message into the *CorrelId* field of a reply message.

The group identifier is usually generated by the queue manager when the first message of a group is put onto a queue. The *MsgSeqNumber* field identifies the position of the message within the group and the *Offset* field identifies the segments within the message.

Where more than one message meets the combined selection criteria, the *MsgDeliverySequence* attribute of the queue determines whether messages are selected in FIFO (first in, first out) or priority order. When messages have equal priority, they are selected in FIFO order. For more information, see "The order in which messages are retrieved from a queue" on page 124.

For an example of an application that uses correlation identifiers, see "The Credit Check sample" on page 409.

Messages that fail to be delivered

When a queue manager is unable to put a message on a queue, you have various options. You can:

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

See Chapter 5, "Handling program errors" on page 53 for more information.

Messages that are backed out

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

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

In MQSeries for MVS/ESA, to make the backout count survive restarts of the queue manager, set the *HardenGetBackout* attribute to MQQA_BACKOUT_HARDENED; otherwise, if the queue manager has to restart, it does not maintain an accurate backout count for each message. Setting the attribute this way adds the penalty of extra processing.

In MQSeries for AS/400, MQSeries for OS/2 Warp, MQSeries for Windows NT, MQSeries for Digital OpenVMS, and MQSeries on UNIX systems, the backout count always survives restarts of the queue manager. Any change to the *HardenGetBackout* attribute is ignored.

For more information on committing and backing out messages, see Chapter 13, "Committing and backing out units of work" on page 185.

Reply-to queue and queue manager

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

- A reply message in response to a request message
- A report message about an unexpected event or expiry
- A report message about a COA or a COD event
- A report message about a PAN or a NAN event

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

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

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

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

Note you can request that a queue manager makes more than one attempt to deliver a message, and you can request that the message is discarded if it fails. If the message, after failing to be delivered, is not to be discarded, the remote queue manager puts the message on its dead-letter (undelivered-message) queue (see "Using the dead-letter (undelivered-message) queue" on page 57).

Message context

Message context information allows the application that retrieves the message to find out about the originator of the message. The retrieving application may want to:

- Check that the sending application has the correct level of authority

- Perform some accounting function so that it can charge the sending application for any work it has to perform

- Keep an audit trail of all the messages it has worked with

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. For more information on how to specify context information, see "Controlling context information" on page 108.

All context information is stored in the eight context fields of the message descriptor. The type of information falls into two categories: identity and origin context information.

Identity context

Identity context information identifies the user of the application that **first** put the message on a queue:

- The queue manager fills the *UserIdentifier* field with a name that identifies the user—the way that the queue manager can do this depends on the environment in which the application is running.

- The queue manager fills the *AccountingToken* field with a token or number that it determined from the application that put the message.

- Applications can use the *ApplIdentityData* field for any extra information that they want to include about the user (for example, an encrypted password).

Suitably authorized applications may set the above fields.

For information on how the queue manager fills the *UserIdentifier* and *AccountingToken* fields, see the descriptions of these fields in the *MQSeries Application Programming Reference* manual.

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

Origin context

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

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

Origin context information is usually supplied by the queue manager. Greenwich Mean Time (GMT) is used for the *PutDate* and *PutTime* fields.

Within MQSeries for OS/2 Warp only, the TZ environment variable is used to calculate the GMT *PutDate* and *PutTime* of a message. For more information on setting this variable, refer to the *MQSeries System Administration* book.

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

Message context

Chapter 4. MQSeries objects

Here is a list of MQSeries objects:

- Queue managers
- Queues
- Namelists (MQSeries for MVS/ESA only)
- Process definitions
- Channels
- Storage classes (MQSeries for MVS/ESA only)

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

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

Queue managers

A *queue manager* supplies an application with MQSeries services. A program must have a connection to a queue manager before it can use the services of that queue manager. A program can make this connection explicitly (using the MQCONN call), or the connection might be made implicitly (this depends on the platform and the environment in which the program is running).

Queues belong to queue managers, but programs can send messages to queues that belong to any queue manager.

Attributes of queue managers

Associated with each queue manager is a set of attributes (or properties) that define its characteristics. Some of the attributes of a queue manager are fixed when it is created; you can change others using the MQSeries commands. You can inquire about the values of **all** the attributes using the MQINQ call.

The *fixed* attributes are:

- The name of the queue manager

- The platform on which the queue manager runs (for example, AS/400)

- The level of system control commands that the queue manager supports

- The maximum priority that you can assign to messages processed by the queue manager

- The name of the queue to which programs can send MQSeries commands

- The identifier of the character set the queue manager uses for character strings when it processes MQI calls (this can be changed in MVS/ESA using the system parameters)
- The maximum length of messages the queue manager can process
- Whether or not the queue manager supports syncpointing when programs put and get messages

The *changeable* attributes are:

- A text description of the queue manager
- The time interval that the queue manager uses to restrict the number of trigger messages
- The name of the queue manager's dead-letter (undelivered-message) queue (useful if this queue is full, but is already set to the maximum size allowed, as it allows the queue manager to start using a different one)
- The name of the queue manager's default transmission queue
- The maximum number of open handles for any one connection
- The enabling and disabling of various categories of event reporting
- The maximum number of uncommitted messages within a unit of work

For a full description of these attributes, see the *MQSeries Application Programming Reference* manual.

Queues

An MQSeries *queue* is a named object on which applications can put messages, and from which applications can get messages. Messages are stored on a queue, so if the putting application is expecting a reply to its message, it is free to do other work while waiting for that reply. Applications access a queue by using the Message Queue Interface (MQI), described in Chapter 6, "Introduction to the Message Queue Interface" on page 65.

Before a message can be put on a queue, the queue must have already been created. A queue is owned by a queue manager, and that queue manager can own many queues. However, each queue must have a name that is unique within that queue manager.

A queue is maintained through a queue manager. Queues are managed physically by their queue managers but this is transparent to an application program.

To create a queue, use the commands described in one of the following manuals:

- The *MQSeries Command Reference* manual
- The *MQSeries for AS/400 Administration Guide*
- The *MQSeries Programmable System Management* book

Alternatively, you can use the MQSeries for MVS/ESA operations and control panels.

With MQSeries for AS/400 V4R2, MQSeries for MVS/ESA, MQSeries for OS/2 Warp, MQSeries for Digital OpenVMS, and MQSeries on UNIX systems, you can also create local queues "dynamically" from your application for temporary jobs,

such as *reply-to* queues (which are not needed after an application ends). For more information, see "Dynamic queues" on page 45.

Before using a queue, you must open the queue, specifying what you want to do with it. For example, you can open a queue:

- For browsing messages only (not retrieving them)
- For retrieving messages (and either sharing the access with other programs, or with exclusive access)
- For putting messages on the queue
- For inquiring about the attributes of the queue
- For setting the attributes of the queue

For a complete list of the options you can specify when you open a queue, see the description of the MQOPEN call in the *MQSeries Application Programming Reference* manual.

Types of queues

The types of queue that MQSeries supports for applications to use are:

Local and remote queues

A queue is known to a program as *local* if it is owned by the queue manager to which the program is connected; the queue is known as *remote* if it is owned by a different queue manager. The important difference between these two types of queue is that you can get messages only from local queues. (You can put messages on both types of queue.)

The queue definition object, created when you define a local queue, will hold the definition information of the queue as well as the physical messages put on the queue. The queue definition object, created when you 'define' a remote queue, will only hold the information necessary for the local queue manager to be able to locate the queue you want your message to go to. This object is known as the *local definition of a remote queue*. All the attributes of the remote queue are held by the queue manager that owns it, because it is a local queue to that queue manager.

Alias queues

To your program, an *alias queue* appears to be a queue, but it is really an MQSeries object that you can use to access another queue. This means that more than one program can work with the same queue, accessing it using different names.

Model and Dynamic queues

A model queue is a template of a queue definition used only when you want to create a dynamic local queue.

You can create a local queue dynamically from an MQSeries program, naming the model queue you wish to use as the template for the queue attributes. You may now, if you wish, change some attributes of the new queue. However, you cannot change the *DefinitionType*. If, for example, you require a permanent queue, you must select a model queue with the definition type set to permanent. Some conversational applications could make use of dynamic queues to hold replies to their queries because they probably would not need to maintain these queues after they have processed the replies.

Types of local queues

Each queue manager can have some local queues that it uses for special purposes:

Transmission queues

A *transmission queue* is a local queue which holds messages destined for a remote queue. The messages are forwarded to their destination queue by MQSeries when a communication program and link are available.

Initiation queues

An *initiation queue* is a local queue on which the queue manager puts a message for the purpose of automatically starting an application when certain conditions (such as more than 10 messages arriving, for example) are met on a local queue.

Dead-letter (undelivered-message) queue

The *dead-letter queue* is a local queue on which the queue manager and applications put messages they cannot deliver. You should plan to process any messages that arrive on this queue.

System command queue

The *system command queue* is a queue to which suitably authorized applications can send MQSeries commands.

System default queues

When you create a queue (other than a dynamic queue), MQSeries uses the queue definitions stored in the *system default queues*.

Channel queues

Channel queues are used for distributed queue management.

Event queues

Event queues hold event messages. These messages are reported by the queue manager or a channel.

These special queues are described in greater detail in the following sections.

Attributes of queues

Some of the attributes of a queue are specified when the queue is defined, and may not be changed afterwards (for example, the type of the queue). Other attributes of queues can be grouped into those that can be changed:

- By the queue manager during the processing of the queue (for example, the current depth of a queue)

- Only by commands (for example, the text description of the queue)

- By applications, using the MQSET call (for example, whether or not put operations are allowed on the queue)

You can find the values of all the attributes using the MQINQ call.

The attributes that are common to more than one type of queue are:

QName

Name of the queue

QType

Type of the queue

QDesc
Text description of the queue

InhibitGet
Whether or not programs are allowed to get messages from the queue (although you can never get messages from remote queues)

InhibitPut
Whether or not programs are allowed to put messages on the queue

DefPriority
Default priority for messages put on the queue

DefPersistence
Default persistence for messages put on the queue

Scope **(***not supported on AS/400 or MVS/ESA***)**
Controls whether an entry for this queue also exists in a cell directory

For a full description of these attributes, see the *MQSeries Application Programming Reference* manual.

Remote queues

To a program, a queue is *remote* if it is owned by a different queue manager to the one to which the program is connected. Where a communication link has been established, it is possible for a program to send a message to a remote queue. A program can never get a message from a remote queue.

When opening a remote queue, to identify the queue you must specify either:

- The name of the local definition that defines the remote queue.

    To create a local definition of a remote queue use the DEFINE QREMOTE command; in MQSeries for AS/400, alternatively use the CRTMQMQ command; in MQSeries for Tandem NSK you can use the MQM screen-based interface.

    From the viewpoint of an application, this is the same as opening a local queue. An application does not need to know if a queue is local or remote, this is sorted out by the queue manager.

- The name of the remote queue manager and the name of the queue as it is known to that remote queue manager.

Local definitions of remote queues have three attributes in addition to the common attributes described in "Attributes of queues" on page 42. These are *RemoteQName* (the name that the queue's owning queue manager knows it by), *RemoteQMgrName* (the name of the owning queue manager), and *XmitQName* (the name of the local transmission queue that is used when forwarding messages to other queue managers). For a fuller description of these attributes, see the *MQSeries Application Programming Reference* manual.

If you use the MQINQ call against the local definition of a remote queue, the queue manager returns the attributes of the local definition only, that is the remote queue name, the remote queue manager name and the transmission queue name, not the attributes of the matching local queue in the remote system.

See also "Transmission queues" on page 46.

Alias queues

An *alias queue* is an MQSeries object that you can use to access another queue. The queue resulting from the resolution of an alias name (known as the base queue) can be either a local queue or the local definition of a remote queue. It can also be either a predefined queue or a dynamic queue, as supported by the platform.

Note: An alias cannot resolve to another alias.

An example of the use of alias queues is for a system administrator to give different access authorities to the base queue name (that is, the queue to which the alias resolves) and to the alias queue name. This would mean that a program or user could be authorized to use the alias queue, but not the base queue.

Alternatively, authorization can be set to inhibit put operations for the alias name, but allow them for the base queue.

In some applications, the use of alias queues means that system administrators can easily change the definition of an alias queue object without having to get the application changed.

MQSeries makes authorization checks against the alias name when programs try to use that name. It does not check that the program is authorized to access the name to which the alias resolves. A program can therefore be authorized to access an alias queue name, but not the resolved queue name.

In addition to the general queue attributes described in "Attributes of queues" on page 42, alias queues have a *BaseQName* attribute. This is the name of the base queue to which the alias name resolves. For a fuller description of this attribute, see the *MQSeries Application Programming Reference* manual.

The *InhibitGet* and *InhibitPut* attributes (see "Attributes of queues" on page 42) of alias queues belong to the alias name. For example, if the alias-queue name ALIAS1 resolves to the base-queue name BASE, inhibitions on ALIAS1 affect ALIAS1 only and BASE is not inhibited. However, inhibitions on BASE also affect ALIAS1.

The *DefPriority* and *DefPersistence* attributes also belong to the alias name. So, for example, you can assign different default priorities to different aliases of the same base queue. Also, you can change these priorities without having to change the applications that use the aliases.

Model queues

A *model queue* is a template of a queue definition, that you use when creating a dynamic queue. You specify the name of a model queue in the *object descriptor* (MQOD) of your MQOPEN call. Using the attributes of the model queue, the queue manager dynamically creates a local queue for you.

You can specify a name (in full) for the dynamic queue, or the stem of a name (for example, ABC) and let the queue manager add a unique part to this, or you can let the queue manager assign a complete unique name for you. If the queue manager assigns the name, it puts it in the MQOD structure.

You can not issue an MQPUT1 call directly to a model queue, however, once a model queue has been opened, you can issue an MQPUT1 to the dynamic queue.

The attributes of a model queue are a subset of those of a local queue. For a fuller description, see the *MQSeries Application Programming Reference* manual.

Dynamic queues

When an application program issues an MQOPEN call to open a model queue, the queue manager dynamically creates an instance of a local queue with the same attributes as the model queue. Depending on the value of the *DefinitionType* field of the model queue, the queue manager creates either a temporary or permanent dynamic queue (See "Creating dynamic queues" on page 100).

Properties of temporary dynamic queues

Temporary dynamic queues have the following properties:

- They hold nonpersistent messages only.

- They are non-recoverable.

- They are deleted when the queue manager is started

- They are deleted when the application that issued the MQOPEN call which resulted in the creation of the queue closes the queue or terminates.

    – If there are any committed messages on the queue, they will be deleted.

    – If there are any uncommitted MQGET, MQPUT, or MQPUT1 calls outstanding against the queue at this time, the queue is marked as being logically deleted, and is only physically deleted (after these calls have been committed) as part of close processing, or when the application terminates.

    – If the queue happens to be in use at this time (by the creating, or another application), the queue is marked as being logically deleted, and is only physically deleted when closed by the last application using the queue.

    – Attempts to access a logically deleted queue (other than to close it) fail with reason code MQRC_Q_DELETED.

    – MQCO_NONE, MQCO_DELETE and MQCO_DELETE_PURGE are all treated as MQCO_NONE when specified on an MQCLOSE call for the corresponding MQOPEN call that created the queue.

Properties of permanent dynamic queues

Permanent dynamic queues have the following properties:

- They hold persistent or nonpersistent messages.

- They are recoverable in the event of system failures.

- They are deleted when an application (not necessarily the one that issued the MQOPEN call which resulted in the creation of the queue) successfully closes the queue using the MQCO_DELETE, or the MQCO_DELETE_PURGE option.

    – A close request with the MQCO_DELETE option fails if there are any messages (committed or uncommitted) still on the queue. A close request with the MQCO_DELETE_PURGE option succeeds even if there are committed messages on the queue (the messages being deleted as part of the close), but fails if there are any uncommitted MQGET, MQPUT, or MQPUT1 calls outstanding against the queue.

- – If the delete request is successful, but the queue happens to be in use (by the creating, or another application), the queue is marked as being logically deleted and is only physically deleted when closed by the last application using the queue.

- They are not deleted if closed by an application (other than the one that issued the MQOPEN call which resulted in the creation of the queue) which is not authorized to delete the queue. Authorization checks are performed against the user identifier (or alternate user identifier if MQOO_ALTERNATE_USER_AUTHORITY was specified) which was used to validate the corresponding MQOPEN call.

- They can be deleted in the same way as a normal queue.

Uses of dynamic queues

You can use dynamic queues for:

- Applications that do not require queues to be retained after the application has terminated.

- Applications that require replies to messages to be processed by another application can dynamically create a reply-to queue by opening a model queue. For example, a client application could:

  1. Create a dynamic queue.

  2. Supply its name in the *ReplyToQ* field of the message descriptor structure of the request message.

  3. Place the request on a queue being processed by a server.

The server could then place the reply message on the reply-to queue. Finally, the client could process the reply, and close the reply-to queue with the delete option.

Recommendations for uses of dynamic queues

You should consider the following points when using dynamic queues:

- In a client-server model, each client should create and use its own dynamic reply-to queue. If a dynamic reply-to queue is shared between more than one client, the deletion of the reply-to queue may be delayed because there is uncommitted activity outstanding against the queue, or because the queue is in use by another client. Additionally, the queue might be marked as being logically deleted, and hence inaccessible for subsequent API requests (other than MQCLOSE).

- If your application environment requires that dynamic queues must be shared between applications, you should ensure that the queue is only closed (with the delete option) when all activity against the queue has been committed. This should be by the last user preferably. This ensures that deletion of the queue is not delayed, and should minimize the period that the queue is inaccessible because it has been marked as being logically deleted.

Transmission queues

When an application sends a message to a remote queue, the local queue manager stores the message in a special local queue, called a *transmission queue*.

A *message channel agent* (channel program) will be associated with the transmission queue and the remote queue manager, and it is this that deals with

the transmitting of the message. When the message has been transmitted, it is deleted from the transmission queue.

The message may have to pass through many queue managers (or *nodes*) on its journey to its final destination. There must be a transmission queue defined at each queue manager along the route, each holding messages waiting to be transmitted to the next node. There can be several transmission queues defined at a particular queue manager. A given transmission queue holds messages whose **next** destination is the same queue manager, although the messages may have different eventual destinations. There may also be several transmission queues for the same remote queue manager, with each one being used for a different type of service, for example.

The queue manager must know the attributes of an application if that application is to start automatically. (For information about this, see Chapter 14, "Starting MQSeries applications using triggers" on page 197). These attributes are defined in the transmission queue definition (for triggered channels) or the process definition object (see "Process definition objects" on page 48).

Initiation queues

An *initiation queue* is a local queue on which the queue manager puts a trigger message when a trigger event occurs on an application queue. A trigger event is an event (for example, more than 10 messages arriving) that an application designer intends the queue manager to use as a cue, or trigger, to start a program that will process the queue. For more information on how triggering works, see Chapter 14, "Starting MQSeries applications using triggers" on page 197.

Dead-letter (undelivered-message) queues

A *dead-letter (undelivered-message) queue* is a local queue on which the queue manager puts messages it cannot deliver.

When the queue manager puts a message on the dead-letter queue, it adds a header to the message. This includes such information as the intended destination of the original message, the reason the queue manager put the message on the dead-letter queue, and the date and time it did this.

Applications can also use the queue for messages they cannot deliver. For more information, see "Using the dead-letter (undelivered-message) queue" on page 57.

System command queues

These queues receive the PCF, MQSC, and CL commands, as supported on your platform, in readiness for the queue manager to action them. In MQSeries for MVS/ESA the queue is known as the **SYSTEM.COMMAND.INPUT.QUEUE** and accepts MQSC commands. On other platforms it is known as the **SYSTEM.ADMIN.COMMAND.QUEUE** and the commands accepted vary by platform. See the *MQSeries Programmable System Management* book for details.

System default queues

The *system default queues* contain the initial definitions of the queues for your system. When you create a new queue, the queue manager copies the definition from the appropriate system default queue.

Namelists

Namelists are supported on MQSeries for MVS/ESA only.

A *namelist* is an MQSeries object that contains a list of queue names. The queues in the list can be of any type and the list can contain queues of more than one type.

You can define and modify namelists using only the commands or operations and control panels of MQSeries for MVS/ESA.

Programs can use the MQI to find out which queues are included in these namelists. The organization of the namelists is the responsibility of the application designer and system administrator.

The attributes of namelists are:

NamelistName
> The name of the namelist. The namelist name must be unique within the queue manager.

NamelistDesc
> Text description of the namelist.

NameCount
> The number of names contained in the list.

Names
> A list of 48-character queue names.

For a full description of these attributes, see the *MQSeries for MVS/ESA Application Programming Reference* manual.

Process definition objects

To allow an application to be started without the need for operator intervention (this is described in Chapter 14, "Starting MQSeries applications using triggers" on page 197), the attributes of the application must be known to the queue manager. These attributes are defined in a *process definition object*.

The *ProcessName* attribute is fixed when the object is created; you can change the others using the MQSeries commands or the MQSeries for MVS/ESA operations and control panels. You can inquire about the values of **all** the attributes using the MQINQ call.

The attributes are:

ProcessName
> A unique name for the process definition

ProcessDesc
> Application-defined field that may be used for a description of the application

ApplType
> The nature of the application (for example, a CICS transaction or AS/400 application)

ApplicId
> A character string identifier for the application (for a CICS transaction, this is a CICS transaction ID)

EnvData
> A user-defined character string that could be used to contain information related to the environment in which the application runs

UserData
> A user-defined character string that could be used to pass user information to the trigger monitor that starts the application

For a full description of these attributes, see the *MQSeries Application Programming Reference* manual.

Channels

A *channel* is a communication link used by distributed applications. There are two categories of channels in MQSeries:

- *Message* channels, which are unidirectional, and transfer messages from one queue manager to another.

- *MQI* channels, which are bidirectional, and transfer MQI calls from an MQSeries client to a queue manager, and responses from a queue manager to an MQSeries client.

These need to be considered when designing your application, but a program will be unaware of MQSeries channel objects. For more information, see the *MQSeries Intercommunication* book and the *MQSeries Clients* book.

Storage classes

Storage classes are supported on MQSeries for MVS/ESA only.

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

For further information about storage classes, see the *MQSeries for MVS/ESA System Management Guide*.

Rules for naming MQSeries objects

An MQSeries queue, process definition, namelist, and channel can all have the same name. However, an MQSeries object cannot have the same name as any other object of the same type. Names in MQSeries are case sensitive.

The character set that can be used for naming all MQSeries objects is as follows:

- Uppercase A–Z

- Lowercase a–z (but there are restrictions on the use of lowercase letters for MVS console support)

  On systems using EBCDIC Katakana you cannot use lowercase characters.

- Numerics 0–9

- Period (.)

- Forward slash (/)

- Underscore (_)

- Percent sign (%)

Notes:

1. Leading or embedded blanks are not allowed.

2. You should also avoid using names with leading or trailing underscores, because they cannot be handled by the MQSeries for MVS/ESA operations and control panels.

3. Any name that is less than the full field length can be padded to the right with blanks. All short names that are returned by the queue manager are always padded to the right with blanks.

4. Any structure to the names (for example, the use of the period or underscore) is not significant to the queue manager.

5. On AS/400 systems lowercase a-z, forward slash (/), and percent (%) are special characters. If you use any of these characters in a name, the name must be enclosed in quotation marks. Lowercase a-z characters are changed to uppercase if the name is not enclosed in quotation marks.

Queue names

The name of a queue has two parts:

- The name of a queue manager
- The local name of the queue as it is known to that queue manager

Each part of the queue name is 48 characters long.

To refer to a local queue, you can omit the name of the queue manager (by replacing it with blank characters or using a leading null character). However, all queue names returned to a program by MQSeries contain the name of the queue manager.

To refer to a remote queue, a program must include the name of the queue manager in the full queue name, or there must be a local definition of the remote queue.

Note that when an application uses a queue name, that name can be either the name of a local queue (or an alias to one) or the name of a local definition of a remote queue, but the application does not need to know which, unless it needs to get a message from the queue (when the queue must be local). When the application opens the queue object, the MQOPEN call performs a name resolution function to determine on which queue to perform subsequent operations. The significance of this is that the application has no built-in dependency on particular queues being defined at particular locations in a network of queue managers.

Therefore if a system administrator relocates queues in the network, and changes their definitions, the applications that use those queues do not need to be changed.

Process definition and namelist names

Process definitions and namelists can have names up to 48 characters long.

Channel names

Channels can have names up to 20 characters long. See the *MQSeries Intercommunication* book for further information on channels.

Reserved object names

Names that start with SYSTEM. are reserved for objects defined by the queue manager.

Naming objects

Chapter 5. Handling program errors

Your application may encounter errors associated with its MQI calls either when it makes a call or when its message is delivered to its final destination:

- Whenever possible, the queue manager returns any errors as soon as an MQI call is made. These are *locally determined errors*.

- When sending messages to a remote queue, errors may not be apparent when the MQI call is made. In this case, the queue manager that identifies the errors reports them by sending another message to the originating program. These are *remotely determined errors*.

This chapter gives advice on how to handle both types of error.

Locally determined errors

The three most common causes of errors that the queue manager can report immediately are:

- Failure of an MQI call; for example, because a queue is full

- An interruption to the running of some part of the system on which your application is dependent; for example, the queue manager

- Messages containing data that cannot be processed successfully

Failure of an MQI call

The queue manager can report immediately any errors in the coding of an MQI call. It does this using a set of predefined return codes. These are divided into completion codes and reason codes.

To show whether or not a call is successful, the queue manager returns a *completion code* when the call completes. There are three completion codes, indicating success, partial completion, and failure of the call. The queue manager also returns a *reason code* which indicates the reason for the partial completion or the failure of the call.

The completion and reason codes for each call are listed with the description of that call in the *MQSeries Application Programming Reference* manual. You will also find in this manual, further information (including some ideas for corrective action) for each completion and reason code, in the chapter on return codes. You should design your programs to handle all the return codes that could arise from each call.

System interruptions

Your application may be unaware of any interruption if the queue manager to which it is connected has to recover from a system failure. However, you must design your application to ensure that your data is not lost if such an interruption occurs.

The methods you can use to make sure that your data remains consistent depends on the platform on which your queue manager is running:

MVS

In the CICS and IMS environments, you can make MQPUT and MQGET calls within units of work that are managed by CICS and IMS. In the batch environment, you can make MQPUT and MQGET calls in the same way, but you must declare syncpoints by using the MQSeries for MVS/ESA CSQBCMT and CSQBBAK calls (see Chapter 13, "Committing and backing out units of work" on page 185).

AS/400

You can make your MQPUT and MQGET calls within units of work that are managed by OS/400 commitment control.

Digital OpenVMS, DOS, OS/2, UNIX systems, Windows NT, and Windows 3.1

In these environments, you can make your MQPUT and MQGET calls in the normal way, but you must declare syncpoints by using the MQCMIT and MQBACK calls (see Chapter 13, "Committing and backing out units of work" on page 185). In the CICS environment, MQCMIT and MQBACK commands are disabled as you can make your MQPUT and MQGET calls within units of work that are managed by CICS.

Tandem NSK

You can make your MQPUT and MQGET calls within units of work that are managed by Tandem's TM/MP product.

You should also use persistent messages for carrying all data you cannot afford to lose. Persistent messages are reinstated on queues if the queue manager has to recover from a failure. With MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT note that an MQGET or MQPUT call within your application will fail at the point of filling up all the log files, with the message MQRC_RESOURCE_PROBLEM. For more information on log files, see the *MQSeries System Administration* Guide for MQSeries for AIX, HP-UX, OS/2, and Windows NT; for other platforms, see the appropriate *System Management Guide*.

If the queue manager is stopped by an operator while an application is running, the quiesce option is normally used. The queue manager enters a quiescing state in which applications can continue to do work, but they should terminate as soon as it is convenient. Small, quick applications can probably ignore the quiescing state and continue until they terminate as normal. Longer running applications, or ones that wait for messages to arrive, should use the *fail if quiescing* option when they use the MQCONN, MQPUT, MQPUT1, and MQGET calls. These options mean that the calls fail when the queue manager quiesces, but the application may still have time to terminate cleanly by issuing calls that ignore the quiescing state. Such applications could also commit, or back out, changes they have made, and then terminate.

If the queue manager is forced to stop (that is, stop without quiescing), applications will receive the MQRC_CONNECTION_BROKEN reason code when they make MQI calls. At this point you must exit the application or, alternatively, on MQSeries

on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT, you can issue an MQDISC call.

Messages containing incorrect data

When you use units of work in your application, if a program cannot successfully process a message that it retrieves from a queue, the MQGET call is backed out. The queue manager maintains a count (in the *BackoutCount* field of the message descriptor) of the number of times this happens. It maintains this count in the descriptor of each message that is affected. This count can provide valuable information about the efficiency of an application. Messages whose backout counts are increasing over time are being repeatedly rejected—you should design your application so that it analyzes the reasons for this and handles such messages accordingly.

In MQSeries for MVS/ESA, to make the backout count survive restarts of the queue manager, set the *HardenGetBackout* attribute to MQQA_BACKOUT_HARDENED; otherwise, if the queue manager has to restart, it does not maintain an accurate backout count for each message. Setting the attribute this way adds the penalty of extra processing.

In MQSeries for AS/400, MQSeries for OS/2 Warp, MQSeries for Windows NT, and MQSeries on UNIX systems, the backout count always survives restarts of the queue manager.

Also, in MQSeries for MVS/ESA, when you remove messages from a queue within a unit of work, you can mark one message so that it is **not** made available again if the unit of work is backed out **by the application**. The marked message is treated as if it has been retrieved under a new unit of work. You mark the message that is to skip backout using the MQGMO_MARK_SKIP_BACKOUT option (in the MQGMO structure) when you use the MQGET call. See "Skipping backout" on page 145 for more information about this technique.

Using report messages for problem determination

The remote queue manager cannot report errors such as failing to put a message on a queue when you make your MQI call, but it can send you a report message to say how it has processed your message.

Within your application you can create (MQPUT) report messages as well as select the option to receive them (in which case they will be sent by either another application or by a queue manager).

Creating report messages

Report messages provide a mechanism for an application to inform another application that it is unable to deal with the message that was sent. However, the *Report* field must initially be analyzed to determine whether or not the application that sent the message is interested in being informed of any problems. Having determined that a report message is required, you have to decide:

- Whether you want to include all the original message (not an option on MVS/ESA), just the first 100 bytes of data, or none of the original message.

- What to do with the original message. You can discard it or let it go to the dead-letter queue.

- Whether the content of the *MsgId* and *CorrelId* fields are needed as well.

Use the *Feedback* field to indicate the reason for the report message being generated. Put your report messages on an application's reply-to queue. Refer to the *MQSeries Application Programming Reference* manual for further information.

Requesting and receiving (MQGET) report messages

When you send a message to another application, you will not be informed of any problems unless you complete the *report* field to indicate the feedback you require. The options available to you are in the *MQSeries Application Programming Reference* manual.

Queue managers always put report messages on an application's reply-to queue and it is recommended that your own applications do the same. When you use the report message facility you must specify the name of your reply-to queue in the message descriptor of your message; otherwise, the MQPUT call fails.

Your application should contain procedures that monitor your reply-to queue and process any messages that arrive on it. Remember that a report message can contain all the original message, the first 100 bytes of the original message, or none of the original message.

The queue manager sets the *Feedback* field of the report message to indicate the reason for the error; for example, the target queue does not exist. Your programs should do the same.

For more information on report messages, see "Report messages" on page 24.

Remotely determined errors

When you send messages to a remote queue, even when the local queue manager has processed your MQI call without finding an error, other factors can influence how your message is handled by a remote queue manager. For example, the queue you are targeting may be full, or may not even exist. If your message has to be handled by other intermediate queue managers on the route to the target queue, any of these could find an error.

Problems delivering a message

When an MQPUT call fails, you have the choice of attempting to put the message on the queue again, returning it to the sender, or putting it on the dead-letter queue.

Each option has its own merits, but you may not want to retry putting a message if the reason that the MQPUT failed was because the destination queue was full. In this instance, putting it on the dead-letter queue allows you to deliver it to the correct destination queue later on.

Retry message delivery

Before the message is put on a dead-letter queue, a remote queue manager attempts to put the message on the queue again if the attributes *MsgRetryCount* and *MsgRetryInterval* have been set for the channel, or if there is a retry exit program for it to use (the name of which is held in the channel attribute *MsgRetryExitId* field).

If the *MsgRetryExitId* field is blank, the values in the attributes *MsgRetryCount* and *MsgRetryInterval* are used.

If the *MsgRetryExitId* field is not blank, the exit program of this name runs. For more information on using your own exit programs, see the *MQSeries Intercommunication* book.

Return message to sender

You return a message to the sender by requesting a report message to be generated to include all of the original message. See "Report messages" on page 24 for details on report message options.

Using the dead-letter (undelivered-message) queue

When a queue manager cannot deliver a message, it attempts to put the message on its dead-letter queue. This queue should be defined when the queue manager is installed.

Your programs can use the dead-letter queue in the same way that the queue manager uses it. You can find the name of the dead-letter queue by opening the queue manager object (using the MQCONN call) and inquiring about the *DeadLetterQName* attribute (using the MQINQ call).

When the queue manager puts a message on this queue, it adds a header to the message, the format of which is described by the dead-letter header (MQDLH) structure, in the *MQSeries Application Programming Reference* manual. This header includes the name of the target queue and the reason the message was put on the dead-letter queue. It must be removed and the problem must be resolved before the message is put on the intended queue. Also, the queue manager changes the *Format* field of the message descriptor (MQMD) to indicate that the message contains an MQDLH structure.

> **MQDLH structure**
>
> You are recommended to add an MQDLH structure to all messages that you put on the dead-letter queue; however, if you intend to use the dead-letter handler provided by certain MQSeries products, you **must** add an MQDLH structure to your messages.

The addition of the header to a message may make the message too long for the dead-letter queue, so you should always make sure that your messages are shorter than the maximum size allowed for the dead-letter queue, by at least the value of the MQ_MSG_HEADER_LENGTH constant. The maximum size of messages allowed on a queue is determined by the value of the *MaxMsgLength* attribute of the queue. For the dead-letter queue, you should make sure that this attribute is set to the maximum allowed by the queue manager. If your application cannot deliver a

message, and the message is too long to be put on the dead-letter queue, follow the advice given in the description of the MQDLH structure.

You need to ensure the dead-letter queue is monitored, and that any messages arriving on it get processed. A dead-letter queue handler is provided by MQSeries on all platforms except MVS/ESA and AS/400. It runs as a batch utility and can be used to perform various actions on selected messages on the dead-letter queue. If you have a queue manager on one of the platforms that does not provide a dead-letter queue handler, you will need to provide your own. The program could be triggered, or run at regular intervals. For further details, see the *MQSeries System Administration* Guide for MQSeries for AIX, HP-UX, OS/2, Sun Solaris, and Windows NT; for other platforms, see the appropriate *System Management Guide*.

If data conversion is necessary, the queue manager converts the header information when you use the MQGMO_CONVERT option on the MQGET call. If the process putting the message is an MCA, the header is followed by all the text of the original message.

You should be aware that messages put on the dead-letter queue may be truncated if they are too long for this queue. A possible indication of this situation is the messages on the dead-letter queue being the same length as the value of the *MaxMsgLength* attribute of the queue.

Dead-letter queue processing

┌──────────────── General-use programming interface ────────────────┐

Dead-letter queue processing is dependent on local system requirements, but you should consider the following when you draw up the specification:

- The message can be identified as having a dead-letter queue header because the value of the format field in the MQMD, is MQFMT_DEAD_LETTER_HEADER.

- In MQSeries for MVS/ESA using CICS, if an MCA puts this message to the dead-letter queue, the *PutApplType* field is MQAT_CICS, and the *PutApplName* field is the *ApplId* of the CICS system followed by the transaction name of the MCA.

- The reason for the message to be routed to the dead-letter queue is contained in the *Reason* field of the dead-letter queue header.

- The dead-letter queue header contains details of the destination queue name and queue manager name.

- The dead-letter queue header contains fields that have to be reinstated in the message descriptor before the message is put to the destination queue. These are:

    1. *Encoding*
    2. *CodedCharSetId*
    3. *Format*

- The message descriptor is the same as PUT by the original application, except for the three fields shown above.

Your dead-letter queue application should do one or more of the following:

- Examine the *Reason* field. A message may have been put by an MCA for the following reasons:
  - The message was longer than the maximum message size for the channel

    The reason will be MQRC_MSG_TOO_BIG_FOR_CHANNEL (or MQRC_MSG_TOO_BIG_FOR_Q_MGR if you are using CICS for distributed queuing on MQSeries for MVS/ESA)
  - The message could not be put to its destination queue

    The reason will be any MQRC_* reason code that can be returned by an **MQPUT** operation
  - A user exit has requested this action

    The reason code will be that supplied by the user exit, or the default MQRC_SUPPRESSED_BY_EXIT
- Try to forward the message to its intended destination, where this is possible.
- Retain the message for a certain length of time before discarding when the reason for the diversion is determined, but not immediately correctable.
- Give instructions to administrators for the correction of problems where these have been determined.
- Discard messages that are corrupted or otherwise not processible.

There are two ways that you deal with the messages you have recovered from the dead-letter queue:

1. If the message is for a local queue, you should:

   - Carry out any code translations required to extract the application data
   - Carry out code conversions on that data if this is a local function
   - Put the resulting message on the local queue with all the detail of the message descriptor restored

2. If the message is for a remote queue, put the message on the queue.

────────────── End of General-use programming interface ──────────────

For information on how undelivered messages are handled in a distributed queuing environment, see the *MQSeries Intercommunication* book.

Remotely determined errors

Part 2. Writing an MQSeries application

Chapter 6. Introduction to the Message Queue Interface

This chapter introduces the features that comprise the Message Queue Interface (MQI).

The remaining chapters in this part of the book describe how to use these features. Detailed descriptions of the calls, structures, data types, return codes, and constants are given in the *MQSeries Application Programming Reference* manual.

What is in the MQI?

The Message Queue Interface comprises the following:

- *Calls* through which programs can access the queue manager and its facilities
- *Structures* that programs use to pass data to, and get data from, the queue manager
- *Elementary data types* for passing data to, and getting data from, the queue manager

MQSeries for MVS/ESA also supplies:

- Two extra calls through which MVS batch programs can commit and back out changes.
- *Namelists*, where each namelist contains a list of queue names.
- *Data definition files* (sometimes known as copy files, macros, include files, and header files) that define the values of constants supplied with MQSeries for MVS/ESA.
- *Stub programs* to link-edit to your applications.
- A suite of sample programs that demonstrate how to use the MQI on the MVS platform. For further information about these samples, see Chapter 31, "Sample programs for MQSeries for MVS/ESA" on page 375.

MQSeries for AS/400 also supplies:

- *Data definition files* (sometimes known as copy files, macros, include files, and header files) that define the values of constants supplied with MQSeries for AS/400.
- A stub program to link-edit to your C/400 applications.
- A suite of sample programs that demonstrate how to use the MQI on the AS/400 platform. For further information about these samples, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

MQSeries for OS/2 Warp, MQSeries for Windows NT, MQSeries for Digital OpenVMS, and MQSeries on UNIX systems also supply:

- Calls through which MQSeries for OS/2 Warp and MQSeries on UNIX systems programs can commit and back out changes.
- *Include files* that define the values of constants supplied on these platforms.
- *Library files* to link your applications.

- A suite of sample programs that demonstrate how to use the MQI on these platforms.
- Sample source and executable code for bindings to external transaction managers.

| MQSeries for Tandem NSK also supplies:

| - *Include files* that define the values of constants supplied with MQSeries for
| Tandem NSK.

| - *Library files* to link your applications.

| - A suite of sample programs that demonstrate how to use the MQI on the
| Tandem NSK platform.

MQSeries for Windows provides a subset of the MQI. For more information, see the following:

- *MQSeries for Windows V2.0 User's Guide*.
- *MQSeries for Windows V2.1 User's Guide*.

Calls

The calls in the MQI can usefully be grouped as follows:

MQCONN, MQCONNX, and MQDISC
Use these calls to connect a program to (with or without options), and disconnect a program from, a queue manager. Writers of CICS for MVS/ESA programs or OS/400 programs do not need to use these calls. However, you can use them if you want your application to be portable to other platforms.

MQOPEN and MQCLOSE
Use these calls to open and close an object, such as a queue.

MQPUT and MQPUT1
Use these calls to put a message on a queue.

MQGET Use this call to browse messages on a queue, or to remove messages from a queue.

MQINQ Use this call to inquire about the attributes of an object.

MQSET Use this call to set some of the attributes of a queue. You cannot set the attributes of other types of object.

MQBEGIN, MQCMIT, and MQBACK
Use these calls to start and commit a global unit of work (where coordination is performed by MQSeries) and indicate to the queue manager that all messages put or retrieved as part of a unit of work since the last syncpoint are to be backed out.

The MQI calls are described fully in the *MQSeries Application Programming Reference* manual.

Syncpoint calls

Syncpoint calls are available as follows:

MQSeries for MVS/ESA calls

MQSeries for MVS/ESA provides the MQCMIT and MQBACK calls. Use these calls in MVS batch programs to tell the queue manager that all the get and put operations since the last syncpoint are to be made permanent (committed) or are to be backed out. To commit and back out changes in other environments:

CICS   Use commands such as EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK.

IMS   Use the IMS syncpoint facilities, such as the GU (get unique) to the IOPCB, CHKP (checkpoint), and ROLB (rollback) calls.

For downward compatibility, the CSQBCMT and CSQBBAK calls are available as synonyms for MQCMIT and MQBACK. These are described fully in the *MQSeries Application Programming Reference* manual.

OS/400 calls

Use the OS/400 COMMIT and ROLLBACK commands, or any other commands or calls that initiate the OS/400 commitment control facilities (for example, EXEC CICS SYNCPOINT).

MQSeries for Tandem NSK calls

The default SYNCPOINT option for the MQPUT and MQGET calls is SYNCPOINT, rather than NO_SYNCPOINT. To use the default (SYNCPOINT) option for MQPUT, MQGET and MQPUT1 operations, the application must have an active TM/MP Transaction that defines the unit of work to be committed.

MQSeries calls on other platforms

The following products provide the MQCMIT and MQBACK calls:

- MQSeries for OS/2 Warp
- MQSeries for Windows
- MQSeries for Windows NT
- MQSeries for Digital OpenVMS
- MQSeries on UNIX systems

Use these calls in programs to tell the queue manager that all the get and put operations since the last syncpoint are to be made permanent (committed) or are to be backed out. To commit and back out changes in the CICS environment, use commands such as EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK.

Data conversion

The MQXCNVC - convert characters call is used only from a data-conversion exit. This call converts message character data from one character set to another.

See the *MQSeries Application Programming Reference* for the syntax used with the MQXCNVC call, and Chapter 11, "Data-conversion exits" on page 157 for guidance on writing and invoking data conversion exits.

Structures

The following structures are used with the MQI calls listed in "Calls" on page 66:

MQBO (Begin options)
> Specifies options for the MQBEGIN call (MQSeries Version 5 products only).

MQCNO (Connect options)
> Specifies options for the MQCONNX call (MQSeries Version 5 products only).

MQDH (Distribution header)
> Describes the data that is present in a message on a transmission queue when that message is a distribution-list message (MQSeries Version 5 products and MQSeries for AS/400 only).

MQGMO (Get message options)
> Specifies options for the MQGET call.

MQMD (Message descriptor)
> Provides control information for a message you are putting on (using MQPUT or MQPUT1), or getting from (using MQGET), a queue.

MQMDE (Message descriptor extension)
> In conjunction with an MQMD Version 1, this contains grouped message and segmentation information that would normally be held in the MQMD Version 2 (MQSeries Version 5 products and MQSeries for AS/400 only).

MQOD (Object descriptor)
> Identifies the object you want to work with when using MQOPEN.

MQOR (Object record)
> Identifies the destinations you want to work with in a distribution list (MQSeries Version 5 products, and MQSeries for AS/400 V4R2 only).

MQPMO (Put message options)
> Specifies options for the MQPUT and MQPUT1 calls.

MQPMR (Put-message record)
> Contains specific information relating to the individual destinations included in a distribution list (MQSeries Version 5 products, and MQSeries for AS/400 V4R2 only).

The following structures are used for special purposes:

MQDLH (Dead-letter header)
> Defines the format of the header of messages put on the dead-letter (undelivered-message) queue (not supported on MQSeries for Windows V2.0).

MQRMH (Reference message header)
> Defines the format of a reference message (MQSeries Version 5 products, and MQSeries for AS/400 only).

MQTM (Trigger message)
> Defines the format of a trigger message.

MQTMC (Trigger message)
> Defines the format of a trigger message as a set of character fields (MQSeries for AS/400 only).

MQTMC2 (Trigger message)

> Defines the format of a trigger message including the queue manager name (MQSeries for MVS/ESA, MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT only).

MQXP (Exit parameter block) structure

> Used to communicate with the API-crossing exit (MQSeries for MVS/ESA only).

MQXQH (Transmission queue header)

> Defines the format of the header that is added to messages put on a transmission queue.

The structures are supplied in data definition files for each of the supported programming languages. MQSeries for MVS/ESA and MQSeries for AS/400 supply files that contain constants for you to use when filling in some of the fields of these structures. For more information on these, see "MQSeries data definitions."

All the structures are described fully in the *MQSeries Application Programming Reference* manual.

Elementary data types

For the C and Visual Basic languages, the MQI provides the following elementary data types or unstructured fields:

MQBYTE	A single byte of data
MQBYTEn	A string of 16, 24, 32, or 64 bytes
MQCHAR	One single-byte character
MQCHARn	A string of 4, 8, 12, 16, 20, 28, 32, 48, 64, 128, or 256 single-byte characters
MQHCONN	A connection handle (this data is 32 bits long)
MQHOBJ	An object handle (this data is 32 bits long)
MQLONG	A 32-bit signed binary integer
PMQLONG	A pointer to data of type MQLONG

These data types are described fully in the *MQSeries Application Programming Reference* manual.

For COBOL, assembler, PL/I, or RPG, use the equivalent declarations shown in the same manuals.

MQSeries data definitions

MQSeries for MVS/ESA supplies data definitions in the form of COBOL copy files, assembler-language macros, a single PL/I include file, and a single C language include file.

MQSeries for AS/400 supplies data definitions in the form of COBOL copy files, RPG copy files, and C language include files.

The data definition files supplied with MQSeries contain:

- Definitions of all the MQSeries constants and return codes
- Definitions of the MQSeries structures and data types
- Constant definitions for initializing the structures
- Function prototypes for each of the calls (for PL/I and the C language only)

For a full description of MQSeries data definition files, see Appendix G, "MQSeries data definition files" on page 509.

MQSeries stub programs and library files

The stub programs and library files provided are listed here, for each platform.

For more information about how to use stub programs and library files when you build an executable application, seePart 3, "Building an MQSeries application" on page 251 . For information about linking to C++ library files, see the *MQSeries Using C*++ book.

MQSeries for MVS/ESA

Before you can run an MQSeries for MVS/ESA program, you must link-edit it to the stub program supplied with MQSeries for MVS/ESA for the environment in which you are running the application. The stub program provides the first stage of the processing of your calls into requests that MQSeries for MVS/ESA can process.

MQSeries for MVS/ESA supplies five stub programs:

CSQBSTUB	Stub program for MVS batch programs
CSQCSTUB	Stub program for CICS programs
CSQQSTUB	Stub program for IMS programs
CSQXSTUB	Stub program for distributed queuing non CICS exits
CSQASTUB	Stub program for data-conversion exits

Note: The CSQXSTUB program is explained in the *MQSeries Intercommunication* book.

Alternatively, you can dynamically call the stub from within your program. This technique is described in "Dynamically calling the MQSeries stub" on page 270.

In IMS, you may also need to use a special language interface module that is supplied by MQSeries.

MQSeries for AS/400

In MQSeries for AS/400, you must bind your ILE C/400 programs and RPG/400 static calls to the supplied AMQZSTUB service program.

If you are using MQSeries for AS/400 V4R2 you can write your applications in C++. To see how to link your C++ applications, and for full details of all aspects of using C++, see the book *MQSeries Using C*++.

MQSeries for OS/2 Warp

In MQSeries for OS/2 Warp, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system:

MQM.LIB	server for 32-bit C
MQIC.LIB	client for C
MQMXA.LIB	static XA interface for C
MQMCICS.LIB	CICS for OS/2 V2 exits for C
MQMCICS3.LIB	Transaction Server for OS/2 Warp, V4 exits
MQMZF.LIB	installable services exits for C
MQICCB16.LIB	client for 16-bit Micro Focus COBOL
MQMCB16.LIB	server for 16-bit Micro Focus COBOL
MQMCBB.LIB	server for 32-bit IBM VisualAge for COBOL
MQMCB32.LIB	server for 32-bit Micro Focus COBOL
MQICCBB.LIB	client for 32-bit IBM VisualAge for COBOL
MQICCB32.LIB	client for 32-bit Micro Focus COBOL
IMQ*.LIB	server for C++

MQSeries for Windows

In MQSeries for Windows, you must link your program to the MQI library files
supplied for the environment in which you are running your application, in addition
to those provided by the operating system:

MQM16.LIB	server for 16-bit C
MQM.LIB	server for 32-bit C
MQM16.LIB	server for 16-bit Visual Basic
MQMSTD.LIB	server for 32-bit Visual Basic

MQSeries for Windows NT

In MQSeries for Windows NT, you must link your program to the MQI library files
supplied for the environment in which you are running your application, in addition
to those provided by the operating system:

MQM.LIB	server for 32-bit C
MQIC.LIB	client for 16-bit C
MQIC32.LIB	client for 32-bit C
MQMXA.LIB	static XA interface for C
MQMCICS.LIB	CICS for Windows NT V2 exits for C
MQMCICS4.LIB	Transaction Server for Windows NT, V4 exits for C
MQMZF.LIB	installable services exits for C
MQMCBB.LIB	server for 32-bit IBM COBOL
MQMCB32	server for 32-bit Micro Focus COBOL
MQICCBB.LIB	client for 32-bit IBM COBOL
MQICCB32	client for 32-bit Micro Focus COBOL
IMQ*.LIB	server for C++
MQMENC.LIB	dynamic XA interface in C for Encina
MQMTUX.LIB	dynamic XA interface in C for Tuxedo

MQSeries for AIX

In MQSeries for AIX, you must link your program to the MQI library files supplied
for the environment in which you are running your application, in addition to those
provided by the operating system.

In a non-threaded application:

libmqm.a	server for C
libmqic.a	client for C
libmqmzf.a	installable service exits for C
libmqmxa.a	XA interface for C
libmqmcbrt.o	MQSeries runtime library for Micro Focus COBOL support
libmqmcb.a	server for COBOL
libmqicb.a	client for COBOL
libimq*.a	client for C++

In a threaded application:

libmqm_r.a	server for C
libmqmzf_r.a	installable service exits for C
libmqmxa_r.a	XA interface for C
libimq*_r.a	client for C++
libmqmxa_r.a	for Encina

MQSeries for AT&T GIS UNIX V2.2

In MQSeries for AT&T GIS UNIX V2.2, you must link your program to the MQI
library files supplied for the environment in which you are running your application,
in addition to those provided by the operating system.

libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C
libmqmzf.so	installable service exits for C
libmqmxa.a	XA interface for C

MQSeries for Digital OpenVMS

In MQSeries for Digital OpenVMS, you must link your program to the MQI library
files supplied for the environment in which you are running your application, in
addition to those provided by the operating system:

mqm.exe	server for C
mqic.exe	client for C
mqmzf.exe	installable service exits for C
mqmxa.exe	XA interface for C
mqcbrt.exe	MQSeries COBOL runtime
mqmcb.exe	server for COBOL
mqicb.exe	client for COBOL

MQSeries for HP-UX

In MQSeries for HP-UX, you must link your program to the MQI library files
supplied for the environment in which you are running your application, in addition
to those provided by the operating system.

In a non-threaded application:

libmqm.sl	server for C
libmqic.sl	client for C
libmqmzf.sl	installable service exits for C
libmqmxa.sl	XA interface for C
libmqmcbrt.o	MQSeries runtime library for Micro Focus COBOL support
libmqmcb.sl	server for COBOL
libmqicb.sl	client for COBOL

In a threaded application:

libmqm_r.sl	server for C
libmqmzf_r.sl	installable service exits for C
libmqmxa_r.sl	XA interface for C

MQSeries for SINIX and DC/OSx

In MQSeries for SINIX and DC/OSx, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

In a non-threaded application:

libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C
libmqmzf.so	installable service exits for C
libmqmxa.a	XA interface for C
libmqmcbrt.o	MQSeries COBOL runtime
libmqmcb.so	server for COBOL
libmqicb.so	client for COBOL

In a threaded application:

libmqm_r.so	for C
libmqmcs_r.so	for C
libmqmcics_r.so	for CICS
libmqmxa_r.a	for XA interface in C

DOS and Windows 3.1 clients

In DOS and Windows 3.1, you must link your program to the MQIC.LIB library file (or imq*vw.lib for C++), followed by the protocol libraries, indicating the protocol you do and do not want.

mqicn.lib	NetBIOS required
mqicdn.lib	NetBIOS not required
mqict.lib	TCP/IP required
mqicdt.lib	TCP/IP not required

MQSeries for SunOS

In MQSeries for SunOS, you must link your program to the MQI library files supplied for the environment in which you are running your application in addition to those provided by the operating system.

libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C
libmqmzf.so	installable service exits for C
libmqmxa.a	XA interface for C
imq*.so	C++

MQSeries for Sun Solaris

In MQSeries for Sun Solaris, you must link your program to the MQI library files supplied for the environment in which you are running your application in addition to those provided by the operating system.

libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C
libmqmzf.so	installable service exits for C
libmqmxa.a	XA interface for C
imq*.so	C++

MQSeries for Tandem NSK

In MQSeries for Tandem NSK, you must link your program to the MQI library files supplied for the environment in which you are running your application in addition to those provided by the operating system.

mqmlibc	for C, nonnative
mqmlibt	for TAL or COBOL, nonnative
mqmlibnc	for native C
mqmlibnt	for native TAL or COBOL

Parameters common to all the calls

There are two types of parameter common to all the calls—these are handles and return codes.

Using connection and object handles

For a program to communicate with a queue manager, the program must have a unique identifier by which it knows that queue manager. This identifier is called a *connection handle.* For CICS and OS/400 programs, the connection handle is always zero. For all other platforms or styles of programs, the connection handle is returned by the MQCONN or MQCONNX call when the program connects to the queue manager. Programs pass the connection handle as an input parameter when they use the other calls.

For a program to work with an MQSeries object, the program must have a unique identifier by which it knows that object. This identifier is called an *object handle*. The handle is returned by the MQOPEN call when the program opens the object to work with it. Programs pass the object handle as an input parameter when they use subsequent MQPUT, MQGET, MQINQ, MQSET, or MQCLOSE calls.

Understanding return codes

A completion code and a reason code are returned as output parameters by each call. These are known collectively as *return codes*.

To show whether or not a call is successful, each call returns a *completion code* when the call is complete. The completion code is usually either MQCC_OK or MQCC_FAILED, showing success and failure, respectively. Some calls can return an intermediate state, MQCC_WARNING, indicating partial success.

Each call also returns a *reason code* that shows the reason for the failure, or partial success, of the call. There are many reason codes, covering such circumstances

as a queue being full, get operations not being allowed for a queue, and a particular queue not being defined for the queue manager. Programs can use the reason code to decide how to proceed. For example, they could prompt the user of the program to make changes to his input data, then make the call again, or they could return an error message to the user.

When the completion code is MQCC_OK, the reason code is always MQRC_NONE.

The completion and reason codes for each call are listed with the description of that call in the *MQSeries Application Programming Reference* manual. In addition, there is further information (including some ideas for corrective action) for each completion and reason code in the chapter on return codes in those manuals.

Specifying buffers

The queue manager refers to buffers only if they are required. If you do not require a buffer on a call or the buffer is zero in length, you can use a null pointer to a buffer.

Always use datalength when specifying the size of the buffer you require.

When you use a buffer to hold the output from a call (for example, to hold the message data for an MQGET call, or the values of attributes queried by the MQINQ call), the queue manager attempts to return a reason code if the buffer you specify is not valid or is in read-only storage. However, it may not be able to return a reason code in some situations.

Programming language considerations

MQSeries provides support for the following programming languages:

- C.
- C++ (MQSeries for AIX, AS/400, HP-UX, OS/2, Sun Solaris, and Windows NT only). See the *MQSeries Using C++* book for information about coding MQSeries programs in C++.
- Visual Basic (MQSeries for Windows only). See the *MQSeries for Windows Version 2.0 User's Guide* and the *MQSeries for Windows Version 2.1 User's Guide* for information about coding MQSeries programs in Visual Basic.
- COBOL.
- Assembler language (MQSeries for MVS/ESA only).
- RPG (MQSeries for AS/400 only).
- PL/I (MQSeries for MVS/ESA, AIX, OS/2 Warp, and Windows NT only).
- TAL (MQSeries for Tandem NSK only).

The call interface, and how you can code the calls in each of these languages, is described in the *MQSeries Application Programming Reference* manual.

MQSeries provides data definition files to assist you with the writing of your applications. For a full description, see Appendix G, "MQSeries data definition files" on page 509.

If you can choose which language to code your programs in, you should consider the maximum length of the messages that your programs will process. If your programs will process only messages of a known maximum length, you can code them in any of the supported programming languages. But if you do not know the maximum length of the messages the programs will have to process, the language you choose will depend on whether you are writing a CICS, IMS, or batch application:

IMS and batch
> Code the programs in C, PL/I, or assembler language to use the facilities these languages offer for obtaining and releasing arbitrary amounts of memory. Alternatively, you could code your programs in COBOL, but use assembler language, PL/I, or C subroutines to get and release storage.

CICS
> Code the programs in any language supported by CICS. The EXEC CICS interface provides the GETMAIN and FREEMAIN calls for managing memory, if necessary.

Coding in C

See Appendix A, "Language compilers and assemblers" on page 433 for the compilers that you can use to process your C programs.

Note the information in the following sections when coding MQSeries programs in C.

Parameters of the MQI calls

Parameters that are *input-only* and of type MQHCONN, MQHOBJ, or MQLONG are passed by value; for all other parameters, the *address* of the parameter is passed by value.

Not all parameters that are passed by address need to be specified every time a function is invoked. Where a particular parameter is not required, a null pointer can be specified as the parameter on the function invocation, in place of the address of the parameter data. Parameters for which this is possible are identified in the call descriptions.

No parameter is returned as the value of the function; in C terminology, this means that all functions return void.

The attributes of the function are defined by the MQENTRY macro variable; the value of this macro variable depends on the environment.

Parameters with undefined data type

The MQGET, MQPUT, and MQPUT1 functions each have one parameter that has an undefined data type, namely the *Buffer* parameter. This parameter is used to send and receive the application's message data.

Parameters of this sort are shown in the C examples as arrays of MQBYTE. It is valid to declare the parameters in this way, but it is usually more convenient to declare them as the particular structure that describes the layout of the data in the message. The function parameter is declared as a pointer-to-void, and so the address of any sort of data can be specified as the parameter on the function invocation.

Data types

All data types are defined by means of the `typedef` statement. For each data type, the corresponding pointer data type is also defined. The name of the pointer data type is the name of the elementary or structure data type prefixed with the letter "P" to denote a pointer. The attributes of the pointer are defined by the MQPOINTER macro variable; the value of this macro variable depends on the environment. The following illustrates how pointer data types are declared:

```
#define MQPOINTER                     /* depends on environment */
...
typedef MQLONG  MQPOINTER PMQLONG;  /* pointer to MQLONG */
typedef MQMD    MQPOINTER PMQMD;    /* pointer to MQMD */
```

Manipulating binary strings

Strings of binary data are declared as one of the MQBYTEn data types. Whenever you copy, compare, or set fields of this type, use the C functions `memcpy`, `memcmp`, or `memset`:

```
#include <string.h>
#include "cmqc.h"

MQMD MyMsgDesc;

memcpy(MyMsgDesc.MsgId,          /* set "MsgId" field to nulls    */
       MQMI_NONE,                /* ...using named constant       */
       sizeof(MyMsgDesc.MsgId));

memset(MyMsgDesc.CorrelId,       /* set "CorrelId" field to nulls */
       0x00,                     /* ...using a different method   */
       sizeof(MQBYTE24));
```

Do not use the string functions strcpy, strcmp, strncpy, or strncmp because these do not work correctly with data declared as MQBYTE24.

Manipulating character strings

When the queue manager returns character data to the application, the queue manager always pads the character data with blanks to the defined length of the field. The queue manager *does not* return null-terminated strings, but you can use them in your input. Therefore, when copying, comparing, or concatenating such strings, use the string functions strncpy, strncmp, or strncat.

Do not use the string functions that require the string to be terminated by a null (strcpy, strcmp, and strcat). Also, do not use the function strlen to determine the length of the string; use instead the sizeof function to determine the length of the field.

Initial values for structures

The include file <cmqc.h> defines various macro variables that may be used to provide initial values for the structures when instances of those structures are declared. These macro variables have names of the form MQxxx_DEFAULT, where MQxxx represents the name of the structure. Use them like this:

```
MQMD   MyMsgDesc = {MQMD_DEFAULT};
MQPMO  MyPutOpts = {MQPMO_DEFAULT};
```

For some character fields, the MQI defines particular values that are valid (for example, for the *StrucId* fields or for the *Format* field in MQMD). For each of the valid values, two macro variables are provided:

- One macro variable defines the value as a string whose length, excluding the implied null, matches exactly the defined length of the field. For example, (the symbol ƀ represents a blank character):

```
#define MQMD_STRUC_ID "MDƀƀ"
#define MQFMT_STRING "MQSTRƀƀƀ"
```

Use this form with the memcpy and memcmp functions.

- The other macro variable defines the value as an array of char; the name of this macro variable is the name of the string form suffixed with "_ARRAY". For example:

```
#define MQMD_STRUC_ID_ARRAY 'M','D','ƀ','ƀ'
#define MQFMT_STRING_ARRAY 'M','Q','S','T','R','ƀ','ƀ','ƀ'
```

Use this form to initialize the field when an instance of the structure is declared with values different from those provided by the MQMD_DEFAULT macro variable.

Initial values for dynamic structures

When a variable number of instances of a structure are required, the instances are usually created in main storage obtained dynamically using the calloc or malloc functions. To initialize the fields in such structures, the following technique is recommended:

1. Declare an instance of the structure using the appropriate MQxxx_DEFAULT macro variable to initialize the structure. This instance becomes the "model" for other instances:

```
MQMD ModelMsgDesc = {MQMD_DEFAULT};
                                    /* declare model instance */
```

The static or auto keywords can be coded on the declaration in order to give the model instance static or dynamic lifetime, as required.

2. Use the calloc or malloc functions to obtain storage for a dynamic instance of the structure:

```
PMQMD InstancePtr;
InstancePtr = malloc(sizeof(MQMD));
                                    /* get storage for dynamic instance */
```

3. Use the memcpy function to copy the model instance to the dynamic instance:

```
memcpy(InstancePtr,&ModelMsgDesc,sizeof(MQMD));
                                    /* initialize dynamic instance */
```

Use from C++

For the C++ programming language, the header files contain the following additional statements that are included only when a C++ compiler is used:

```
#ifdef __cplusplus
  extern "C" {
#endif

/* rest of header file */

#ifdef __cplusplus
  }
#endif
```

Coding in COBOL

See Appendix A, "Language compilers and assemblers" on page 433 for the compilers that you can use to process your COBOL programs.

Note the information in the following sections when coding MQSeries programs in COBOL.

Named constants

In this book, the names of constants are shown containing the underscore character (_) as part of the name. In COBOL, you must use the hyphen character (-) in place of the underscore.

Constants that have character-string values use the single quotation mark character (') as the string delimiter. To make the compiler accept this character, use the compiler option APOST.

The copy file CMQV contains declarations of the named constants as level-10 items. To use the constants, declare the level-01 item explicitly, then use the COPY statement to copy in the declarations of the constants:

```
 WORKING-STORAGE SECTION.
 01  MQM-CONSTANTS.
     COPY CMQV.
```

However, this method causes the constants to occupy storage in the program even if they are not referred to. If the constants are included in many separate programs within the same run unit, multiple copies of the constants will exist—this may result in a significant amount of main storage being used. You can avoid this situation by adding the GLOBAL clause to the level-01 declaration:

```
* Declare a global structure to hold the constants
 01  MQM-CONSTANTS GLOBAL.
     COPY CMQV.
```

This causes storage to be allocated for only *one* set of constants within the run unit; the constants, however, can be referred to by *any* program within the run unit, not just the program that contains the level-01 declaration.

Coding in System/390 assembler language

This is supported on MVS/ESA only.

See Appendix A, "Language compilers and assemblers" on page 433 for the assemblers that you can use to process your assembler-language programs.

Note the information in the following sections when coding MQSeries for MVS/ESA programs in assembler language.

Names
In this book, the names of parameters in the descriptions of calls, and the names of fields in the descriptions of structures are shown in mixed case. In the assembler-language macros supplied with MQSeries, all names are in uppercase.

Using the MQI calls
The MQI is a call interface, so assembler-language programs must observe the OS linkage convention. In particular, before they issue an MQI call, assembler-language programs must point register R13 at a save area of at least 18 full words. This save area is to provide storage for the called program. It stores the registers of the caller before their contents are destroyed, and restores the contents of the caller's registers on return.

Note: This is of particular importance for CICS assembler-language programs that use the DFHEIENT macro to set up their dynamic storage, but that choose to override the default DATAREG from R13 to other registers. When the CICS Resource Manager Interface receives control from the stub, it saves the current contents of the registers at the address to which R13 is pointing. Failing to reserve a proper save area for this purpose gives unpredictable results, and will probably cause an abend in CICS.

Declaring constants
Most constants are declared as equates in macro CMQA. However, the following constants cannot be defined as equates, and these are not included when you call the macro using default options:

    MQACT_NONE
    MQCI_NONE
    MQFMT_NONE
    MQFMT_ADMIN
    MQFMT_COMMAND_1
    MQFMT_COMMAND_2
    MQFMT_DEAD_LETTER_HEADER
    MQFMT_EVENT
    MQFMT_IMS
    MQFMT_IMS_VAR_STRING
    MQFMT_PCF
    MQFMT_STRING
    MQFMT_TRIGGER
    MQFMT_XMIT_Q_HEADER
    MQMI_NONE

To include them, add the keyword EQUONLY=NO when you call the macro.

CMQA is protected against multiple declaration, so you can include it many times. However, the keyword EQUONLY takes effect only the first time the macro is included.

Specifying the name of a structure

To allow more than one instance of a structure to be declared, the macro that generates the structure prefixes the name of each field with a user-specifiable string and an underscore character (_). Specify the string when you invoke the macro. If you do not specify a string, the macro uses the name of the structure to construct the prefix:

```
* Declare two object descriptors
         CMQODA                  Prefix used="MQOD_" (the default)
MY_MQOD  CMQODA                  Prefix used="MY_MQOD_"
```

The structure declarations in the *MQSeries Application Programming Reference* manual show the default prefix.

Specifying the form of a structure

The macros can generate structure declarations in one of two forms, controlled by the DSECT parameter:

DSECT=YES   An assembler-language DSECT instruction is used to start a new data section; the structure definition immediately follows the DSECT statement. No storage is allocated, so no initialization is possible. The label on the macro invocation is used as the name of the data section; if no label is specified, the name of the structure is used.

DSECT=NO   Assembler-language DC instructions are used to define the structure at the current position in the routine. The fields are initialized with values, which you can specify by coding the relevant parameters on the macro invocation. Fields for which no values are specified on the macro invocation are initialized with default values.

DSECT=NO is assumed if the DSECT parameter is not specified.

Controlling the listing

You can control the appearance of the structure declaration in the assembler-language listing by means of the LIST parameter:

LIST=YES   The structure declaration appears in the assembler-language listing.

LIST=NO   The structure declaration does not appear in the assembler-language listing. This is assumed if the LIST parameter is not specified.

Specifying initial values for fields

You can specify the value to be used to initialize a field in a structure by coding the name of that field (without the prefix) as a parameter on the macro invocation, accompanied by the value required.

For example, to declare a message descriptor structure with the *MsgType* field initialized with MQMT_REQUEST, and the *ReplyToQ* field initialized with the string MY_REPLY_TO_QUEUE, you could use the following code:

```
MY_MQMD        CMQMDA          MSGTYPE=MQMT_REQUEST,             X
               REPLYTOQ=MY_REPLY_TO_QUEUE
```

If you specify a named constant (or equate) as a value on the macro invocation, you must use the CMQA macro to define the named constant. You must not enclose in single quotation marks (' ') values that are character strings.

Writing reenterable programs

MQSeries uses its structures for both input and output. If you want your program to remain reenterable, you should:

1. Define working storage versions of the structures as DSECTs, or define the structures inline within an already-defined DSECT. Then copy the DSECT to storage that is obtained using:

   - For batch and TSO programs, the STORAGE or GETMAIN MVS assembler macros

   - For CICS, the working storage DSECT (DFHEISTG) or the EXEC CICS GETMAIN command

   To correctly initialize these working storage structures, copy a constant version of the corresponding structure to the working storage version.

   Note: The MQMD and MQXQH structures are each more than 256 bytes long. To copy these structures to storage, you will have to use the MVCL assembler instruction.

2. Reserve space in storage by using the LIST form (`MF=L`) of the CALL macro. When you use the CALL macro to make an MQI call, use the EXECUTE form (`MF=E`) of the macro, using the storage reserved earlier, as shown in the example under "Using CEDF." For more examples of how to do this, see the assembler language sample programs as shipped with MQSeries.

Use the assembler language RENT option to help you determine if your program is reenterable.

For information on writing reenterable programs, see the *MVS/ESA Application Development Guide: Assembler Language Programs*, GC28-1644.

Using CEDF

If you want to use the CICS Execution Diagnostic Facility (CEDF) supplied transaction to help you to debug your program, you must add the `,VL` keyword to each `CALL` statement, for example:

```
CALL MQCONN,(NAME,HCONN,COMPCODE,REASON),MF=(E,PARMAREA),VL
```

The above example is reenterable assembler-language code where `PARMAREA` is an area in the working storage you specified.

Coding in RPG

RPG is supported on AS/400 only.

See Appendix A, "Language compilers and assemblers" on page 433 for the compilers that you can use to process your RPG programs.

Note the following when coding MQSeries for AS/400 programs in RPG:

- In this book, the MQI calls are described using their individual names. In RPG, you must make all calls to the single name QMQM, and specify the particular function you require by coding an additional parameter that precedes the normal parameters for that call. The names you use for this additional parameter are shown in Appendix B, "MQI names in RPG" on page 437.

- The structure declarations supplied with MQSeries for AS/400 do not contain DS statements. This allows you to declare either a single data structure or a multiple-occurrence data structure. To do the latter, use the DS statement followed by the /COPY statement to copy in the remainder of the declaration. For example:

```
I*..1....:....2....:....3....:....4....:....5....:....6....:....7
I* Declare an MQMD data structure with 5 occurrences
IMYMD         DS                          5
I/COPY CMQMDR
```

- The *buffer* field, as declared in the AMQ1GET4 sample program, has a maximum size limit of 256 bytes. This limit can be increased to:

   - 9999 bytes by declaring *buffer* as a structure
   - 32 KB if *buffer* is declared as a physical file

   This is explained in the *MQSeries for AS/400 Application Programming Reference (RPG)* manual for the MQPUT, MQPUT1, and MQGET calls.

In this book, the parameters of calls, the names of data types, the fields of structures, and the names of constants are described using their long names. In RPG, these names are abbreviated to six or fewer uppercase characters. For example, the field *MsgType* becomes *MDMT* in RPG. The short RPG names for these names are shown in Appendix B, "MQI names in RPG" on page 437.

Coding in PL/I

This is supported on AIX, MVS/ESA, OS/2 Warp, and Windows NT only.

See Appendix A, "Language compilers and assemblers" on page 433 for the compilers that you can use to process your PL/I programs.

Note the information in the following sections when coding MQSeries for MVS/ESA programs in PL/I.

Structures

Structures are declared with the BASED attribute, and so do not occupy any storage unless the program declares one or more instances of a structure.

An instance of a structure can be declared by using the `like` attribute, for example:

```
dcl my_mqmd       like MQMD;  /* one instance */
dcl my_other_mqmd like MQMD;  /* another one */
```

The structure fields are declared with the INITIAL attribute; when the `like` attribute is used to declare an instance of a structure, that instance inherits the initial values defined for that structure. Thus it is necessary to set only those fields where the value required is different from the initial value.

PL/I is not sensitive to case, and so the names of calls, structure fields, and constants can be coded in lowercase, uppercase, or mixed case.

Named constants

The named constants are declared as macro variables; as a result, named constants which are not referenced by the program do not occupy any storage in the compiled procedure. However, the compiler option which causes the source to be processed by the macro preprocessor must be specified when the program is compiled.

All of the macro variables are character variables, even the ones which represent numeric values. Although this may seem counter intuitive, it does not result in any data-type conflict after the macro variables have been substituted by the macro processor, for example:

```
%dcl MQMD_STRUC_ID char;
%MQMD_STRUC_ID = '''MD '''; 

%dcl MQMD_VERSION_1 char;
%MQMD_VERSION_1 = '1';
```

Coding in TAL

TAL is supported on MQSeries for Tandem NSK only.

See Appendix A, "Language compilers and assemblers" on page 433 for the compilers that you can use to process your TAL programs.

Note the following when coding MQSeries for Tandem NSK programs in TAL:

- The MQI library (bound into the application process) does not open $RECEIVE and does not open $TMP (TM/MP transaction pseudo-file) itself, so you may code your application to use these features.

- The MQI library uses a SERVERCLASS_SEND_() call in initial communication with the Queue Manager. While connected, it maintains two process file opens (with the LINKMON process and a Local Queue Manager Agent) and a small number of disk file opens (fewer than 10).

MVS batch considerations

MVS batch programs that call the MQI can be in either supervisor or problem state. However, they must meet the following conditions:

- They must be in task mode, not service request block (SRB) mode.

- They must be in Primary address space control (ASC) mode (not Access Register ASC mode).

- They must not be in cross-memory mode. The primary address space number (ASN) must be equal to the secondary ASN and the home ASN.

- No MVS locks can be held.

- There can be no function recovery routines (FRRs) on the FRR stack.

- Any program status word (PSW) key can be in force for the MQCONN call (provided the key is compatible with using storage that is in the TCB key), but subsequent calls that use the connection handle returned by MQCONN:

  – Must have the same PSW key that was used on the MQCONN call

– Must have parameters accessible (for write, where appropriate) under the same PSW key

– Must be issued under the same task (TCB), but not in any subtask of the task

• They can be in either 24-bit or 31-bit addressing mode. However, if 24-bit addressing mode is in force, parameter addresses must be interpreted as valid 31-bit addresses.

If any of these conditions is not met, a program check may occur. In some cases the call will fail and a reason code will be returned.

UNIX signal handling

Each MQI function sets up its own signal handler for the signals:

SIGSEGV
SIGBUS
SIGALRM

Users' handlers for these are suspended for the duration of the call. Other signals can be caught in the normal way by user-written handlers. If you do not install a handler, the default actions (for example, ignore, core dump, or exit) are left in place.

If you are writing multi-threaded applications (for MQSeries for AIX, MQSeries for HP-UX, MQSeries for SINIX and DC/OSx, and MQSeries for Sun Solaris), you should note that there is only one process-wide signal handler for each signal. MQSeries alters this signal handler appropriately when it makes calls to MQSeries. If an application thread alters the signal handlers for the above signals while another thread is making a call to MQSeries, the new handler may pass signals inside the MQSeries code. Therefore, application threads should ***not*** establish signal handlers for signals while there is any possibility that another thread of the same process is also making an MQSeries call.

If an MQSeries signal handler gets a signal for a thread that is not currently in MQSeries code, it attempts to find the handler established by the application (before it called MQSeries) and pass the signal to that handler. However, it may not be possible to pass all aspects of the call to the signal handler.

While you are in a signal handler, you cannot call an MQI function. If you do, while another MQI function is active, MQRC_CALL_IN_PROGRESS is returned. If you do, while no other MQI function is active, it is likely to fail because of the operating system restrictions on which calls can be issued from within a handler.

In the case of C++ destructor methods, which may be called automatically during program exit, you may not be able to stop the MQI functions from being called. Therefore, ignore any errors about MQRC_CALL_IN_PROGRESS. If a signal handler calls exit(), MQSeries backs out uncommitted messages in syncpoint as normal and closes any open queues.

MQI functions do not return the code EINTR or any equivalent to application programs. If a signal occurs during an MQI call, and the handler calls 'return', the call continues to run as if the signal had not happened. In particular, MQGET cannot be interrupted by a signal to return control immediately to the application. If

you want to break out of an MQGET, set the queue to GET_DISABLED; alternatively, use a loop around a call to MQGET with a finite time expiry (MQGMO_WAIT_INTERVAL), and use your signal handler to set a flag which breaks the loop.

Chapter 7. Connecting and disconnecting a queue manager

To use MQSeries programming services, a program must have a connection to a queue manager. The way this connection is made depends on the platform and the environment in which the program is operating:

MVS batch, MQSeries on UNIX systems, MQSeries for Digital OpenVMS, MQSeries for OS/2 Warp, MQSeries for Tandem NSK, MQSeries for Windows, and MQSeries for Windows NT

Programs that run in these environments can use the MQCONN MQI call to connect to, and the MQDISC call to disconnect from, a queue manager. Alternatively, MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT can use the MQCONNX call. This chapter describes how writers of such programs should use these calls.

IMS   The IMS control region is connected to one or more queue managers when it starts. This connection is controlled by IMS commands. (For information on how to control the IMS adapter of MQSeries for MVS/ESA, see the *MQSeries for MVS/ESA System Management Guide*.) However, writers of message queuing IMS programs must use the MQCONN MQI call to specify the queue manager to which they want to connect. They can use the MQDISC call to disconnect from that queue manager. This chapter describes how writers of such programs should use these calls. Before the IMS adapter processes a message for another user following a Get Unique call from the IOPCB, or one implied by a checkpoint call, the adapter ensures that the application closes handles and disconnects from the queue manager.

CICS for MVS/ESA and CICS/MVS

CICS for MVS/ESA and CICS/MVS programs do not need to do any work to connect to a queue manager because the CICS system itself is connected. This connection is usually made automatically at initialization, but you can also use the CKQC transaction, which is supplied with MQSeries for MVS/ESA. CKQC is discussed in the *MQSeries for MVS/ESA System Management Guide*.

Note: CICS for MVS/ESA and CICS/MVS programs can also use the MQI connect and disconnect calls (MQCONN and MQDISC). You may want to do this so that you can port these applications to non-CICS environments with a minimum of recoding. Be warned, though, that these calls *always* complete successfully in a CICS environment. This means that the return code may not reflect the true state of the connection to the queue manager.

CICS for Windows NT and Open Systems

These programs do not need to do any work to connect to a queue manager because the CICS system itself is connected. Therefore, only one connection at a time is supported. CICS applications must issue an MQCONN call to obtain a connection handle, and should issue an MQDISC call before they exit.

AS/400   The connection to a queue manager can be made explicitly using the MQCONN call, or it can be made implicitly in the first successful MQOPEN or MQPUT1 call. If you want to use your application on other platforms as well as AS/400, it is a good idea to use the MQCONN call.

If you do use the MQCONN call, you must also use the MQDISC call. The performance of your application is better if you connect and disconnect (either explicitly or implicitly) as infrequently as possible.

If you use the MQPUT1 call and rely on an implicit connection, your application is implicitly disconnected when the call completes. When an application connects implicitly using an MQOPEN call, it is disconnected implicitly when it issues the last MQCLOSE call.

Connecting to a queue manager using the MQCONN call

In general, you can connect either to a specific queue manager, or to the default queue manager. In both cases, the queue manager you connect to must be *local* to the task. This means that it must belong to the same system. For MQSeries for MVS/ESA, in the batch environment, the default queue manager is specified in the CSQBDEFV module. For MQSeries for OS/2 Warp, MQSeries for Windows NT, MQSeries for Digital OpenVMS, MQSeries for Tandem NSK, and MQSeries on UNIX systems, it is specified in the mqs.ini file. MQSeries for Windows allows only one queue manager to run at a time; it uses the running queue manager as its default.

In the IMS environment, the queue manager must be connected to the IMS control region and to the dependent region that the program uses. The default queue manager is specified in the CSQQDEFV module when MQSeries for MVS/ESA is installed.

With the CICS on Open Systems environment, and Transaction Server for Windows NT and AIX, the queue manager must be defined as an XA resource to CICS.

On AS/400, there can be only one queue manager on the same system as an application.

As input to MQCONN, you must supply a queue manager name.

To connect to the default queue manager, call MQCONN, specifying a name consisting entirely of blanks or starting with a null (X'00') character.

Within MQSeries on UNIX systems, an application must be authorized for it to successfully connect to a queue manager. For more information, see the *MQSeries System Administration* Guide for MQSeries for AIX, HP-UX, and Sun Solaris; for other platforms, see the appropriate *System Management Guide*.

The output from MQCONN is:

- A connection handle
- A completion code
- A reason code

You will need to use the connection handle on subsequent MQI calls.

If the reason code indicates that the application is already connected to that queue manager, the connection handle that is returned is the same as the one that was returned when the application first connected. So the application probably should not issue the MQDISC call in this situation because the calling application will expect to remain connected.

The scope of the connection handle is the same as that for the object handle (see "Opening objects using the MQOPEN call" on page 94).

Descriptions of the parameters are given in the description of the MQCONN call in the *MQSeries Application Programming Reference* manual.

The MQCONN call fails if the queue manager is in a quiescing state when you issue the call, or if the queue manager is shutting down.

Scope of MQCONN

Within MQSeries on UNIX systems, MQSeries for Digital OpenVMS, MQSeries for OS/2 Warp, MQSeries for Windows, and MQSeries for Windows NT, the scope of an MQCONN call is the thread that issued it. That is, the connection handle returned from an MQCONN call is valid only within the thread which issued the call. Only one call may be made at any one time using the handle. If it is used from a different thread, it will be rejected as invalid. If you have multiple threads in your application that each wish to use MQSeries calls, then they must each individually issue MQCONN.

Each thread can connect to a different queue manager on OS/2 and Windows NT, but not on UNIX.

If your application is running as a client, it may connect to more than one queue manager within a thread. This does not apply if your application is not running as a client.

OS/2 has a limit of 4095 active threads in a system. However, the default is 64. This value may be controlled by the THREADS=xxxx parameter in CONFIG.SYS. Limitations on the number of concurrent MQCONN calls that can be made within a system are dependent on this value, although other factors to consider are disk space availability for the swapper.dat file and shared memory availability.

On MQSeries for Windows, the scope of an MQCONN call is the application process.

Connecting to a queue manager using the MQCONNX call

| *This is not supported on AS/400, MVS/ESA, MQSeries for Tandem NSK, and Digital OpenVMS*.

The MQCONNX call is similar to the MQCONN call, but includes options to control the way that the call actually works.

As input to MQCONNX, you must supply a queue manager name. The output from MQCONNX is:

- A connection handle
- A completion code
- A reason code

You will need to use the connection handle on subsequent MQI calls.

A description of all of the parameters of MQCONNX is given in the *MQSeries Application Programming Reference* manual. The `Options` field allows you to set STANDARD_BINDING or FASTPATH_BINDING:

MQCNO_STANDARD_BINDING

By default, MQCONNX (like MQCONN) implies two threads where the MQSeries application and the local queue manager agent run in separate processes. The MQSeries application performs the MQSeries operation and the local queue manager agent performs the application operation. This is defined by the MQCNO_STANDARD_BINDING option on the MQCONNX call.

Note: This default maintains the integrity of the queue manager (that is, it makes the queue manager immune to errant programs), but impairs the performance of the MQI calls.

MQCNO_FASTPATH_BINDING

Trusted applications imply that the MQSeries application and the local queue manager agent become the same process. Since the agent process no longer needs to use an interface to access the queue manager, these applications become an extension of the queue manager. This is defined by the MQCNO_FASTPATH_BINDING option on the MQCONNX call.

You need to link trusted applications to the threaded MQSeries libraries. For instructions on how to set up an MQSeries application to run as trusted, see the *MQSeries Application Programming Reference* manual.

Note: This option compromises the integrity of the queue manager as there is no protection from overwriting its storage. This also applies if the application contains errors which can be exposed to messages and other data in the queue manager too. These issues must be considered before using this option.

Restrictions

The following restrictions apply to trusted applications:

- On MQSeries on UNIX systems, it is necessary to use mqm as the effective userID and groupID for all MQI calls. You may change these IDs before making a non-MQI call requiring authentication (for example, opening a file), but you **must** change it back to mqm before making the next MQI call.

- On MQSeries on UNIX systems, trusted applications must run in threaded processes but only one thread can be connected at a time.

- On MQSeries for OS/2 Warp and MQSeries for Windows NT, a thread within a trusted application cannot connect to a queue manager while another thread in the same process is connected to a different queue manager.

- You must explicitly disconnect trusted applications from the queue manager.

- You must stop trusted applications before ending the queue manager with the endmqm command.

- You must not use asynchronous signals and timer interrupts (such as `sigkill`) with MQCNO_FASTPATH_BINDING.

- On MQSeries for AIX, trusted applications cannot be compiled using the PL/I programming language.

- On MQSeries for AIX, there are restrictions on the use of shared memory segments:

MQSeries uses a single "shmat()" command to connect to shared memory resources. However, on AIX, one process cannot attach to more than 10 memory segments.

MQSeries uses two additional shared memory segments for trusted applications, reducing the amount of shared storage available. Therefore, it is important that your applications do not connect to too many shared segments, causing a failure in the application code.

Here is a breakdown of the memory segments:

Segment	Use	
0	Reserved for AIX	
1	Reserved for AIX	
2	Stack and heap	
3	CICS	MQSeries (trusted applications only)
4	DB2 and DT/6000	
5		
6		
7		MQSeries (trusted applications only)
8		MQSeries
9	CICS	
A	CICS	
B	CICS	
C	DB2	
D	Reserved for AIX	
E	Reserved for AIX	
F	Reserved for AIX	

This also implies that trusted applications cannot use the maxdata binder option to specify a greater user data area: this conflicts with the queue manager use of shared memory within the application process as it causes the program data to be placed in shared memory segment 3.

Environment variable

On MQSeries for OS/2 Warp, MQSeries for Windows NT, and MQSeries on UNIX systems, the environment variable, MQ_CONNECT_TYPE, can be used in combination with the type of binding specified in the *Options* field. This environment variable allows you to execute the application with the STANDARD_BINDING if any problems occur with the FASTPATH_BINDING. If the environment variable is specified, it should have the value FASTPATH or STANDARD to select the type of binding required. However, the FASTPATH binding is used only if the connect option is appropriately specified as shown in Table 1:

Table 1. Environment variable

MQCONNX	Environment variable	Result
STANDARD	UNDEFINED	STANDARD
FASTPATH	UNDEFINED	FASTPATH
STANDARD	STANDARD	STANDARD
FASTPATH	STANDARD	STANDARD
STANDARD	FASTPATH	STANDARD
FASTPATH	FASTPATH	FASTPATH

So, to run a trusted application, either:

1. Specify the MQCNO_FASTPATH_BINDING option on the MQCONNX call and the FASTPATH environment variable,

or

2. Specify the MQCNO_FASTPATH_BINDING option on the MQCONNX call and leave the environment variable undefined.

If neither MQCNO_STANDARD_BINDING nor MQCNO_FASTPATH_BINDING is specified, you can use MQCNO_NONE, which defaults to MQCNO_STANDARD_BINDING.

Disconnecting programs from a queue manager using MQDISC

When a program that has connected to a queue manager using the MQCONN call has finished all interaction with the queue manager, it must break the connection using the MQDISC call.

However, in MQSeries for AS/400, if the connection was made implicitly, the MQDISC call is optional. Also, on MVS/ESA for CICS applications, the call is optional.

After MQDISC is called, the connection handle (*Hconn*) is no longer valid, and you cannot issue any further MQI calls until you call MQCONN again. MQDISC does an implicit MQCLOSE for any objects that are still open using this handle.

In MQSeries for AS/400, when you sign off from the operating system, an implicit MQDISC call is made.

As input to the MQDISC call, you must supply the connection handle (*Hconn*) that was returned by MQCONN when you connected to the queue manager.

The output from this call is a completion code and a reason code, with the connection handle set to the value MQHC_UNUSABLE_HCONN.

Descriptions of the parameters are given in the description of the MQDISC call in the *MQSeries Application Programming Reference* manual.

Authority checking

The MQCLOSE and MQDISC calls usually perform no authority checking. In the normal course of events a job which has the authority to open or connect to an MQSeries object will close or disconnect from that object. Even if the authority of a job that has connected to, or opened an MQSeries object is revoked, the MQCLOSE and MQDISC calls are accepted.

Chapter 8. Opening and closing objects

To perform any of the following operations, you must first *open* the relevant
MQSeries object:

- Put messages on a queue
- Get (browse or retrieve) messages from a queue
- Set the attributes of an object
- Inquire about the attributes of any object

Use the MQOPEN call to open the object, using the options of the call to specify
what you want to do with the object. The only exception is if you want to put a
single message on a queue, then close the queue immediately. In this case, you
can bypass the "opening" stage by using the MQPUT1 call (see "Putting one
message on a queue using the MQPUT1 call" on page 110).

Before you open an object using the MQOPEN call, you must connect your
program to a queue manager. This is explained in detail, for all environments, in
Chapter 7, "Connecting and disconnecting a queue manager" on page 87.

There are four types of MQSeries object that can be opened:

- Queue
- Namelist (MQSeries for MVS/ESA only)
- Process definition
- Queue manager

You open all of these objects in a similar way using the MQOPEN call. For more
information about MQSeries objects, see Chapter 4, "MQSeries objects" on
page 39.

You can open the same object more than once, and each time you get a new
object handle. You might want to browse messages on a queue using one handle,
and remove messages from the same queue using another handle. This saves
using up resources to close and reopen the same object. You can also open a
queue for browsing **and** removing messages at the same time.

Moreover, you can open multiple objects with a single MQOPEN and close them
using MQCLOSE. See "Distribution lists" on page 111 for information about how to
do this.

When you attempt to open an object, the queue manager checks that you are
authorized to open that object for the options you specify in the MQOPEN call.

Objects are closed automatically when a program disconnects from the queue
manager. In the IMS environment, disconnection is forced when a program starts
processing for a new user following a GU (get unique) IMS call. On the AS/400
platform, objects are closed automatically when a job ends.

It is good programming practice to close objects you have opened. Use the
MQCLOSE call to do this.

The remaining sections of this chapter describe how to:

- Open objects using the MQOPEN call
- Create dynamic queues using model queues (MQSeries for AS/400, MQSeries for MVS/ESA, MQSeries for OS/2 Warp, MQSeries for Tandem NSK, MQSeries for Windows NT, and MQSeries on UNIX systems only)
- Open remote queues
- Close objects using the MQCLOSE call

Opening objects using the MQOPEN call

As input to the MQOPEN call, you must supply:

- A connection handle. For CICS and AS/400 applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN call. For other programs, always use the connection handle returned by the MQCONN call.
- A description of the object you want to open, using the object descriptor structure (MQOD).
- One or more options that control the action of the call.

The output from MQOPEN is:

- An object handle that represents your access to the object. Use this on input to any subsequent MQI calls.
- A modified object-descriptor structure, if you are creating a dynamic queue (and it is supported on your platform).
- A completion code.
- A reason code.

Scope of an object handle

The scope of an object handle is the same as the scope of a connection handle, however there are variations between platforms:

CICS
: In a CICS program, you can use the handle only within the same CICS task from which you made the MQOPEN call.

IMS and MVS batch
: In the IMS and batch environments, you can use the handle within the same task, but not within any subtasks.

AS/400
: In an AS/400 program, you can use the handle only within the same job from which you made the MQOPEN call.

MQSeries for OS/2 Warp
: In the MQSeries for OS/2 Warp environment, you can use the same handle within the same thread.

MQSeries for Windows NT
: In the MQSeries for Windows NT environment, you can use the same handle within the same thread.

MQSeries on UNIX systems
: In these environments, you can use the same handle within the same thread.

DOS   In the DOS environment, there are no restrictions on where you can use the handle.

Windows 3.1   In the Windows 3.1 environment, you can use the handle in the same Windows 3.1 instance.

Descriptions of the parameters of the MQOPEN call are given in the *MQSeries Application Programming Reference* manual.

The following sections describe the information you must supply as input to MQOPEN.

Identifying objects (the MQOD structure)

Use the MQOD structure to identify the object you want to open. This structure is an input parameter for the MQOPEN call. (The structure is modified by the queue manager when you use the MQOPEN call to create a dynamic queue.) The following sections give you help on filling in the fields of MQOD. There is a description of the structure in the *MQSeries Application Programming Reference* manual.

StrucId

This identifies the structure as an object descriptor. It is a 4-character field. Always specify MQOD_STRUC_ID.

Version

This describes the version number of the structure. MQOD_VERSION_1 can open only a single queue, whereas MQOD_VERSION_2 can open one or more queues (thus creating a distribution list). The default is MQOD_VERSION_1.

See "Distribution lists" on page 111 for a description of distribution lists.

ObjectType

This specifies the type of the object you want to open.

When you use the MQOPEN call to create a distribution list, **do not** use a model queue in the *ObjectType* field; this fails with reason code MQRC_Q_TYPE_ERROR.

ObjectName

This identifies the name of the object to open.

This supplies the name of a model queue in the *ObjectName* field when you are using the MQOPEN call to create a dynamic queue. The name of the queue created will be returned to you in this field. For more information, see "Creating dynamic queues" on page 100.

When you use the MQOPEN call to create a distribution list, the *ObjectName* field must be blank.

When opening the queue manager object (which you would do only if you want to inquire about its attributes), specify blanks or a leading null (X'00') character in the *ObjectName* field.

ObjectQMgrName

This identifies the name of the queue manager on which the object you want to open is defined. To identify the local queue manager, specify blanks or a leading null (X'00') character.

When you use the MQOPEN call to create a distribution list, the *ObjectQMgrName* field must be blank.

DynamicQName
>   This specifies the name that you want to give if you are filling in the MQOD
>   structure to create a dynamic queue (see "Creating dynamic queues" on
>   page 100).
>
>   This field is ignored if the name you put in *ObjectName* is not a model queue.

AlternateUserId
>   This nominates an alternate user identifier that is to be used to test for authority
>   to open the object. "Alternate" means a user identifier other than the one under
>   which the application is running.
>
>   The use of this field is discussed in "MQOPEN option for alternate user
>   authority" on page 100.

For information about using the MQOD structure for distribution lists, see Using the
MQOD structure under "Distribution lists" on page 111.

Name resolution

When you open an MQSeries queue, the MQOPEN call performs a name
resolution function on the queue name you specify. This determines on which
queue the queue manager performs subsequent operations. This means that when
you specify the name of an alias queue or a remote queue in your object descriptor
(MQOD), the call resolves the name either to a local queue or to a transmission
queue. See Table 2 on page 97 for an overview of the name resolution process.
Note that the name you supply in *ObjectQMgrName* is resolved **before** that in
ObjectName.

Table 2 on page 97 also shows how you can use a local definition of a remote
queue to define an alias for the name of a queue manager. This allows you to
select which transmission queue is used when you put messages on a remote
queue, so you could, for example, use a single transmission queue for messages
destined for many remote queue managers. This technique is described further in
the *MQSeries Intercommunication* book.

Table 2. Resolving queue names when using MQOPEN

Input to MQOD		Resolved names		
ObjectQMgrName	*ObjectName*	*ObjectQMgrName*	*ObjectName*	**Transmission queue**
Blank or local queue manager	Local queue	Local queue manager	Input *ObjectName*	Not applicable
Blank or local queue manager	Model queue	Local queue manager	Generated name	Not applicable
Blank or local queue manager	Alias queue	Perform name resolution again with *ObjectQMgrName* set to the name of the local queue manager, and input *ObjectName* set to the *BaseQName* in the alias queue definition object		
Blank or local queue manager	Local definition of a remote queue	*RemoteQMgrName* in the remote queue definition object	*RemoteQName* in the remote definition object	Name of *XmitQName* attribute, if non-blank; otherwise *RemoteQMgrName* in the remote queue definition object
Name of a local transmission queue	(Not resolved)	Input *ObjectQMgrName*	Input *ObjectName*	Input *ObjectQMgrName*
Queue manager alias definition (and *RemoteQMgrName* in the alias queue definition object is the local queue manager)	(Not resolved)	Perform name resolution again with *ObjectQMgrName* set to the name of the local queue manager		
Queue manager alias definition (and *RemoteQMgrName* in the alias queue definition object is **not** the local queue manager)	(Not resolved)	*RemoteQMgrName* in the alias queue definition object	Input *ObjectName*	Name of *XmitQName* attribute, if non-blank; otherwise *RemoteQMgrName*
Queue manager is not the name of any local object	(Not resolved)	*RemoteQMgrName*	Input *ObjectName*	*DefXmitQName* attribute of the queue manager[3]

Notes:

1. *BaseQName* is the name of the base queue from the definition of the alias queue.

2. *RemoteQName* is the name of the remote queue from the local definition of the remote queue.

3. *RemoteQMgrName* is the name of the remote queue manager from the local definition of the remote queue.

---

[3] where *DefXmitQName* is supported

4. *XmitQName* is the name of the transmission queue from the local definition of the remote queue.

Opening an alias queue also opens the base queue to which the alias resolves, and opening a remote queue also opens the transmission queue. Therefore you cannot delete either the queue you specify or the queue to which it resolves while the other one is open.

There is more information on name resolution in a distributed queuing environment in the *MQSeries Intercommunication* book.

Using the options of the MQOPEN call

In the `Options` parameter of the MQOPEN call, you must choose one or more options to control the access you are given to the object you are opening. With these options you can:

- Open a queue to allow you to put messages on it
- Open a queue to allow you to get (browse or get) messages from it
- Open an object to allow you to inquire about and set its attributes (but you can set the attributes of queues only)
- Associate context information with a message
- Nominate an alternate user identifier to be used for security checks
- Control the call if the queue manager is in a quiescing state

MQOPEN option for putting messages

To open a queue in order to put messages on it, use the MQOO_OUTPUT option.

MQOPEN option for browsing messages

To open a queue so that you can *browse* the messages on it, use the MQOPEN call with the MQOO_BROWSE option. This creates a *browse cursor* that the queue manager uses to identify the next message on the queue. For more information, see "Browsing messages on a queue" on page 149.

Notes:

1. You cannot browse messages on a remote queue. Therefore you cannot open a remote queue using the MQOO_BROWSE option.

2. You cannot specify this option when opening a distribution list. For further information about distribution lists, see "Distribution lists" on page 111.

MQOPEN options for removing messages

There are three options that control the opening of a queue in order to remove messages from it. You can use only one of them in any MQOPEN call. These options define whether your program has exclusive or shared access to the queue. *Exclusive access* means that, until you close the queue, only you can remove messages from it. If another program attempts to open the queue to remove messages, its MQOPEN call fails. *Shared access* means that more than one program can remove messages from the queue.

The most advisable approach is to accept the type of access that was intended for the queue when the queue was defined. The queue definition involved the setting

of the *Shareability* and the *DefInputOpenOption* attributes. To accept this access, use the MQOO_INPUT_AS_Q_DEF option. Refer to Table 3 to see how the setting of these attributes affects the type of access you will be given when you use this option.

Table 3. How queue attributes and options of the MQOPEN call affect access to queues				
Queue attributes		**Type of access with MQOPEN options**		
Shareability	*DefInputOpenOption*	**AS_Q_DEF**	**SHARED**	**EXCLUSIVE**
SHAREABLE	SHARED	shared	shared	exclusive
SHAREABLE	EXCLUSIVE	exclusive	shared	exclusive
NOT_SHAREABLE*	SHARED*	exclusive	exclusive	exclusive
NOT_SHAREABLE	EXCLUSIVE	exclusive	exclusive	exclusive
Note: * Although you can define a queue to have this combination of attributes, the default input open option is overridden by the shareability attribute.				

Alternatively:

- If you know that your application can work successfully even if other programs can remove messages from the queue at the same time, use the MQOO_INPUT_SHARED option. Table 3 shows how, in some cases you will be given exclusive access to the queue, even with this option.

- If you know that your application can work successfully only if other programs are prevented from removing messages from the queue at the same time, use the MQOO_INPUT_EXCLUSIVE option.

Notes:

1. You cannot remove messages from a remote queue. Therefore you cannot open a remote queue using any of the MQOO_INPUT_⋆ options.

2. You cannot specify this option when opening a distribution list. For further information, see "Distribution lists" on page 111.

MQOPEN options for setting and inquiring about attributes

To open a queue so that you can set its attributes, use the MQOO_SET option. You cannot set the attributes of any other type of object (see Chapter 12, "Inquiring about and setting object attributes" on page 181).

To open an object so that you can inquire about its attributes, use the MQOO_INQUIRE option.

Note: You cannot specify this option when opening a distribution list.

MQOPEN options relating to message context

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

The options allow you to differentiate between context information that relates to the **user** who originated the message, and that which relates to the **application** that originated the message. Also, you can opt to set the context information when

you put the message on the queue, or you can opt to have the context taken automatically from another queue handle.

For more information about the subject of message context, see "Message context" on page 35.

MQOPEN option for alternate user authority
This is not supported on MQSeries for Windows.

When you attempt to open an object using the MQOPEN call, the queue manager checks that you have the authority to open that object. If you are not authorized, the call fails.

However, server programs may want the queue manager to check the authorization of the user on whose behalf they are working, rather than the server's own authorization. To do this, they must use the MQOO_ALTERNATE_USER_AUTHORITY option of the MQOPEN call, and specify the alternate user ID in the *AlternateUserId* field of the MQOD structure. Typically, the server would get the user ID from the context information in the message it is processing.

MQOPEN option for queue manager quiescing
This is not supported on MQSeries for Windows.

In the CICS environment, if you use the MQOPEN call when the queue manager is in a quiescing state, the call always fails. In other MVS environments, AS/400, OS/2, Windows NT, and in UNIX systems environments, the call fails when the queue manager is quiescing only if you use the MQOO_FAIL_IF_QUIESCING option of the MQOPEN call.

Creating dynamic queues

You should use a dynamic queue for those cases where you do not need the queue after your application ends. For example, you may want to use a dynamic queue for your "reply-to" queue. You specify the name of the reply-to queue in the *ReplyToQ* field of the MQMD structure when you put a message on a queue (see "Defining messages using the MQMD structure" on page 104).

To create a dynamic queue, you use a template known as a model queue, together with the MQOPEN call. You create a model queue using the MQSeries commands or the operations and control panels. The dynamic queue you create takes the attributes of the model queue.

When you call MQOPEN, specify the name of the model queue in the *ObjectName* field of the MQOD structure. When the call completes, the *ObjectName* field is set to the name of the dynamic queue that is created. Also, the *ObjectQMgrName* field is set to the name of the local queue manager.

There are three ways to specify the name of the dynamic queue you create:

- Give the full name you want in the *DynamicQName* field of the MQOD structure.

- Specify a prefix (less than 33 characters) for the name, and allow the queue manager to generate the rest of the name. This means that the queue manager generates a unique name, but you still have some control (for

example, you may want each user to use a certain prefix, or you may want to give a special security classification to queues with a certain prefix in their name). To use this method, specify an asterisk (⋆) for the last non-blank character of the *DynamicQName* field. Do not specify a single asterisk (⋆) for the dynamic queue name.

- Allow the queue manager to generate the full name. To use this method, specify an asterisk (⋆) in the first character position of the *DynamicQName* field.

For more information about these methods, see the description of the *DynamicQName* field in the *MQSeries Application Programming Reference* manual.

There is more information on dynamic queues in "Dynamic queues" on page 45.

Opening remote queues

A remote queue is a queue owned by a queue manager other than the one to which the application is connected.

To open a remote queue, use the MQOPEN call as for a local queue, but there are two ways you can specify the name of the queue:

1. In the *ObjectName* field of the MQOD structure, specify the name of the remote queue as known to the *local* queue manager.

2. In the *ObjectName* field of the MQOD structure, specify the name of the remote queue, as known to the *remote* queue manager. In the *ObjectQMgrName* field, specify either:

   - The name of the transmission queue that has the same name as the remote queue manager.

   - The name of an alias queue object that resolves to the transmission queue that has the same name as the remote queue manager.

   This tells the queue manager the destination of the message as well as the transmission queue it needs to be put on to get there.

3. If *DefXmitQname* is supported, in the *ObjectName* field of the MQOD structure, specify the name of the remote queue name as known by the *remote* queue manager.

In both methods, only local names are validated when you call MQOPEN; the last check is for the existence of the transmission queue to be used.

These two methods are the ones that are summarized in Table 2 on page 97.

Closing objects using the MQCLOSE call

To close an object, you use the MQCLOSE call. If the object is a queue, you should note the following:

- There is no need to empty a temporary dynamic queue before you close it.

  When you close a temporary dynamic queue, the queue is deleted, along with any messages that may still be on it. This is true even if there are uncommitted MQGET, MQPUT, or MQPUT1 calls outstanding against the queue.

- In MQSeries for MVS/ESA, if you have any MQGET requests with an MQGMO_SET_SIGNAL option outstanding for that queue, they are canceled.
- If you opened the queue using the MQOO_BROWSE option, your browse cursor is destroyed.

In MQSeries for AS/400, if the first MQOPEN call made an implicit connection to the queue manager, the final MQCLOSE call makes an implicit disconnection.

Closure is unrelated to syncpoint, so you can close queues before or after syncpoint.

As input to the MQCLOSE call, you must supply:

- A connection handle. Use the same connection handle used to open it, or alternatively, for CICS and AS/400 applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero).
- The handle of the object you want to close. Get this from the output of the MQOPEN call.
- MQCO_NONE in the *Options* field (unless you are closing a permanent dynamic queue).
- The control option to determine whether the queue manager should delete the queue even if there are still messages on it (when closing a permanent dynamic queue).

The output from MQCLOSE is:

- A completion code.
- A reason code.
- The object handle, reset to the value MQHO_UNUSABLE_HOBJ.

Descriptions of the parameters of the MQCLOSE call are given in the *MQSeries Application Programming Reference* manual.

Chapter 9. Putting messages on a queue

Use the MQPUT call to put messages on the queue. You can use MQPUT repeatedly to put many messages on the same queue, following the initial MQOPEN call. Call MQCLOSE when you have finished putting all your messages on the queue.

If you want to put a single message on a queue and close the queue immediately afterwards, you can use the MQPUT1 call. MQPUT1 performs the same functions as the following sequence of calls:

- MQOPEN
- MQPUT
- MQCLOSE

Generally however, if you have more than one message to put on the queue, it is more efficient to use the MQPUT call. This depends on the size of the message and the platform you are working on.

The sections in this chapter describe how to:

- Put messages on a local queue
- Put messages on a remote queue
- Control context information
- Put a single message on a queue using the MQPUT1 call
- Put a single message to multiple destinations using a single MQPUT call

The final section mentions some cases where put calls might fail.

Putting messages on a local queue using the MQPUT call

As input to the MQPUT call, you must supply:

- A connection handle (HCONN).

- A queue handle (HObj).

- A description of the message you want to put on the queue. This is in the form of a message descriptor structure (MQMD).

- Control information, in the form of a put-message options structure (MQPMO).

- The length of the data contained within the message (MQLONG).

- The message data itself.

The output from the MQPUT call is

- A reason code (MQLONG)
- A completion code (MQLONG)

If the call completes successfully, it also returns your options structure and your message descriptor structure. The call modifies your options structure to show the name of the queue and the queue manager to which the message was sent. If you request that the queue manager generates a unique value for the identifier of the message you are putting (by specifying binary zero in the *MsgId* field of the MQMD structure), the call inserts the value in the *MsgId* field before returning this structure to you. This value must be reset before you issue another MQPUT.

There is a description of the MQPUT call in the *MQSeries Application Programming Reference* manual.

The following sections describe the information you must supply as input to the MQPUT call.

Specifying handles

For the connection handle (*Hconn*) in CICS on MVS/ESA and AS/400 applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or you can use the connection handle returned by the MQCONN call. For other applications, always use the connection handle returned by the MQCONN call.

Whatever environment you are working in, use the same queue handle (*Hobj*) that is returned by the MQOPEN call.

Defining messages using the MQMD structure

The message descriptor structure (MQMD) is an input/output parameter for the MQPUT and MQPUT1 calls. You use it to define the message you are putting on a queue.

Note: You must reset the *MsgId* and *CorrelId* to null prior to putting a new message in order to ensure they are unique. The values in these fields are returned on a successful MQPUT. However, if you set the *Version* field of the MQMD structure to 2, you can use the MQMO_MATCH_MSG_ID and MQMO_MATCH_CORREL_ID flags instead of resetting.

There is an introduction to the message properties that MQMD describes in Chapter 3, "MQSeries messages" on page 21, and there is a description of the structure itself in the *MQSeries Application Programming Reference* manual.

Specifying options using the MQPMO structure

You use the MQPMO (Put Message Option) structure to pass options to the MQPUT and MQPUT1 calls.

The following sections give you help on filling in the fields of this structure. There is a description of the structure in the *MQSeries Application Programming Reference* manual.

The fields of the structure include:

- *StrucId*
- *Version*
- *Options*
- *Context*
- *ResolvedQName*
- *ResolvedQMgrName*

These fields are described below.

StrucId

This identifies the structure as a put-message options structure. This is a 4-character field. Always specify MQPMO_STRUC_ID.

Version

This describes the version number of the structure. The default is MQPMO_VERSION_1. If you enter MQPMO_VERSION_2, you can use distribution lists (see on page 111). If you enter MQPMO_CURRENT_VERSION, your application is set to always use the most recent level.

Options

This controls the following:

- Whether the put operation is included in a unit of work
- How much context information is associated with a message
- Where the context information is taken from
- Whether the call fails if the queue manager is in a quiescing state
- Generation of a new message identifier and correlation identifier
- The order in which messages and segments are put on a queue

If you leave the *Options* field set to the default value (MQPMO_NONE), the message you put has default context information associated with it.

Also, the way that the call operates with syncpoints is determined by the platform. The syncpoint control default is 'yes' in MVS/ESA; for other platforms, it is 'no'.

Context

This states the name of the queue handle that you want context information to be copied from (if requested in the *Options* field).

For an introduction to message context, see "Message context" on page 35. For information about using the MQPMO structure to control the context information in a message, see "Controlling context information" on page 108.

ResolvedQName

This contains the name (after resolution of any alias name) of the queue that was opened to receive the message. This is an output field.

ResolvedQMgrName

This contains the name (after resolution of any alias name) of the queue manager that owns the queue in *ResolvedQName*. This is an output field.

The MQPMO can also accommodate fields required for distribution lists (see on page 111). If you wish to use this facility, Version 2 of the MQPMO structure is used. This includes the following fields:

Version

This field describes the version number of the structure. For distribution lists, you are required to specify MQPMO_VERSION_2.

RecsPresent

This field contains the number of Put Message Records (MQPMR) and Response Records (MQRR) present.

The value you enter can be the same as the number of Object Records provided at MQOPEN. However, if the value is less than the number of Object Records provided on the MQOPEN call (or if no Put Message Records are

provided), the values of the queues that are not defined are taken from the default values provided by the message descriptor. Also, if the value is greater than the number of Object Records provided, the excess Put Message Records are ignored.

You are recommended to do one of the following:

- If you want to receive a report or reply from each destination, enter the same value as appears in the MQOR structure and use MQPMRs containing *MsgId* fields. Either initialize these *MsgId* fields to zeros or specify MQPMO_NEW_MSG_ID.

  When you have put the message to the queue, *MsgId* values that the queue manager has created become available in the MQPMRs; you can use these to identify which destination is associated with each report or reply.

- If you do not want to receive reports or replies, choose one of the following:

  1. If you want to identify destinations that fail immediately, you may still want to enter the same value in the *RecsPresent* field as appears in the MQOR structure and provide MQRRs to identify these destinations. Do not specify any MQPMRs.

  2. If you do not want to identify failed destinations, enter zero in the *RecsPresent* field and do not provide MQPMRs nor MQRRs.

  Note: If you are using MQPUT1, the number of Response Record Pointers and Response Record Offsets must be zero.

  For a full description of Put Message Records (MQPMR) and Response Records (MQRR), see the *MQSeries Application Programming Reference* manual.

PutMsgRecFields
This indicates which fields are present in each Put Message Record (MQPMR). For a list of these fields, see "Using the MQPMR structure" on page 116.

PutMsgRecOffset **and** *PutMsgRecPtr*
Pointers (typically in C) and offsets (typically in COBOL) are used to address the Put Message Records (see "Using the MQPMR structure" on page 116 for an overview of the MQPMR structure).

Use the *PutMsgRecPtr* field to specify a pointer to the first Put Message Record, or the *PutMsgRecOffset* field to specify the offset of the first Put Message Record. This is the offset from the start of the MQPMO. Depending on the *PutMsgRecFields* field, enter a non-null value for either *PutMsgRecOffset* or *PutMsgRecPtr*.

ResponseRecOffset **and** *ResponseRecPtr*
You also use pointers and offsets to address the Response Records (see "Using the MQRR structure" on page 114 for further information about Response Records).

Use the *ResponseRecPtr* field to specify a pointer to the first Response Record, or the *ResponseRecOffset* field to specify the offset of the first Response Record. This is the offset from the start of the MQPMO structure. Enter a non-null value for either *ResponseRecOffset* or *ResponseRecPtr*.

Note: If you are using MQPUT1 to put messages to a distribution list, *ResponseRecPtr* must be null or zero and *ResponseRecOffset* must be zero.

Additional information for putting to a distribution list (see "Distribution lists" on page 111) is provided in Version 2 of the Put Message Option structure (MQPMR). This is described in the *MQSeries Application Programming Reference*.

The data in your message

Give the address of the buffer that contains your data in the *Buffer* parameter of the MQPUT call. You can include anything in the data in your messages. The amount of data in the messages, however, affects the performance of the application that is processing them.

The maximum size of the data is determined by:

- The *MaxMsgLength* attribute of the queue manager
- The *MaxMsgLength* attribute of the queue on which you are putting the message
- The size of any message header added by MQSeries (including the Dead-letter header, MQDLH and the Distribution list header, MQDH)

The *MaxMsgLength* attribute of the queue manager holds the size of message that the queue manager can process. This has a default of 4 MB (1 MB=1048576 bytes). To determine the value of this attribute, use the MQINQ call on the queue manager object. For large messages, you can change this value (not on AS/400).

The *MaxMsgLength* attribute of a queue determines the maximum size of message you can put on the queue. If you attempt to put a message with a size larger than the value of this attribute, your MQPUT call fails. If you are putting a message on a remote queue, the maximum size of message that you can successfully put is determined by the *MaxMsgLength* attribute of the remote queue, of any intermediate transmission queues that the message is put on along the route to its destination, and of the channels used.

For an MQPUT operation, the size of the message must be smaller than or equal to the *MaxMsgLength* attribute of both the queue and the queue manager. The values of these attributes are independent, but you are recommended to set the *MaxMsgLength* of the queue to a value less than or equal to that of the queue manager.

MQSeries adds header information to messages in the following circumstances:

- When you put a message on a remote queue, MQSeries adds an MQXQH structure to the message. This structure includes the name of the destination queue and its owning queue manager.
- If MQSeries cannot deliver a message to a remote queue, it attempts to put the message on the dead-letter (undelivered-message) queue. It adds an MQDLH structure to the message. This structure includes the name of the destination queue and the reason the message was put on the dead-letter (undelivered-message) queue.
- If you want to send a message to multiple destination queues, MQSeries adds an MQDH header to the message. This describes the data that is present in a message, belonging to a distribution list, on a transmission queue. This point should be considered when choosing an optimum value for the maximum message length.

These structures are described in the *MQSeries Application Programming Reference* manual.

If your messages are of the maximum size allowed for these queues, the addition of these headers means that the put operations fail because the messages are now too big. To reduce the possibility of the put operations failing:

- Make the size of your messages smaller than the *MaxMsgLength* attribute of the transmission and dead-letter (undelivered-message) queues. Allow at least the value of the MQ_MSG_HEADER_LENGTH constant (more for large distribution lists).

- Make sure that the *MaxMsgLength* attribute of the dead-letter (undelivered-message) queue is set to the same as the *MaxMsgLength* of the queue manager that owns the dead-letter queue.

The queue managers attributes and the message queuing constants are described in the *MQSeries Application Programming Reference* manual.

For information on how undelivered messages are handled in a distributed queuing environment, see the *MQSeries Intercommunication* book.

Putting messages on a remote queue

When you want to put a message on a remote queue (that is, a queue owned by a queue manager other than the one to which your application is connected) rather than a local queue, the only extra consideration is how you specify the name of the queue when you open it. This is described in "Opening remote queues" on page 101. There is no change to how you use the MQPUT or MQPUT1 call for a local queue.

For more information on using remote and transmission queues, see the *MQSeries Intercommunication* book.

Controlling context information

To control context information, you use the *Options* field in the MQPMO structure.

If you don't, the queue manager will overwrite context information that may already be in the message descriptor with the identity and context information it has generated for your message. This is the same as specifying the MQPMO_DEFAULT_CONTEXT option. You may want this default context information when you create a new message (for example, when processing user input from an inquiry screen).

If you want no context information associated with your message, use the MQPMO_NO_CONTEXT option.

Passing identity context

In general, programs should pass identity context information from message to message around an application until the data reaches its final destination. Programs should change the origin context information each time they change the data. However, applications that want to change or set any context information must have the appropriate level of authority. The queue manager checks this authority when the applications open the queues; they must have authority to use the appropriate context options for the MQOPEN call.

If your application gets a message, processes the data from the message, then puts the changed data into another message (possibly for processing by another application), the application should pass the identity context information from the original message to the new message. You can allow the queue manager to create the origin context information.

To save the context information from the original message, you must use the MQOO_SAVE_ALL_CONTEXT option when you open the queue for getting the message. This is in addition to any other options you use with the MQOPEN call. Note, however, that you cannot save context information if you only browse the message.

When you create the second message, you must:

- Open the queue using the MQOO_PASS_IDENTITY_CONTEXT option (in addition to the MQOO_OUTPUT option).

- In the *Context* field of the put-message options structure, give the handle of the queue from which you saved the context information.

- In the *Options* field of the put-message options structure, specify the MQPMO_PASS_IDENTITY_CONTEXT option.

Passing all context

If your application gets a message, and puts the message data (unchanged) into another message, the application should pass both the identity and the origin context information from the original message to the new message. An example of an application that might do this is a message mover, which moves messages from one queue to another.

Follow the same procedure as for passing identity context, except you use the MQOPEN option MQOO_PASS_ALL_CONTEXT and the put-message option MQPMO_PASS_ALL_CONTEXT.

Setting identity context

If you want to set the identity context information for a message, leaving the queue manager to set the origin context information:

- Open the queue using the MQOO_SET_IDENTITY_CONTEXT option.

- Put the message on the queue, specifying the MQPMO_SET_IDENTITY_CONTEXT option. In the message descriptor, specify whatever identity context information you require.

Setting all context

If you want to set both the identity and the origin context information for a message,

- Open the queue using the MQOO_SET_ALL_CONTEXT option.

- Put the message on the queue, specifying the MQPMO_SET_ALL_CONTEXT option. In the message descriptor, specify whatever identity and origin context information you require.

Appropriate authority is needed for each type of context setting.

Putting one message on a queue using the MQPUT1 call

Use the MQPUT1 call when you want to close the queue immediately after you have put a single message on it. For example, a server application is likely to use the MQPUT1 call when it is sending a reply to each of the different queues.

MQPUT1 is functionally equivalent to calling MQOPEN followed by MQPUT, followed by MQCLOSE. The only difference in the syntax for the MQPUT and MQPUT1 calls is that for MQPUT you must specify an object handle, whereas for MQPUT1 you must specify an object descriptor structure (MQOD) as defined in MQOPEN (see "Identifying objects (the MQOD structure)" on page 95). This is because you need to give information to the MQPUT1 call about the queue it has to open, whereas when you call MQPUT, the queue must already be open.

As input to the MQPUT1 call, you must supply:

- A connection handle.

- A description of the object you want to open. This is in the form of an object descriptor structure (MQOD).

- A description of the message you want to put on the queue. This is in the form of a message descriptor structure (MQMD).

- Control information in the form of a put-message options structure (MQPMO).

- The length of the data contained within the message (MQLONG).

- The address of the message data.

The output from MQPUT1 is:

- A completion code
- A reason code

If the call completes successfully, it also returns your options structure and your message descriptor structure. The call modifies your options structure to show the name of the queue and the queue manager to which the message was sent. If you request that the queue manager generates a unique value for the identifier of the message you are putting (by specifying binary zero in the *MsgId* field of the MQMD structure), the call inserts the value in the *MsgId* field before returning this structure to you.

Note: You cannot use MQPUT1 with a model queue name; however, once a model queue has been opened, you can issue an MQPUT1 to the dynamic queue.

The six input parameters for MQPUT1 are:

Hconn
  This is a connection handle. For CICS and AS/400 applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN call. For other programs, always use the connection handle returned by the MQCONN call.

ObjDesc
  This is an object descriptor structure (MQOD).

  In the *ObjectName* and *ObjectQMgrName* fields, give the name of the queue on which you want to put a message, and the name of the queue manager that owns this queue.

  The *DynamicQName* field is ignored for the MQPUT1 call because it cannot use model queues.

  Use the *AlternateUserId* field only if you want to nominate an alternate user identifier that is to be used to test authority to open the queue.

MsgDesc
  This is a message descriptor structure (MQMD). As with the MQPUT call, use this structure to define the message you are putting on the queue.

PutMsgOpts
  This is a put-message options structure (MQPMO). Use it as you would for the MQPUT call (see "Specifying options using the MQPMO structure" on page 104).

  When the *Options* field is set to zero, the queue manager uses your own user ID when it performs tests for authority to access the queue. Also, the queue manager ignores any alternate user identifier given in the *AlternateUserId* field of the MQOD structure.

BufferLength
  This is the length of your message.

Buffer
  This is the buffer area that contains the text of your message.

There is a description of the MQPUT1 call in the *MQSeries Application Programming Reference* manual.

Distribution lists

These are supported on MQSeries Version 5 products, and MQSeries for AS/400.

Distribution lists allow you to put a message to multiple destinations in a single MQPUT or MQPUT1 call. Multiple queues can be opened using a single MQOPEN and a message can then be put to each of those queues using a single MQPUT. Some generic information from the MQI structures used for this process can be superseded by specific information relating to the individual destinations included in the distribution list.

When an MQOPEN call is issued, generic information is taken from the Object Descriptor (MQOD). If you specify MQOD_VERSION_2 in the *Version* field and a value greater than zero in the *RecsPresent* field, the *Hobj* can be defined as a handle of a list (of one or more queues) rather than of a queue. In this case, specific information is given through the object records (MQORs), which give details of destination (that is, *ObjectName* and *ObjectQMgrName*).

The object handle (*Hobj*) is passed to the MQPUT call, allowing you to put to a list rather than to a single queue.

When a message is put on the queues (MQPUT), generic information is taken from the Put Message Option structure (MQPMO) and the Message Descriptor (MQMD). Specific information is given in the form of Put Message Records (MQPMRs).

Response Records (MQRR) can receive a completion code and reason code specific to each destination queue.

Note: There may be an impact on portability if pointers are used in these languages on the following platforms:

Language	Platform
Visual Basic	Windows '95, Windows NT, Windows 3.1
COBOL	AS/400
RPG	AS/400

Figure 5 shows how distribution lists work:
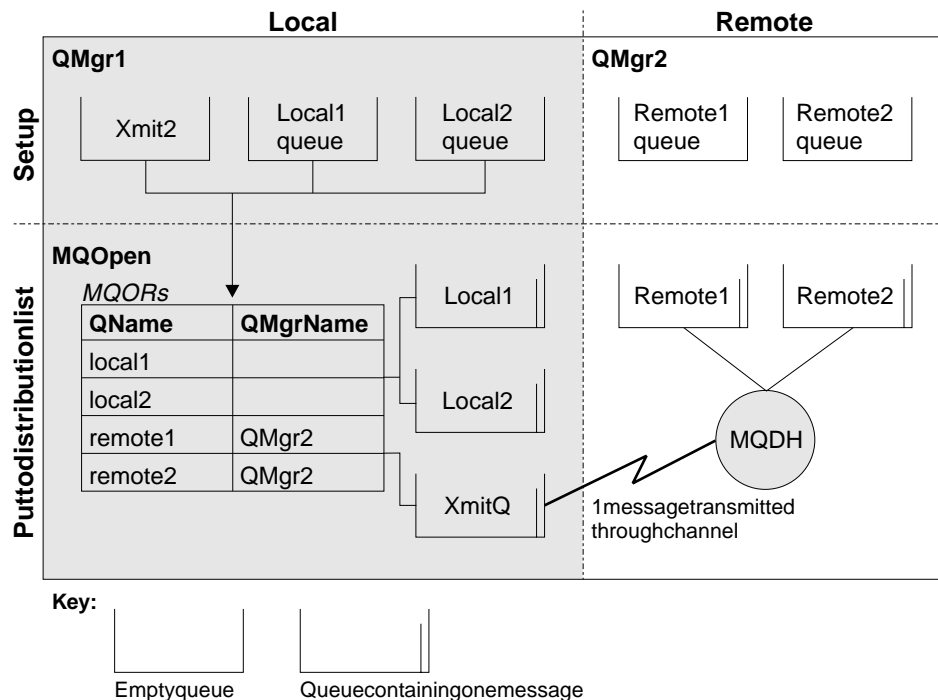
*Figure 5. How distribution lists work. This diagram shows that **one** message is transmitted through the channel and can be put on more than one remote queue.*

Opening distribution lists

Use the MQOPEN call to open a distribution list, and use the options of the call to specify what you want to do with the list.

As input to MQOPEN, you must supply:

- A connection handle (see Chapter 9, " Putting messages on a queue" on page 103 for a description)
- Generic information in the Object Descriptor structure (MQOD)

- The name of each queue you want to open, using the Object Record structure (MQOR)

The output from MQOPEN is:

- An object handle that represents your access to the distribution list
- A generic completion code
- A generic reason code
- Response Records (optional), containing a completion code and reason for each destination

Using the MQOD structure

Use the MQOD structure to identify the queues you want to open. To define a distribution list, you must specify MQOD_VERSION_2 in the *Version* field, a value greater than zero in the *RecsPresent* field, and MQOT_Q in the *ObjectType* field. See the *MQSeries Application Programming Reference* manual for a description of all the fields of the MQOD structure.

Using the MQOR structure

An MQOR structure must be provided for each destination. The structure contains the destination queue and queue manager names. The *ObjectName* and *ObjectQMgrName* fields in the MQOD are not used for distribution lists. There must be one or more object records. If the *ObjectQMgrName* is left blank, the local queue manager is used. See the *MQSeries Application Programming Reference* manual for further information about these fields.

You can specify the destination queues in two ways:

- By using the offset field *ObjectRecOffset*.

  In this case, the application should declare its own structure containing an MQOD structure, followed by the array of MQOR records (with as many array elements as are needed), and set *ObjectRecOffset* to the offset of the first element in the array from the start of the MQOD. Care must be taken to ensure that this offset is correct.

  Use of built-in facilities provided by the programming language is recommended, if these are available in all of the environments in which the application must run. The following illustrates this technique for the COBOL programming language:

  ```
  01  MY-OPEN-DATA.
      02 MY-MQOD.
         COPY CMQODV.
      02 MY-MQOR-TABLE OCCURS 100 TIMES.
         COPY CMQORV.
      MOVE LENGTH OF MY-MQOD TO MQOD-OBJECTRECOFFSET.
  ```

  Alternatively, the constant MQOD_CURRENT_LENGTH can be used if the programming language does not support the necessary built-in facilities in all of the environments concerned. The following illustrates this technique:

```
01  MY-MQ-CONSTANTS.
    COPY CMQV.
01  MY-OPEN-DATA.
    02 MY-MQOD.
       COPY CMQODV.
    02 MY-MQOR-TABLE OCCURS 100 TIMES.
       COPY CMQORV.
    MOVE MQOD-CURRENT-LENGTH TO MQOD-OBJECTRECOFFSET.
```

However, this will work correctly only if the MQOD structure and the array of MQOR records are contiguous; if the compiler inserts skip bytes between the MQOD and the MQOR array, these must be added to the value stored in *ObjectRecOffset*.

Using *ObjectRecOffset* is recommended for programming languages that do not support the pointer data type, or that implement the pointer data type in a way that is not portable to different environments (for example, the COBOL programming language).

- By using the pointer field *ObjectRecPtr*.

  In this case, the application can declare the array of MQOR structures separately from the MQOD structure, and set *ObjectRecPtr* to the address of the array. The following illustrates this technique for the C programming language:

```
MQOD MyMqod;
MQOR MyMqor[100];
MyMqod.ObjectRecPtr = MyMqor;
```

  Using *ObjectRecPtr* is recommended for programming languages that support the pointer data type in a way that is portable to different environments (for example, the C programming language).

Whichever technique is chosen, one of *ObjectRecOffset* and *ObjectRecPtr* must be used; the call fails with reason code MQRC_OBJECT_RECORDS_ERROR if both are zero, or both are nonzero.

Using the MQRR structure
These structures are destination specific as each Response Record contains a *CompCode* and *Reason* field for each queue of a distribution list. You must use this structure to enable you to distinguish where any problems lie.

For example, if you receive a reason code of MQRC_MULTIPLE_REASONS and your distribution list contains five destination queues, you will not know which queues the problems apply to if you do not use this structure. However, if you have a completion code and reason code for each destination, you can locate the errors more easily.

See the *MQSeries Application Programming Reference* manual for further information about the MQRR structure.

Figure 6 shows how you can open a distribution list in C:

MQOD 2   n ptr   ptr

MQOR

0	Q	QMgr
	Q	QMgr
n-1	Q	QMgr

MQRR

0	CompCode	Reason
	CompCode	Reason
n-1	CompCode	Reason

Figure 6. Opening a distribution list in C. The MQOD uses pointers to the MQOR and MQRR structures.

Figure 7 shows how you can open a distribution list in COBOL:

MQOD 2   n offset offset Q QMgr Q QMgr CompCode Reason CompCode R

Figure 7. Opening a distribution list in COBOL. The MQOD uses offsets in COBOL.

Using the MQOPEN options

The following options can be specified when opening a distribution list:

- MQOO_OUTPUT
- MQOO_FAIL_IF_QUIESCING (optional)
- MQOO_ALTERNATE_USER_AUTHORITY (optional)
- MQOO_*_CONTEXT (optional)

See Chapter 8, "Opening and closing objects" on page 93 for a description of these options.

Putting messages to a distribution list

To put messages to a distribution list, you can use MQPUT or MQPUT1. As input, you must supply:

- A connection handle (see Chapter 9, " Putting messages on a queue" on page 103 for a description).

- An object handle. If a distribution list is opened using MQOPEN, the *Hobj* allows you only to put to the list.

- A message descriptor structure (MQMD). See the *MQSeries Application Programming Reference* for a description of this structure.

- Control information in the form of a put-message option structure (MQPMO). See "Specifying options using the MQPMO structure" on page 104 for information about filling in the fields of the MQPMO structure.

- Control information in the form of Put Message Records (MQPMR).

- The length of the data contained within the message (MQLONG).

- The message data itself.

The output is:

- A completion code
- A reason code
- Response Records (optional)

Using the MQPMR structure

This structure is optional and gives destination-specific information for some fields that you may want to identify differently from those already identified in the MQMD. For a description of these fields, see the *MQSeries Application Programming Reference* manual.

The content of each record depends on the information given in the *PutMsgRecFields* field of the MQPMO. For example, in the sample program AMQSPTL0.C (see "The Distribution List sample program" on page 327 for a description) showing the use of distribution lists, the sample chooses to provide values for *MsgId* and *CorrelId* in the MQPMR. This section of the sample program looks like this:

```
typedef struct
{
MQBYTE24 MsgId;
MQBYTE24 CorrelId;
} PutMsgRec;...
/**********************
MQLONG PutMsgRecFields=MQPMRF_MSG_ID | MQPMRF_CORREL_ID;
```

This implies that *MsgId* and *CorrelId* are provided for each destination of a distribution list. The Put Message Records are provided as an array.

Using MQPUT1

If you are using MQPUT1, consider the following:

1. The values of the *ResponseRecOffset* and *ResponseRecPtr* fields must be null or zero.

2. The Response Records, if required, must be addressed from the MQOD.

Some cases where the put calls fail

If certain attributes of a queue are changed using the FORCE option on a command during the interval between you issuing an MQOPEN and an MQPUT call, the MQPUT call fails and returns the MQRC_OBJECT_CHANGED reason code. The queue manager marks the object handle as being no longer valid. This also happens if the changes are made while an MQPUT1 call is being processed, or if the changes apply to any queue to which the queue name resolves. The attributes that affect the handle in this way are listed in the description of the MQOPEN call in the *MQSeries Application Programming Reference* manual. If your call returns the MQRC_OBJECT_CHANGED reason code, close the queue, reopen it, then try to put a message again.

If put operations are inhibited for a queue on which you are attempting to put messages (or any queue to which the queue name resolves), the MQPUT or MQPUT1 call fails and returns the MQRC_PUT_INHIBITED reason code. You may be able to put a message successfully if you attempt the call at a later time, if the design of the application is such that other programs change the attributes of queues regularly.

Further, if the queue that you are trying to put your message on is full, the MQPUT or MQPUT1 call fails and returns MQRC_Q_FULL.

If a dynamic queue (either temporary or permanent) has been deleted, MQPUT calls using a previously acquired object handle fail and return the MQRC_Q_DELETED reason code. In this situation it is good practice to close the object handle as it is no longer any use to you.

In the case of distribution lists, multiple completion codes and reason codes can occur in a single request. These cannot be handled using only the *CompCode* and *Reason* output fields on MQOPEN and MQPUT.

When distribution lists are used to put messages to multiple destinations, the Response Records contain the specific *CompCode* and *Reason* for each destination. If you receive a completion code of MQCC_FAILED, no message is put on any destination queue successfully. If the completion code is MQCC_WARNING, the message is successfully put on one or more of the destination queues. If you receive a return code of MQRC_MULTIPLE_REASONS, the reason codes are not all the same for every destination. Therefore, it is recommended to use the MQRR structure so that you can determine which queue or queues caused an error and the reasons for each.

Put call failure

Chapter 10. Getting messages from a queue

You can get messages from a queue in two ways:

1. You can **remove** a message from the queue so that other programs can no longer see it.

2. You can **copy** a message, leaving the original message on the queue. This is known as *browsing*. You can easily remove the message once you have browsed it.

In both cases, you use the MQGET call, but first your application must be connected to the queue manager, and you must use the MQOPEN call to open the queue (for input, browse, or both). These operations are described in Chapter 7, "Connecting and disconnecting a queue manager" on page 87 and Chapter 8, "Opening and closing objects" on page 93.

When you have opened the queue, you can use the MQGET call repeatedly to browse or remove messages on the same queue. Call MQCLOSE when you have finished getting all the messages you want from the queue.

This chapter tells you how to get messages from a queue using the MQGET call, and includes:

- Information on the MQGET call itself
- The order in which messages are retrieved
- Getting a particular message
- Waiting for messages
- Handling large messages
- Getting notification when a message arrives (signaling)
- Avoiding loops by skipping backout
- Converting messages from different systems
- Browsing messages
- Why MQGET might fail

Getting messages from a queue using the MQGET call

The MQGET call gets a message from an open local queue. It cannot get a message from a queue on another system.

As input to the MQGET call, you must supply:

- A connection handle.

- A queue handle.

- A description of the message you want to get from the queue. This is in the form of a message descriptor (MQMD) structure.

- Control information in the form of a Get Message Options (MQGMO) structure.

- The size of the buffer you have assigned to hold the message (MQLONG).

- The address of the storage in which the message must be put.

The output from MQGET is:

- A reason code

- A completion code

- The message in the buffer area you specified, if the call completes successfully

- Your options structure, modified to show the name of the queue from which the message was retrieved

- Your message descriptor structure, with the contents of the fields modified to describe the message that was retrieved

- The length of the message (MQLONG)

There is a description of the MQGET call in the *MQSeries Application Programming Reference* manual.

The following sections describe the information you must supply as input to the MQGET call.

Specifying connection handles

For CICS on MVS/ESA and AS/400 applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN call. For other applications, always use the connection handle returned by the MQCONN call.

Use the queue handle (*Hobj*) that is returned when you call MQOPEN.

Describing messages using the MQMD structure and the MQGET call

To identify the message you want to get from a queue, use the message descriptor structure (MQMD). This is an input/output parameter for the MQGET call. There is an introduction to the message properties that MQMD describes in Chapter 3, "MQSeries messages" on page 21, and there is a description of the structure itself in the *MQSeries Application Programming Reference* manual.

If you know which message you want to get from the queue, see "Getting a particular message" on page 133.

If you do not specify a particular message, MQGET retrieves the **first** message in the queue."The order in which messages are retrieved from a queue" on page 124 describes how the priority of a message, the *MsgDeliverySequence* attribute of the queue, and the MQGMO_LOGICAL_ORDER option determine the order of the messages in the queue.

Note: If you want to use MQGET more than once (for example, to step through the messages in the queue), you can set the *MsgId* and *CorrelId* fields of this structure to null after each call. This prevents the the call from filling these fields with the identifiers of the message that was retrieved, and therefore having messages with the same identifiers as the previous message.

However, if you want to group your messages, the *GroupId* should be the same for messages in the same group, so that the call will look for a message having the same identifiers as the previous message in order to make up the whole group.

Specifying MQGET options using the MQGMO structure

The MQGMO structure is an input/output variable for passing options to the MQGET call.

The following sections give you help on filling in some of the fields of this structure. There is a description of the structure in the *MQSeries Application Programming Reference* manual.

StrucId

StrucId is a 4-character field used to identify the structure as a get-message options structure. Always specify MQGMO_STRUC_ID.

Version

Version describes the version number of the structure. MQGMO_VERSION_1 is the default. If you wish to use the Version 2 fields or retrieve messages in logical order, specify MQGMO_VERSION_2. MQGMO_CURRENT_VERSION sets your application to use the most recent level.

Options

Within your code, you can select the options in any order as each option is represented by a bit in the *Options* field.

The *Options* field controls:

- Whether the MQGET call waits for a message to arrive on the queue before it completes (see "Waiting for messages" on page 142)

- Whether the get operation is included in a unit of work

- Whether a nonpersistent message is retrieved outside syncpoint, allowing fast messaging

- In MQSeries for MVS/ESA, whether the message retrieved is marked as skipping backout (see "Skipping backout" on page 145)

- Whether the message is removed from the queue, or merely browsed

- Whether to select a message by using a browse cursor or by other selection criteria

- Whether the call succeeds even if the message is longer than your buffer

- In MQSeries for MVS/ESA, whether to allow the call to complete, but set a signal to indicate that you want to be notified when a message arrives

- Whether the call fails if the queue manager is in a quiescing state

- On MVS/ESA, whether the call fails if the connection is in a quiescing state

- Whether application message data conversion is required (see "Application data conversion" on page 148)

- On MQSeries Version 5 products and MQSeries for AS/400, the order in which messages and segments are retrieved from a queue

- On MQSeries Version 5 products and MQSeries for AS/400, whether complete, logical messages only are retrievable

- On MQSeries Version 5 products and MQSeries for AS/400, whether messages in a group can be retrieved only when **all** messages in the group are available

- On MQSeries Version 5 products and MQSeries for AS/400, whether segments in a logical message can be retrieved only when **all** segments in the logical message are available

If you leave the *Options* field set to the default value (MQGMO_NO_WAIT), the MQGET call operates this way:

- If there is no message matching your selection criteria on the queue, the call does not wait for a message to arrive, but completes immediately. Also, in MQSeries for MVS/ESA, the call does not set a signal requesting notification when such a message arrives.

- The way that the call operates with syncpoints is determined by the platform:

Platform	Under syncpoint control
AS/400	No
UNIX systems	No
MVS/ESA	YES
OS/2	No
Tandem NSK	YES
Windows NT	No
Windows	No

- In MQSeries for MVS/ESA, the message retrieved is not marked as skipping backout.

- The selected message is removed from the queue (not browsed).

- No application message data conversion is required.

- The call fails if the message is longer than your buffer.

WaitInterval

The *WaitInterval* field specifies the maximum time (in milliseconds) that the MQGET call waits for a message to arrive on the queue when you use the MQGMO_WAIT option. If no message arrives within the time specified in *WaitInterval*, the call completes and returns a reason code showing that there was no message that matched your selection criteria on the queue.

In MQSeries for MVS/ESA, if you use the MQGMO_SET_SIGNAL option, the *WaitInterval* field specifies the time for which the signal is set.

For more information on these options, see "Waiting for messages" on

Signal1

Signal1 is supported on MQSeries for MVS/ESA, MQSeries for Tandem NSK, and MQSeries for Windows Version 2.1 only.

If you have chosen to use the MQGMO_SET_SIGNAL option to request that your application is notified when a suitable message arrives, you must specify the type of signal in the *Signal1* field. In MQSeries on all other platforms the *Signal1* field is reserved and its value is not significant.

Signal2
On MQSeries for Windows Version 2.1 this specifies an identifier for the signal message. The *Signal2* field is reserved on all other platforms and its value is not significant.

For more information, see "Signaling" on page 143.

ResolvedQName
ResolvedQName is an output field in which the queue manager returns the name of the queue (after resolution of any alias) from which the message was retrieved.

MatchOptions
MatchOptions controls the selection criteria for MQGET.

GroupStatus
GroupStatus indicates whether the message you have retrieved is in a group.

SegmentStatus
SegmentStatus indicates whether the item you have retrieved is a segment of a logical message.

Segmentation
Segmentation indicates whether segmentation is allowed for the message retrieved.

Specifying the size of the buffer area

In the *BufferLength* parameter of the MQGET call, specify the size of the buffer area you want to use to hold the message data that you retrieve. There are three ways to decide how big this should be:

1. You may already know what length of messages to expect from this program. If so, specify a buffer of this size.

   However, you can use the MQGMO_ACCEPT_TRUNCATED_MSG option in the MQGMO structure if you want the MQGET call to complete even if the message is too big for the buffer. In this case:

   • The buffer is filled with as much of the message as it can hold

   • The call returns a warning completion code

   • The message is removed from the queue (discarding the remainder of the message), or the browse cursor is advanced (if you are browsing the queue)

   • The real length of the message is returned in *DataLength*

   Without this option, the call still completes with a warning, but it does not remove the message from the queue (or advance the browse cursor).

2. Estimate a size for the buffer (or even specify a size of zero bytes) and ***do not*** use the MQGMO_ACCEPT_TRUNCATED_MSG option. If the MQGET call fails (for example, because the buffer is too small), the length of the message is returned in the *DataLength* parameter of the call. (The buffer is still filled with as much of the message as it can hold, but the processing of the call is not completed.) Store the *MsgId* of this message, then repeat the MQGET call, specifying a buffer area of the correct size, and the *MsgId* you noted from the first call.

If your program is serving a queue that is also being served by other programs, one of those other programs may remove the message you want before your program can issue another MQGET call. Your program could waste time searching for a message that no longer exists. To avoid this, first browse the queue until you find the message you want, specifying a *BufferLength* of zero and using the MQGMO_ACCEPT_TRUNCATED_MSG option. This positions the browse cursor under the message you want. You can then retrieve the message by calling MQGET again, specifying the MQGMO_MSG_UNDER_CURSOR option. If another program removes the message between your browse and removal calls, your second MQGET fails immediately (without searching the whole queue), because there is no message under your browse cursor.

3. The *MaxMsgLength* queue attribute determines the maximum length of messages accepted for that queue and the *MaxMsgLength* queue manager attribute determines the maximum length of messages accepted for that queue manager. If you do not know what length of message to expect, you can inquire about the *MaxMsgLength* attribute (using the MQINQ call), then specify a buffer of this size.

For further information about the *MaxMsgLength* attribute, see "Increasing the maximum message length" on page 136.

The order in which messages are retrieved from a queue

You have control over the order in which you retrieve messages from a queue. This section looks at the options:

Priority

A program can assign a priority to a message when it puts the message on a queue (see "Message priorities" on page 31). Messages of equal priority are stored in a queue in order of arrival, not the order in which they are committed.

The queue manager maintains queues either in strict FIFO (first in, first out) sequence, or in FIFO within priority sequence. This depends on the setting of the *MsgDeliverySequence* attribute of the queue. When a message arrives on a queue, it is inserted immediately following the last message that has the same priority.

Programs can either get the first message from a queue, or they can get a particular message from a queue, ignoring the priority of those messages. For example, a program may want to process the reply to a particular message that it sent earlier. For more information, see "Getting a particular message" on page 133.

If an application puts a sequence of messages on a queue, another application can retrieve those messages in the same order that they were put, provided:

- The messages all have the same priority

- The messages were all put within the same unit of work, or all put outside a unit of work

- The queue is local to the putting application

If these conditions are not met, and the applications depend on the messages being retrieved in a certain order, the applications must either include sequencing

information in the message data, or establish a means of acknowledging receipt of a message before the next one is sent.

On MQSeries for MVS/ESA, the queue attribute, *IndexType*, can be used to increase the speed of MQGET operations on the queue. For more information, see "Type of index" on page 135.

Logical and physical ordering

Logical and physical ordering is supported on MQSeries Version 5 products only.

Messages on queues can occur (within each priority level) in *physical* or *logical* order:

Order	Meaning
Physical	This is the order in which messages arrive on a queue.
Logical	This is when all of the messages and segments within a group are in their logical sequence, adjacent to each other, in the position determined by the physical position of the first item belonging to the group.

For a description of groups, messages, and segments, see "Message groups" on page 32. These physical and logical orders may differ because:

- Groups can arrive at a destination at similar times from different applications, therefore losing any distinct physical order.

- Even within a single group, messages may get out of order due to rerouting or delay of some of the messages in the group.

For example, the logical order might look like Figure 8:
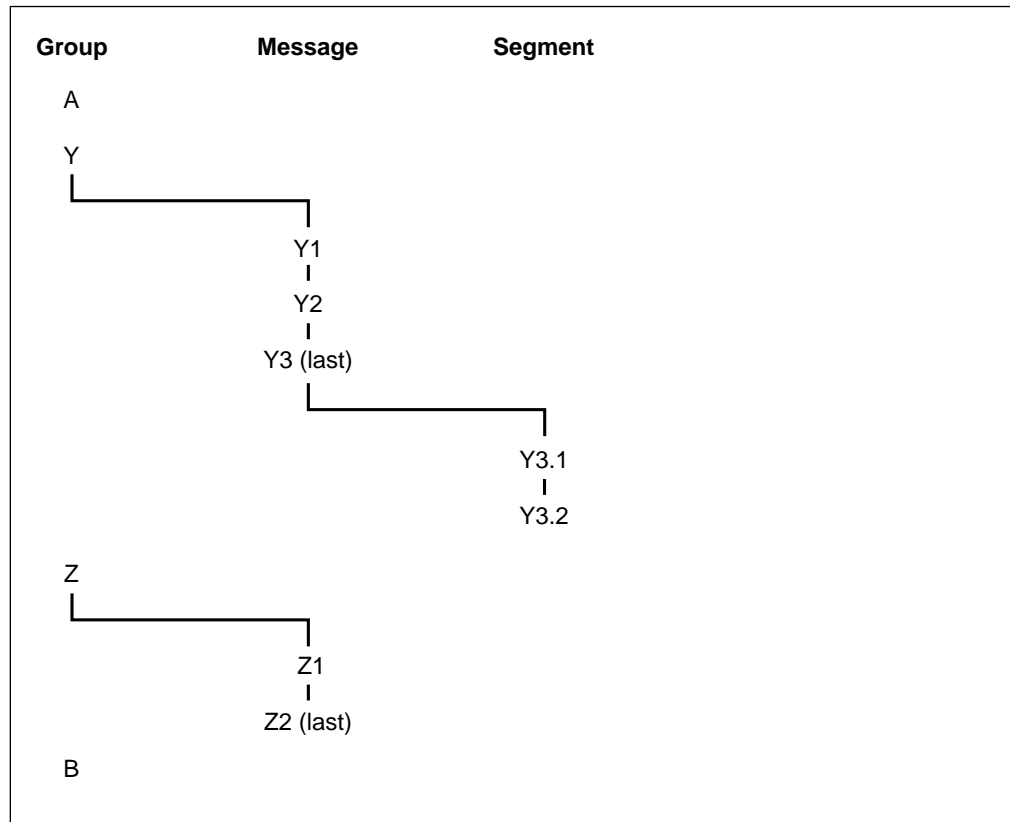
Figure 8. Logical order on a queue

These messages would appear in the following order on a queue:

1. Message A (not in a group)
2. Logical message 1 of group Y
3. Logical message 2 of group Y
4. Segment 1 of (last) logical message 3 of group Y
5. (Last) segment 2 of (last) logical message 3 of group Y
6. Logical message 1 of group Z
7. (Last) logical message 2 of group Z
8. Message B (not in a group)

The physical order, however, might be entirely different. As stated on page 125, the physical position of the *first* item within each group determines the logical position of the whole group. For example, if groups Y and Z arrived at similar times, and message 2 of group Z overtook message 1 of the same group, the physical order would look like Figure 9:
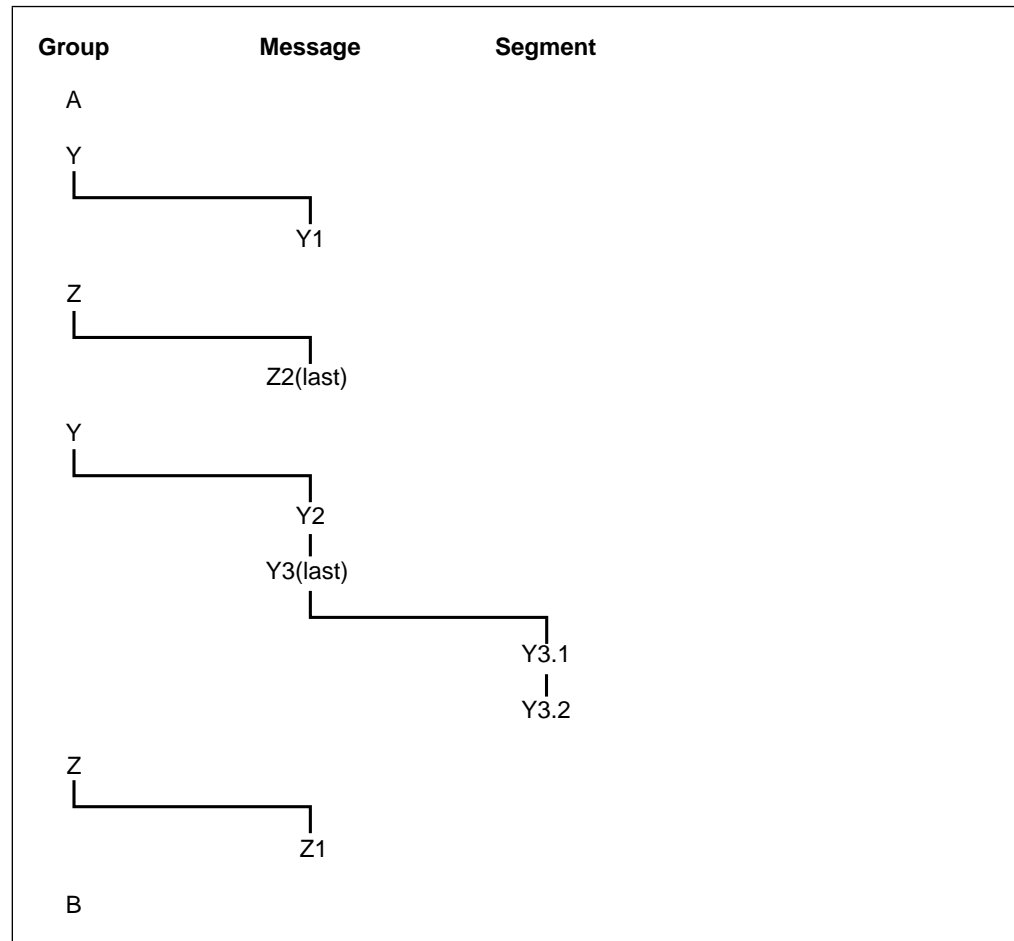
Figure 9. Physical order on a queue

These messages appear in the following order on the queue:

1. Message A (not in a group)
2. Logical message 1 of group Y
3. Logical message 2 of group Z
4. Logical message 2 of group Y
5. Segment 1 of (last) logical message 3 of group Y
6. (Last) segment 2 of (last) logical message 3 of group Y
7. Logical message 1 of group Z
8. Message B (not in a group)

When getting messages, you can specify MQGMO_LOGICAL_ORDER to retrieve messages in logical rather than physical order.

If you issue an MQGET call with MQGMO_BROWSE_FIRST and MQGMO_LOGICAL_ORDER, subsequent MQGETs with MQGMO_BROWSE_NEXT must also specify this option. Conversely, if the MQGET with MQGMO_BROWSE_FIRST does not specify MQGMO_LOGICAL_ORDER, neither must the following MQGETs with MQGMO_BROWSE_NEXT.

The group and segment information that the queue manager retains for MQGET calls that browse messages on the queue is separate from the group and segment information that the queue manager retains for MQGET calls that remove

messages from the queue. When MQGMO_BROWSE_FIRST is specified, the queue manager ignores the group and segment information for browsing, and scans the queue as though there were no current group and no current logical message.

Note: Special care is needed if an MQGET call is used to browse **beyond the end** of a message group (or logical message not in a group) when MQGMO_LOGICAL_ORDER is not specified. For example, if the last message in the group happens to **precede** the first message in the group on the queue, using MQGMO_BROWSE_NEXT to browse beyond the end of the group, specifying MQMO_MATCH_MSG_SEQ_NUMBER with *MsgSeqNumber* set to 1 (to find the first message of the next group) would return again the first message in the group already browsed. This could happen immediately, or a number of MQGET calls later (if there are intervening groups).

The possibility of an infinite loop can be avoided by opening the queue **twice** for browse:

- Use the first handle to browse only the first message in each group.
- Use the second handle to browse only the messages within a specific group.
- Use the MQMO_* options to move the second browse cursor to the position of the first browse cursor, before browsing the messages in the group.
- Do not use the MQGMO_BROWSE_NEXT browse beyond the end of a group.

For further information about this, see the *MQSeries Application Programming Reference* manual.

For most applications you will probably choose either logical or physical ordering when browsing. However, if you want to switch between these modes, remember that when you first issue a browse with MQGMO_LOGICAL_ORDER, your position within the logical sequence is established.

If the first item within the group is not present at this time, the group you are in is not considered to be part of the logical sequence.

Once the browse cursor is within a group, it can continue within the same group, even if the first message is removed. Initially though, you can never move into a group using MQGMO_LOGICAL_ORDER where the first item is not present.

Grouping logical messages

There are two main reasons for using logical messages in a group:

- The messages may need to be processed in the correct order
- Each of the messages in a group may need to be processed in a related way.

In either case, retrieval of the entire group must be carried out by the same getting application instance.

For example, assume that the group consists of four logical messages. The putting application looks like this:

```
    PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

    MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
    MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
    MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
    MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP

    MQCMIT
```

The getting application chooses not to start processing any group until all of the
messages within it have arrived. MQGMO_ALL_MSGS_AVAILABLE is therefore
specified for the first message in the group; the option is ignored for subsequent
messages within the group.

Once the first logical message of the group is retrieved,
MQGMO_LOGICAL_ORDER is used to ensure that the remaining logical messages
of the group are retrieved in order.

So, the getting application looks like this:
```
    /* Wait for the first message in a group, or a message not in a group */
    GMO.Options = MQGMO_SYNCPOINT | MQGMO_WAIT
                | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_LOGICAL_ORDER
    do while ( GroupStatus == MQGS_MSG_IN_GROUP )
       MQGET
       /* Process each remaining message in the group */
       ...

    MQCMIT
```

For further examples of grouping messages, see "Application segmentation of
logical messages" on page 138 and "Putting and getting a group that spans units
of work."

Putting and getting a group that spans units of work
In the previous case, messages or segments cannot start to leave the node (if its
destination is remote) or start to be retrieved until all of the group has been put and
the unit of work is committed. This may not be what you want if it takes a long
time to put the whole group, or if queue space is limited on the node. To overcome
this, the group can be put in several units of work.

If the group is put within multiple units of work, it is possible for some of the group
to commit even when a failure of the putting application occurs. The application
must therefore save status information, committed with each unit of work, which it
can use after a restart to resume an incomplete group. The simplest place to
record this information is in a STATUS queue. If a complete group has been
successfully put, the STATUS queue is empty.

If segmentation is involved, the logic is similar. In this case, the StatusInfo must
include the *Offset*.

MQGET retrieval sequence

Here is an example of putting the group in several units of work:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

/* First UOW */

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

/* Next and subsequent UOWs */
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

/* Last UOW */
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
MQCMIT
```

If all the units of work have been committed, the entire group has been put successfully, and the STATUS queue is empty. If not, the group must be resumed at the point indicated by the status information. MQPMO_LOGICAL_ORDER cannot be used for the first put, but can thereafter.

Restart processing looks like this:

```
MQGET (StatusInfo from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if (Reason == MQRC_NO_MSG_AVAILABLE)
   /* Proceed to normal processing */
   ...

else
   /* Group was terminated prematurely */
   Set GroupId, MsgSeqNumber in MQMD to values from Status message
   PMO.Options = MQPMO_SYNCPOINT
   MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP

   /* Now normal processing is resumed.
      Assume this is not the last message */
   PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT
   MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
   MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
   StatusInfo = GroupId,MsgSeqNumber from MQMD
   MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
   MQCMIT
```

From the getting application, you may want to start processing the messages in a group before the whole group has arrived. This improves response times on the

messages within the group, and also means that storage is not required for the entire group.

For recovery reasons, each message must be retrieved within a unit of work. However, in order to realize the above benefits, several units of work must be used for each group of messages.

As with the corresponding putting application, this requires status information to be recorded somewhere atomically as each unit of work is committed. Again, the simplest place to record this information is on a STATUS queue. If a complete group has been successfully processed, the STATUS queue is empty.

Note:

> For intermediate units of work, you can avoid the MQGETs from the STATUS queue by specifying that each MQPUT to the status queue is a segment of a message (that is, by setting the MQMF_SEGMENT flag), instead of putting a complete new message for each unit of work. In the last unit of work, a final segment is put to the status queue specifying MQMF_LAST_SEGMENT, and then the status information is cleared with an MQGET specifying MQGMO_COMPLETE_MSG.

> During restart processing, instead of using a single MQGET to get a possible status message, browse the status queue with MQGMO_LOGICAL_ORDER until you reach the last segment (that is, until no further segments are returned). In the first unit of work after restart, also specify the offset explicitly when putting the status segment.

In the following example, we consider only messages within a group. It is assumed that the application's buffer is always large enough to hold the entire message, whether or not the message has been segmented. MQGMO_COMPLETE_MSG is therefore specified on each MQGET. The same principles apply if segmentation is involved (in this case, the StatusInfo must include the *Offset*).

For simplicity, we are assuming that a maximum of 4 messages should be retrieved within a single UOW:

```
msgs = 0     /* Counts messages retrieved within UOW */
/* Should be no status message at this point */

/* Retrieve remaining messages in the group */
do while ( GroupStatus == MQGS_MSG_IN_GROUP )

   /* Process up to 4 messages in the group */
   GMO.Options = MQGMO_SYNCPOINT | MQGMO_WAIT
                 | MQGMO_LOGICAL_ORDER
   do while ( (GroupStatus == MQGS_MSG_IN_GROUP) && (msgs < 4) )
      MQGET
      msgs = msgs + 1
      /* Process this message */
      ...
   /* end while */

   /* Have retrieved last message or 4 messages  */
   /* Update status message if not last in group */
   MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
   if ( GroupStatus == MQGS_MSG_IN_GROUP )
      StatusInfo = GroupId,MsgSeqNumber from MQMD
      MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
   MQCMIT
   msgs = 0
/* end while */

if ( msgs > 0 )
   /* Come here if there was only 1 message in the group */
   MQCMIT
```

If all of the units of work have been committed, then the entire group has been retrieved successfully, and the STATUS queue is empty. If not, then the group must be resumed at the point indicated by the status information. MQGMO_LOGICAL_ORDER cannot be used for the first retrieve, but can thereafter.

Restart processing looks like this:

```
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if (Reason == MQRC_NO_MSG_AVAILABLE)
   /* Proceed to normal processing */
   ...

else
   /* Group was terminated prematurely */
   /* The next message on the group must be retrieved by matching
      the sequence number and group id with those retrieved from the
      status information. */
   GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT | MQGMO_WAIT
   MQGET GMO.MatchOptions = MQMO_MATCH_GROUP_ID | MQMO_MATCH_MSG_SEQ_NUMBER,
         MQMD.GroupId       = value from Status message,
         MQMD.MsgSeqNumber = value from Status message plus 1
   msgs = 1
   /* Process this message */
   ...

   /* Now normal processing is resumed */
   /* Retrieve remaining messages in the group */
   do while ( GroupStatus == MQGS_MSG_IN_GROUP )

      /* Process up to 4 messages in the group */
      GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT | MQGMO_WAIT
                  | MQGMO_LOGICAL_ORDER
      do while ( (GroupStatus == MQGS_MSG_IN_GROUP) && (msgs < 4) )
         MQGET
         msgs = msgs + 1
         /* Process this message */
         ...

      /* Have retrieved last message or 4 messages  */
      /* Update status message if not last in group */
      MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
      if ( GroupStatus == MQGS_MSG_IN_GROUP )
         StatusInfo = GroupId,MsgSeqNumber from MQMD
         MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
      MQCMIT
      msgs = 0
```

Getting a particular message

To get a particular message from a queue, use the *MsgId* and *CorrelId* fields of the MQMD structure. Note, however, that applications can explicitly set these fields, so the values you specify may not identify a unique message. Table 4 on page 134 shows which message is retrieved for the possible settings of these fields. These fields are ignored on input if you specify MQGMO_MSG_UNDER_CURSOR in the *GetMsgOpts* parameter of the MQGET call.

Table 4. Using message and correlation identifiers		
To retrieve ...	*MsgId*	*CorrelId*
First message in the queue	MQMI_NONE	MQCI_NONE
First message that matches *MsgId*	Nonzero	MQCI_NONE
First message that matches *CorrelId*	MQMI_NONE	Nonzero
First message that matches both *MsgId* and *CorrelId*	Nonzero	Nonzero

In each case, **first** means the first message that satisfies the selection criteria (unless MQGMO_BROWSE_NEXT is specified, when it means the **next** message in the sequence satisfying the selection criteria).

On return, the MQGET call sets the *MsgId* and *CorrelId* fields to the message and correlation identifiers (respectively) of the message returned (if any).

If you set the *Version* field of the MQMD structure to 2, you can use the *GroupId*, *MsgSeqNumber*, and *Offset* fields. Table 5 shows which message is retrieved for the possible settings of these fields.

Table 5. Using the group identifier	
To retrieve ...	**Match options**
First message in the queue	MQMO_NONE
First message that matches *MsgId*	MQMO_MATCH_MSG_ID
First message that matches *CorrelId*	MQMO_MATCH_CORREL_ID
First message that matches *GroupId*	MQMO_MATCH_GROUP_ID
First message that matches *MsgSeqNumber*	MQMO_MATCH_MSG_SEQ_NUMBER
First message that matches *Offset*	MQMO_MATCH_OFFSET

Notes:

1. MQMO_XXX implies that the *XXX* field in the MQMD structure is set to the value to be matched.

2. The MQMO flags can be used in combination. For example, MQMO_MATCH_GROUP_ID, MQMO_MATCH_MSG_SEQ_NUMBER, and MQMO_MATCH_OFFSET can be used together to give the segment identified by the *GroupId*, *MsgSeqNumber*, and *Offset* fields.

3. If you specify MQGMO_LOGICAL_ORDER, the message you are trying to retrieve is affected because the option depends on state information controlled for the queue handle. For information about this, see "Logical and physical ordering" on page 125 and the *MQSeries Application Programming Reference* manual.

Notes:

1. The MQGET call usually retrieves the first message from a queue. If you specify a particular message when you use the MQGET call, the queue manager has to search the queue until it finds that message. This can affect the performance of your application.

2. If you are using Version 2 of the MQMD structure, you can use the MQMO_MATCH_MSG_ID and MQMO_MATCH_CORREL_ID flags. This avoids having to reset the *MsgId* and *CorrelId* fields between MQGETs.

On MQSeries for MVS/ESA, the queue attribute, *IndexType*, can be used to increase the speed of MQGET operations on the queue. For more information, see "Type of index" on page 135.

Type of index

This is supported on MQSeries for MVS/ESA only.

The queue attribute, *IndexType*, specifies the type of index that the queue manager maintains in order to increase the speed of MQGET operations on the queue.

You have three options:

Value	Description
NONE	No index is maintained. Use this when messages are retrieved sequentially (see "Priority" on page 124).
MSGID	An index of message identifiers is maintained. Use this when messages are retrieved using the *MsgId* field as a selection criterion on the MQGET call (see "Getting a particular message" on page 133).
CORRELID	An index of correlation identifiers is maintained. Use this when messages are retrieved using the *CorrelId* field as a selection criterion on the MQGET call (see "Getting a particular message" on page 133).

Notes:

1. If you are indexing using the MSGID option or CORRELID option, set the relative *MsgId* or *CorrelId* parameters in the MQMD. It is **not** beneficial to set both.

2. Indexes are ignored when browsing messages on a queue (see on page 149).

3. Avoid queues (indexed by *MsgId* or *CorrelId*) containing many messages (that is, thousands) as this affects restart time. (This does not apply to nonpersistent messages as they are deleted at restart.)

For a full description of the *IndexType* attribute, see the *MQSeries Application Programming Reference* manual. For conditions needed to change the *IndexType* attribute, see the *MQSeries Command Reference* manual.

Handling large messages

This is supported on MQSeries Version 5 products, and MQSeries for AS/400 only.

Messages can be too large for the application, queue, or queue manager. MQSeries provides three ways of dealing with large messages:

1. Increase the queue and queue manager *MaxMsgLength* attributes (not AS/400).

2. Use segmented messages. (Messages can be segmented by either the application or the queue manager.)

3. Use reference messages.

Each of these approaches is described in the remainder of this section.

Increasing the maximum message length

The *MaxMsgLength* queue manager attribute defines the maximum length of a message that can be handled by a queue manager. Similarly, the *MaxMsgLength* queue attribute is the maximum length of a message that can be handled by a queue. The **default** maximum message length supported depends on the environment you are working in, as shown under "Message properties and protocols" on page 530.

If you are handling large messages, you can alter these attributes independently. (not AS/400) The attribute value can be set between 32768 bytes and 100 MB.

After changing one or both of the *MaxMsgLength* attributes, restart your applications and channels to ensure that the changes take effect. When these changes are made, the message length must be less than or equal to both the queue and the queue manager *MaxMsgLength* attributes. However, existing messages may be longer than either attribute.

If the message is too big for the queue, MQRC_MSG_TOO_BIG_FOR_Q is returned. Similarly, if the message is too big for the queue manager, MQRC_MSG_TOO_BIG_FOR_Q_MGR is returned. However, these return codes can be avoided by segmenting your messages (see "Message segmentation" for full details of this).

This method of handling large messages is easy and convenient. However, consider the following factors before using it:

- Uniformity among queue managers is reduced. The maximum size of message data is determined by the *MaxMsgLength* for each queue (including transmission queues) on which the message will be put. This value is often defaulted to the queue manager's *MaxMsgLength*, especially for transmission queues. This makes it difficult to predict whether a message is too large when it is to travel to a remote queue manager.

- Usage of system resources is increased. For example, applications need larger buffers, and on some platforms, there may be increased usage of shared storage. Note that queue storage should be affected only if actually required for larger messages.

- Channel batching is affected. A large message still counts as just one message towards the batch count but needs longer to transmit, thereby increasing response times for other messages.

Message segmentation

Increasing the maximum message length as discussed on page 136 has some negative implications. Also, it could still result in the message being too large for the queue or queue manager. In these cases, a message can be segmented. For information about segments, see "Message groups" on page 32.

The next sections look at common uses for segmenting messages. For putting and destructively getting, it is assumed that the MQPUT or MQGET calls **always** operate within a unit of work. It is strongly recommended that this technique is always used, to reduce the possibility of incomplete groups being present in the network. Single-phase commit by the queue manager is assumed, but of course other coordination techniques are equally valid.

Also, in the getting applications, it is assumed that if multiple servers are processing the same queue, each server executes similar code, so that one server never fails to find a message or segment that it expects to be there (because it had specified MQGMO_ALL_MSGS_AVAILABLE or MQGMO_ALL_SEGMENTS_AVAILABLE earlier).

Segmentation and reassembly by queue manager

This is the simplest scenario, in which one application puts a message to be retrieved by another. The message may be large: not too large for either the putting or the getting application to handle in a single buffer, but possibly too large for the queue manager or a queue on which the message is to be put.

The only changes necessary for these applications are for the putting application to authorize the queue manager to perform segmentation if necessary,

```
PMO.Options = (existing options)
MQPUT MD.MsgFlags = MQMF_SEGMENTATION_ALLOWED
```

and for the getting application to ask the queue manager to reassemble the message if it has been segmented:

```
GMO.Options = MQGMO_COMPLETE_MSG | (existing options)
MQGET
```

The application buffer must be large enough to contain the reassembled message (unless the MQGMO_ACCEPT_TRUNCATED_MSG option is included).

If data conversion is necessary it may have to be done by the getting application specifying MQGMO_CONVERT. This should be straightforward because the data conversion exit is presented with the complete message. Attempting to do data conversion in a sender channel will not be successful if the message is segmented, and the format of the data is such that the data-conversion exit cannot carry out the conversion on incomplete data.

Application segmentation

This example shows how to segment a single large message

Application segmentation is used for two main reasons:

1. Queue-manager segmentation alone is not adequate because the message is too large to be handled in a single buffer by the applications.

2. Data conversion must be performed by sender channels, and the format is such that the putting application needs to stipulate where the segment boundaries are to be in order for conversion of an individual segment to be possible.

However, if data conversion is not an issue, or if the getting application always uses MQGMO_COMPLETE_MSG, queue-manager segmentation can also be allowed by specifying MQMF_SEGMENTATION_ALLOWED. In our example, the application segments the message into 4 segments:

```
            PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

        MQPUT MD.MsgFlags = MQMF_SEGMENT
        MQPUT MD.MsgFlags = MQMF_SEGMENT
        MQPUT MD.MsgFlags = MQMF_SEGMENT
        MQPUT MD.MsgFlags = MQMF_LAST_SEGMENT

        MQCMIT
```

If you do not use MQPMO_LOGICAL_ORDER, the application must set the *Offset* and the length of each segment. In this case, logical state is not maintained automatically.

The getting application cannot, or chooses not to, guarantee to have a buffer that will hold any reassembled message. It must therefore be prepared to process segments individually.

For messages that are segmented, this application does not want to start processing one segment until all of the segments that constitute the logical message are present. MQGMO_ALL_SEGMENTS_AVAILABLE is therefore specified for the first segment. If you specify MQGMO_LOGICAL_ORDER and there is a current logical message, MQGMO_ALL_SEGMENTS_AVAILABLE is ignored.

Once the first segment of a logical message has been retrieved, MQGMO_LOGICAL_ORDER is used to ensure that the remaining segments of the logical message are retrieved in order.

No consideration is given to messages within different groups. If such messages do occur, they are processed in the order in which the first segment of each message appears on the queue.

```
    GMO.Options = MQGMO_SYNCPOINT | MQGMO_LOGICAL_ORDER
                | MQGMO_ALL_SEGMENTS_AVAILABLE | MQGMO_WAIT
    do while ( SegmentStatus == MQSS_SEGMENT )
       MQGET
       /* Process each remaining segment of the logical message */
       ...

    MQCMIT
```

Application segmentation of logical messages
The messages must be maintained in logical order in a group, and some or all of them may be so large that they require application segmentation.

In our example, a group of four logical messages is to be put. All but the third message are large, and require segmentation which is performed by the putting application:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP        | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP        | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP        | MQMF_LAST_SEGMENT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP        | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP        | MQMF_LAST_SEGMENT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP

MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_LAST_SEGMENT

MQCMIT
```

In the getting application, MQGMO_ALL_MSGS_AVAILABLE is specified on the first MQGET. This means that no messages or segments of a group are retrieved until the entire group is available. When the first physical message of a group has been retrieved, MQGMO_LOGICAL_ORDER is used to ensure that the segments and messages of the group are retrieved in order:

```
GMO.Options = MQGMO_SYNCPOINT | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_MESSAGES_AVAILABLE | MQGMO_WAIT

do while ( (GroupStatus    != MQGS_LAST_MSG_IN_GROUP) ||
             (SegmentStatus != MQGS_LAST_SEGMENT) )
   MQGET
   /* Process a segment or complete logical message.  Use the GroupStatus
      and SegmentStatus information to see what has been returned */
   ...

MQCMIT
```

Note: If you specify MQGMO_LOGICAL_ORDER and there is a current group, MQGMO_ALL_MSGS_AVAILABLE is ignored.

Putting and getting a segmented message that spans units of work

You can put and get a segmented message that spans a unit of work in a similar way to "Putting and getting a group that spans units of work" on page 129.

Reference messages

This method allows a large object to be transferred from one node to another without the need for the object to be stored on MQSeries queues at either the source or the destination nodes. This is of particular benefit where the data already exists in another form, for example, for mail applications.

To do this, you need to specify a message exit at both ends of a channel. For information on how to do this, see the *MQSeries Intercommunication* book.

MQSeries defines the format of a reference message header (MQRMH). See the *MQSeries Application Programming Reference* manual for a description of this. This is recognized by means of a defined format name and may or may not be followed by actual data.

Handling large messages

To initiate transfer of a large object, an application can put a message consisting of a reference message header with no data following it. As this message leaves the node, the message exit retrieves the object in an appropriate way and appends it to the reference message. It then returns the message (now larger than before) to the sending Message Channel Agent for transmission to the receiving MCA.

Another message exit is configured at the receiving MCA. When this message exit sees one of these messages, it creates the object using the object data that was appended and passes on the reference message *without* it. The reference message can now be received by an application and this application knows that the object (or at least the portion of it represented by this reference message) has been created at this node.

The maximum amount of object data that a sending message exit can append to the reference message is limited by the negotiated maximum message length for the channel. The exit can only return a single message to the MCA for each message that it is passed, so the putting application can put several messages to cause one object to be transferred. Each message must identify the *logical* length and offset of the object that is to be appended to it. However, in cases where it is not possible to know the total size of the object or the maximum size allowed by the channel, the sending message exit can be designed so that the putting application just puts a single message, and the exit itself puts the next message on the transmission queue when it has appended as much data as it can to the message it has been passed.

Before using this method of dealing with large messages, consider the following:

- The MCA and the message exit run under an MQSeries user ID. The message exit (and therefore, the user ID) needs to access the object to either retrieve it at the sending end or create it at the receiving end; this may only be feasible in cases where the object is universally accessible. This raises a security issue.

- If the reference message with bulk data appended to it must travel through several queue managers before reaching its destination, the bulk data *is* present on MQSeries queues at the intervening nodes. However, no special support or exits need to be provided in these cases.

- Designing your message exit is made difficult if rerouting or dead-letter queuing is allowed. In these cases, the portions of the object may arrive out of order.

- When a reference message arrives at its destination, the receiving message exit creates the object. However, this is not synchronized with the MCA's unit of work, so if the batch is backed out, another reference message containing this same portion of the object will arrive in a later batch, and the message exit may attempt to recreate the same portion of the object. If the object is, for example, a series of database updates, this might be unacceptable. If so, the message exit must keep a log of which updates have been applied; this may require the use of an MQSeries queue.

- Depending on the characteristics of the object type, the message exits and applications may need to cooperate in maintaining use counts, so that the object can be deleted when it is no longer needed. An instance identifier may also be required; a field is provided for this in the reference message header (see the *MQSeries Application Programming Reference* manual).

- If a reference message is put as a distribution list, the object must be retrievable for each resulting distribution list or individual destination at that

node. You may need to maintain use counts. Also consider the possibility that a given node may be the final node for some of the destinations in the list, but an intermediate node for others.

- Bulk data is not normally converted. This is because conversion takes place **before** the message exit is invoked. For this reason, conversion should not be requested on the originating sender channel. If the reference message passes through an intermediate node, the bulk data is converted when sent from the intermediate node, if requested.

- Reference messages cannot be segmented.

Using the MQRMH and MQMD structures
See the *MQSeries Application Programming Reference* manual for a description of the fields in the reference message header and the message descriptor.

In the MQMD structure, the *Format* field must be set to MQFMT_REF_MSG_HEADER. The MQHREF format, when requested on MQGET, is converted automatically by MQSeries along with any bulk data that follows.

Here is an example of the use of the *DataLogicalOffset* and *DataLogicalLength* fields of the MQRMH:

A putting application might put a reference message with:

- No physical data
- *DataLogicalLength* = 0 (this message represents the entire object)
- *DataLogicalOffset* = 0.

Assuming that the object is 70,000 bytes long, the sending message exit sends the first 40,000 bytes along the channel in a reference message containing:

- 40,000 bytes of physical data following the MQRMH
- *DataLogicalLength* = 40,000
- *DataLogicalOffset* = 0 (from the start of the object).

It then places another message on the transmission queue containing:

- No physical data

- *DataLogicalLength* = 0 (to the end of the object). You could specify a value of 30,000 here.

- *DataLogicalOffset* = 40,000 (starting from this point).

When this message exit is seen by the sending message exit, the remaining 30,000 bytes of data is appended, and the fields are set to:

- 30,000 bytes of physical data following the MQRMH
- *DataLogicalLength* = 30,000
- *DataLogicalOffset* = 40,000 (starting from this point).

The MQRMHF_LAST flag is also set.

For a description of the sample programs provided for the use of reference messages, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

Waiting for messages

If you want a program to wait until a message arrives on a queue, specify the MQGMO_WAIT option in the *Options* field of the MQGMO structure. Use the *WaitInterval* field of the MQGMO structure to specify the maximum time (in milliseconds) that you want an MQGET call to wait for a message to arrive on a queue.

If the message does not arrive within this time, the MQGET call completes with the MQRC_NO_MSG_AVAILABLE reason code.

You can specify an unlimited wait interval using the constant MQWI_UNLIMITED in the *WaitInterval* field. However, events outside your control could cause your program to wait for a long time, so use this constant with caution. IMS applications should not specify an unlimited wait interval because this would prevent the IMS system terminating. (When IMS terminates, it requires all dependent regions to end.) Instead, IMS applications should specify a finite wait interval; then, if the call completes without retrieving a message after that interval, issue another MQGET call with the wait option.

In the Windows 3.1 environment, while your application is waiting for an MQGET to return, MQSeries will still recover Windows messages to allow the application and the rest of Windows to function normally. You must ensure that your code that processes Windows program messages does not assume that the MQGET call returns data to the application immediately. If it attempts to access data that is not yet available, errors can easily occur. Also, if you attempt to make other MQI calls while the MQGET call is waiting, MQRC_CALL_IN_PROGRESS is returned to show that another call is busy.

Note: If more than one program is waiting on the same shared queue to **remove** a message, only one program is activated by a message arriving. However, if more than one program is waiting to browse a message, all the programs can be activated. For more information, see the description of the *Options* field of the MQGMO structure in the *MQSeries Application Programming Reference* manual.

If the state of the queue or the queue manager changes before the wait interval expires, the following actions occur:

- If the queue manager enters the quiescing state, and you used the MQGMO_FAIL_IF_QUIESCING option, the wait is canceled and the MQGET call completes with the MQRC_Q_MGR_QUIESCING reason code. Without this option, the call remains waiting.

- On MVS/ESA, if the connection (for a CICS or IMS application) enters the quiescing state, and you used the MQGMO_FAIL_IF QUIESCING option, the wait is canceled and the MQGET call completes with the MQRC_CONN_QUIESCING reason code. Without this option, the call remains waiting.

- If the queue manager is forced to stop, or is canceled, the MQGET call completes with either the MQRC_Q_MGR_STOPPING or the MQRC_CONNECTION_BROKEN reason code.

- If the attributes of the queue (or a queue to which the queue name resolves) are changed so that get requests are now inhibited, the wait is canceled and the MQGET call completes with the MQRC_GET_INHIBITED reason code.

- If the attributes of the queue (or a queue to which the queue name resolves) are changed in such a way that the FORCE option is required, the wait is canceled and the MQGET call completes with the MQRC_OBJECT_CHANGED reason code.

If you want your application to wait on more than one queue, use the signal facility of MQSeries for MVS/ESA (see "Signaling"). For more information about the circumstances in which these actions occur, see the *MQSeries Application Programming Reference* manual.

Signaling

| *Signaling is supported on MQSeries for MVS/ESA, MQSeries for Tandem NSK, and MQSeries for Windows Version 2.1 only*.

Signaling is an option on the MQGET call to allow the operating system to notify (or *signal*) a program when an expected message arrives on a queue. This is similar to the "get with wait" function described on page 142 because it allows your program to continue with other work while waiting for the signal. However, if you use signaling, you can free the application thread and rely on the operating system to notify the program when a message arrives.

To set a signal

To set a signal, do the following in the MQGMO structure that you use on your MQGET call:

1. Set the MQGMO_SET_SIGNAL option in the *Options* field.

2. Set the maximum life of the signal in the *WaitInterval* field. This sets the length of time (in milliseconds) for which you want MQSeries to monitor the queue. Use the MQWI_UNLIMITED value to specify an unlimited life.

   Note: IMS applications should not specify an unlimited wait interval because this would prevent the IMS system from terminating. (When IMS terminates, it requires all dependent regions to end.) Instead, IMS applications should examine the state of the ECB at regular intervals (see step 3).

   A program can have signals set on several queue handles at the same time:

| 3. On MQSeries for Tandem NSK specify an application tag in the *Signal1* field.
| This can be used by an application to associate the IPC notification message
| (see "When the message arrives" on page 144) with a particular MQGET call.

   On MQSeries for Windows Version 2.1, specify the handle of the window to which you want the signal sent in the *Signal1* field.

   On MQSeries for MVS/ESA, specify the address of the *Event Control Block* (ECB) in the *Signal1* field. This notifies you of the result of your signal. The ECB storage must remain available until the queue is closed.

4. On MQSeries for Windows Version 2.1, specify an identifier for the signal message in the *Signal2* field. This specifies the Windows message that you receive when a suitable message arrives. Use a RegisterWindow message to find a suitable identifier.

   Note: You cannot use the MQGMO_SET_SIGNAL option in conjunction with the MQGMO_WAIT option.

When the message arrives

When a suitable message arrives, the following occurs:

| - On MQSeries for Tandem NSK An Inter-Process Communication (IPC)
| message is sent to the $RECEIVE queue of the process that made the
| MQGET call.

- On MQSeries for Windows Version 2.1, MQSeries sends a Windows message (identified in step 4) to the window you specified in your *Signal1* field. It also puts a completion code in the WPARAM field of the Windows message.

- On MQSeries for MVS/ESA, a completion code is returned to the ECB.

The completion code describes one of the following:

- The message you set the signal for has arrived on the queue. The message is not reserved for the program that requested a signal, so the program must issue an MQGET call again to get the message.

  Note: Another application could get the message in the time between you receiving the signal and you issuing another MQGET call.

- The wait interval you set has expired and the message you set the signal for did not arrive on the queue. MQSeries has canceled the signal.

- The signal has been canceled. This happens, for example, if the queue manager stops or the attribute of the queue is changed so that MQGET calls are no longer allowed.

When a suitable message is already on the queue, the MQGET call completes in the same way as an MQGET call without signaling. Also, if an error is detected immediately, the call completes and the return codes are set.

When the call is accepted and no message is immediately available, control is returned to the program so that it can continue with other work. None of the output fields in the message descriptor are set, but the *CompCode* and *Reason* parameters are set to MQCC_WARNING and MQRC_SIGNAL_REQUEST_ACCEPTED, respectively.

For information on what MQSeries can return to your application when it makes an MQGET call using signaling, see the *MQSeries Application Programming Reference* manual.

On MQSeries for MVS/ESA, if the program has no other work to do while it is waiting for the ECB to be posted, it can wait for the ECB using:

- For a CICS for MVS/ESA program, the EXEC CICS WAIT EXTERNAL command
- For a CICS/MVS program, the DFHKC macro
- For batch and IMS programs, the MVS WAIT macro

If the state of the queue or the queue manager changes while the signal is set (that is, the ECB has not yet been posted), the following actions occur:

- If the queue manager enters the quiescing state, and you used the MQGMO_FAIL_IF_QUIESCING option, the signal is canceled. The ECB is posted with the MQEC_Q_MGR_QUIESCING completion code. Without this option, the signal remains set.

- If the queue manager is forced to stop, or is canceled, the signal is canceled. The signal is delivered with the MQEC_WAIT_CANCELED completion code.

- If the attributes of the queue (or a queue to which the queue name resolves) are changed so that get requests are now inhibited, the signal is canceled. The signal is delivered with the MQEC_WAIT_CANCELED completion code.

Notes:

1. If more than one program has set a signal on the same shared queue to remove a message, only one program is activated by a message arriving. However, if more than one program is waiting to browse a message, all the programs can be activated. The rules that the queue manager follows when deciding which applications to activate are the same as those for waiting applications: for more information, see the description of the *Options* field of the MQGMO structure in the *MQSeries Application Programming Reference* manual.

2. If there is more than one MQGET call waiting for the same message, with a mixture of wait and signal options, each waiting call is considered equally. For more information, see the description of the *Options* field of the MQGMO structure in the *MQSeries Application Programming Reference* manual.

3. Under some conditions, it is possible both for an MQGET call to retrieve a message and for a signal (resulting from the arrival of the same message) to be delivered. This means that when your program issues another MQGET call (because the signal was delivered), there could be no message available. You should design your program to test for this situation.

For information about how to set a signal, see the description of the MQGMO_SET_SIGNAL option and the *Signal1* field in the *MQSeries Application Programming Reference* manual.

Skipping backout

Skipping backout is supported on MQSeries for MVS/ESA only.

As part of a unit of work, an application program can issue one or more MQGET calls to get messages from a queue. If the application program detects an error, it can back out the unit of work. This restores all the resources updated during that unit of work to the state they were in before the unit of work started, and reinstates the messages retrieved by the MQGET calls.

Once reinstated, these messages are available to subsequent MQGET calls issued by the application program. In many cases, this does not cause a problem for the application program. However, in cases where the error leading to the backout cannot be circumvented, having the message reinstated on the queue can cause the application program to enter an 'MQGET–error–backout' loop.

To avoid this problem, specify the MQGMO_MARK_SKIP_BACKOUT option on the MQGET call. This marks the MQGET request as not being involved in application-initiated backout; that is, it should not be backed out. Use of this option means that when a backout occurs, updates to other resources are backed out as required, but the marked message is treated as if it had been retrieved under a new unit of work. The application program can then perform exception handling, such as informing the originator that the message has been discarded, and then

Skipping backout

commit the new unit of work, causing the message to be removed from the queue. If the new unit of work is backed out (for any reason) the message is reinstated on the queue.

Within a unit of work, there can be only one MQGET request marked as skipping backout; however, there can be several other messages that are not marked as skipping backout. Once a message has been marked as skipping backout, any further MQGET calls within the unit of work that specify MQGMO_MARK_SKIP_BACKOUT will fail with reason code MQRC_SECOND_MARK_NOT_ALLOWED.

Note: The marked message only skips backout if the unit of work containing it is terminated by an application request to back it out. If the unit of work is backed out for any other reason, the message is backed out on to the queue in the same way that it would be if it was not marked to skip backout.
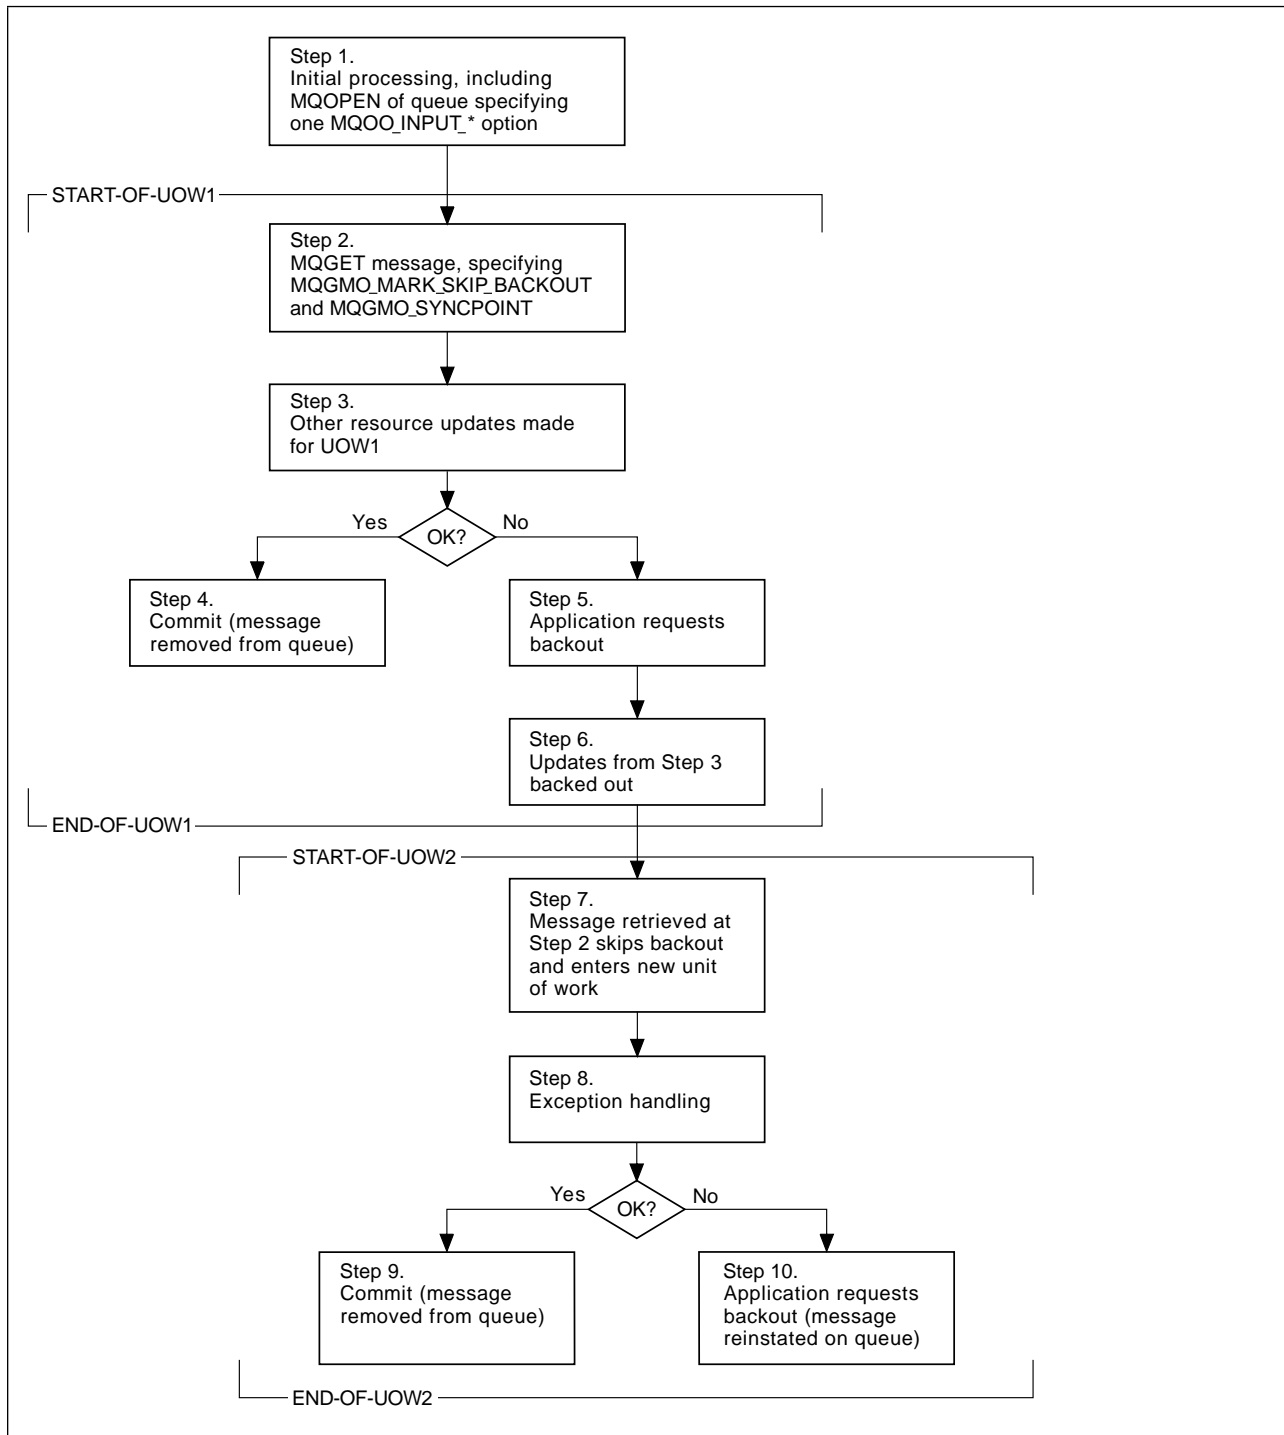
Figure 10. Skipping backout using MQGMO_MARK_SKIP_BACKOUT

Figure 10 illustrates a typical sequence of steps that an application program might contain when an MQGET request is required to skip backout:

Step 1   Initial processing occurs within the transaction, including an MQOPEN call to open the queue (specifying one of the MQOO_INPUT_* options in order to get messages from the queue in Step 2).

Step 2   MQGET is called, with MQGMO_SYNCPOINT and MQGMO_MARK_SKIP_BACKOUT. MQGMO_SYNCPOINT is required because MQGET must be within a unit of work for

MQGMO_MARK_SKIP_BACKOUT to be effective. In Figure 10 this unit of work is referred to as UOW1.

Step 3   Other resource updates are made as part of UOW1. These may include further MQGET calls (issued without MQGMO_MARK_SKIP_BACKOUT).

Step 4   All updates from Steps 2 and 3 complete as required. The application program commits the updates and UOW1 ends. The message retrieved in Step 2 is removed from the queue.

Step 5   Some of the updates from Steps 2 and 3 do not complete as required. The application program requests that the updates made during these steps are backed out.

Step 6   The updates made in Step 3 are backed out.

Step 7   The MQGET request made in Step 2 skips backout and becomes part of a new unit of work, UOW2.

Step 8   UOW2 performs exception handling in response to UOW1 being backed out. (For example, an MQPUT call to another queue, indicating that a problem occurred that caused UOW1 to be backed out.)

Step 9   Step 8 completes as required, the application program commits the activity, and UOW2 ends. As the MQGET request is part of UOW2 (see Step 7), this commit causes the message to be removed from the queue.

Step 10 Step 8 does not complete as required and the application program backs out UOW2. Because the get message request is part of UOW2 (see Step 7), it too is backed out and reinstated on the queue. It is now available to further MQGET calls issued by this or another application program (in the same way as any other message on the queue).

Application data conversion

When necessary, MCAs convert the message descriptor data into the required character set and encoding. Either end of the link (that is, the local MCA or the remote MCA) may do the conversion.

When an application puts messages on a queue, the local queue manager adds control information to the message descriptors to facilitate the control of the messages when they are processed by queue managers and MCAs. Depending on the environment, the message header data fields will be created in the character set and encoding of the local system.

When you move messages between systems, it is necessary, on some occasions, to convert the application data into the character set and encoding required by the receiving system. This can be done either from within application programs on the receiving system or by the MCAs on the sending system. If data conversion is supported on the receiving system, it is recommended to use application programs to convert the application data, rather than depending on the conversion having already occurred at the sending system.

Application data is converted within an application program when the MQGMO_CONVERT option is specified in the `Options` field of the MQGMO structure passed to an MQGET call, and **all** of the following are true:

- The *CodedCharSetId* or *Encoding* fields set in the MQMD structure associated with the message on the queue differ from the *CodedCharSetId* or *Encoding* fields set in the MQMD structure specified on the MQGET call.

- The *Format* field in the MQMD structure associated with the message is not MQFMT_NONE.

- The *BufferLength* specified on the MQGET call is not zero.

- The message data length is not zero.

- The queue manager supports conversion between the *CodedCharSetId* and *Encoding* fields specified in the MQMD structures associated with the message and the MQGET call. See the *MQSeries Application Programming Reference* for details of the coded character set identifiers and machine encodings supported.

- The queue manager supports conversion of the message format. If the *Format* field of the MQMD structure associated with the message is one of the built-in formats, the queue manager is able to convert the message. If the *Format* is not one of the built-in formats, you need to write a data-conversion exit to convert the message.

If the sending MCA is to convert the data, the CONVERT(YES) keyword must be specified on the definition of each sender or server channel for which conversion is required. If the data conversion fails, the message is sent to the DLQ at the sending queue manager and the *Feedback* field of the MQDLH structure indicates the reason. If the message cannot be put on the DLQ, the channel will close and the unconverted message will remain on the transmission queue. Data conversion within applications rather than at sending MCAs avoids this situation.

As a general rule, data in the message that is described as "character" data by the built-in format or data-conversion exit is converted from the coded character set used by the message to that requested, and "numeric" fields are converted to the encoding requested.

For further details of the conversion processing conventions used when converting the built-in formats, and for information about writing your own data-conversion exits, see Chapter 11, "Data-conversion exits" on page 157. See also the *MQSeries Application Programming Reference* for information about the language support tables and about the supported machine encodings.

Browsing messages on a queue

To use the MQGET call to browse the messages on a queue:

1. Call MQOPEN to open the queue for browsing, specifying the MQOO_BROWSE option.

2. To browse the first message on the queue, call MQGET with the MQGMO_BROWSE_FIRST option. To find the message you want, you can call MQGET repeatedly with the MQGMO_BROWSE_NEXT option to step through many messages.

   You **must** set the *MsgId* and *CorrelId* fields of the MQMD structure to null after each MQGET call in order to see all messages.

3. Call MQCLOSE to close the queue.

The browse cursor

When you open (MQOPEN) a queue for browsing, the call establishes a browse cursor for use with MQGET calls that use one of the browse options. You can think of the browse cursor as a logical pointer that is positioned before the first message on the queue.

You can have more than one browse cursor active (from a single program) by issuing several MQOPEN requests for the same queue.

When you call MQGET for browsing, use one of the following options in your MQGMO structure:

MQGMO_BROWSE_FIRST   Gets a copy of the first message that satisfies the conditions specified in your MQMD structure.

MQGMO_BROWSE_NEXT   Gets a copy of the next message that satisfies the conditions specified in your MQMD structure.

In both cases, the message remains on the queue.

When you open a queue, the browse cursor is positioned logically just before the first message on the queue. This means that if you make your MQGET call immediately after your MQOPEN call, you can use the MQGMO_BROWSE_NEXT option to browse the first message; you do not have to use the MQGMO_BROWSE_FIRST option.

The order in which messages are copied from the queue is determined by the *MsgDeliverySequence* attribute of the queue. (For more information, see "The order in which messages are retrieved from a queue" on page 124.)

Queues in FIFO (first in, first out) sequence

The first message in a queue in this sequence is the message that has been on the queue the longest.

Use MQGMO_BROWSE_NEXT to read the messages sequentially in the queue. You will see any messages put to the queue while you are browsing, as a queue in this sequence will have messages placed at the end. When the cursor has recognized that it has reached the end of the queue, the browse cursor will stay where it is and return with MQRC_NO_MSG_AVAILABLE. You may then either leave it there waiting for further messages or reset it to the beginning of the queue with a MQGMO_BROWSE_FIRST call.

Queues in priority sequence

The first message in a queue in this sequence is the message that has been on the queue the longest and has the highest priority at the time the MQOPEN call is issued.

Use MQGMO_BROWSE_NEXT to read the messages in the queue.

The browse cursor will point to the next message, working from the priority of the first message to finish with the message at the lowest priority. It will browse any messages put to the queue during this time as long as they are of equal or lower priority to the message identified by the current browse cursor.

Any messages put to the queue of higher priority can only be browsed by:

- Opening the queue for browse again, at which point a new browse cursor is established
- Using the MQGMO_BROWSE_FIRST option

Uncommitted messages

An uncommitted message is never visible to a browse, the browse cursor skips passed it. Not until the unit-of-work is committed can the messages within it be browsed.

Change to queue sequence

If the message delivery sequence is changed from priority to FIFO while there are messages on the queue, the order of the messages already queued is not changed. Messages added to the queue subsequently take the default priority of the queue.

Browsing messages when message length unknown

To browse a message when you do not know the size of the message, and you do not wish to use the *MsgId*, *CorrelId*, or *GroupId* fields to locate the message, you can use the MQGMO_BROWSE_MSG_UNDER_CURSOR option (not supported on MVS/ESA):

1. Issue an MQGET with:

   - Either the MQGMO_BROWSE_FIRST or MQGMO_BROWSE_NEXT option
   - The MQGMO_ACCEPT_TRUNCATED_MSG option
   - Buffer length zero

   Note: If another program is likely to get the same message, consider using the MQGMO_LOCK option as well. MQRC_TRUNCATED_MSG_ACCEPTED should be returned.

2. Use the returned *DataLength* to allocate the storage needed.

3. Issue an MQGET with the MQGMO_BROWSE_MSG_UNDER_CURSOR.

The message pointed to is the last one that was retrieved; the browse cursor will not have moved. You can choose to either lock the message using the MQGMO_LOCK option, or unlock a locked message using MQGMO_UNLOCK option.

The call fails if no MQGET with either the MQGMO_BROWSE_FIRST or MQGMO_BROWSE_NEXT options has been issued successfully since the queue was opened.

Restriction

MQGMO_LOCK and MQGMO_UNLOCK are not available on MQSeries for Tandem NSK and MQSeries for MVS/ESA.

Removing a message you have browsed

You can remove from the queue a message you have already browsed provided you have opened the queue for removing messages as well as for browsing. (You must specify one of the MQOO_INPUT_⋆ options, as well as the MQOO_BROWSE option, on your MQOPEN call.)

To remove the message, call MQGET again, but in the *Options* field of the MQGMO structure, specify MQGMO_MSG_UNDER_CURSOR. In this case, the MQGET call ignores the *MsgId*, *CorrelId*, and *GroupId* fields of the MQMD structure.

In the time between your browsing and removal steps, another program may have removed messages from the queue, including the message under your browse cursor. In this case, your MQGET call returns a reason code to say that the message is not available.

Browsing messages in logical order

| *Browsing messages in logical order is supported on MQSeries Version 5 products and MQSeries for AS/400 only*.

"Logical and physical ordering" on page 125 discusses the difference between the logical and physical order of messages on a queue. This distinction is particularly important when browsing a queue, because, in general, messages are not being deleted and browse operations do not necessarily start at the beginning of the queue. If an application browses through the various messages of one group (using logical order), it is important that logical order should be followed to reach the start of the next group, since the last message of one group may occur physically **after** the first message of the next group. The MQGMO_LOGICAL_ORDER option ensures that logical order is followed when scanning a queue.

MQGMO_ALL_MSGS_AVAILABLE (or MQGMO_ALL_SEGMENTS_AVAILABLE) needs to be used with care for browse operations. Consider the case of logical messages with MQGMO_ALL_MSGS_AVAILABLE. The effect of this is that a logical message is available only if all of the remaining messages in the group are also present. If they are not, the message is passed over. This can mean that when the missing messages arrive subsequently, they will not be noticed by a browse-next operation.

For example, if the following logical messages are present,

```
Logical message 1 (not last) of group 123
Logical message 1 (not last) of group 456
Logical message 2 (last)     of group 456
```

and a browse function is issued with MQGMO_ALL_MSGS_AVAILABLE, the first logical message of group 456 is returned, leaving the browse cursor on this logical message. If the second (last) message of group 123 now arrives,

```
Logical message 1 (not last) of group 123
Logical message 2 (last)     of group 123
Logical message 1 (not last) of group 456  <=== browse cursor
Logical message 2 (last)     of group 456
```

and the same browse-next function is issued, it will not be noticed that group 123 is now complete, because the first message of this group is *before* the browse cursor.

In some cases (for example, if messages are retrieved destructively when the group is present in its entirety), it may be acceptable to use MQGMO_ALL_MSGS_AVAILABLE together with MQGMO_BROWSE_FIRST. Otherwise, the browse scan must be repeated in order to take note of newly arrived messages that have been missed; just issuing MQGMO_WAIT together with MQGMO_BROWSE_NEXT and MQGMO_ALL_MSGS_AVAILABLE does not take account of them. (This also happens to higher-priority messages that might arrive after scanning the messages is complete.)

The next sections look at browsing examples that deal with unsegmented messages; segmented messages follow similar principles.

Browsing messages in groups

In this example, the application browses through each message on the queue, in logical order.

Messages on the queue may either be grouped or not. For grouped messages, the application does not want to start processing any group until all of the messages within it have arrived. MQGMO_ALL_MSGS_AVAILABLE is therefore specified for the first message in the group; for subsequent messages in the group, this option is unnecessary.

MQGMO_WAIT is used in this example. However, although the wait can be satisfied if a new group arrives, for the reasons in "Browsing messages in logical order" on page 152, it will not be satisfied if the browse cursor has already passed the first logical message in a group, and the remaining messages now arrive. Nevertheless, waiting for a suitable interval ensures that the application does not constantly loop while waiting for new messages or segments.

MQGMO_LOGICAL_ORDER is used throughout, to ensure that the scan is in logical order. This contrasts with the destructive MQGET example, where because each group is being removed, MQGMO_LOGICAL_ORDER is not used when looking for the first (or only) message in a group.

It is assumed that the application's buffer is always large enough to hold the entire message, whether or not the message has been segmented. MQGMO_COMPLETE_MSG is therefore specified on each MQGET.

The following gives an example of browsing logical messages in a group:

```
/* Browse the first message in a group, or a message not in a group */
GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
            | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_WAIT
MQGET GMO.MatchOptions = MQMO_MATCH_MSG_SEQ_NUMBER, MD.MsgSeqNumber = 1
/* Examine first or only message */
...

GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
   MQGET
   /* Examine each remaining message in the group */
   ...
```

The above group is repeated until MQRC_NO_MSG_AVAILABLE is returned.

Browsing and retrieving destructively
In this example, the application browses each of the logical messages within a group, before deciding whether to retrieve that group destructively.

The first part of this example is similar to the previous one. However in this case, having browsed an entire group, we may decide to go back and retrieve it destructively.

As each group is removed in this example, MQGMO_LOGICAL_ORDER is not used when looking for the first or only message in a group.

The following gives an example of browsing and then retrieving destructively:

```
GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
            | MQGMO_ALL_MESSAGES_AVAILABE | MQGMO_WAIT
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
   MQGET
   /* Examine each remaining message in the group (or as many as
      necessary to decide whether or not to get it destructively) */
   ...

if ( we want to retrieve the group destructively )

   if ( GroupStatus == ' ' )
      /* We retrieved an ungrouped message */
      GMO.Options = MQGMO_MSG_UNDER_CURSOR | MQGMO_SYNCPOINT
      MQGET GMO.MatchOptions = 0
      /* Process the message */
      ...

   else
      /* We retrieved one or more messages in a group.  The browse cursor */
      /* will not normally be still on the first in the group, so we have */
      /* to match on the GroupId and MsgSeqNumber = 1.                     */
      /* Another way, which works for both grouped and ungrouped messages,*/
      /* would be to remember the MsgId of the first message when it was   */
      /* browsed, and match on that.                                       */
      GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT
      MQGET GMO.MatchOptions = MQMO_MATCH_GROUP_ID
                             | MQMO_MATCH_MSG_SEQ_NUMBER,
          (MQMD.GroupId       = value already in the MD)
           MQMD.MsgSeqNumber = 1
      /* Process first or only message */
      ...

      GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT
                  | MQGMO_LOGICAL_ORDER
      do while ( GroupStatus == MQGS_MSG_IN_GROUP )
         MQGET
         /* Process each remaining message in the group */
         ...
```

Some cases where the MQGET call fails

If certain attributes of a queue are changed using the FORCE option on a command between issuing an MQOPEN and an MQGET call, the MQGET call fails and returns the MQRC_OBJECT_CHANGED reason code. The queue manager marks the object handle as being no longer valid. This also happens if the changes apply to any queue to which the queue name resolves. The attributes that affect the handle in this way are listed in the description of the MQOPEN call in the *MQSeries Application Programming Reference* manual. If your call returns the MQRC_OBJECT_CHANGED reason code, close the queue, reopen it, then try to get a message again.

If get operations are inhibited for a queue from which you are attempting to get messages (or any queue to which the queue name resolves), the MQGET call fails and returns the MQRC_GET_INHIBITED reason code. This happens even if you are using the MQGET call for browsing. You may be able to get a message successfully if you attempt the MQGET call at a later time, if the design of the application is such that other programs change the attributes of queues regularly.

If a dynamic queue (either temporary or permanent) has been deleted, MQGET calls using a previously acquired object handle fail and return the MQRC_Q_DELETED reason code.

MQGET failure

Chapter 11. Data-conversion exits

Data-conversion exits are not supported on MQSeries for Windows.

The Message Descriptor of a message is created by your application when you do an MQPUT. As MQSeries needs to be able to understand the contents of the MQMD regardless of the platform it is created on, it is converted automatically by the system.

Application data, however, is not converted automatically. If character data is being exchanged between platforms where the *CodedCharSetId* and *Encoding* fields differ, for example, between ASCII and EBCDIC, it is the responsibility of the application to arrange for conversion of the message. Application data conversion may be performed by the queue manager itself or by a user exit program, referred to as a *data-conversion exit*. This chapter discusses the data-conversion exit facility that MQSeries provides.

Control may be passed to the data-conversion exit during an MQGET call. This avoids converting across different platforms before reaching the final destination. However, if the final destination is a platform that does not support data conversion on the MQGET, you must specify CONVERT(YES) on the sender channel that sends the data to its final destination. This ensures that MQSeries converts the data during transmission. In this case, your data-conversion exit must reside on the system where the sender channel is defined.

The MQGET call can be issued directly by an application. Set the *CodedCharSetId* field of the MQMD to MQCCSI_DEFAULT to pick up the default CCSID of the queue manager. This ensures that MQSeries knows the correct target CCSID.

The conditions required for the data-conversion exit to be called are defined for the MQGET call in the *MQSeries Application Programming Reference* manual.

For a description of the parameters that are passed to the data-conversion exit and detailed usage notes see the MQDATACONVEXIT call and the MQDXP structure in the *MQSeries Application Programming Reference* manual.

Programs that convert application data between different machine encodings and CCSIDs must conform to the MQSeries data conversion interface (DCI).

Invoking the data-conversion exit

A data-conversion exit is a user-written exit that receives control during the processing of an MQGET call. The exit is invoked if the following are true:

- The MQGMO_CONVERT option is specified on the MQGET call.

- The *CodedCharSetId* or *Encoding* fields in the MQMD structure associated with the message on the queue differ from the *CodedCharSetId* or *Encoding* fields in the MQMD structure specified on the MQGET call (see the language support tables in the *MQSeries Intercommunication* book).

- The *Format* field in the MQMD structure associated with the message is not MQFMT_NONE (MQFMT_STRING indicates that the message consists entirely of character data).

- The *BufferLength* specified on the MQGET call is not zero.

- The message data length is not zero.

- Either the message format is not one that can be handled by one of the built-in conversion routines, or its format can be handled by one of the built-in conversion routines but the routine is unable to convert the message itself. The conversion routines supplied with the product **always** attempt to convert the built-in format messages first; user-written routines are called only if these product-supplied routines fail to convert.

There are some other conditions, described fully in the usage notes of the MQDATACONVEXIT call in the *MQSeries Application Programming Reference* manual.

See the *MQSeries Application Programming Reference* manual for details of the MQGET call. Data-conversion exits cannot use MQI calls, other than MQXCNVC.

A new copy of the exit is loaded when an application attempts to retrieve the first message that uses that *Format* since the application connected to the queue manager. A new copy may also be loaded at other times if the queue manager has discarded a previously-loaded copy.

The data-conversion exit runs in an environment similar to that of the program which issued the MQGET call. As well as user applications, the program can be an MCA (message channel agent) sending messages to a destination queue manager that does not support message conversion. The environment includes address space and user profile, where applicable. The exit cannot compromise the queue manager's integrity, since it does not run in the queue manager's environment.

In a client-server environment, the exit is loaded at the server, and conversion takes place there.

Data conversion on MVS/ESA

On MVS/ESA, you must also be aware of the following:

- Exit programs can only be written in assembler language.

- They must be reentrant, and capable of running anywhere in storage.

- They must restore the environment on exit to that at entry, and must free any storage obtained.

- They must not WAIT, or issue ESTAEs or SPIEs.

- They are normally invoked as if by MVS LINK in:

    - Non-authorized problem program state
    - Primary address space control mode
    - Non cross-memory mode
    - Non access-register mode
    - 31 bit addressing mode
    - TCB-PRB mode

- Data conversion is not supported by the CICS/MVS R212 adapter. It is only supported by the CICS/ESA adapter.

- When used by a CICS application, the exit is invoked by EXEC CICS LINK, and should conform to the CICS programming conventions. The parameters

are passed by pointers (addresses) in the CICS communication area (COMMAREA).

Although not recommended, user exit programs can also make use of CICS API calls, with the following caution:

– Do not issue syncpoints, as the results could influence units of work declared by the MCA.

– Do not update any resources controlled by a resource manager other than MQSeries for MVS/ESA, including those controlled by CICS for MVS/ESA.

- For distributed queuing without CICS, the exit is loaded from the data set referenced by the CSQXLIB DD statement. In other environments, the exit is loaded from the same place as application programs.

- For distributed queuing using CICS, data-conversion exits are not supported.

Conversion processing

The queue manager performs the following actions if the MQGMO_CONVERT option is specified on the MQGET call, and there is a message to be returned to the application:

1. If one or more of the following is true, no conversion of the message data is necessary:

   - The *CodedCharSetId* and *Encoding* values in the control information in the message are identical to those in the *MsgDesc* parameter.

   - The length of the application message data is zero.

   - The length of the *Buffer* parameter is zero.

   However, although there is no message data to convert, the *CodedCharSetId* and *Encoding* values in the *MsgDesc* parameter are set to the values supplied on the MQGET call, unless the length of the *Buffer* parameter is zero and the MQGMO_ACCEPT_TRUNCATED_MSG option was **not** specified on the MQGET call. In this case, the *CodedCharSetId* and *Encoding* values in the *MsgDesc* parameter are set to values in the control information in the message.

The following steps are performed only if the *CodedCharSetId* or *Encoding* value in the control information in the message differs from that in the *MsgDesc* parameter, and there is data to be converted:

2. If the *Format* field in the control information in the message has the value MQFMT_NONE, the message is returned unconverted, with completion code MQCC_WARNING and reason code MQRC_FORMAT_ERROR.

   In all other cases conversion processing continues.

3. The message is removed from the queue and placed in a temporary buffer which is the same size as the *Buffer* parameter. For browse operations, the message is copied into the temporary buffer, instead of being removed from the queue.

4. If the message has to be truncated to fit in the buffer, the following is done:

   - If the MQGMO_ACCEPT_TRUNCATED_MSG option was *not* specified, the message is returned unconverted, with completion code MQCC_WARNING and reason code MQRC_TRUNCATED_MSG_FAILED.

- If the MQGMO_ACCEPT_TRUNCATED_MSG option *was* specified, the completion code is set to MQCC_WARNING, the reason code is set to MQRC_TRUNCATED_MSG_ACCEPTED, and conversion processing continues.

5. If the message can be accommodated in the buffer without truncation, or the MQGMO_ACCEPT_TRUNCATED_MSG option was specified, the following is done:

- If the format is a built-in format, the buffer is passed to the queue-manager's data-conversion service.

- If the format is not a built-in format, the buffer is passed to a user-written exit which has the same name as the format. If the exit cannot be found, the message is returned unconverted, with completion code MQCC_WARNING and reason code MQRC_FORMAT_ERROR.

If no error occurs, the output from the data-conversion service or from the exit is the converted message, plus the completion code and reason code to be returned to the application issuing the MQGET call.

6. If the conversion is successful, the queue manager returns the converted message to the application. In this case, the completion code and reason code returned will usually be one of the following combinations:

Completion code	Reason code
MQCC_OK	MQRC_NONE
MQCC_WARNING	MQRC_TRUNCATED_MSG_ACCEPTED

If the conversion fails (for whatever reason), the queue manager returns the unconverted message to the application, with the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter set to the values in the control information in the message, and with completion code MQCC_WARNING. See "Processing conventions" for possible reason codes.

If the local MCA is converting the message and the conversion fails, the message is sent to the dead-letter queue at the sending queue manager. If the message cannot be sent to the DLQ, the channel is closed and the unconverted message remains in the transmission queue.

Processing conventions

When you convert a built-in format, the queue manager follows the processing conventions described below. It is recommended that user-written exits should also follow these conventions, although this is not enforced by the queue manager. The built-in formats converted by the queue manager are:

```
MQFMT_ADMIN
MQFMT_COMMAND_1
MQFMT_COMMAND_2
MQFMT_DEAD_LETTER_HEADER
MQFMT_DIST_HEADER (not supported on MQSeries for MVS/ESA)
MQFMT_EVENT
MQFMT_IMS
MQFMT_IMS_VAR_STRING
MQFMT_MD_EXTENSION (not supported on MQSeries for MVS/ESA)
MQFMT_PCF
MQFMT_REF_MSG_HEADER
MQFMT_STRING
```

MQFMT_TRIGGER
MQFMT_XMIT_Q_HEADER

1. If the message expands during conversion, and exceeds the size of the *Buffer* parameter, the following is done:

   - If the MQGMO_ACCEPT_TRUNCATED_MSG option was *not* specified, the message is returned unconverted, with completion code MQCC_WARNING and reason code MQRC_CONVERTED_MSG_TOO_BIG.

   - If the MQGMO_ACCEPT_TRUNCATED_MSG option *was* specified, the message is truncated, the completion code is set to MQCC_WARNING, the reason code is set to MQRC_TRUNCATED_MSG_ACCEPTED, and conversion processing continues.

2. If truncation occurs (either before or during conversion), it is possible for the number of valid bytes returned in the *Buffer* parameter to be **less than** the length of the buffer.

   This can occur, for example, if a 4-byte integer or a DBCS character straddles the end of the buffer. The incomplete element of information is not converted, and so those bytes in the returned message do not contain valid information. This can also occur if a message that was truncated before conversion, shrinks during conversion.

   If the number of valid bytes returned is less than the length of the buffer, the unused bytes at the end of the buffer are set to nulls.

3. If an array or string straddles the end of the buffer, as much of the data as possible is converted; only the particular array element or DBCS character which is incomplete is not converted – preceding array elements or characters are converted.

4. If truncation occurs (either before or during conversion), the length returned for the *DataLength* parameter is the length of the unconverted message before truncation.

5. The data returned to the application is never partially converted; either all of the data returned is converted, or none of it is. For example, if the integers in the data can be converted, but the character strings cannot (because the character-set identifier is not recognized), none of the data is converted.

6. If the *CodedCharSetId* or *Encoding* fields in the control information of the message being retrieved, or in the *MsgDesc* parameter, specify values which are undefined or not supported, the queue manager may ignore the error if the undefined or unsupported value does not need to be used in converting the message.

   For example, if the *Encoding* field in the message specifies an unsupported float encoding, but the message contains only integer data, or contains floating-point data which does not require conversion (because the source and target float encodings are identical), the error may or may not be diagnosed.

   If the error is diagnosed, the message is returned unconverted, with completion code MQCC_WARNING and one of the MQRC_SOURCE_*_ERROR or MQRC_TARGET_*_ERROR reason codes (as appropriate); the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to the values in the control information in the message.

If the error is not diagnosed and the conversion completes successfully, the values returned in the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are those specified by the application issuing the MQGET call.

7. In all cases, if the message is returned to the application unconverted the completion code is set to MQCC_WARNING, and the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to the values appropriate to the unconverted data. This is done for MQFMT_NONE also.

   The *Reason* parameter is set to a code that indicates why the conversion could not be carried out, unless the message also had to be truncated; reason codes related to truncation take precedence over reason codes related to conversion. (To determine if a truncated message was converted, check the values returned in the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter.)

   When an error is diagnosed, either a specific reason code is returned, or the general reason code MQRC_NOT_CONVERTED. The reason code returned depends on the diagnostic capabilities of the underlying data-conversion service.

8. If completion code MQCC_WARNING is returned, and more than one reason code is relevant, the order of precedence is as follows:

   a. The following reason takes precedence over all others:

      MQRC_TRUNCATED_MSG_ACCEPTED

   b. Next in precedence is the following:

      MQRC_FORMAT_ERROR

   c. The order of precedence within this final group is not defined:

      MQRC_CONVERTED_MSG_TOO_BIG
      MQRC_NOT_CONVERTED
      MQRC_SOURCE_CCSID_ERROR
      MQRC_SOURCE_DECIMAL_ENC_ERROR
      MQRC_SOURCE_FLOAT_ENC_ERROR
      MQRC_SOURCE_INTEGER_ENC_ERROR
      MQRC_TARGET_CCSID_ERROR
      MQRC_TARGET_DECIMAL_ENC_ERROR
      MQRC_TARGET_FLOAT_ENC_ERROR
      MQRC_TARGET_INTEGER_ENC_ERROR

9. On completion of the MQGET call:

   • The following reason code indicates that the message was converted successfully:

      MQRC_NONE

   • The following reason code indicates that the message *may* have been converted successfully (check the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter to find out):

      MQRC_TRUNCATED_MSG_ACCEPTED

   • All other reason codes indicate that the message was not converted.

The following processing is specific to the built-in formats; it is not applicable to user-defined formats:

10. With the exception of the format MQFMT_STRING, none of the built-in formats can be converted from or to double-byte character sets (DBCS); only single-byte character sets (SBCS) can be used with these formats.

    If DBCS character sets are specified with these formats, the message is returned unconverted, with completion code MQCC_WARNING and reason code MQRC_SOURCE_CCSID_ERROR or MQRC_TARGET_CCSID_ERROR, as appropriate.

11. If the message data for a built-in format is truncated, fields within the message which contain lengths of strings, or counts of elements or structures, are **not** adjusted to reflect the length of the data actually returned to the application; the values returned for such fields within the message data are the values applicable to the message **prior to truncation**.

    When processing messages such as a truncated MQFMT_ADMIN message, care must be taken to ensure that the application does not attempt to access data beyond the end of the data returned.

12. If the format name is MQFMT_DEAD_LETTER_HEADER, the message data begins with an MQDLH structure, and this is followed by zero or more bytes of application message data. The format, character set, and encoding of the application message data are defined by the *Format*, *CodedCharSetId*, and *Encoding* fields in the MQDLH structure at the start of the message. Because the MQDLH structure and application message data can have different character sets and encodings, it is possible for one, other, or both of the MQDLH structure and application message data to require conversion.

    The queue manager converts the MQDLH structure first, as necessary. If conversion is successful, or the MQDLH structure does not require conversion, the queue manager checks the *CodedCharSetId* and *Encoding* fields in the MQDLH structure to see if conversion of the application message data is required. If conversion **is** required, the queue manager invokes the user-written exit with the name given by the *Format* field in the MQDLH structure, or performs the conversion itself (if *Format* is the name of a built-in format).

    If the MQGET call returns a completion code of MQCC_WARNING, and the reason code is one of those indicating that conversion was not successful, one of the following applies:

    - The MQDLH structure could not be converted. In this case the application message data will not have been converted either.

    - The MQDLH structure was converted, but the application message data was not.

    The application can examine the values returned in the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter, and those in the MQDLH structure, in order to determine which of the above applies.

13. If the format name is MQFMT_XMIT_Q_HEADER, the message data begins with an MQXQH structure, and this is followed by zero or more bytes of application message data. The format, character set, and encoding of the application message data are defined by the *Format*, *CodedCharSetId*, and *Encoding* fields in the MQMD structure contained *within* the MQXQH.

    If the MQGMO_CONVERT option is specified for an MQFMT_XMIT_Q_HEADER message, the application message data is converted, but the data in the MQXQH structure is not. Also, if the message

data begins with an MQXQH structure followed by MQDH, the distribution header (MQDH) and the record arrays (MQOR and MQPMR) are not converted. On return from the MQGET call:

- The values of the *Format*, *CodedCharSetId*, and *Encoding* fields in the *MsgDesc* parameter describe the data in the MQXQH structure, and **not** the application message data; the values will therefore be the same as those in the control information in the message, and **not** the values specified by the application that issued the MQGET call.

  The effect of this is that an application which repeatedly gets messages from a transmission queue with the MQGMO_CONVERT option specified must reset the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter to the values desired for the application message data, prior to each MQGET call.

- The values of the *Format*, *CodedCharSetId*, and *Encoding* fields in the MQMD structure within the MQXQH structure describe the application message data. If conversion is successful, the values will be the same as those specified in the *MsgDesc* parameter by the application that issued the MQGET call.

This facility is primarily intended for use by message channel agents when they get messages from transmission queues.

Conversion of report messages

A report message can contain varying amounts of application message data, according to the report options specified by the sender of the original message. In particular, a report message can contain either:

1. No application message data

2. Some of the application message data from the original message (this occurs when the sender of the original message specifies MQRO_★_WITH_DATA and the message is longer than 100 bytes)

3. All of the application message data from the original message (this occurs when the sender of the original message specifies MQRO_★_WITH_FULL_DATA, or specifies MQRO_★_WITH_DATA and the message is 100 bytes or shorter)

When the queue manager or message channel agent generates a report message, it copies the format name from the original message into the *Format* field in the control information in the report message. The format name in the report message may therefore imply a length of data which is different from the length actually present in the report message (cases 1 and 2 above).

If the MQGMO_CONVERT option is specified when the report message is retrieved:

- For case 1 above, the data-conversion exit will not be invoked (because the report message will have no data).

- For case 3 above, the format name correctly implies the length of the message data.

- But for case 2 above, the data-conversion exit will be invoked to convert a message which is **shorter** than the length implied by the format name.

In addition, the reason code passed to the exit will usually be MQRC_NONE (that is, the reason code will not indicate that the message has been truncated). This happens because the message data was truncated by the **sender** of the report message, and not by the receiver's queue manager in response to the MQGET call.

Because of these possibilities, the data-conversion exit should **not** use the format name to deduce the length of data passed to it; instead the exit should check the length of data provided, and be prepared to convert less data than the length implied by the format name. If the data can be converted successfully, completion code MQCC_OK and reason code MQRC_NONE should be returned by the exit. The length of the message data to be converted is passed to the exit as the *InBufferLength* parameter.

Writing a data-conversion exit program

For MVS/ESA, you must write data-conversion exits in assembler language. For other platforms, it is recommended that you use the C programming language.

To help you create a data-conversion exit program, the following are supplied:

- A skeleton source file

- A convert characters call

- A utility that creates a fragment of code that performs data conversion on data type structures

  This utility takes C input only. On MVS/ESA, it produces assembler code.

These are described below.

For the procedure for writing the programs see:

Skeleton source file

These can be used as your starting point when writing a data-conversion exit program. The files supplied are listed in Table 6 on page 166:

Table 6. Skeleton source files	
Platform	**File**
AS/400	QMQMSAMP/QCSRC(AMQSVFC4)
Digital OpenVMS	AMQSVFCX.C
OS/2	AMQSVFC2.C
MVS/ESA	CSQ4BAX8 (1) CSQ4BAX9 (2) CSQ4CAX9 (3)
Tandem NSK	amqsvfcn
UNIX systems	amqsvfcx.c
Windows NT	amqsvfcn.c
Notes:	
1. Illustrates the MQXCVNC call. 2. A wrapper for the code fragments generated by the utility for use in all environments except CICS. 3. A wrapper for the code fragments generated by the utility for use in the CICS environment.	

Convert characters call

The MQXCNVC (Convert characters) call may be used from within a data-conversion exit program to convert character message data from one character set to another. For certain multibyte character sets (for example, UCS2 character sets), the appropriate options must be used.

No other MQI calls can be made from within the exit; an attempt to make such a call fails with reason code MQRC_CALL_IN_PROGRESS.

See the *MQSeries Application Programming Reference* for further information on the MQXCNVC call and appropriate options.

Utility for creating conversion-exit code

The commands for creating conversion-exit code are:

AS/400
CVTMQMDTA (Convert MQSeries Data Type)

OS/2, Digital OpenVMS, Tandem NSK, and Windows NT, UNIX systems
crtmqcvx (Create MQSeries conversion-exit)

MVS/ESA
CSQUCVX

The command for your platform produces a fragment of code that performs data conversion on data type structures, for use in your data-conversion exit program. The command takes a file containing one or more C language structure definitions. On MVS/ESA, it then generates a data set containing assembler code fragments and conversion functions. On other platforms, it generates a file with a C function to convert each structure definition. The utility requires access to the LE/370 run-time libraries.

Invoking the CSQUCVX utility on MVS/ESA

11 shows an example of the JCL used to invoke the CSQUCVX utility.

```
//CVX     EXEC PGM=CSQUCVX
//STEPLIB  DD DISP=SHR,DSN=thlqual.SCSQANLE
//         DD DISP=SHR,DSN=thlqual.SCSQMVRy
//         DD DISP=SHR,DSN=thlqual.SCSQAUTH
//         DD DISP=SHR,DSN=le370.library
//SYSPRINT DD SYSOUT=*
//CSQUINP  DD DISP=SHR,DSN=MY.MQSERIES.FORMATS(MSG1)
//CSQUOUT  DD DISP=OLD,DSN=MY.MQSERIES.EXITS(MSG1)
```

Figure 11. Sample JCL used to invoke the CSQUCVX utility

Where y is 1 if you are using the IBM TCP/IP feature, and 2 if you are using the Interlink SNS/TCP access feature.

Data definition statements

The CSQUCVX utility requires DD statements with the following DDnames:

SYSPRINT

This specifies a data set or print spool class for reports and error messages.

CSQUINP

This specifies the sequential data set containing the definitions of the data structures to be converted.

CSQUOUT

This specifies the sequential data set where the conversion code fragments are to be written. The logical record length (LRECL) must be 80 and the record format (RECFM) must be FB.

Error messages in OS/2, Windows NT, and UNIX systems

The **crtmqcvx** command returns messages in the range AMQ7953 through AMQ7970. These are listed in the *MQSeries Messages* book for MQSeries Version 5 products. For other platforms, see the appropriate *System Management Guide* for your platform.

There are two main types of error:

- Major errors, such as syntax errors, when processing cannot continue.

  A message is displayed on the screen giving the line number of the error in the input file. The output file may have been partially created.

- Other errors when a message is displayed stating that a problem has been found but parsing of the structure can continue.

  The output file has been created and contains error information on the problems that have occurred. This error information is prefixed by #error so that the code produced will not be accepted by any compiler without intervention to rectify the problems.

Valid syntax

Your input file for the utility must conform to the C language syntax. If you are unfamiliar with C, refer to "Example of valid syntax for the input data set."

In addition, you must be aware of the following rules:

- typedef is recognized only before the struct keyword.
- A structure tag is required on your structure declarations.
- Empty square brackets [ ] may be used to denote a variable length array or string at the end of a message.
- Multi-dimensional arrays and arrays of strings are not supported.
- The following additional data types are recognized:

      MQBYTE
      MQCHAR
      MQSHORT
      MQLONG

  MQCHAR fields are code page converted, but MQBYTE is left untouched. If the encoding is different, MQSHORT and MQLONG are converted accordingly.

- The following should **not** be used:

      float
      double
      pointers
      bit-fields

  This is because the utility for creating conversion-exit code does not provide the facility to convert these data types. To overcome this, you can write your own routines and call them from the exit.

Other points to note:

- Do not use sequence numbers in the input data set.
- If there are fields for which you want to provide your own conversion routines, declare them as MQBYTE, and then replace the generated CMQXCFBA macros with your own conversion code.

Example of valid syntax for the input data set

```
struct TEST { MQLONG    SERIAL_NUMBER;
              MQCHAR    ID[5];
              MQSHORT   VERSION;
              MQBYTE    CODE[4];
              MQLONG    DIMENSIONS[3];
              MQCHAR    NAME[24];
            } ;
```

This corresponds to the following declarations in the other programming languages:

COBOL

```
              10 TEST.
                15 SERIAL-NUMBER  PIC S9(9) BINARY.
                15 ID            PIC X(5).
                15 VERSION       PIC S9(4) BINARY.
              * CODE IS NOT TO BE CONVERTED
                15 CODE          PIC X(4).
                15 DIMENSIONS    PIC S9(9) BINARY OCCURS 3 TIMES.
                15 NAME          PIC X(24).
```

System/390 assembler: *Supported on MVS/ESA only*

```
TEST          EQU *
SERIAL_NUMBER DS  F
ID            DS  CL5
VERSION       DS  H
CODE          DS  XL4
DIMENSIONS    DS  3F
NAME          DS  CL24
```

PL/I: *Supported on AIX, MVS/ESA, OS/2 Warp, and Windows NT only*

```
DCL 1 TEST,
      2 SERIAL_NUMBER  FIXED BIN(31),
      2 ID             CHAR(5),
      2 VERSION        FIXED BIN(15),
      2 CODE           CHAR(4),       /* not to be converted */
      2 DIMENSIONS(3)  FIXED BIN(31),
      2 NAME           CHAR(24);
```

Writing a data-conversion exit program for MQSeries for AS/400

Follow these steps:

1. Name your message. The name must fit in the *Format* field of the MQMD. The *Format* name should not have leading embedded blanks, and trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long. Remember to use this name each time you send a message (our example uses the name Format).

2. Create a structure to represent your message. See "Valid syntax" on page 168 for an example.

3. Run this structure through the CVTMQMDTA command to create a code fragment for your data-conversion exit.

   The functions generated by the CVTMQMDTA command use macros that are shipped in the file QMQM/H(AMQSVMHA). These macros are written assuming that all structures are packed; they should be amended if this is not the case.

4. Take a copy of the supplied skeleton source file, QMQMSAMP/QCSRC(AMQSVFC4) and rename it. (Our example uses the name EXIT_MOD.)

5. Find the following comment boxes in the source file and insert code as described:

   a. Towards the bottom of the source file, a comment box starts with:

      ```
      /* Insert the functions produced by the data-conversion exit */
      ```

Here, insert the code fragment generated in step 3.

b. Near the middle of the source file, a comment box starts with:

```
/* Insert calls to the code fragments to convert the format's */
```

This is followed by a commented-out call to the function `ConverttagSTRUCT`.

Change the name of the function to the name of the function you added in step 5a above. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

c. Near the top of the source file, a comment box starts with:

```
/* Insert the function prototypes for the functions produced by */
```

Here, insert the function prototype statements for the functions added in step 5a above.

If the message contains character data, the generated code will call MQXCNVC; this can be resolved by binding the service programs AMQVSTUB and AMQZSTUB.

6. Compile the source module, EXIT_MOD, using the CRTCMOD command.

7. Create/link the program:

```
CRTPGM PGM(library/Format)  +
       MODULE(library/EXIT_MOD)  +
       BNDSRVPGM(QMQM/AMQVSTUB)  +
       BNDSRVPGM(QMQM/AMQZSTUB)  +
       ACTGRP(QMQM)  +
       USRPRF(*USER)
```

8. Place the output in the library list for the MQSeries job. It is recommended that, for production, data-conversion exit programs be stored in QSYS.

Notes:

1. If CVTMQMDTA uses packed structures, all MQSeries applications must use the _Packed qualifier.

2. Data-conversion exit programs must be re-entrant.

3. MQXCNVC is the **only** MQI call that may be issued from a data-conversion exit.

4. The exit program should be compiled with the user profile compiler option set to *USER, so that the exit runs with the authority of the user.

Writing a data-conversion exit for MQSeries for OS/2 Warp

Follow these steps:

1. Name your message. The name must fit in the *Format* field of the MQMD. The *Format* name should not have leading embedded blanks, and trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long.

A .DEF file called AMQSVFC2.DEF is also supplied in the samples directory, <drive:\directory>\MQM\TOOLS\C\SAMPLES. Take a copy of this file and rename it, for example, to MYFORMAT.DEF. Make sure that the name of the DLL being created and the name specified in MYFORMAT.DEF are the same. Overwrite the name `FORMAT1` in MYFORMAT.DEF with the new format name.

Remember to use this name each time you send a message.

2. Create a structure to represent your message. See "Valid syntax" on page 168 for an example.

3. Run this structure through the CRTMQCVX command to create a code fragment for your data-conversion exit.

   The functions generated by the CRTMQCVX command use macros which are written assuming that all structures are packed; they should be amended if this is not the case.

4. Take a copy of the supplied skeleton source file, AMQSVFC2.C, renaming it to the name of your message that you decided on in step 1 (that is, MYFORMAT.C in this example).

   AMQSVFC2.C is in <drive:\directory>\MQM\TOOLS\C\SAMPLES (where <drive:\directory> was specified at installation).

   The skeleton includes a sample header file AMQSVMHA.H in the same directory. Make sure that your include path points to this directory to pick up this file.

   The AMQSVMHA.H file contains macros that are used by the code generated by the CRTMQCVX command. If the structure to be converted contains character data, then these macros call MQXCNVC.

5. Find the following comment boxes in the source file and insert code as described:

   a. Towards the bottom of the source file, a comment box starts with:

      ```
      /* Insert the functions produced by the data-conversion exit */
      ```

      Here, insert the code fragment generated in step 3.

   b. Near the middle of the source file, a comment box starts with:

      ```
      /* Insert calls to the code fragments to convert the format's */
      ```

      This is followed by a commented-out call to the function `ConverttagSTRUCT`.

      Change the name of the function to the name of the function you added in step 5a above. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

   c. Near the top of the source file, a comment box starts with:

      ```
      /* Insert the function prototypes for the functions produced by */
      ```

      Here, insert the function prototype statements for the functions added in step 5a above.

6. Resolve this call by linking the routine with the library MQMVX.LIB, in the directory <drive:\directory>\MQM\TOOLS\LIB.

7. Create the following command file:

   ```
   icc /Ge-   \
   /I<drive:\directory>\mqm\tools\c\include MYFORMAT.C   \
   <drive:\directory>\mqm\tools\lib\mqm.lib MYFORMAT.DEF  \
   <drive:\directory>\mqm\tools\lib\mqmvx.lib
   ```

   where `<drive:\directory>` is specified at installation.

   Issue the command file to compile your exit as a DLL file.

8. Place the output in the \mqm\exits subdirectory. The path used to look for the data-conversion exits is given in the qm.ini file as DefaultExitPath. This path is set for each queue manager and the exit will only be looked for in that path or paths.

Notes:

1. If CVTMQCVX uses packed structures, all MQSeries applications must be compiled in this way.

2. Data-conversion exit programs must be re-entrant.

3. MQXCNVC is the *only* MQI call that may be issued from a data-conversion exit.

Writing a data-conversion exit program for MQSeries for MVS/ESA

Follow these steps:

1. Take the supplied source skeleton CSQ4BAX9 (for non-CICS environments) or CSQ4CAX9 (for CICS) as your starting point.

2. Run the CSQUCVX utility.

3. Follow the instructions in the prolog of CSQ4BAX9 or CSQ4CAX9 to incorporate the routines generated by the CSQUCVX utility, in the order that the structures occur in the message you want to convert.

4. The utility assumes that the data structures are not packed, that the implied alignment of the data is honored, and that the structures start on a full-word boundary, with bytes being skipped as required (as between ID and VERSION in the "Example of valid syntax for the input data set" on page 168). If the structures are packed, you will need to omit the CMQXCALA macros that are generated. You are therefore strongly recommended to declare your structures in such a way that all fields are named and no bytes are skipped; in the "Example of valid syntax for the input data set" on page 168, you would add a field "MQBYTE DUMMY;" between ID and VERSION.

5. The supplied exit returns an error if the input buffer is shorter than the message format to be converted. Although the exit converts as many complete fields as possible, the error causes an unconverted message to be returned to the application. If you want to allow short input buffers to be converted as far as possible, including partial fields, change the TRUNC= value on the CSQXCDFA macro to YES: no error is returned, so the application receives a converted message. The application is responsible for handling the truncation.

6. Add any other special processing code that you need.

7. Rename the program to your data format name.

8. Compile and link-edit your program like a batch application program (unless it is for use with CICS applications). The macros in the code generated by the utility are in the library, **thlqual**.SCSQMACS.

   If the message contains character data, the generated code will call MQXCNVC. If your exit uses this call, link-edit it with the exit stub program CSQASTUB. The stub is language and environment independent. Alternatively, you can load the stub dynamically using the dynamic call name CSQXCNVC. See "Dynamically calling the MQSeries stub" on page 270 for more information.

Place the link-edited module in your application load library, and in a data set that is referenced by the CSQXLIB DD statement of your task procedure started by your channel initiator.

9. If the exit is for use by CICS applications, compile and link-edit it like a CICS application program, including CSQASTUB if required. Place it in your CICS application program library. Define the program to CICS in the usual way, specifying EXECKEY(CICS) in the definition.

Note: Although the LE/370 runtime libraries are needed for running the CSQUCVX utility (see step 2 on page 172), they are not needed for link-editing or running the data-conversion exit itself (see steps 8 on page 172 and 9).

See "Writing MQSeries-IMS bridge applications" on page 228 for information about data conversion within the MQSeries-IMS bridge.

Writing a data-conversion exit for MQSeries for Tandem NSK

Dynamically bound libraries are not supported by MQSeries for Tandem NSK. Data conversion exits (and channel exits) are implemented by including statically bound stub functions in the MQSeries libraries and executables that can be replaced using the REPLACE bind option.

A data conversion exit **must** be called DATACONVEXIT (see sample AMQSVFCN), and can be bound into the chosen executable (or library) using the TACL macro BEXITE.

Note: This procedure modifies the target executable; you are recommended to make a back-up copy of the target executable or library before using the macro.

Exit functions, once compiled, must be bound directly into the target executable or library to be accessible by MQSeries. The following TACL macro is used for this purpose:

BEXITE   Usage: `BEXITE target-executable-or-library`
            `source-exit-file-or-library`

For example, to bind the sample data conversion exit into the sample MQSGETA, follow these steps:

1. Compile the exit function DATACONVEXIT (CSAMP AMQSVFCN).

2. Compile the get application (CSAMP AMQSGET0).

3. Bind the get application (BSAMP AMQSGET).

4. Bind the exit function into the get application (BEXITE AMQSGET AMQSVFCO).

Alternatively, if all applications are to have this data conversion exit, the following steps would create both a user library and an application with the exit bound in:

1. Compile the exit function DATACONVEXIT (CSAMP AMQSVFCN).

2. Compile the get application (CSAMP AMQSGET0).

3. Bind the exit function into the user library (BEXITE ZMQSLIB.MQMLIBC AMQSVFCO).

4. Bind the get application with the modified library (BSAMP AMQSGET).

| If the data conversion exit is to be used by channels processing within MQSeries, it
| must also be bound into the caller executable by the system administrator. For
| example:

| `BEXITE ZMQSEXE.MQMCACAL AMQSVFCO`

| Use the TACL macro BDCXALL to bind the data conversion exit into all required
| MQSeries processes. For example:

| `BDCXALL source-exit-file-or-library`

| # Reusing data-conversion exit programs

| In other MQSeries Version 2 products, a data-conversion exit is required for each
| application-defined format to be supported. The data-conversion exit programs are
| named according to the *Format* value (from MQMD) of the message to be
| converted. The format for which conversion is being requested can be determined
| from the *Format* field of the *MsgDesc* parameter. The appropriate data-conversion
| exit program can therefore be invoked from MQDATACONVEXIT(). The
| parameters supplied to MQDATACONVEXIT() can be supplied to the invoked
| data-conversion function.

Writing a data-conversion exit for MQSeries on UNIX systems and Digital OpenVMS

For SINIX and DC/OSx, data-conversion exits must not use DCE.

> **Note**
>
> The name **mqmtop** is used in this section to represent the name of the
> installation directory.

Follow these steps:

1. Name your message. The name must fit in the *Format* field of the MQMD, and
   be in uppercase, for example, MYFORMAT. The *Format* name should not
   have leading embedded blanks, and trailing blanks are ignored. The object's
   name must have no more than eight non-blank characters because the *Format*
   is only eight characters long. Remember to use this name each time you send
   a message.

2. Create a structure to represent your message. See "Valid syntax" on page 168
   for an example.

3. Run this structure through the **crtmqcvx** command to create a code fragment
   for your data-conversion exit.

   The functions generated by the **crtmqcvx** command use macros which are
   written assuming that all structures are packed; they should be amended if this
   is not the case.

4. Take a copy of the supplied skeleton source file, amqsvfcx.c, renaming it to the
   name of your message that you decided on in step 1 (that is, MYFORMAT.C).

   The skeleton includes a sample header file amqsvmha.h in the directory
   /**mqmtop**/mqm/inc. Make sure that your include path points to this directory to
   pick up this file.

The amqsvmha.h file contains macros that are used by the code generated by the **crtmqcvx** command. If the structure to be converted contains character data, then these macros call MQXCNVC.

5. Find the following comment boxes in the source file and insert code as described:

   a. Towards the bottom of the source file, a comment box starts with:

      ```
      /* Insert the functions produced by the data-conversion exit */
      ```

      Here, insert the code fragment generated in step 3.

   b. Near the middle of the source file, a comment box starts with:

      ```
      /* Insert calls to the code fragments to convert the format's */
      ```

      This is followed by a commented-out call to the function `ConverttagSTRUCT`.

      Change the name of the function to the name of the function you added in step 5a above. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

   c. Near the top of the source file, a comment box starts with:

      ```
      /* Insert the function prototypes for the functions produced by */
      ```

      Here, insert the function prototype statements for the functions added in step 5a above.

6. Resolve this call by linking the routine with the library libmqm. For DCE-threaded programs the routine must be linked with the library libmqm_r.

7. Compile your exit as a shared library, using MQStart as the entry point. To do this, see "Compiling data-conversion exits on UNIX" on page 177, or "Compiling data-conversion exits on Digital OpenVMS" on page 176.

8. Place the output in the default system directory, /var/mqm/exits, to ensure that it can be loaded when required. The path used to look for the data-conversion exits is given in the qm.ini file. This path can be set for each queue manager and the exit is only looked for in that path or paths.

Notes:

1. If **crtmqcvx** uses packed structures, all MQSeries applications must be compiled in this way.

2. Data-conversion exit programs must be re-entrant.

3. MQXCNVC is the *only* MQI call that may be issued from a data-conversion exit.

UNIX environment

There are two environments to consider: non threaded and threaded.

Non-threaded environment

The loadable object must have its name in upper case, for example MYFORMAT. The libmqm library should be used to resolve the calls to MQXCNVC.

Threaded environment

In addition to creating the data-conversion exit for the basic environment, another is required in the threaded environment. This loadable object must be followed by _r (on AIX and HP-UX) and _d (on Sun Solaris) to indicate that it is a DCE-threaded version. The libmqm_r library (on AIX and HP-UX) and the lmqmcs_d library (on Sun Solaris) should be used to resolve the calls to MQXCNVC. Note that both loadable objects (non-threaded and threaded) are required for a threading environment.

Note: If the data-conversion exits are in a mixed non-threaded and threaded environment, the calling environment is detected and the appropriate object loaded. The shared object should be placed in /var/mqm/exits to ensure it can be loaded when required.

Compiling data-conversion exits on Digital OpenVMS

The names of the routines which are called by the data-conversion exit must be made universal.

```
$ CC /INCLUDE_DIRECTORY=MQS_INCLUDE AMQSVFCX.C
$ LINK /SYS$SHARE:[SYSLIB]MYFORMAT AMQSVFCX.OBJ,MYFORMAT/OPTIONS
```

The contents of MYFORMAT.OPT vary depending on which platform you are working on:

On Alpha:

```
SYS$SHARE:MQM/SHAREABLE
SYS$SHARE:MQMCS/SHAREABLE
SYMBOL_VECTOR=(MQSTART=PROCEDURE)
```

On VAX:

```
SYS$SHARE:MQM/SHAREABLE
SYS$SHARE:MQMCS/SHAREABLE
UNIVERSAL=MQSTART
```

If you are using threaded applications linked with the pthread library, you must also build a second copy of the data-conversion exit with the thread options and libraries:

```
$ CC /INCLUDE_DIRECTORY=MQS_INCLUDE AMQSVFCX.C
$ LINK /SYS$SHARE:[SYSLIB]MYFORMAT AMQSVFCX.OBJ,MYFORMAT/OPTIONS
```

Again, the contents of MYFORMAT.OPT vary depending on which platform you are working on:

On Alpha:

```
SYS$SHARE:MQM_R/SHAREABLE
SYS$SHARE:MQMCS_R/SHAREABLE
SYS$SHARE:CMA$OPEN_RTL.EXE/SHAREABLE
SYMBOL_VECTOR-(MQSTART=PROCEDURE)
```

On VAX:

```
SYS$SHARE:MQM_R/SHAREABLE
SYS$SHARE:MQMCS_R/SHAREABLE
SYS$SHARE:CMA$OPEN_RTL.EXE/SHAREABLE
UNIVERSAL=MQSTART
```

Compiling data-conversion exits on UNIX

The following sections give examples of how to compile a data conversion exit on the UNIX platforms.

On all platforms, the entry point to the module is MQStart.

On AIX

```
$ cc -c -I/usr/lpp/mqm/inc MYFORMAT.C
$ ld MYFORMAT.o -e MQStart -o MYFORMAT -bM:SRE -H512 -T512 -lmqm -lc
$ cp MYFORMAT /var/mqm/exits
```

If you are using threaded application linked with the pthreads library, you must build a second copy of the conversion exit with the thread options and libraries.

```
$ cc_r -c -I/usr/lpp/mqm/inc MYFORMAT.C
$ ld MYFORMAT.o -e MQStart -o MYFORMAT_r -bM:SRE -H512      \
-T512 -lpthreads -lc_r -lmqm_r
$ cp MYFORMAT_r /var/mqm/exits
```

On AT&T GIS UNIX

```
$ cc -c -K PIC -I/mqmtop/inc MYFORMAT.C
$ ld -G MYFORMAT.O -o MYFORMAT
$ cp MYFORMAT /mqmtop/lib
```

On HP-UX

```
$ CC -c -Aa +z -I/mqmtop/inc MYFORMAT.C
$ ld -b MYFORMAT.o -o MYFORMAT -L /mqmtop/lib -lmqm +IMQStart
$ cp MYFORMAT /var/mqm/exits
```

If you are using threaded applications linked with the pthreads library, you must build a second copy of the conversion exit with the thread options and libraries.

```
$ CC -c -Aa +z -I/mqmtop/inc -D_REENTRANT MYFORMAT.C
$ ld -b MYFORMAT.o -o MYFORMAT_r -L /mqmtop/lib      \
-lmqm_r -lcma -lc_r +IMQStart
$ cp MYFORMAT_r /var/mqm/exits
```

On SINIX

```
$ cc -c -K PIC -I/mqmtop/inc -lmproc -lext MYFORMAT.C
$ ld -G MYFORMAT.O  -o MYFORMAT
$ cp MYFORMAT /mqmtop/lib
```

On DC/OSx

```
$ cc -c -K PIC -I/mqmtop/inc -liconv -lmproc -lext MYFORMAT.C
$ ld -G MYFORMAT.O  -o MYFORMAT
$ cp MYFORMAT /mqmtop/lib
```

---

4 This platform has become NCR UNIX SVR4 MP-RAS, R3.0.

On SunOS

```
$ acc -c -sys5 -PIC -I/usr/mqm/inc MYFORMAT.C
$ /bin/ld MYFORMAT.O -o MYFORMAT -Bdynamic -L /usr/mqm/lib -lmqm
$ cp MYFORMAT /usr/mqm/lib
```

On Sun Solaris

If your application uses no threading calls or Posix V10 threading calls:

```
cc -c -KPIC -I/mqmtop/inc MYFORMAT.C

ld -G /opt/SUNWspro/SC4.0/lib/crt1.o
/opt/SUNWspro/SC4.0/lib/crti.o
/opt/SUNWspro/SC4.0/lib/crtn.o
/opt/SUNWspro/SC4.0/lib/values-xt.o
MYFORMAT.o -o MYFORMAT -lmqm -lthread -lsocket -lc -lnsl -ldl

cp MYFORMAT /var/mqm/exits
```

If your application requires DCE threading (for example, if it is a CICS application):

```
cc -c -KPIC -I/mqmtop/inc MYFORMAT.C

ld -G /opt/SUNWspro/SC4.0/lib/crt1.o
/opt/SUNWspro/SC4.0/lib/crti.o
/opt/SUNWspro/SC4.0/lib/crtn.o
/opt/SUNWspro/SC4.0/lib/values-xt.o
MYFORMAT.o -o MYFORMAT_d -ldce -lnsl -lthread -lm -lsocket
-lmqmcs_d -lmqm -lc -ldl

cp MYFORMAT /var/mqm/exits
```

Note: The SC4.0 directory name varies depending on the release of compiler.

If you want to run applications using both the Posix V10-threaded and the DCE-threaded variants on a single queue manager:

1. Build a Posix V10 type of data-conversion exit. Name it MYFORMAT and place it in the appropriate exit directory.

2. Build a DCE-threaded type of data-conversion exit. Name it MYFORMAT_d and place it in the appropriate exit directory.

Two object files are generated; one of which loads the MYFORMAT data-conversion exit, and the other of which loads the MYFORMAT_d data-conversion exit.

Writing a data-conversion exit for MQSeries for Windows NT

Follow these steps:

1. Name your message. The name must fit in the *Format* field of the MQMD. The *Format* name should not have leading embedded blanks, and trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long.

   A .DEF file called amqsvfcn.def is also supplied in the samples directory, <drive:\directory>\MQM\TOOLS\C\SAMPLES. Take a copy of this file and rename it, for example, to MYFORMAT.DEF. Make sure that the name of the

DLL being created and the name specified in MYFORMAT.DEF are the same. Overwrite the name FORMAT1 in MYFORMAT.DEF with the new format name.

Remember to use this name each time you send a message.

2. Create a structure to represent your message. See "Valid syntax" on page 168 for an example.

3. Run this structure through the CRTMQCVX command to create a code fragment for your data-conversion exit.

The functions generated by the CVTMQCVX command use macros which are written assuming that all structures are packed; they should be amended if this is not the case.

4. Take a copy of the supplied skeleton source file, amqsvfcn.c. Give the copy the name assigned to the message in step 1 (that is, MYFORMAT).

amqsvfcn.c is in <drive:\directory>\MQM\TOOLS\C\SAMPLES (where <drive:\directory> was specified at installation).

The skeleton includes a sample header file amqsvmha.h in the same directory. Make sure that your include path points to this directory to pick up this file.

The amqsvmha.h file contains macros that are used by the code generated by the CRTMQCVX command. If the structure to be converted contains character data, then these macros call MQXCNVC.

5. Find the following comment boxes in the source file and insert code as described:

   a. Towards the bottom of the source file, a comment box starts with:

      ```
      /* Insert the functions produced by the data-conversion exit */
      ```

      Here, insert the code fragment generated in step 3.

   b. Near the middle of the source file, a comment box starts with:

      ```
      /* Insert calls to the code fragments to convert the format's */
      ```

      This is followed by a commented-out call to the function `ConverttagSTRUCT`.

      Change the name of the function to the name of the function you added in step 5a above. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

   c. Near the top of the source file, a comment box starts with:

      ```
      /* Insert the function prototypes for the functions produced by */
      ```

      Here, insert the function prototype statements for the functions added in step 5a above.

6. Resolve this call by linking the routine with the library MQMVX.LIB, in the directory <drive:\directory>\MQM\TOOLS\LIB.

7. Create the following command file:

   ```
   cl -I <drive:\directory>\mqm\tools\c\include -Tp   \
   MYFORMAT.C -LD -DEFAULTLIB \
   <drive:\directory>\mqm\tools\lib\mqm.lib  \
   <drive:\directory>\mqm\tools\lib\mqmvx.lib   \
   MYFORMAT.DEF
   ```

   where `<drive:\directory>` is specified at installation,

   Issue the command file to compile your exit as a DLL file.

8. Place the output in the mqm\exits subdirectory. The path used to look for the data-conversion exits is given in the qm.ini file. This path can be set for each queue manager and the exit is only looked for in that path or paths.

Notes:

1. If CVTMQCVX uses packed structures, all MQSeries applications must be compiled in this way.

2. Data-conversion exit programs must be re-entrant.

3. MQXCNVC is the *only* MQI call that may be issued from a data-conversion exit.

Chapter 12. Inquiring about and setting object attributes

Attributes are the properties that define the characteristics of an MQSeries object. They affect the way that an object is processed by a queue manager. The attributes of each type of MQSeries object are described in detail in the *MQSeries Application Programming Reference* manual.

Some attributes are set when the object is defined, and can be changed only by using the MQSeries commands; an example of such an attribute is the default priority for messages put on a queue. Other attributes are affected by the operation of the queue manager and may change over time; an example is the current depth of a queue.

You can inquire about the current values of all these attributes using the MQINQ call. The MQI also provides an MQSET call with which you can change some queue attributes. You cannot use the MQI calls to change the attributes of any other type of object; instead you must use:

For MQSeries for MVS/ESA
> The ALTER operator commands (or the DEFINE commands with the REPLACE option), which are described in the *MQSeries Command Reference*.

For MQSeries for AS/400
> The CHGMQMx CL commands, which are described in the *MQSeries for AS/400 Administration Guide*, or you can use the MQSC facility.

For MQSeries for Tandem NSK
> in MQSeries for Tandem NSK you can use the MQM screen-based interface. which is described in the *MQSeries for Tandem NonStop Kernel System Management Guide*, or you can use the MQSC facility.

For MQSeries for all other platforms
> The MQSC facility, described in the *MQSeries Command Reference*.

Note: The names of the attributes of objects are shown in this book in the form that you use them with the MQINQ and MQSET calls. When you use MQSeries commands to define, alter, or display the attributes, you must identify the attributes using the keywords shown in the descriptions of the commands in the above books.

Both the MQINQ and the MQSET calls use arrays of selectors to identify those attributes you want to inquire about or set. There is a selector for each attribute you can work with. The selector name has a prefix, determined by the nature of the attribute:

MQCA_   These selectors refer to attributes that contain character data (for example, the name of a queue).

MQIA_   These selectors refer to attributes that contain either numeric values (such as *CurrentQueueDepth*, the number of messages on a queue) or a constant value (such as *SyncPoint*, whether or not the queue manager supports syncpoints).

Before you use the MQINQ or MQSET calls your application must be connected to the queue manager, and you must use the MQOPEN call to open the object for setting or inquiring about attributes. These operations are described in Chapter 7, "Connecting and disconnecting a queue manager" on page 87 and Chapter 8, "Opening and closing objects" on page 93.

Inquiring about the attributes of an object

Use the MQINQ call to inquire about the attributes of any type of MQSeries object.

As input to this call, you must supply:

- A connection handle.

- An object handle.

- The number of selectors.

- An array of attribute selectors, each selector having the form MQCA_⋆ or MQIA_⋆. Each selector represents an attribute whose value you want to inquire about, and each selector must be valid for the type of object that the object handle represents. You can specify selectors in any order.

- The number of integer attributes that you are inquiring about. Specify zero if you are not inquiring about integer attributes.

- The length of the character attributes buffer in *CharAttrLength*. This must be at least the sum of the lengths required to hold each character attribute string. Specify zero if you are not inquiring about character attributes.

The output from MQINQ is:

- A set of integer attribute values copied into the array. The number of values is determined by *IntAttrCount*. If either *IntAttrCount* or *SelectorCount* is zero, this parameter is not used.

- The buffer in which character attributes are returned. The length of the buffer is given by the *CharAttrLength* parameter. If either *CharAttrLength* or *SelectorCount* is zero, this parameter is not used.

- A completion code. If the completion code gives a warning, this means that the call completed only partially. In this case, you should examine the reason code.

- A reason code. There are three partial-completion situations:

  - The selector does not apply to the queue type
  - There is not enough space allowed for integer attributes
  - There is not enough space allowed for character attributes

  If more than one of these situations arise, the first one that applies is returned.

Note: The values returned by the call are a snapshot of the selected attributes. The attributes can change before your program acts on the returned values.

There is a description of the MQINQ call in the *MQSeries Application Programming Reference* manual.

Some cases where the MQINQ call fails

If you open an alias to inquire about its attributes, you are returned the attributes of the alias queue (the MQSeries object used to access another queue) only, not those of the base queue. However, the definition of the base queue to which the alias resolves is also opened by the queue manager, and if another program changes the usage of the base queue in the interval between your MQOPEN and MQINQ calls, your MQINQ call fails and returns the MQRC_OBJECT_CHANGED reason code. The call also fails if the attributes of the alias queue object are changed.

Similarly, when you open a remote queue to inquire about its attributes, you are returned the attributes of the local definition of the remote queue only.

If you specify one or more selectors that are not valid for the type of queue about whose attributes you are inquiring, the MQINQ call completes with a warning and sets the output as follows:

- For integer attributes, the corresponding elements of *IntAttrs* are set to MQIAV_NOT_APPLICABLE.

- For character attributes, the corresponding portions of the *CharAttrs* string are set to asterisks.

If you specify one or more selectors that are not valid for the type of object about whose attributes you are inquiring, the MQINQ call fails and returns the MQRC_SELECTOR_ERROR reason code.

It is not possible to call MQINQ to look at a model queue. You will have to use either the MQSC facility or use the commands available on your platform.

Setting queue attributes

You can set only the following queue attributes using the MQSET call:

- *InhibitGet* (but not for remote queues)
- *DistList*
- *InhibitPut*
- *TriggerControl*
- *TriggerType*
- *TriggerDepth*
- *TriggerMsgPriority*
- *TriggerData*

The MQSET call has the same parameters as the MQINQ call. However for MQSET, all parameters except the completion code and reason code are input parameters. There are no partial-completion situations.

Note: You cannot use the MQI to set the attributes of MQSeries objects other than queues.

There is a description of the MQSET call in the *MQSeries Application Programming Reference* manual.

Using MQSET

Chapter 13. Committing and backing out units of work

This chapter describes how to commit, and back out, recoverable get and put operations that have occurred in a unit of work. But before a discussion on how to do this, let us define the following terms:

- Commit
- Back out
- Syncpoint coordination
- Syncpoint
- Unit of work
- Single-phase commit
- Two-phase commit

If you are familiar with these transaction processing terms, you can skip to "Syncpoint considerations in MQSeries applications" on page 186.

When a program puts a message on a queue within a unit of work, that message is made visible to other programs only when the program *commits* the unit of work. To commit a unit of work, **all** updates must be successful to preserve data integrity. If the program detects an error and decides that the put operation should not be made permanent, it can *back out* the unit of work. When a program performs a back out, MQSeries restores the message on the queue. How the program performs the commit and back out operations depends on the environment in which the program is running.

Similarly, when a program gets a message from a queue within a unit of work, that message remains on the queue until the program commits the unit of work, but the message is not available to be retrieved by other programs. The message is permanently deleted from the queue when the program commits the unit of work. If the program backs out the unit of work, MQSeries restores the queue to the state it was in before the program performed the get operation.

Changes to queue attributes (either by the MQSET call or by commands) are not affected by the committing or backing out of units of work.

Syncpoint coordination is the process by which units of work are either committed or backed out with data integrity.

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

A *single-phase commit* process is one in which a program can commit updates to a queue without coordinating its changes with other resource managers.

A *two-phase commit* process is one in which updates that a program has made to MQSeries queues can be coordinated with updates to other resources (for example, databases under the control of DB2). Under such a process, updates to **all** resources are committed or backed out together.

185

To help handle units of work, MQSeries provides the *BackoutCount* attribute. This is incremented each time a message, within a unit of work, is backed out. If the message repeatedly causes the unit of work to abend, the value of the *BackoutCount* finally exceeds that of the *BackoutThreshold*. This value is set when the queue is defined. In this situation, the application can choose to remove the message from the unit of work and put it onto another queue, as defined in *BackoutRequeueQName*. When the message is moved, the unit of work can commit.

Syncpoint considerations in MQSeries applications

Commit Protocol	Transaction Manager
Two-phase	MQSeries for AIX
	MQSeries for OS/2 Warp
	MQSeries for HP-UX
	MQSeries for Sun Solaris
	MQSeries for Tandem NSK
	MQSeries for Windows NT
	CICS for MVS/ESA
	CICS on Open Systems
	Transaction Server for Windows NT
	IMS/ESA
	Other external coordinators using the X/Open XA interface
Single-phase	MQSeries for AS/400
	MQSeries for OS/2 Warp
	MQSeries for Tandem NSK
	CICS for Windows NT V2.0
	MQSeries for Windows
	MQSeries for Windows NT
	MQSeries on UNIX systems
	MQSeries for Digital OpenVMS
	CICS for OS/2
	MVS batch

Note: For further details on external interfaces see "Interfaces to external syncpoint managers" on page 192, and the XA documentation *X/Open CAE Specifications Distributed Transaction Processing: The XA Specification*.

Transaction managers (such as CICS, IMS, Encina, Tuxedo, and Top End) can participate in two-phase commit, coordinated with other recoverable resources. This means that the queuing functions provided by MQSeries can be brought within the scope of a unit of work, managed by the transaction manager.

Samples shipped with MQSeries show the use of MQSeries coordinating XA-compliant databases. For further information about these samples, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

In your MQSeries application, you can specify on every put and get call whether or not you want the call to be under syncpoint control. To make a put operation operate under syncpoint control, use the MQPMO_SYNCPOINT value in the *Options* field of the MQPMO structure when you call MQPUT. For a get operation, use the MQGMO_SYNCPOINT value in the *Options* field of the MQGMO structure. If you do not explicitly choose an option, the default action depends on the platform. The syncpoint control default on MVS/ESA and Tandem NSK is 'yes'; for all other platforms, it is 'no'.

If a program issues the MQDISC call while there are uncommitted requests, an implicit syncpoint occurs. If the program ends abnormally, an implicit backout occurs. On MVS/ESA, an implicit syncpoint occurs if the program ends normally without first calling MQDISC.

For MQSeries for MVS/ESA programs, you can use the MQGMO_MARK_SKIP_BACKOUT option to specify that a message should not be backed out if backout occurs (in order to avoid an 'MQGET-error-backout' loop). For information about using this option, see "Skipping backout" on page 145.

Syncpoints in CICS for MVS/ESA and CICS/MVS applications

In a CICS for MVS/ESA and CICS/MVS application you establish a syncpoint by using the EXEC CICS SYNCPOINT command. To back out all changes to the previous syncpoint, you can use the EXEC CICS SYNCPOINT ROLLBACK command. For more information, see the *CICS Application Programming Reference* manual.

If other recoverable resources are also involved in the unit of work, the queue manager (in conjunction with the CICS syncpoint manager) participates in a two-phase commit protocol; otherwise, the queue manager performs a single-phase commit process.

If a CICS application issues the MQDISC call, no implicit syncpoint is taken. If the application closes down normally, any open queues are closed and an implicit commit occurs. If the application closes down abnormally, any open queues are closed and an implicit backout occurs.

Syncpoints in IMS applications

In an IMS application, you establish a syncpoint by using IMS calls such as GU (get unique) to the IOPCB and CHKP (checkpoint). To back out all changes since the previous checkpoint, you can use the IMS ROLB (rollback) call. For more information, see the following books:

- *IMS/ESA Version 4 Application Programming: DL/I Calls*
- *IMS/ESA Version 4 Application Programming: Design Guide*
- *IMS/ESA Version 5 Application Programming: Database Manager*
- *IMS/ESA Version 5 Application Programming: Design Guide*

The queue manager (in conjunction with the IMS syncpoint manager) participates in a two-phase commit protocol if other recoverable resources are also involved in the unit of work.

All open handles are closed by the IMS adapter at a syncpoint (except in a nonmessage batch-oriented BMP). This is because a different user could initiate the next unit of work and MQSeries security checking is performed when the MQCONN and MQOPEN calls are made, not when the MQPUT or MQGET calls are made. The handles are closed at the beginning of the MQI call following the IMS call which initiated the syncpoint.

If you have not installed IMS APAR PN83757, handles are also closed after a ROLB call unless you are running IMS Version 3 or are running a nonmessage BMP.

If an IMS application (either a BMP or an MPP) issues the MQDISC call, open queues are closed but no implicit syncpoint is taken. If the application closes down normally, any open queues are closed and an implicit commit occurs. If the application closes down abnormally, any open queues are closed and an implicit backout occurs.

Syncpoints in MVS batch applications

For batch applications, you can use the MQCMIT and MQBACK calls. For downward compatibility, CSQBCMT and CSQBBAK are available as synonyms.

Committing changes using the MQCMIT call

As input, you must supply the connection handle (*Hconn*), which is returned by the MQCONN call.

The output from MQCMIT is a completion code and a reason code. The call completes with a warning if the syncpoint was completed but the queue manager backed out the put and get operations since the previous syncpoint.

Successful completion of the MQCMIT call indicates to the queue manager that the application has reached a syncpoint and that all put and get operations made since the previous syncpoint are to be made permanent.

There is a description of the MQCMIT call in the *MQSeries Application Programming Reference* manual.

Backing out changes using the MQBACK call

As input, you must supply a connection handle (*Hconn*). Use the handle that is returned by the MQCONN call.

The output from MQBACK is a completion code and a reason code.

It indicates to the queue manager that the application has reached a syncpoint and that all gets and puts that have been made since the last syncpoint are to be backed out.

There is a description of the MQBACK call in the *MQSeries Application Programming Reference* manual.

Syncpoints in MQSeries for AS/400 applications

To start the AS/400 commitment control facilities, use the STRCMTCTL command.

Note: The default value of *Commitment definition scope* is *ACTGRP. This must be defined as *JOB for MQSeries for AS/400.

All MQSeries for AS/400 code runs in a single, named activation group: QMQM.

If you call MQPUT, MQPUT1 or MQGET, specifying MQPMO_SYNCPOINT or MQGMO_SYNCPOINT, when MQSeries for AS/400 is not registered as an API commitment resource inside your commitment definition, MQSeries for AS/400 adds itself to the definition. This is typically the first such call in a job. While there are any API commitment resources registered under a particular commitment definition, you cannot end commitment control for that definition.

MQSeries for AS/400 removes its registration, as an API commitment resource, when you disconnect from the queue manager, provided there are no pending MQI operations in the current unit of work. The precise mechanism used depends on how you connected to the queue manager:

- If you connected to the queue manager explicitly, using the CCTMQM command, you must issue the DSCMQM command to disconnect from the queue manager.

- If you connected to the queue manager explicitly, using the MQCONN call, you must issue an MQDISC call to disconnect from the queue manager.

- If you connected to the queue manager implicitly, by opening an object using the MQOPEN call without a previous MQCONN call, you must use the MQCLOSE call for each of the opened objects, to disconnect from the queue manager.

If you disconnect from the queue manager while there are pending MQPUT, MQPUT1 or MQGET operations in the current unit of work, MQSeries for AS/400 remains registered as an API commitment resource in order that it is notified of the next commit or rollback. When the next syncpoint is reached, MQSeries for AS/400 commits or rolls back the changes as required.

If you attempt to issue an ENDCMTCTL command for that commitment definition, message `CPF8355` is issued, indicating that pending changes were active. This message also appears in the job log when the job ends. In order to avoid this situation, ensure that you commit or rollback all pending MQSeries for AS/400 operations, and that you disconnect from the queue manager. Thus, using commands COMMIT or ROLLBACK and DSCMQM before command ENDCMTCTL, should enable end-commitment-control to complete successfully.

To commit or rollback (otherwise backout) your unit of work, use one of the programming languages that supports the commitment control:

- CL commands: COMMIT and ROLLBACK
- ILE C Programming Functions: _Rcommit and _Rrollback

Syncpoints in CICS for AS/400 applications

MQSeries for AS/400 participates in CICS for AS/400 units of work. You can use the MQI within a CICS for AS/400 application to put and get messages inside the current unit of work.

You can use the EXEC CICS SYNCPOINT command to establish a syncpoint that includes the MQSeries for AS/400 operations. To back out all changes up to the previous syncpoint, you can use the EXEC CICS SYNCPOINT ROLLBACK command.

If you use MQPUT, MQPUT1, or MQGET with the MQPMO_SYNCPOINT, or MQGMO_SYNCPOINT, option set in a CICS for AS/400 application, you cannot log off CICS for AS/400 until MQSeries for AS/400 has removed its registration as an API commitment resource. Therefore, you should commit or back out any pending put or get operations before you disconnect from the queue manager. This will allow you to log off CICS for OS/400.

Syncpoints in MQSeries for OS/2 Warp, MQSeries for Windows NT, MQSeries for Digital OpenVMS, and MQSeries on UNIX systems

Syncpoint support operates on two types of units of work:

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 procedure.

A *global* unit of work is one in which resources belonging to other resource managers, such as XA-compliant databases, are also updated. For full integrity, a two-phase commit procedure must be used. MQSeries can either coordinate such units of work itself or it can be coordinated by another XA-compliant transaction manager such as IBM CICS. CICS for OS/2, Transaction Server for OS/2 Warp, and CICS for Windows NT all provide a single-phase commit procedure in which MQSeries can participate.

Local units of work

Units of work that involve only the queue manager are called *local* units of work. Syncpoint coordination is provided by the queue manager itself (internal coordination) using a single-phase commit process.

To start a local unit of work, the application issues MQGET, MQPUT, or MQPUT1 requests specifying the appropriate syncpoint option. The unit of work is committed using MQCMIT or rolled back using MQBACK. However, the unit of work also ends when the connection between the application and the queue manager is broken, whether intentionally or unintentionally.

If an application disconnects (MQDISC) from a queue manager while a unit of work is still active, the unit of work is committed. If, however, the application terminates without disconnecting, the unit of work is rolled back as the application is deemed to have terminated abnormally.

Global units of work

Use global units of work when you also need to include updates to resources belonging to other resource managers. Here the coordination may be internal or external to the queue manager:

Internal syncpoint coordination

Queue manager coordination of global units of work is supported on MQSeries Version 5 products only.

Here, the coordination is performed by MQSeries. To start a global unit of work, the application issues the MQBEGIN call. This allows the queue manager to coordinate its own updates in addition to those made by other resource managers.

As input to the MQBEGIN call, you must supply the connection handle (*Hconn*), which is returned by the MQCONN call. This represents the connection to the MQSeries queue manager.

Again, the application issues MQGET, MQPUT, or MQPUT1 requests specifying the appropriate 'syncpoint' option. This means that MQBEGIN can be used to initiate a global unit of work that updates local resources, resources belonging to other resource managers, or both. Updates made to resources belonging to other resouce managers are made using the API of that resource manager. However, it is not possible to use the MQI to update queues that belong to other queue managers. MQCMIT or MQBACK must be issued before starting further units of work (local or global).

The global unit of work is committed using MQCMIT; this initiates a two-phase commit of all the resource managers involved in the unit of work. A two-phase commit process is used whereby resource managers (for example, XA-compliant database managers such as DB2 and Oracle) are firstly all asked to prepare to commit. Only if all are prepared are they asked to commit. If any resource manager signals that it cannot commit, each is asked to back out instead. Alternatively, MQBACK can be used to roll back the updates of all the resource managers.

If an application disconnects (MQDISC) while a global unit of work is still active, the unit of work is committed. If, however, the application terminates without disconnecting, the unit of work is rolled back as the application is deemed to have terminated abnormally.

The output from MQBEGIN is a completion code and a reason code.

When MQBEGIN is used to start a global unit of work, all the external resource managers that have been configured with the queue manager are included. However, the call starts a unit of work but completes with a warning if:

- There are no participating resource managers (that is, no resource managers have been configured with the queue manager)

or

- One or more resource managers are not available.

In these cases, the unit of work can include updates to only those resource managers that were available when the unit of work was started.

If one of the resource managers is unable to commit its updates, all of the resource managers are instructed to roll back their updates, and MQCMIT completes with a warning. In unusual circumstances (typically, operator intervention), an MQCMIT call may fail if some resource managers commit their updates but others roll them back; the work is deemed to have completed with a 'mixed' outcome. Such occurrences are diagnosed in the error log of the queue manager so remedial action may be taken.

An MQCMIT of a global unit of work succeeds if all of the resource managers involved commit their updates.

For a description of the MQBEGIN call, see the *MQSeries Application Programming Reference*.

External syncpoint coordination

This is when a syncpoint coordinator other than MQSeries has been selected, for example CICS, Encina, or Tuxedo. In this situation, MQSeries for OS/2 Warp, MQSeries for Windows NT, and MQSeries on UNIX systems register their interest in the outcome of the unit of work with the syncpoint coordinator in order that they can commit or roll back any uncommitted get or put operations as required. The external syncpoint coordinator determines whether one- or two-phase commitment protocols are provided.

You can also use the Transaction Monitor supplied with the Version 5 products for external syncpoint coordination. You cannot use a transaction monitor with a client application.

When an external coordinator is used (that is, when you are in XA mode), MQCMIT, MQBACK, and MQBEGIN may not be issued. Calls to these functions fail with the reason code MQRC_ENVIRONMENT_ERROR.

The way in which an externally coordinated unit of work is started is dependent on the interface provided by the syncpoint coordinator. There may or may not be an explicit call required. If an explicit call is required, and you issue an MQPUT call specifying the MQPMO_SYNCPOINT option when a unit of work is not started, the completion code MQRC_SYNCPOINT_NOT_AVAILABLE is returned.

The scope of the unit of work is determined by the syncpoint coordinator. The state of the connection between the application and the queue manager affects the success or failure of MQI calls that an application issues, not the state of the unit of work. It is, for example, possible for an application to disconnect and reconnect to a queue manager during an active unit of work and perform further MQGET and MQPUT operations inside the same unit of work (a *pending disconnect*).

Interfaces to external syncpoint managers

MQSeries on UNIX systems MQSeries for OS/2 Warp, and MQSeries for Windows NT support coordination of transactions by external syncpoint managers which utilize the X/Open XA interface. This support is available only on server configurations. The interface is not available to client applications.

Some XA transaction managers (not CICS on Open Systems or Encina) require that each XA resource manager supplies its name. This is the string called name in the XA switch structure. The resource manager for MQSeries on UNIX systems is named "MQSeries_XA_RMI". For further details on XA interfaces refer to XA

documentation *X/Open CAE Specifications Distributed Transaction Processing: The XA Specification*.

In an XA configuration, MQSeries on UNIX systems, MQSeries for OS/2 Warp, and MQSeries for Windows NT fulfil the role of an XA Resource Manager. An XA syncpoint coordinator can manage a set of XA Resource Managers, and synchronize the commit or backout of transactions in both Resource Managers. This is how it works for a statically registered resource manager:

1. An application notifies the syncpoint coordinator that it wishes to start a transaction.

2. The syncpoint coordinator issues a call to any resource managers that it knows of, to notify them of the current transaction.

3. The application issues a call (for example MQGET in syncpoint) to the resource manager that is associated with the current transaction.

4. The application requests that the syncpoint coordinator either commit or roll back the transaction.

5. The syncpoint coordinator issues a call to each resource manager using two-phase commit protocols to complete the transaction as requested.

XA requires each Resource Manager to provide a structure called an *XA Switch*. This structure declares the capabilities of the Resource Manager, and the functions that are to be called by the syncpoint coordinator.

There are two versions of this structure:

MQRMIXASwitch   Static XA resource management

MQRMIXASwitchDynamic   Dynamic XA resource management

The structure is found in the following libraries:

mqmxa.lib   OS/2 and Windows NT XA library for Static resource management

mqmenc.lib   AIX, HP-UX, Sun Solaris, and Windows NT Encina XA library for Dynamic resource management

libmqmxa.a   UNIX systems XA library (non-threaded) for both Static and Dynamic resource management

libmqmxa_r.a   UNIX systems XA library (threaded) for both Static and Dynamic resource management

The method that must be used to link them to an XA syncpoint coordinator is defined by the coordinator, and you will need to consult the documentation provided by that coordinator to determine how to enable MQSeries to cooperate with your XA syncpoint coordinator.

The *xa_info* structure that is passed on any *xa_open* call by the syncpoint coordinator should be the name of the queue manager that is to be administered. This takes the same form as the queue manager name passed to MQCONN, and may be blank if the default queue manager is to be used.

Restrictions

- On OS/2, all functions declared in the XA switch are declared as _System functions.

- On Windows NT, all functions declared in the XA switch are declared as _cdecl functions.

- Only one queue manager may be administered by an external syncpoint coordinator at a time. This is due to the fact that the coordinator has an effective connection to each queue manager, and is therefore subject to the rule that only one connection is allowed at a time.

- All applications that are run using the syncpoint coordinator can connect only to the queue manager that is administered by the coordinator because they are already effectively connected to that queue manager. They must issue MQCONN to obtain a connection handle and must issue MQDISC before they exit. Alternatively, they can use the CICS user exit 15 for CICS for OS/2 V2 and V3, and CICS for Windows NT V2, or the exit UE014015 for Transaction Server for Windows NT V4 and CICS on Open Systems.

The features not implemented are:

- Association migration
- Asynchronous calls

Because Transaction Server V4 is 32 bit, changes are required to the source of CICS user exits. The supplied samples have been updated to work with Transaction Server V4 as shown in Table 7.

User exit	CICS V2 source	CICS V2 dll	TS V4 source	TS V4 dll
exit 15	amqzsc52.c	faaexp15.dll	amqzsc53.c	faaex315.dll
exit 17	amqzsc72.c	faaexp17.dll	amqzsc73.c	faaex317.dll

Table 7. Linking MQSeries for OS/2 Warp with CICS Version 3 applications

For Transaction Server V4, the supplied user exits faaex315.dll and faaex317.dll should be renamed to the standard names faaexp15.dll and faaexp17.dll.

Syncpoints in MQSeries for Tandem NSK applications

When using MQSeries for Tandem NSK transaction management is performed under the control of Tandem's TM/MP product, rather than by MQSeries itself.

The effects of this difference are:

- The default SYNCPOINT option for the MQPUT and MQGET calls is SYNCPOINT, rather than NO_SYNCPOINT.

- To use the default (SYNCPOINT) option for MQPUT, MQGET, and MQPUT1 operations, the application must have an active TM/MP Transaction that defines the unit of work to be committed. An application initiates a TM/MP transaction by calling the BEGINTRANSACTION() function. All MQPUT, MQPUT1, and MQGET operations performed by the application while this transaction is active are within the same unit of work (transaction). Any other database operations

performed by the application are also within this UOW. Note that there are system-imposed limits on the number and size of messages that can be written and deleted within a single TM/MP transaction. When the application has completed the UOW, the TM/MP transaction is ended (the UOW is committed) using the ENDTRANSACTION() function. If any error is encountered, the application can cancel the TM/MP transaction (backout the UOW) using the ABORTTRANSACTION() function. Consequently, the standard Version 2 functions MQCMIT() and MQBACK() are not supported on this product. If they are called, an error is returned.

* If an application uses the NO_SYNCPOINT option for MQPUT, MQGET, and MQPUT1 operations, MQSeries starts a TM/MP transaction itself, performs the queuing operation, and commits the transaction before returning to the application. Each operation is therefore performed in its own UOW and, once complete, cannot be backed out by the application using TM/MP.

* A TM/MP transaction does not need to be active for MQI calls other than MQGET, MQPUT, and MQPUT1.

* Because TM/MP can cause previously performed MQGET, MQPUT, and MQPUT1 operations to be backed out without notification, the current queue-depth and input-and-output-open counts of queues can become inaccurate. The **cleanqm** control command runs in the background to correct such inaccuracies, and each MQOPEN call corrects the value of these attributes on a queue if necessary. However, applications should be coded to be resilient to inaccuracies in these quantities, especially in an environment that may involve backed-out transactions.

* The back-out count attribute cannot be maintained in the same way as on standard Version 2 implementations. Also, the harden backout count attribute is not used.

* The MQRC_SYNCPOINT_LIMIT_REACHED reason code is used by MQSeries for Tandem NSK V2.2 to inform an application that the system-imposed limit on the number of I/O operations within a single TM/MP transaction has been reached. If the application specified the SYNCPOINT option, it should cancel the transaction (back out the UOW) and retry with a smaller number of operations in that UOW.

* The MQRC_UOW_CANCELED reason code informs the application that the UOW (TM/MP transaction) has been canceled, either by the system itself (TM/MP imposes some system-wide resource-usage thresholds that will cause this), by user action, or by the initiator of the transaction itself.

General XA support

This is not supported on Digital OpenVMS or Tandem NSK

An XA switch load module is provided to enable you to link CICS with MQSeries on UNIX systems. Additionally, sample source code files are provided to enable you to develop the XA switches for other transaction messages. The names of the switch load modules provided are:

General XA support

Table 8. Essential Code for CICS applications		
Description	**C (source)**	**C (exec) - add one of the following to your XAD.Stanza**
XA initialization routine	amqzscix.c	amqzsc - CICS for AIX Version 2.1, amqzsc - Transaction Server for AIX, Version 4, amqzsc - CICS for HP 9000, amqzsc - CICS for Siemens Nixdorf SINIX Version 2.2, amqzsc - CICS for Solaris Version 2.1.1
	amqzscin.c	mqmc4swi - Transaction Server for Windows NT, Version 4

Chapter 14. Starting MQSeries applications using triggers

Triggering is not supported on MQSeries for Windows.

Some MQSeries applications that serve queues run continuously, so they are always available to retrieve messages that arrive on the queues. However, this may not be desirable when the number of messages arriving on the queues is unpredictable. In this case, applications could be consuming system resources even when there are no messages to retrieve.

MQSeries provides a facility that enables an application to be started automatically when there are messages available to retrieve. This facility is known as *triggering*.

This chapter describes:

- What triggering is
- The prerequisites for using triggering
- The conditions that must exist for a trigger event to occur
- How to control trigger events
- How to design an application that uses triggered queues
- The properties of trigger messages
- What happens when triggering does not work

What is triggering?

The queue manager defines certain conditions as constituting "trigger events". If triggering is enabled for a queue and a trigger event occurs, the queue manager sends a *trigger message* to a queue called an *initiation queue*. The presence of the trigger message on the initiation queue indicates that a trigger event has occurred.

Trigger messages generated by the queue manager are not persistent. This has the effect of reducing logging (thereby improving performance), and minimizing duplicates during restart, so improving restart time.

The program which processes the initiation queue is called a *trigger-monitor application*, and its function is to read the trigger message and take appropriate action, based on the information contained in the trigger message. Normally this action would be to start some other application to process the queue which caused the trigger message to be generated. From the point of view of the queue manager, there is nothing special about the trigger-monitor application—it is simply another application that reads messages from a queue (the initiation queue).

If triggering is enabled for a queue, you have the option to create a *process-definition object* associated with it. This object contains information about the application that processes the message which caused the trigger event. If the process definition object is created, the queue manager extracts this information and places it in the trigger message, for use by the trigger-monitor application. The name of the process definition associated with a queue is given by the *ProcessName* local-queue attribute. Each queue can specify a different process definition, or several queues can share the same process definition.

Triggering

| On MQSeries Version 5 products, and MQSeries for AS/400 in the case of triggering a channel, you do not need to create a process definition object; the transmission queue definition is used instead. When a trigger event occurs, the transmission queue definition contains information about the application that processes the message which caused the event. Again, when the queue manager generates the trigger message, it extracts this information and places it in the trigger message.

Support in an MQSeries client environment

Triggering is supported by MQSeries clients in the UNIX systems, Digital OpenVMS, OS/2, Windows 3.1, Windows 95, and Windows NT environments. An application running in a client environment is the same as one running in a full MQSeries environment, only you link it with the client libraries. However the trigger monitor and the application to be started must both be in the same environment as one another.

Triggering involves:

Application queue
An *application queue* is a local queue, which, when it has triggering set on and when the conditions are met, requires that trigger messages are written.

Process Definition
An application queue can have a *process definition object* associated with it that holds details of the application that will get messages from the application queue. (See "Process definition objects" on page 48 for a list of attributes.)

On MQSeries Version 5 products, the process definition object is optional in the case of triggering channels.

Transmission queue
The transmission queue holds the name of the channel to be triggered. This can replace the process definiton for triggering channels, but is used only when a process definition is not created.

Trigger event
A *trigger event* is an event that causes a trigger message to be generated by the queue manager. This is usually a message arriving on an application queue, but it can also occur at other times (see "Conditions for a trigger event" on page 204). MQSeries has a range of options to allow you to control the conditions that cause a trigger event (see "Controlling trigger events" on page 208).

Trigger message
The queue manager creates a *trigger message* when it recognizes a trigger event (see "Conditions for a trigger event" on page 204). It copies into the trigger message information about the application to be started. This information comes from the application queue and the process definition object associated with the application queue. Trigger messages have a fixed format (see "Format of trigger messages" on page 216).

Initiation queue
An *initiation queue* is a local queue on which the queue manager puts trigger messages. A queue manager can own more than one initiation queue, and each one is associated with one or more application queues.

Trigger monitor

A *trigger monitor* is a continuously-running program that serves one or more initiation queues. When a trigger message arrives on an initiation queue, the trigger monitor retrieves the message. The trigger monitor uses the information in the trigger message. It issues a command to start the application that is to retrieve the messages arriving on the application queue, passing it information contained in the trigger message header, which includes the name of the application queue. (For more information, see "Trigger monitors" on page 212.)

To understand how triggering works, consider Figure 12, which is an example of trigger type FIRST, (MQTT_FIRST).
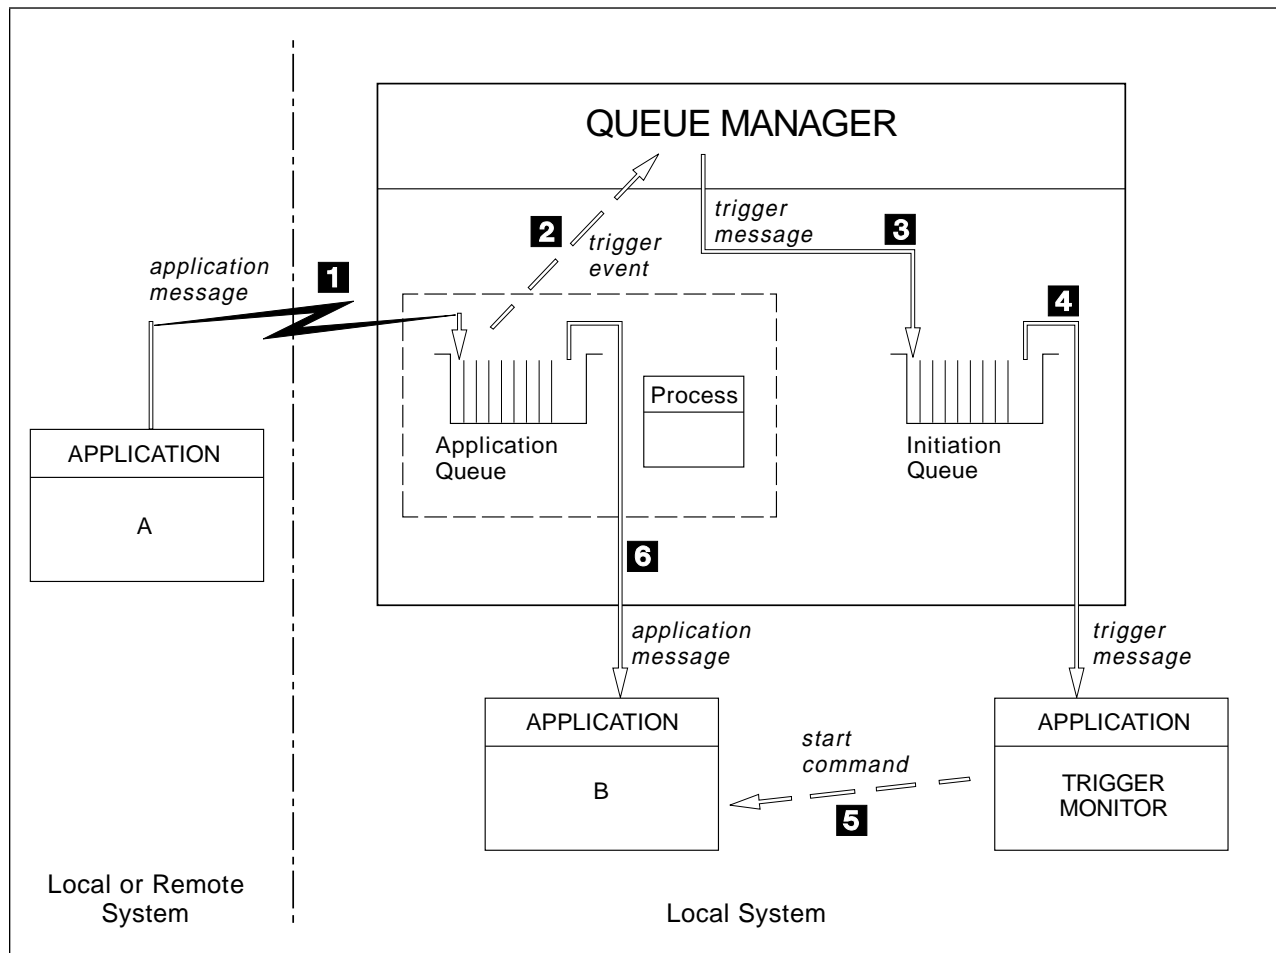
Figure 12. Flow of application and trigger messages

In Figure 12, the sequence of events is:

1. Application A, which can be either local or remote to the queue manager, puts a message on the application queue. Note that no application has this queue open for input. However this fact is only relevant to trigger type FIRST and DEPTH.

2. The queue manager checks to see if the conditions are met under which it has to generate a trigger event. They are, and a trigger event is generated, passing on information held within the associated process definition object.

3. The queue manager creates a trigger message and puts it on the initiation queue associated with this application queue, but only if an application (trigger monitor) has the initiation queue open for input.

4. The trigger monitor retrieves the trigger message from the initiation queue.

5. The trigger monitor issues a command to start program B (the server application).

6. Application B opens the application queue and retrieves the message.

Notes:

1. If the application queue is open for input, by any program, and has triggering set for FIRST or DEPTH, no trigger event will occur - it's not needed.

2. If the initiation queue is not open for input, the queue manager will not generate any trigger messages, it will wait until an application opens the initiation queue for input.

3. Only use type FIRST or DEPTH when using triggering for channels.

So far, the relationship between the queues within triggering has been only on a one to one basis. Consider Figure 13 on page 201.

Figure 13. Relationship of queues within triggering

An application queue has a process definition object associated with it that holds details of the application that will process the message. The queue manager places the information in the trigger message, so only one initiation queue is necessary. The trigger monitor extracts this information from the trigger message and starts the relevant application to deal with the message on each application queue.

On MQSeries Version 5 products, in the case of triggering a channel, the process definition object is optional. The transmission queue definition can determine the channel to be triggered.

Prerequisites for triggering

Before your application can take advantage of triggering, follow the steps below:

1. Either:

    a. create an initiation queue for your application queue. For example:

    ```
    DEFINE QLOCAL (initiation.queue) REPLACE      +
           LIKE (SYSTEM.DEFAULT.LOCAL.QUEUE)      +
           DESCR ('initiation queue description')
    ```

    or

    b. determine the name of a local queue that already exists and can be used by your application, and specify its name in the *InitiationQName* field of the application queue.

    You can think of this task as associating the initiation queue with the application queue. A queue manager can own more than one initiation queue—you may want some of your application queues to be served by different programs, in which case you could use one initiation queue for each serving program, although you do not have to. Here is an example of how to create an application queue:

    ```
    DEFINE QLOCAL (application.queue) REPLACE   +
    LIKE (SYSTEM.DEFAULT.LOCAL.QUEUE)           +
    DESCR ('appl queue description')     +
    INITQ ('initiation.queue')           +
    PROCESS ('process.name')             +
    TRIGGER                                     +
    TRIGTYPE (FIRST)
    ```

2. If you are triggering an application, create a process definition object to contain information relating to the application that is to serve your application queue. For example:

    ```
    DEFINE PROCESS (process.name) +
           REPLACE +
           DESCR ('process description') +
           APPLTYPE ('CICS') +
           APPLICID ('CKSG') +
           USERDATA ('EXAMPLE.CHANNEL')
    ```

    Here is an extract from an MQSeries for AS/400 CL program that creates a process definition object:

```
        /*   Queue used by AMQSINQA                               */
                CRTMQMQ    QNAME('SYSTEM.SAMPLE.INQ')          +
                           QTYPE(*LCL)  REPLACE(*YES)          +
                                                               +
                           TEXT('queue for AMQSINQA')          +
                           SHARE(*YES)             /* Shareable   */+
                           DFTMSGPST(*YES)/* Persistent messages OK */+
                                                               +
                           TRGENBL(*YES)  /* Trigger control on    */+
                           TRGTYPE(*FIRST)/* Trigger on first message*/+
                           PRCNAME('SYSTEM.SAMPLE.INQPROCESS')     +
                           INITQNAME('SYSTEM.SAMPLE.TRIGGER')

        /*   Process definition                                   */
                CRTMQMPRC  PRCNAME('SYSTEM.SAMPLE.INQPROCESS')     +
                           REPLACE(*YES)                       +
                                                               +
                           TEXT('trigger process for AMQSINQA')    +
                           ENVDATA('JOBPTY(3)') /* Submit parameter */+
                           APPID('AMQSINQA')    /* Program name    */
```

When the queue manager creates a trigger message, it copies information from the attributes of the process definition object into the trigger message. This step is optional in the case of triggering channels.

Platform	To create a process definition object
UNIX systems, Digital OpenVMS, OS/2, Windows NT	Use DEFINE PROCESS or use SYSTEM.DEFAULT.PROCESS and modify using ALTER PROCESS
MVS/ESA	Use DEFINE PROCESS (see sample code in step 2), or use the operations and control panels.
OS/400	Use a CL program like the extract from AMQSDEF4 in step 2.

3. If you are triggering a channel on a release other than MQSeries Version 5 products, or MQSeries for AS/400 V4R2, you need to create a process definition. Create a transmission queue definition and specify the *ProcessName* attribute as blanks. The *TrigData* attribute can contain the name of the channel to be triggered or it can be left blank. When the queue manager creates a trigger message, it copies information from the *TrigData* attribute of the transmission queue definition, not the trigger message.

4. If you have created a process definition object, associate your application queue with the application that is to serve that queue by naming the process definition object in the *ProcessName* attribute of the queue.

Platform	Use commands
UNIX systems, Digital OpenVMS, OS/2, Windows NT	ALTER QLOCAL
MVS/ESA	ALTER QLOCAL
AS/400	CHGMQMQ

5. Start instances of the trigger monitors (or trigger servers in MQSeries for AS/400) that are to serve the initiation queues you have defined. See "Trigger monitors" on page 212 for more information.

If you wish to be aware of any undelivered trigger messages, make sure your queue manager has a dead-letter (undelivered-message) queue defined. Specify the name of the queue in the *DeadLetterQName* queue manager field.

You can then set the trigger conditions you require, using the attributes of the queue object that defines your application queue. For more information on this, see "Controlling trigger events" on page 208.

Conditions for a trigger event

The queue manager creates a trigger message when the following conditions are satisfied:

1. A message is **put** on a queue.

2. The message has a priority greater than or equal to the threshold trigger priority of the queue. This priority is set in the *TriggerMsgPriority* local queue attribute—if it is set to zero, any message qualifies.

3. The number of messages on the queue with priority greater than or equal to *TriggerMsgPriority* was previously, depending on *TriggerType*:

   • Zero (for trigger type MQTT_FIRST)
   • Any number (for trigger type MQTT_EVERY)
   • *TriggerDepth* minus 1 (for trigger type MQTT_DEPTH)

   Note: The queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist. Consequently an application may be started when there are no messages for it to retrieve because the messages on the queue have not been committed. In this situation, you are strongly recommended to consider using the wait option, and relating the *WaitInterval* to the number of messages in the unit of work.

4. For triggering of type FIRST or DEPTH, no program has the application queue open for removing messages (that is, the *OpenInputCount* local queue attribute is zero).

5. On MQSeries for MVS/ESA, if the application queue is one with a *Usage* attribute of MQUS_NORMAL, get requests for it are not inhibited (that is, the *InhibitGet* queue attribute is MQQA_GET_ALLOWED). Also, on MQSeries for non-MVS/ESA platforms, if the application queue is one with a *Usage* attribute of MQUS_XMITQ, get requests for it are not inhibited.

6. Either:

   • The *ProcessName* local queue attribute for the queue is not blank, and the process definition object identified by that attribute has been created.

   or

   • The *ProcessName* local queue attribute for the queue is all blank, but the queue is a transmission queue. In this case, the trigger message contains attributes with the following values:

     ProcessName: blanks
     TriggerData: trigger data

> > *ApplType*: MQAT_UNKNOWN
> > *ApplId*: blanks
> > *EnvData*: blanks
> > *UserData*: blanks

> **Note:** As the process definition is optional, the *TriggerData* attribute may also contain the name of the channel to be started. This option is available on MQSeries for AS/400, OS/2, HP-UX, AIX, Sun Solaris, and Windows NT only.

7. An initiation queue has been created, and has been specified in the *InitiationQName* local queue attribute. Also:

   - Get requests are not inhibited for the initiation queue (that is, the *InhibitGet* queue attribute is MQQA_GET_ALLOWED).

   - Put requests must not be inhibited for the initiation queue (that is, the *InhibitPut* queue attribute must be MQQA_PUT_ALLOWED).

   - The *Usage* attribute of the initiation queue must be MQUS_NORMAL.

   - In environments where dynamic queues are supported, the initiation queue must not be a dynamic queue that has been marked as logically deleted.

8. A trigger monitor currently has the initiation queue open for removing messages (that is, the *OpenInputCount* local queue attribute is greater than zero).

9. The trigger control (*TriggerControl* local queue attribute) for the application queue is set to MQTC_ON. To do this, set the *trigger* attribute when you define your queue, or use the ALTER QLOCAL command.

10. The trigger type (*TriggerType* local queue attribute) is not MQTT_NONE.

    If all of the above required conditions are met, and the message that caused the trigger condition is put as part of a unit of work, the trigger message does not become available for retrieval by the trigger monitor application until the unit of work completes, whether the unit of work is committed **or** backed out.

11. A suitable message is placed on the queue, for a *TriggerType* of MQTT_FIRST or MQTT_DEPTH, and the queue:

    - Was not previously empty (MQTT_FIRST)

    or

    - Had *TriggerDepth* or more messages (MQTT_DEPTH)

    and conditions 2 through 10 (excluding 3) are satisfied, if in the case of MQTT_FIRST a sufficient interval (*TriggerInterval* queue-manager attribute) has elapsed since the last trigger message was written for this queue.

    This is to allow for a queue server that ends before processing all of the messages on the queue. The purpose of the trigger interval is to reduce the number of duplicate trigger messages that are generated.

12. The only application serving a queue issues an MQCLOSE call, for a *TriggerType* of MQTT_FIRST or MQTT_DEPTH, and there is at least:

    - One (MQTT_FIRST)

    or

    - *TriggerDepth* (MQTT_DEPTH)

    messages on the queue of sufficient priority (condition 2 on page 204), and conditions 6 through 10 are also satisfied.

This is to allow for a queue server that issues an MQGET call, finds the queue empty, and so ends; however, in the interval between the MQGET and the MQCLOSE calls, one or more messages arrive.

Notes:

a. If the program serving the application queue does not want to retrieve all the messages, this can cause a closed loop. Each time the program closes the queue, the queue manager creates another trigger message which causes the trigger monitor to start the server program again.

b. If the program serving the application queue backs out its get request (or if the program abends) before it closes the queue, the same happens.

c. To prevent such a loop occurring, you could use the *BackoutCount* field of MQMD to detect messages that are repeatedly backed out. For more information, see "Messages that are backed out" on page 34.

13. The following conditions are satisfied using MQSET or a command:

a.

- *TriggerControl* is changed to MQTC_ON

or

- *TriggerControl* is already MQTC_ON and the value of either *TriggerType*, *TriggerMsgPriority*, or *TriggerDepth* (if relevant) is changed,

and there is at least:

- One (MQTT_FIRST or MQTT_EVERY)

or

- *TriggerDepth* (MQTT_DEPTH)

messages on the queue of sufficient priority (condition 2 on page 204), and conditions 4 through 10 on page 205 (excluding 8) are also satisfied.

This is to allow for an application or operator changing the triggering criteria, when the conditions for a trigger to occur are already satisfied.

b. The *InhibitPut* queue attribute of an initiation queue changes from MQQA_PUT_INHIBITED to MQQA_PUT_ALLOWED, and there is at least:

- One (MQTT_FIRST or MQTT_EVERY)

or

- *TriggerDepth* (MQTT_DEPTH)

messages of sufficient priority (condition 2 on page 204) on any of the queues for which this is the initiation queue, and conditions 4 through 10 on page 205 are also satisfied. (One trigger message is generated for each such queue satisfying the conditions.)

This is to allow for trigger messages not being generated because of the MQQA_PUT_INHIBITED condition on the initiation queue, but this condition now having been changed.

c. The *InhibitGet* queue attribute of an application queue changes from MQQA_GET_INHIBITED to MQQA_GET_ALLOWED, and there is at least:

- One (MQTT_FIRST or MQTT_EVERY)

or

- *TriggerDepth* (MQTT_DEPTH)

messages of sufficient priority (condition 2 on page 204) on the queue, and conditions 4 through 10 on page 205, excluding 5, are also satisfied.

This allows applications to be triggered only when they are able to retrieve messages from the application queue.

d. A trigger-monitor application issues an MQOPEN call for input from an initiation queue, and there is at least:

- One (MQTT_FIRST or MQTT_EVERY)

or

- *TriggerDepth* (MQTT_DEPTH)

messages of sufficient priority (condition 2 on page 204) on any of the application queues for which this is the initiation queue, and conditions 4 through 10 on page 205 (excluding 8) are also satisfied, and no other application has the initiation queue open for input (one trigger message is generated for each such queue satisfying the conditions).

This is to allow for messages arriving on queues while the trigger monitor is not running, and for the queue manager restarting and trigger messages (which are nonpersistent) being lost.

Note: From step 12 (where trigger messages are generated as a result of some event other than a message arriving on the application queue), the trigger message is not put as part of a unit of work. Also, if the *TriggerType* is MQTT_EVERY, and if there are one or more messages on the application queue, only one trigger message is generated.

14. MSGDLVSQ is set correctly.

If you set MSGDLVSQ=FIFO, messages are delivered to the queue in a First In First Out basis. The priority of the message is ignored and the default priority of the queue is assigned to the message. If *TriggerMsgPriority* is set to a higher value than the default priority of the queue, no messages are triggered. If *TriggerMsgPriority* is set equal to or lower than the default priority of the queue, triggering occurs for type FIRST, EVERY, and DEPTH. For information about these types, see the description of the *TriggerType* field under "Controlling trigger events" on page 208.

If you set MSGDLVSQ=PRIORITY and the message priority is equal to or greater than the *TriggerMsgPriority* field, messages only **count** towards a trigger event. In this case, again triggering occurs for type FIRST, EVERY, and DEPTH.

As an example, if you put 100 messages of lower priority than the *TriggerMsgPriority*, the effective queue depth for triggering purposes is still zero. If you then put another message on the queue, but this time the priority is greater than or equal to the *TriggerMsgPriority*, the effective queue depth increases from zero to one and the condition for *TriggerType* FIRST is satisfied.

Controlling trigger events

You control trigger events using some of the attributes that define your application queue. You can enable and disable triggering, and you can select the number or priority of the messages that count toward a trigger event. There is a full description of these attributes in the *MQSeries Application Programming Reference* manual.

The relevant attributes are:

TriggerControl

Use this attribute to enable and disable triggering for an application queue.

TriggerMsgPriority

The minimum priority that a message must have for it to count toward a trigger event. If a message of priority less than *TriggerMsgPriority* arrives on the application queue, the queue manager ignores the message when it determines whether to create a trigger message. If *TriggerMsgPriority* is set to zero, all messages count toward a trigger event.

TriggerType

In addition to the trigger type NONE (which disables triggering just like setting the *TriggerControl* to OFF), you can use the following trigger types to set the sensitivity of a queue to trigger events:

EVERY   A trigger event occurs every time a message arrives on the application queue. Use this type of trigger if you want a serving program to process only one message, then end.

FIRST   A trigger event occurs only when the number of messages on the application queue changes from zero to one. Use this type of trigger if you want a serving program to start when the first message arrives on a queue, continue until there are no more messages to process, then end. Also see "Special case of trigger type FIRST" on page 209.

DEPTH   A trigger event occurs only when the number of messages on the application queue reaches the value of the *TriggerDepth* attribute. A typical use of this type of triggering is for starting a program when all the replies to a set of requests are received.

> **Triggering by depth**
>
> With triggering by depth, the queue manager disables triggering (using the *TriggerControl* attribute) after it creates a trigger message. Your application must reenable triggering itself (by using the MQSET call) after this has happened.

The action of disabling triggering is not under syncpoint control, so triggering cannot be reenabled simply by backing out a unit of work. If a program backs out a put request that caused a trigger event, or if the program abends, you must reenable triggering by using the MQSET call or the ALTER QLOCAL command.

TriggerDepth
> The number of messages on a queue that causes a trigger event when using triggering by depth.

The conditions that must be satisfied for a queue manager to create a trigger message are described in "Conditions for a trigger event" on page 204.

Example of the use of trigger type EVERY

Consider an application that generates requests for motor insurance. The application might send request messages to a number of insurance companies, specifying the same reply-to queue each time. It could set a trigger of type EVERY on this reply-to queue so that each time a reply arrives, the reply could trigger an instance of the server to process the reply.

Example of the use of trigger type FIRST

Consider an organization with a number of branch offices that each transmit details of the day's business to the head office. They all do this at the same time, at the end of the working day, and at the head office there is an application that processes the details from all the branch offices. The first message to arrive at the head office could cause a trigger event which starts this application. This application would continue processing until there are no more messages on its queue.

Example of the use of trigger type DEPTH

Consider a travel agency application that creates a single request to confirm a flight reservation, to confirm a reservation for a hotel room, to rent a car, and to order some travelers' checks. The application could separate these items into four request messages, sending each to a separate destination. It could set a trigger of type DEPTH on its reply-to queue (with the depth set to the value 4), so that it is restarted only when all four replies have arrived.

If another message (possibly from a different request) arrives on the reply-to queue before the last of the four replies, the requesting application is triggered early. To avoid this, when DEPTH triggering is being used to collect multiple replies to a request, you should always use a new reply-to queue for each request.

Special case of trigger type FIRST

With trigger type FIRST, if there is already a message on the application queue when another message arrives, the queue manager does not usually create another trigger message. However, the application serving the queue might not actually open the queue (for example, the application might end, possibly because of a system problem). If an incorrect application name has been put into the process definition object, the application serving the queue will not pick up any of the messages. In these situations, if another message arrives on the application queue, there is no server running to process this message (and any other messages on the queue).

To deal with this, the queue manager creates another trigger message if another message arrives on the application queue, but only if a predefined time interval has elapsed since the queue manager created the last trigger message for that queue. This time interval is defined in the queue manager attribute *TriggerInterval*. Its default value is 999 999 999 milliseconds.

You should consider the following points when deciding on a value for the trigger interval to be used in your application:

- If *TriggerInterval* is set to a low value, trigger type FIRST might behave like trigger type EVERY (this depends on the rate that messages are being put onto the application queue, which in turn may depend on other system activity). This is because, if the trigger interval is very small, another trigger message is generated each time a message is put onto the application queue, even though the trigger type is FIRST, not EVERY. (Trigger type FIRST with a trigger interval of zero is equivalent to trigger type EVERY.)

- If a unit of work is backed out (see "Trigger messages and units of work") and the trigger interval has been set to a high value (or the default value), one trigger message is generated when the unit of work is backed out. However, if you have set the trigger interval to a low value or to zero (causing trigger type FIRST to behave like trigger type EVERY) many trigger messages can be generated. If the unit of work is backed out, all the trigger messages are still made available. The number of trigger messages generated depends on the trigger interval, the maximum number being reached when trigger interval has been set to zero.

Designing an application that uses triggered queues

You have seen how to set up, and control, triggering for your applications. Here are some tips you should consider when you design your application.

Trigger messages and units of work

Trigger messages created because of trigger events that are not part of a unit of work are put on the initiation queue, outside any unit of work, with no dependence on any other messages, and are available for retrieval by the trigger monitor immediately.

Trigger messages created because of trigger events that *are* part of a unit of work are put on the initiation queue as part of the same unit of work. Trigger monitors cannot retrieve these trigger messages until the unit of work completes. This applies whether the unit of work is committed or backed out.

If the queue manager fails to put a trigger message on an initiation queue, it will be put on the dead-letter (undelivered-message) queue.

Note: The queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist.

With triggering of type FIRST or DEPTH, trigger messages are made available even if the unit of work is backed out so that a trigger message is always available when the required conditions are met. For example, consider a put request within a unit of work for a queue that is triggered with trigger type FIRST. This causes the queue manager to create a trigger message. If another put request occurs, from another unit of work, this does not cause another trigger event because the number of messages on the application queue has now changed from one to two, which does not satisfy the conditions for a trigger event. Now if the first unit of work is backed out, but the second is committed, a trigger message is still created.

However, this does mean that trigger messages are sometimes created when the conditions for a trigger event are *not* satisfied. Applications that use triggering

must always be prepared to handle this situation. It is recommended that you use the wait option with the MQGET call, setting the *WaitInterval* to a suitable value.

Getting messages from a triggered queue

When you design applications that use triggering, you must be aware that there may be a delay between a program being started by a trigger monitor, and other messages becoming available on the application queue. This can happen when the message that causes the trigger event is committed before the others.

To allow time for messages to arrive, always use the wait option when you use the MQGET call to remove messages from a queue for which trigger conditions are set. The *WaitInterval* should be sufficient to allow for the longest reasonable time between a message being put and that put call being committed. If the message is arriving from a remote queue manager, this time is affected by:

- The number of messages that are put before being committed
- The speed and availability of the communication link
- The sizes of the messages

For an example of a situation where you should use the MQGET call with the wait option, consider the same example we used when describing units of work. This was a put request within a unit of work for a queue that is triggered with trigger type FIRST. This event causes the queue manager to create a trigger message. If another put request occurs, from another unit of work, this does not cause another trigger event because the number of messages on the application queue has not changed from zero to one. Now if the first unit of work is backed out, but the second is committed, a trigger message is still created. So the trigger message is created at the time the first unit of work is backed out. If there is a significant delay before the second message is committed, the triggered application may need to wait for it.

With triggering of type DEPTH, a delay can occur even if all relevant messages are eventually committed. Suppose that the *TriggerDepth* queue attribute has the value 2. When two messages arrive on the queue, the second causes a trigger message to be created. However, if the second message is the first to be committed, it is at that time the trigger message becomes available. The trigger monitor starts the server program, but the program can retrieve only the second message until the first one is committed. So the program may need to wait for the first message to be made available.

You should design your application so that it terminates if no messages are available for retrieval when your wait interval expires. If one or more messages arrive subsequently, you should rely on your application being retriggered to process them. This method prevents applications being idle, and unnecessarily using resources.

Trigger monitors

To a queue manager, a trigger monitor is like any other application that serves a queue. However, a trigger monitor serves initiation queues.

A trigger monitor is usually a continuously-running program. When a trigger message arrives on an initiation queue, the trigger monitor retrieves that message. It uses information in the message to issue a command to start the application that is to process the messages on the application queue.

The trigger monitor must pass sufficient information to the program it is starting so that the program can perform the right actions on the right application queue.

A channel initiator is an example of a special type of trigger monitor for message channel agents. In this situation however, you must use either trigger type FIRST or DEPTH.

MQSeries for MVS/ESA trigger monitors

The following trigger monitor is provided for CICS for MVS/ESA:

CKTI
You need to start one instance of CKTI for each initiation queue (see the *MQSeries for MVS/ESA System Management Guide* for information on how to do this). CKTI passes the MQTM structure of the trigger message to the program it starts by EXEC CICS START TRANSID. The started program gets this information by using the EXEC CICS RETRIEVE command. A program can use the EXEC CICS RETRIEVE command with the RTRANSID option to determine how the program was started; if the value returned is CKTI, the program was started by MQSeries for MVS/ESA. For an example of how to use CKTI, see the source code supplied for module CSQ4CVB2 in the Credit Check sample application supplied with MQSeries for MVS/ESA. See "The Credit Check sample" on page 409 for a full description.

The following trigger monitor is provided for IMS/ESA:

CSQQTRMN
You need to start one instance of CSQQTRMN for each initiation queue (see the *MQSeries for MVS/ESA System Management Guide* for information on how to do this). CSQQTRMN passes the MQTMC2 structure of the trigger message to the programs it starts.

MQSeries for AS/400 trigger monitors

The following are provided:

AMQSTRG4
This is a trigger monitor that submits an OS/400 job for the process that is to be started, but this means there is a processing overhead associated with each trigger message.

AMQSERV4
This is a trigger server. For each trigger message, this server runs the command for the process in its own job, and can call CICS transactions.

Both the trigger monitor and the trigger server pass an MQTMC structure to the programs they start. For a description of this structure, see the *MQSeries for*

AS/400 Application Programming Reference manuals. Both of these samples are delivered in both source and executable forms.

MQSeries for OS/2 Warp, Digital OpenVMS, Tandem NSK, UNIX systems, and Windows NT trigger monitors

The following trigger monitors are provided for the server environment:

amqstrg0   This is a sample trigger monitor that provides a subset of the function provided by **runmqtrm**. See Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313 for more information on amqstrg0.

runmqtrm   **runmqtrm** [-*m QMgrName*] [-*q InitQ*] is the command. The default is SYSTEM.DEFAULT.INITIATION.QUEUE on the default queue manager. It calls programs for the appropriate trigger messages. This trigger monitor supports the default application type.

The command string passed by the trigger monitor to the operating system is built as follows:

1. The *ApplId* from the relevant PROCESS definition (if created)

2. The MQTMC2 structure, enclosed in quotation marks

3. The *EnvData* from the relevant PROCESS definition (if created)

where *ApplId* is the name of the program to run - as it would be entered on the command line.

The parameter passed is the MQTMC2 character structure. A command string is invoked which has this string, exactly as provided, in 'quotation marks', in order that the system command will accept it as one parameter.

The trigger monitor will not look to see if there is another message on the initiation queue until the completion of the application it has just started. If the application has a lot of processing to do, this may mean that the trigger monitor cannot keep up with the number of trigger messages arriving. You have two options:

- Have more trigger monitors running
- Run the started applications in the background

If you choose to have more trigger monitors running you have control over the maximum number of applications that can run at any one time. If you choose to run applications in the background, there is no restriction imposed by MQSeries. on the number of applications that can run.

To run the started application in the background under OS/2, or Windows NT, within the *ApplId* field you must prefix the name of your application with a START command. For example:

```
START AMQSECHA /B
```

To run the started application in the background on UNIX systems, you must put an '&' at the end of the *EnvData* of the PROCESS definition.

The following trigger monitors are provided for the MQSeries client:

runmqtmc This is the same as **runmqtrm** except that it links with the MQSeries client libraries.

For CICS:

The following trigger monitor is provided for CICS:

amqltmc0 The CICS Trigger monitor works in the same fashion as the standard trigger monitor, **runmqtrm**, but you run it in a different way and it triggers CICS transactions.

It is supplied as a CICS program and you must define it with a 4-character transaction name. Enter the 4-character name to start the trigger monitor, when it will use the default queue manager (as named in the qm.ini file), and the SYSTEM.CICS.INITIATION.QUEUE.

If you want to use a different queue manager or queue, you must build the trigger monitor the MQTMC2 structure: this requires you to write a program using the EXEC CICS START call, because the structure is too long to add as a parameter. Then, pass the MQTMC2 structure as data to the START request for the trigger monitor.

When you use the MQTMC2 structure, you only need to supply the *StrucId*, *Version*, *QName*, and *QMgrName* parameters to the trigger monitor as it does not reference any other fields.

Messages are read from the initiation queue and used to start CICS transactions, using EXEC CICS START, assuming the APPL_TYPE in the trigger message is MQAT_CICS. The reading of messages from the initiation queue is performed under CICS syncpoint control.

Messages are generated when the monitor has started and stopped as well as when an error occurs. These messages are sent to the CSMT transient data queue.

Here are the available versions and appropriate use of the trigger monitor:

Version	Use
amqltmc0	CICS for OS/2 Version 2
	CICS for Windows NT Version 2
	Transaction Server for AIX, Version 4
amqltmc3	Transaction Server for OS/2, Version 4
amqltmc4	Transaction Server for Windows NT, Version 4

If you need a trigger monitor for other environments, you need to write a program that can process the trigger messages that the queue manager puts on the initiation queues. Such a program should:

1. Use the MQGET call to wait for a message to arrive on the initiation queue.

2. Examine the fields of the MQTM structure of the trigger message to find the name of the application to start and the environment in which it runs.

3. Issue an environment-specific start command. For example, in MVS batch, submit a job to the internal reader.

4. Convert the MQTM structure to the MQTMC2 structure if required.

5. Pass either the MQTMC2 or MQTM structure to the started application. This may contain user data.

6. Associate with your application queue the application that is to serve that queue. You do this by naming the process definition object (if created) in the *ProcessName* attribute of the queue.

   Use DEFINE QLOCAL or ALTER QLOCAL. On AS/400 you can also use CRTMQMQ or CHGMQMQ.

For more information on the trigger monitor interface, see the *MQSeries Application Programming Reference* manual:

Properties of trigger messages

The following sections describe some other properties of trigger messages.

Persistence and priority of trigger messages

Trigger messages are not persistent as there is no requirement for them to be so. The conditions for generating triggering events are persistent, hence trigger messages will be generated whenever these conditions are met. In the event that a trigger message is lost, the continued existence of the application message on the application queue will guarantee that the queue manager will generate a trigger message as soon as all the conditions are met.

If a unit of work is rolled-back, any trigger messages it generated will always be delivered.

Trigger messages take the default priority of the initiation queue.

Queue manager restart and trigger messages

Following the restart of a queue manager, when an initiation queue is next opened for input, a trigger message may be put to this initiation queue if an application queue associated with it has messages on it, and is defined for triggering.

Trigger messages and changes to object attributes

Trigger messages are created according to the values of the trigger attributes in force at the time of the trigger event. If the trigger message is not made available to a trigger monitor until later (because the message that caused it to be generated was put within a unit of work), any changes to the trigger attributes in the meantime have no effect on the trigger message. In particular, disabling triggering does not prevent a trigger message being made available once it has been created. Also, the application queue may no longer exist at the time the trigger message is made available.

Format of trigger messages

The format of a trigger message is defined by the MQTM structure. This has the following fields, which the queue manager fills when it creates the trigger message, using information in the object definitions of the application queue and of the process associated with that queue:

StrucId   The structure identifier.

Version   The version of the structure.

QName   The name of the application queue on which the trigger event occurred. When the queue manager creates a trigger message, it fills this field using the *QName* attribute of the application queue.

ProcessName The name of the process definition object that is associated with the application queue. When the queue manager creates a trigger message, it fills this field using the *ProcessName* attribute of the application queue.

TriggerData A free-format field for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the *TriggerData* attribute of the application queue.

ApplType   The type of the application that the trigger monitor is to start. When the queue manager creates a trigger message, it fills this field using the *ApplType* attribute of the process definition object identified in *ProcessName*.

ApplId   A character string that identifies the application that the trigger monitor is to start. When the queue manager creates a trigger message, it fills this field using the *ApplId* attribute of the process definition object identified in *ProcessName*. When you use an MQSeries for MVS/ESA-supplied trigger monitor (CKTI or CSQQTRMN) the *ApplId* attribute of the process definition object is a CICS or IMS transaction identifier.

EnvData   A character field containing environment-related data for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the *EnvData* attribute of the process definition object identified in *ProcessName*. The MQSeries for MVS/ESA-supplied trigger monitors (CKTI or CSQQTRMN) do not use this field, but other trigger monitors may choose to use it.

UserData   A character field containing user data for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the *UserData* attribute of the process definition object identified in *ProcessName*.

There is a full description of the trigger monitor structures in the *MQSeries Application Programming Reference* manual.

When triggering does not work

A program is not triggered if the trigger monitor cannot start the program or the queue manager cannot deliver the trigger message.

If a trigger message is created but cannot be put on the initiation queue (for example, because the queue is full or the length of the trigger message is greater

Triggering failure

than the maximum message length specified for the initiation queue), the trigger message is put instead on the dead-letter (undelivered-message) queue.

If the put operation to the dead-letter queue cannot complete successfully, the trigger message is discarded and a warning message is sent to the console (MVS) or to the system operator (AS/400), or put on the error log. Note that putting these messages to the dead-letter (undelivered-message) queue will not cause trigger events for the the DLQ on OS/400.

Putting the trigger message on the dead-letter queue may generate a trigger message for that queue. This second trigger message is discarded if it adds a message to the dead-letter queue.

If the program is triggered successfully but abends before it gets the message from the queue, use a trace utility (for example, CICS AUXTRACE if the program is running under CICS) to find out the cause of the failure.

How CKTI detects errors

If the CKTI trigger monitor in MQSeries for MVS/ESA detects an error in the structure of a trigger message, or if it cannot start a program, it puts the trigger message on the dead-letter (undelivered-message) queue. CKTI adds a dead-letter header structure (MQDLH) to the trigger message. It uses a feedback code in the *Reason* field of this structure to explain why it put the message on the dead-letter (undelivered-message) queue.

An instance of CKTI stops serving an initiation queue if it attempts to get a trigger message from the queue and finds that the attributes of the queue have changed since it last accessed that queue. The attributes could have been changed by another program, or by an operator using the commands or operations and control panels of MQSeries. CKTI produces an error message, which includes a reason code, explaining the action it has taken.

How CSQQTRMN detects errors

If the CSQQTRMN trigger monitor in MQSeries for MVS/ESA detects an error in the structure of a trigger message, or if it cannot start a program, it puts the trigger message on the dead-letter (undelivered-message) queue and sends a diagnostic message to a user specified LTERM (the default is MASTER). CSQQTRMN adds a dead-letter header structure (MQDLH) to the trigger message. It uses a feedback code in the *Reason* field of this structure to explain why it put the message on the dead-letter (undelivered-message) queue. If any other errors are detected, CSQQTRMN sends a diagnostic message to the specified LTERM, and then terminates.

How RUNMQTRM detects errors

If the RUNMQTRM trigger monitor in MQSeries for OS/2 Warp and MQSeries on UNIX systems detects an error in either the:

- Structure of a trigger message
- Application type is unsupported

or it either:

- Cannot start a program
- Detects a data-conversion error

Triggering failure

it puts the trigger message on the dead-letter (undelivered-message) queue, having added a dead-letter header structure (MQDLH) to the message. It uses a feedback code in the *Reason* field of this structure to explain why it put the message on the dead-letter (undelivered-message) queue.

Chapter 15. Applications on MQSeries for MVS/ESA

MQSeries for MVS/ESA applications can be made up from programs that run in many different environments. This means they can take advantage of the facilities available in more than one environment. This chapter explains the MQSeries facilities available to programs running in each of the supported environments.

Environment-dependent MQSeries for MVS/ESA functions

The main differences to be considered between MQSeries functions in the environments in which MQSeries for MVS/ESA runs are:

- MQSeries for MVS/ESA supplies the following trigger monitors:
  - CKTI for use in the CICS environment
  - CSQQTRMN for use in the IMS environment

  You must write your own module to start applications in other environments.

- Syncpointing using two-phase commit is supported in the CICS and IMS environments. Single-phase commit is supported in the MVS environment.

- For the batch and IMS environments, the MQI provides calls to connect programs to, and to disconnect them from, a queue manager. Programs can connect to more than one queue manager.

- A CICS system can connect to only one queue manager. This can be made to happen when CICS is initiated if the subsystem name is defined in the CICS system startup job. The MQI connect and disconnect calls are tolerated, but have no effect, in the CICS environment.

- The API-crossing exit allows a program to intervene in the processing of all MQI calls. This exit is available in the CICS environment only.

- In CICS on multiprocessor systems, some performance advantage is gained because MQI calls can be executed under multiple MVS TCBs. For more information, see the *MQSeries for MVS/ESA System Management Guide*.

These features are summarized in Table 9.

Table 9. MVS/ESA environmental features				
	CICS for MVS/ESA	**CICS/MVS**	**IMS**	**Batch/TSO**
Trigger monitor supplied	Yes	Yes	Yes	No
Two-phase commit	Yes	Yes	Yes	No
Single-phase commit	Yes	No	No	Yes
Connect/disconnect MQI calls	Tolerated	Tolerated	Yes	Yes
API-crossing exit	Yes	Yes	No	No

Program debugging facilities

MQSeries for MVS/ESA provides a trace facility that you can use to debug your programs in all environments. Additionally, in the CICS environment you can use:

- The CICS Execution Diagnostic Facility (EDF)
- The CICS Trace Control Transaction (CETR)
- The MQSeries for MVS/ESA API-crossing exit

On the MVS platform, you can use any available interactive debugging tool that is supported by the programming language you are using.

All these tools are discussed further in the *MQSeries for MVS/ESA System Management Guide*.

Syncpoint support

The synchronization of the start and end of units of work is necessary in a transaction processing environment so that transaction processing can be used safely. This is fully supported by MQSeries for MVS/ESA in the CICS and IMS environments. Full support means cooperation between resource managers so that units of work can be committed or backed out in unison, under control of CICS or IMS. Examples of resource managers are DB2, CICS File Control, IMS, and MQSeries for MVS/ESA.

MVS batch applications can use MQSeries for MVS/ESA calls to give a single-phase commit facility. This means that an application-defined set of queue operations can be committed, or backed out, without reference to other resource managers.

Recovery support

If the connection between a queue manager and a CICS or IMS system is broken during a transaction, some units of work may not be backed out successfully. However, these units of work are resolved by the queue manager (under the control of the syncpoint manager) when its connection with the CICS or IMS system is reestablished.

The MQSeries for MVS/ESA interface with the application environment

To allow applications running in different environments to send and receive messages through a message queuing network, MQSeries for MVS/ESA provides an *adapter* for each of the environments it supports. These adapters are the interface between the application programs and an MQSeries for MVS/ESA subsystem. They allow the programs to use the MQI.

The batch adapter

The *batch adapter* provides access to MQSeries for MVS/ESA resources for programs running in:

- Task (TCB) mode
- Problem or Supervisor state
- Primary address space control mode

The programs must not be in cross-memory mode.

Connections between application programs and MQSeries for MVS/ESA are at the task level. The adapter provides a single connection thread from an application task control block (TCB) to MQSeries for MVS/ESA. At any one time, a task (represented by a TCB) can maintain only one MQSeries connection. If you want your programs to maintain multiple connections, you can do this by attaching additional TCBs. However, MQSeries for MVS/ESA does not coordinate syncpointing or resource sharing between TCBs.

The adapter supports a single-phase commit protocol for changes made to resources owned by MQSeries for MVS/ESA; it does not support multiphase-commit protocols.

The CICS adapter

A CICS system can have only one connection to an MQSeries for MVS/ESA queue manager, and this connection is managed by the MQSeries for MVS/ESA *CICS adapter*. The CICS adapter provides access to MQSeries for MVS/ESA resources for CICS programs. In addition to providing access to the MQI calls, the adapter provides:

- A trigger monitor (or task initiator) program that can start programs automatically when certain trigger conditions on a queue are met. For more information, see Chapter 14, "Starting MQSeries applications using triggers" on page 197.

- An API-crossing exit that can be invoked before and after each MQI call. For more information, see "The API-crossing exit for MVS/ESA" on page 223.

- A trace facility to help you when debugging programs. For more information, see the *MQSeries for MVS/ESA System Management Guide*.

- Facilities that allow the MQI calls to be executed under multiple MVS TCBs. For more information, see the *MQSeries for MVS/ESA System Management Guide*.

The adapter supports a two-phase commit protocol for changes made to resources owned by MQSeries for MVS/ESA, with CICS acting as the syncpoint coordinator.

The CICS adapter also supplies facilities (for use by system programmers and administrators) for managing the CICS-MQSeries for MVS/ESA connection, and for collecting task and connection statistics. These facilities are described in the *MQSeries for MVS/ESA System Management Guide*.

The IMS adapter

The *IMS adapter* provides access to MQSeries for MVS/ESA resources for

- On-line message processing programs (MPPs)
- Interactive Fast Path programs (IFPs)
- Batch message processing programs (BMPs)

To use these resources, the programs must be running in task (TCB) mode and problem state; they must not be in cross-memory mode or access-register mode.

The adapter provides a connection thread from an application task control block (TCB) to MQSeries. The adapter supports a two-phase commit protocol for changes made to resources owned by MQSeries for MVS/ESA, with IMS acting as the syncpoint coordinator.

The adapter also provides a trigger monitor program that can start programs automatically when certain trigger conditions on a queue are met. For more information, see Chapter 14, "Starting MQSeries applications using triggers" on page 197.

If you are writing batch DL/I programs, follow the guidance given in this book for MVS batch programs.

Writing MVS OpenEdition applications

The batch adapter supports queue manager connections from Batch and TSO address spaces:

If we consider a Batch address space, the adapter supports connections from multiple TCBs within that address space as follows:

- Each TCB can connect to multiple queue managers via the MQCONN call (but a TCB can only have one instance of a connection to a particular queue manager at any one time).

- Multiple TCBs can connect to the same queue manager (but the queue manager handle returned on any MQCONN call is bound to the issuing TCB and cannot be used by any other TCB).

MVS OpenEdition supports two types of pthread_create call:

1. Heavyweight threads, run one per TCB, that is ATTACHed and DETACHed at thread start and end by MVS.

2. Mediumweight threads, run one per TCB, but the TCB can be one of a pool of long-running TCBs. The onus is on the application to perform all necessary application clean up, since, if it is connected to a server, the default thread termination that may be provided by the server at Task (TCB) termination, will **not** always be driven.

Lightweight threads are not supported. (If an application creates permanent threads which do their own dispatching of work requests, then the **application** is responsible for cleaning up any resources before starting the next work request.)

MQSeries for MVS/ESA supports MVS OpenEdition threads via the Batch Adapter as follows:

1. Heavyweight threads are fully supported as Batch connections.

   Each thread runs in its own TCB which is ATTACHed and DETACHed at thread start and end. Should the thread end before issuing an MQDISC call, then MQSeries for MVS/ESA performs its standard task clean up which includes committing any outstanding unit of work if the thread terminated normally, or backing it out if the thread terminated abnormally.

2. Mediumweight threads are fully supported but if the TCB is going to be reused by another thread, then the application must ensure that an MQDISC call, preceded by either MQCMIT or MQBACK, is issued prior to the next thread start. This implies that if the application has established a Program Interrupt Handler, and the application then abends, then the Interrupt Handler should issue MQCMIT and MQDISC calls before reusing the TCB for another thread.

Again, lightweight threads are not supported.

Note: Threading models do **not** support access to common MQSeries resources from multiple threads.

The API-crossing exit for MVS/ESA

Product-sensitive programming interface

An exit is a point in IBM-supplied code where you can run your own code. MQSeries for MVS/ESA provides an *API-crossing exit* that you can use to intercept calls to the MQI, and to monitor or modify the function of the MQI calls. This section describes how to use the API-crossing exit, and describes the sample exit program that is supplied with MQSeries for MVS/ESA.

> **Note**
>
> The API-crossing exit is invoked only by the CICS adapter of MQSeries for MVS/ESA. The exit program runs in the CICS address space.

Using the API-crossing exit

You could use the API-crossing exit to:

- Operate additional security checks by examining the contents of each message before and after each MQI call

- Replace the queue name supplied in the message with another queue name

- Cancel the call and either issue a return code of 0 to simulate a successful call, or another value to indicate that the call was not performed

- Monitor the use of MQI calls in an application

- Gather statistics

- Modify input parameters on specific calls

- Modify the results of specific calls

Defining the exit program

Before the exit can be used, an exit program load module must be available when the CICS adapter connects to MQSeries for MVS/ESA. The exit program is a CICS program that must be named CSQCAPX and reside in a library in the DFHRPL concatenation. CSQCAPX must be defined in the CICS system definition file (CSD), and the program must be enabled.

When CSQCAPX is loaded, a confirmation message is written to the CKQC adapter control panel or to the console. If the program cannot be loaded, a diagnostic message is displayed.

How the exit is invoked

When enabled, the API-crossing exit is invoked:

- By *all* applications that use the CICS adapter of MQSeries for MVS/ESA

- For the following MQI calls:

  - MQCLOSE
  - MQGET
  - MQINQ
  - MQOPEN
  - MQPUT
  - MQPUT1
  - MQSET

- Every time one of these MQI calls is made

- Both before *and* after a call

This means that using the API-crossing exit degrades the performance of MQSeries for MVS/ESA, so plan your use of it carefully.

The exit program can be invoked once *before* a call is executed, and once *after* the call is executed. On the before type of exit call, the exit program can modify any of the parameters on the MQI call, suppress the call completely, or allow the call to be processed. If the call is processed, the exit is invoked again after the call has completed.

Note: The exit program is not recursive. Any MQI calls made inside the exit do not invoke the exit program for a second time.

Communicating with the exit program

After it has been invoked, the exit program is passed a parameter list in the CICS communication area pointed to by a field called DFHEICAP. The CICS Exec Interface Block field EIBCALEN shows the length of this area. The structure of this communication area is defined in the CMQXPA assembler-language macro that is supplied with MQSeries for MVS/ESA :

```
*
MQXP_COPYPLIST         DSECT
                       DS  0D          Force doubleword alignment
MQXP_PXPB              DS   AL4         Pointer to exit parameter block
MQXP_PCOPYPARM         DS 11AL4         Copy of original plist
*
                       ORG  MQXP_PCOPYPARM
MQXP_PCOPYPARM1        DS   AL4         Copy of 1st parameter
MQXP_PCOPYPARM2        DS   AL4         Copy of 2nd parameter
MQXP_PCOPYPARM3        DS   AL4         Copy of 3rd parameter
MQXP_PCOPYPARM4        DS   AL4         Copy of 4th parameter
MQXP_PCOPYPARM5        DS   AL4         Copy of 5th parameter
MQXP_PCOPYPARM6        DS   AL4         Copy of 6th parameter
MQXP_PCOPYPARM7        DS   AL4         Copy of 7th parameter
MQXP_PCOPYPARM8        DS   AL4         Copy of 8th parameter
MQXP_PCOPYPARM9        DS   AL4         Copy of 9th parameter
MQXP_PCOPYPARM10       DS   AL4         Copy of 10th parameter
MQXP_PCOPYPARM11       DS   AL4         Copy of 11th parameter
*
MQXP_COPYPLIST_LENGTH  EQU  *-MQXP_PXPB
                       ORG  MQXP_PXPB
MQXP_COPYPLIST_AREA    DS   CL(MQXP_COPYPLIST_LENGTH)
*
```

Field *MQXP_PXPB* points to the exit parameter block, MQXP.

Field *MQXP_PCOPYPARM* is an array of addresses of the call parameters. For example, if the application issues an MQI call with parameters P1,P2,or P3, the communication area contains:

*P*XPB,*P*P1,*P*P2,*P*P3

where *P* denotes a pointer (address) and *X*PB is the exit parameter block.

Writing your own exit program

You can use the sample API-crossing exit program (CSQCAPX) that is supplied with MQSeries for MVS/ESA as a framework for your own program. This is described on page 227.

When writing an exit program, to find the name of an MQI call issued by an application, examine the *ExitCommand* field of the MQXP structure. To find the number of parameters on the call, examine the *ExitParmCount* field. You can use the 16-byte *ExitUserArea* field to store the address of any dynamic storage that the application obtains. This field is retained across invocations of the exit and has the same life time as a CICS task.

Your exit program can suppress execution of an MQI call by returning MQXCC_SUPPRESS_FUNCTION or MQXCC_SKIP_FUNCTION in the *ExitResponse* field. To allow the call to be executed (and the exit program to be reinvoked after the call has completed), your exit program must return MQXCC_OK.

When invoked after an MQI call, an exit program can inspect and modify the completion and reason codes set by the call.

Usage notes

Here are some general points you should bear in mind when writing your exit program:

- For performance reasons, you should write your program in assembler language. If you write it in any of the other languages supported by MQSeries for MVS/ESA, you must provide your own data definition file.

- Link-edit your program as AMODE(31) and RMODE(ANY).

- To define the exit parameter block to your program, use the assembler-language macro, CMQXPA.

- If you are using the CICS for MVS/ESA storage protection feature, your program must run in CICS execution key. That is, you must specify EXECKEY(CICS) when defining both your exit program and any programs to which it passes control. For information about CICS exit programs and the CICS storage protection facility, see the *CICS Customization Guide*.

- Your program can use all the APIs (for example, IMS, DB2, and CICS) that a CICS task-related user exit program can use. It can also use any of the MQI calls except MQCONN and MQDISC. However, any MQI calls within the exit program do not invoke the exit program a second time.

- Your program can issue EXEC CICS SYNCPOINT or EXEC CICS SYNCPOINT ROLLBACK commands. However, these commands commit or roll back **all** the updates done by the task up to the point that the exit was used, and so their use is not recommended.

- Your program must end by issuing an EXEC CICS RETURN command. It must not transfer control with an XCTL command.

- Exits are written as extensions to the MQSeries for MVS/ESA code. You must take great care that your exit does not disrupt any MQSeries for MVS/ESA programs or transactions that use the MQI. These are usually indicated with a prefix of "CSQ" or "CK".

- If CSQCAPX is defined to CICS, the CICS system will attempt to load the exit program when CICS connects to MQSeries for MVS/ESA. If this attempt is successful, message CSQC301I is sent to the CKQC panel or to the system console. If the load is unsuccessful (for example, if the load module does not exist in any of the libraries in the DFHRPL concatenation), message CSQC315 is sent to the CKQC panel or to the system console.

- Because the parameters in the communication area are addresses, the exit program must be defined as local to the CICS system (that is, not as a remote program).

The sample API-crossing exit program, CSQCAPX

The sample exit program is supplied as an assembler-language program. The source file (CSQCAPX) is supplied in the library **thlqual**.SCSQASMS (where **thlqual** is the high-level qualifier used by your installation). This source file includes pseudocode that describes the program logic.

The sample program contains initialization code and a layout that you can use when writing your own exit programs.

The sample shows how to:

- Set up the exit parameter block

- Address the call and exit parameter blocks

- Determine for which MQI call the exit is being invoked

- Determine whether the exit is being invoked before or after processing of the MQI call

- Put a message on a CICS temporary storage queue

- Use the macro DFHEIENT for dynamic storage acquisition to maintain reentrancy

- Use DFHEIBLK for the CICS exec interface control block

- Trap error conditions

- Return control to the caller

Design of the sample exit program

The sample exit program writes messages to a CICS temporary storage queue (CSQ1EXIT) to show the operation of the exit. The messages show whether the exit is being invoked before or after the MQI call. If the exit is invoked after the call, the message contains the completion code and reason code returned by the call. The sample uses named constants from the CMQXPA macro to check on the type of entry (that is, before or after the call).

The sample does not perform any monitoring function, but simply places time-stamped messages into a CICS queue indicating the type of call it is processing. This provides an indication of the performance of the MQI, as well as the proper functioning of the exit program.

Note: The sample exit program issues six EXEC CICS calls for each MQI call that is made while the program is running. If you use this exit program, MQSeries for MVS/ESA performance is degraded.

Preparing and using the sample exit

The sample exit is supplied in source form only.

- For CICS for MVS/ESA, when you update the CICS system definition (CSD) data set, the definitions you need are in the member **thlqual**.SCSQPROC(CSQ4B100).

- For CICS/MVS, the definition for the exit program is provided in PPT format in CSQCPPT in **thlqual**.SCSQMACS. This should be renamed DFHPPTMQ. The table should be assembled and then migrated (using the DFHCSDUP utility) to the CSD of your CICS system. If this is all done at installation time, you can use RDO to redefine the API exit program definitions.

> **Note:** The definitions use a suffix of MQ. If this suffix is already used in your enterprise, this must be changed before the assembly stage.

If you use the default CICS program definitions supplied, the exit program CSQCAPX is installed in a ***disabled*** state. This is because using the exit program can produce a significant reduction in performance.

To activate the API-crossing exit temporarily:

1. Issue the command CEMT S PROGRAM(CSQCAPX) ENABLED from the CICS master terminal.

2. Run the CKQC transaction, and use option 3 in the Connection pull-down to alter the status of the API-crossing exit to 'Enabled'.

If you want to run MQSeries for MVS/ESA with the API-crossing exit permanently enabled, do one of the following:

- For CICS for MVS/ESA do one of the following:

  – Alter the CSQCAPX definition in member CSQ4B100, changing STATUS(DISABLED) to STATUS(ENABLED). You can update the CICS CSD definition using the CICS-supplied batch program DFHCSDUP.

  – Alter the CSQCAPX definition in the CSQCAT1 group by changing the status from DISABLED to ENABLED.

- For CICS/MVS, alter the CSQCAPX definition in member CSQCPPT in **thlqual**.SCSQMACS changing PGMSTAT(DISABLED) to PGMSTAT(ENABLED). You can update the CICS CSD definition using the CICS-supplied batch program DFHCSDUP.

In both cases you must reinstall the group. You can do this by cold-starting your CICS system or by using the CICS CEDA transaction to reinstall the group while CICS is running.

Note: Using CEDA may cause an error if any of the entries in the group are currently in use.

└─────────── End of Product-sensitive programming interface ───────────┘

Writing MQSeries-IMS bridge applications

This section discusses writing applications to exploit the MQSeries-IMS bridge. The following topics are discussed:

- "How the MQSeries-IMS bridge deals with messages" on page 229
- "Writing your program" on page 232
- "Triggering" on page 234

For information about the MQSeries-IMS bridge, see the *MQSeries for MVS/ESA System Management Guide*.

How the MQSeries-IMS bridge deals with messages

When you use the MQSeries-IMS bridge to send messages to an IMS application, you need to construct your messages in a special format. You must also put your messages on MQSeries queues that have been defined with a storage class that specifies the XCF group and member name of the target IMS system.

The MQSeries-IMS bridge accepts the following types of message:

- Messages containing IMS transaction data and an MQIIH structure (described in the *MQSeries Application Programming Reference* manual):

```
MQIIH LLZZ<trancode><data>[LLZZ<data>][LLZZ<data>]
```

  Notes:

  1. The square brackets, [ ], represent optional multi-segments.

  2. The *Format* field of the MQMD structure must be set to MQFMT_IMS to use the MQIIH structure.

- Messages containing IMS transaction data but no MQIIH structure:

```
LLZZ<trancode><data>  \
[LLZZ<data>][LLZZ<data>]
```

MQSeries validates the message data to ensure that the sum of the LL bytes is equal to the message length after the MQIIH structure (if it is present).

When the MQSeries-IMS bridge gets messages from the OTMA queues, it processes them as follows:

- If the message contains IMS transaction data and an MQIIH structure the bridge verifies the MQIIH (see the *MQSeries Application Programming Reference* manual) and puts the message on to the appropriate IMS queue. The transaction code is specified in the input message. If this is an LTERM, IMS replies with a DFS1288E message. If the transaction code represents a command, IMS executes the command.

- If the message contains IMS transaction data, but no MQIIH structure, the IMS bridge makes the following assumptions:

  – The transaction code is in bytes 5 through 12 of the user data
  – The transaction is in non-conversational mode
  – The transaction is in commit mode 0 (commit-then-send)
  – The *Format* in the MQMD is used as the *MFSMapName* (on input)
  – The security mode is MQISS_CHECK

  The reply message is also built without an MQIIH structure, taking the *Format* for the MQMD from the *MFSMapName* of the IMS output.

The MQ-IMS bridge uses one or two tpipes for each MQ queue:

- A synchronous tpipe is used for all messages using Commit mode 0 (COMMIT_THEN_SEND) (these show with SYN in the status field of the IMS /DIS TMEMBER client TPIPE xxxx command)

- An asynchronous tpipe is used for all messages using Commit mode (SEND_THEN_COMMIT)

The tpipes are created by MQ when they are first used. An asynchronous tpipe exist until IMS is restarted. Synchronous tpipes exist until IMS is cold started.

Mapping MQSeries messages to IMS transaction types

Table 10. Mapping MQSeries messages to IMS transaction types

MQSeries message type	Commit-then-send (mode 0) - uses synchronous IMS Tpipes	Send-then-commit (mode 1) - uses asynchronous IMS Tpipes
Persistent MQSeries messages	• Recoverable full function transactions • Irrecoverable transactions are rejected by IMS	• Fastpath transactions • Conversational transactions • Full function transactions
Nonpersistent MQSeries messages	• Irrecoverable full function transactions • Recoverable transactions are rejected by IMS	• Fastpath transactions • Conversational transactions • Full function transactions

Note: IMS commands cannot use persistent MQSeries messages with commit mode 0. See the *IMS/ESA Open Transaction Manager Access User's Guide* for more information.

If the message cannot be put to the IMS queue

If the message cannot be put to the IMS queue, the following action is taken by MQSeries:

- If a message cannot be put to the IMS queue because the message is invalid, the message is put to the dead-letter queue and a message is sent to the system console.

- If the message is valid, but is rejected by IMS with a sense code of 001A and a DFS message, MQSeries puts the original message to the dead-letter queue, and puts the DFS message to the reply-to queue. If MQSeries is unable to put the DFS message to the reply-to queue, it is put to the dead-letter queue.

- If the negative acknowledgement (NAK) from IMS represents a message error, an error message is sent to the system console, and the MQSeries message is put to the dead-letter queue.

Note: In the circumstances listed above, if MQSeries is unable to put the message to the dead-letter queue for any reason, the message is returned to the originating MQSeries queue. An error message is sent to the system console, and no further messages are sent using the tpipe associated with that queue until the problem with the dead-letter queue has been resolved.

To resend the messages, do **one** of the following:

1. Stop and restart the Tpipes in IMS corresponding to the queue
2. Alter the queue to GET(DISABLED), and again to GET(ENABLED)
3. Stop and restart the IMS OTMA
4. Stop and restart your MQSeries subsystem

If the NAK received from IMS represents anything else, the MQSeries message is returned to the originating queue, MQSeries stops processing the queue, and an error message is sent to the system console.

If an exception report message is required, the bridge puts it to the reply-to queue with the authority of the originator. If the message cannot be put to the queue, the report message is put to the dead-letter queue with the authority of the bridge. If it cannot be put to the DLQ, it is discarded.

Reply messages from IMS

Reply messages from IMS are put onto the reply-to queue specified in the original message. If the message cannot be put onto the reply-to queue, it is put onto the dead-letter queue using the authority of the bridge. If the message cannot be put onto the dead-letter queue, a negative acknowledgement is sent to IMS to say that the message cannot be received. Responsibility for the message is then returned to IMS. If you are using commit mode 0, messages from that Tpipe are not sent to the bridge, and remain on the IMS queue; that is, no further messages are sent until restart. If you are using commit mode 1, other work can continue.

If the reply has an MQIIH structure, its format type is MQFMT_IMS; if not, its format type is specified by the IMS MOD name used when inserting the message.

Using alternate response PCBs: If your IMS application uses alternate response PCBs, invoking these applications through the MQSeries-IMS bridge will cause the IMS pre-routing and destination resolution exits to be called. See the *MQSeries for MVS/ESA System Management Guide* for information about these exit programs.

Message segmentation

IMS transactions may be defined as expecting single- or multi-segment input. The originating MQSeries application must construct the user input following the MQIIH structure as one or more LLZZ-data segments. All segments of an IMS message must be contained in a single MQSeries message sent with a single MQPUT.

The maximum length of any one LLZZ-data segment is defined by IMS/OTMA (32764 bytes). The total MQSeries message length is the sum of the LL bytes, plus the length of the MQIIH structure.

All the segments of the reply are contained in a single MQSeries message.

Data conversion

The MQSeries-IMS bridge converts messages to the coded character set and encoding of the local queue manager as required, using both *built-in* formats and user exit programs. This means that you can send messages to an IMS application using the MQSeries-IMS bridge from any MQSeries platform.

The conversion (including the calling of any necessary exits) is performed by the distributed queuing facility when it puts a message to a destination queue that has XCF information defined for its storage class. Any exits needed must be available to the distributed queuing facility in the data set referenced by the CSQXLIB DD statement.

Note: Messages arriving through the CICS distributed queuing facility are not converted.

If there are conversion errors, the message is put to the queue unconverted; this results eventually in it being treated as an error by the MQSeries-IMS bridge, because the bridge cannot recognize the header format. If a conversion error occurs, an error message is sent to the MVS console.

See Chapter 11, "Data-conversion exits" on page 157 for detailed information about data conversion in general.

Sending messages to the MQSeries-IMS bridge: To ensure that conversion is performed correctly, you must tell the queue manager what the format of the message is. If the message has an MQIIH structure, the *Format* in the MQMD must be set to the built-in format MQFMT_IMS, and the *Format* in the MQIIH must be set to the name of the format that describes your message data. If there is no MQIIH, set the *Format* in the MQMD to your format name.

If your data (other than the LLZZs) is all character data (MQCHAR), use as your format name (in the MQIIH or MQMD, as appropriate) the built-in format MQFMT_IMS_VAR_STRING. Otherwise, use your own format name, in which case you must also provide a data-conversion exit for your format. The exit must handle the conversion of the LLZZs in your message, in addition to the data itself (but it does not have to handle any MQIIH at the start of the message).

If you use this format, the *MFSMapName* passed to IMS is MQFMT_IMS_VAR_STRING. If your application makes use of *MFSMapName*, you are recommended to use messages with the MQFMT_IMS instead.

Receiving messages from the MQSeries-IMS bridge: If an MQIIH structure is present on the original message that you are sending to IMS, one is also present on the reply message.

To ensure your reply is converted correctly, follow these steps:

- If you have an MQIIH structure on your original message, specify the format you want for your reply message in the MQIIH *ReplytoFormat* field of the original message. This value is placed in the MQIIH *Format* field of the reply message.

- If you do not have an MQIIH structure on your original message, specify the format you want for the reply message as the MFS MOD name in the IMS application's ISRT to the IOPCB.

- Specify CONVERT(YES) on the sender channel between your MQSeries for MVS/ESA system and your destination MQSeries system.

Writing your program

The coding required to handle IMS transactions through MQSeries is platform-specific. However, there are several points to be borne in mind when your application handles IMS screen formatting information.

In IMS, your application can modify certain 3270 screen behavior, for example, highlighting a field which has had invalid data entered. This type of information is communicated by adding a two byte attribute field to the IMS message for each screen field needing to be modified by the program.

Thus, if you are coding an application to mimic a 3270, you need to take account of these fields when building or receiving messages.

You may need to code information in your program to process:

- Which key is pressed (Enter, PF1....)

- Where the cursor is when the message is passed to your application

- Whether the attribute fields have been set by the IMS application

  - High/normal/zero intensity

- Color
- Whether IMS is expecting the field back the next time enter is pressed
- Whether the IMS application has used null characters (X'3F') in any fields.

If your IMS message contains only character data (apart from the LLZZ-data segment), and you are using an MQIIH structure, set the MQMD format to MQFMT_IMS and the MQIIH format to MQFMT_IMS_VAR_STRING.

If your IMS message contains only character data (apart from the LLZZ-data segment), and you are **not** using an MQIIH structure, set the MQMD format to MQFMT_IMS_VAR_STRING and ensure that your IMS application specifies MODname MQFMT_IMS_VAR_STRING when replying.

If your IMS message contains binary, packed, or floating point data (apart from the LLZZ-data segment), you will need to code your own data-conversion routines. Refer to the *IMS/ESA Application Programming: Transaction Manager* manual for information about IMS screen formatting.

Dealing with unsolicited messages from IMS

You need to write pre-routing and destination resolution exits to handle unsolicited messages from IMS. See the *MQSeries for MVS/ESA System Management Guide* for information about these exit programs.

Unsolicited messages can create new tpipes. For example if an existing IMS transaction switched to a new LTERM (for example PRINT01) but the implementation required that the output be delivered through OTMA; a new tpipe (called PRINT01 in this example) would be created. By default this will be an asynchronous tpipe. If the implementation requires the message to be recoverable the destination resolution exit Output flag must be set, see the *IMS Customization Guide* for more information.

Writing MQSeries applications to invoke IMS conversational transactions

When you write an application which will invoke an IMS conversation, you should bear the following in mind:

- You must include an MQIIH structure with your application message.

- You must set the *CommitMode* in MQIIH to MQICM_SEND_THEN_COMMIT.

- To invoke a new conversation, set *TranState* in MQIIH to MQITS_NOT_IN_CONVERSATION.

- To invoke second and subsequent steps of a conversation, set *TranState* to MQITS_IN_CONVERSATION, and set *TranInstanceId* to the value of that field returned in the previous step of the conversation.

- There is no easy way in IMS to find the value of a *TranInstanceId*, should you lose the original message sent from IMS.

- The application must check the *TranState* of messages from IMS to check whether the IMS transaction has terminated the conversation.

- You can use /EXIT to end a conversation. You must also quote the *TranInstanceId*, set *TranState* to MQITS_IN_CONVERSATION, and use the MQSeries queue on which the conversation is being carried out.

- You cannot use /HOLD or /REL to hold or release a conversation.

• Conversations invoked through the MQSeries-IMS bridge are terminated if IMS is restarted.

Triggering

The MQSeries-IMS bridge does not support trigger messages.

If you define an initiation queue that uses a storage class with XCF parameters, messages put to that queue are rejected when they get to the bridge.

Writing IMS applications using MQSeries

This section discusses the following subjects that you should consider when using MQSeries in IMS applications:

• "Syncpoints in IMS applications"
• "MQI calls in IMS applications"

Syncpoints in IMS applications

In an IMS application, you establish a syncpoint by using IMS calls such as GU (get unique) to the IOPCB and CHKP (checkpoint). To back out all changes since the previous checkpoint, you can use the IMS ROLB (rollback) call. For more information, see the following books:

• *IMS/ESA Application Programming: Transaction Manager*
• *IMS/ESA Application Programming: Design Guide*

The queue manager is a participant in a two-phase commit protocol; the IMS syncpoint manager is the coordinator.

All open handles are closed by the IMS adapter at a syncpoint (except in a batch-oriented BMP). This is because a different user could initiate the next unit of work and MQSeries security checking is performed when the MQCONN and MQOPEN calls are made, not when the MQPUT or MQGET calls are made.

Handles are also closed after a ROLB call unless you are running IMS Version 3 or are running a batch-oriented BMP.

If an IMS application (either a BMP or an MPP) issues the MQDISC call, open queues are closed but no implicit syncpoint is taken. If the application closes down normally, any open queues are closed and an implicit commit occurs. If the application closes down abnormally, any open queues are closed and an implicit backout occurs.

MQI calls in IMS applications

This section covers the use of MQI calls in the following types of IMS applications:

• "Server applications" on page 235
• "Enquiry applications" on page 238

Server applications

Here is an outline of the MQI server application model:

```
 Initialize/Connect
 .
 Open queue for input shared
 .
 Get message from MQSeries queue
 .
 Do while Get does not fail
     .
     If expected message received
       Process the message
     Else
       Process unexpected message
     End if
     .
     Commit
     .
     Get next message from MQSeries queue
     .
 End do
 .
 Close queue/Disconnect
 .
END
```

Sample program CSQ4ICB3 shows the implementation, in C/370, of a BMP using this model. The program establishes communication with IMS first, and then with MQSeries:

```
main()
----
    Call InitIMS
    If IMS initialization successful
       Call InitMQM
       If MQSeries initialization successful
          Call ProcessRequests
          Call EndMQM
       End-if
    End-if

Return
```

The IMS initialization determines whether the program has been called as a message-driven or a batch-oriented BMP and controls MQSeries queue manager connection and queue handles accordingly:

```
InitIMS
-------
Get the IO, Alternate and Database PCBs
Set MessageOriented to true

Call ctdli to handle status codes rather than abend
If call is successful (status code is zero)
   While status code is zero
      Call ctdli to get next message from IMS message queue
      If message received
         Do nothing
         Else if no IOPBC
            Set MessageOriented to false
            Initialize error message
            Build 'Started as batch oriented BMP' message
            Call ReportCallError to output the message
         End-if
         Else if response is not 'no message available'
            Initialize error message
            Build 'GU failed' message
                Call ReportCallError to output the message
                Set return code to error
         End-if
      End-if
   End-while
Else
   Initialize error message
   Build 'INIT failed' message
   Call ReportCallError to output the message
   Set return code to error
End-if

Return to calling function
```

The MQSeries initialization performs queue manager connection opens the queues.
In a Message-driven BMP this is called after each IMS syncpoint is taken; in a
batch-oriented BMP, this is only called during program start-up:

```
InitMQM
-------
   Connect to the queue manager
   If connect is successful
      Initialize variables for the open call
      Open the request queue
      If open is not successful
         Initialize error message
         Build 'open failed' message
         Call ReportCallError to output the message
         Set return code to error
      End-if
   Else
      Initialize error message
      Build 'connect failed' message
      Call ReportCallError to output the message
      Set return code to error
   End-if

   Return to calling function
```

The implementation of the server model in an MPP is influenced by the fact that the MPP processes a single unit of work per invocation. This is because, when a syncpoint (GU) is taken, the connection and queue handles are closed and the next IMS message is delivered. This limitation can be partially overcome by one of the following:

- **Processing many messages within a single unit-of-work**

  This involves:

  - Reading a message
  - Processing the required updates
  - Putting the reply

  in a loop until all messages have been processed or until a set maximum number of messages has been processed, at which time a syncpoint is taken.

  Only certain types of application (for example, a simple database update or inquiry) can be approached in this way. Although the MQI reply messages can be put with the authority of the originator of the MQI message being handled, the security implications of any IMS resource updates need to be addressed carefully.

- **Processing one message per invocation of the MPP and ensuring multiple scheduling of the MPP to process all available messages.**

  Use the MQSeries IMS trigger monitor program (CSQQTRMN) to schedule the MPP transaction when there are messages on the MQSeries queue and no applications serving it.

  If the MPP is started by the trigger monitor the queue manager name and queue name are be passed to the program, as shown in the following COBOL code extract:

```
      * Data definition extract
       01  WS-INPUT-MSG.
           05 IN-LL1                 PIC S9(3) COMP.
           05 IN-ZZ1                 PIC S9(3) COMP.
           05 WS-STRINGPARM          PIC X(1000).
       01  TRIGGER-MESSAGE.
           COPY CMQTMC2L.
      *
      * Code extract
       GU-IOPCB SECTION.
           MOVE SPACES TO WS-STRINGPARM.
           CALL 'CBLTDLI' USING GU,
                                IOPCB,
                                WS-INPUT-MSG.
           IF IOPCB-STATUS = SPACES
               MOVE WS-STRINGPARM TO MQTMC.
      *     ELSE handle error
      *
      * Now use the queue manager and queue names passed
           DISPLAY 'MQTMC-QMGRNAME          ='
                   MQTMC-QMGRNAME OF MQTMC '='.
           DISPLAY 'MQTMC-QNAME             ='
                   MQTMC-QNAME    OF MQTMC '='.
```

The server model, which is expected to be a long running task, is better supported in a batch processing region, although the BMP cannot be triggered using CSQQTRMN.

Enquiry applications

A typical MQSeries application initiating an inquiry or update works as follows:

- Gather data from the user
- Put one or more MQSeries messages
- Get the reply messages (you might have to wait for them)
- Provide a response to the user

Because messages put on to MQSeries queues do not become available to other MQSeries applications until they are committed, they must either be put out of syncpoint, or the IMS application must be split into two transactions.

If the inquiry involves putting a single message, it is acceptable to use the "no syncpoint" option; however, if the inquiry is more complex, or resource updates are involved, you might get consistency problems if failure occurs if you don't use syncpoint.

To overcome this, IMS MPP transactions using MQI calls could be split using a program-to-program message switch; see the *IMS/ESA Application Programming: Data Communication* manual for information about this. This would allow an inquiry program to be implemented in an MPP :

```
 Initialize first program/Connect
 .
 Open queue for output
 .
 Put inquiry to MQSeries queue
 .
 Switch to second MQSeries program, passing necessary data in save
 pack area (this commits the put)
 .
END
 .
 .
 Initialize second program/Connect
 .
 Open queue for input shared
 .
 Get results of inquiry from MQSeries queue
 .
 Return results to originator
 .
END
```

Chapter 16. Lotus Notes

MQSeries Version 5 products include an add-in task that runs on a Lotus Notes server. This allows Lotus Notes applications to communicate with MQSeries applications. The add-in task recognizes and interprets:

- Requests that a Lotus Notes application wants to send to MQSeries

- Messages from an MQSeries application sent in reply, to update a Lotus Notes document

To achieve this, the add-in task uses information stored in a Lotus Notes database, called the link database.

This chapter describes:

- How MQSeries and Lotus Notes work together
- The databases needed within Lotus Notes
- The queues needed within MQSeries
- The role of the link database and an example to create an entry

Lotus Notes document and MQSeries message flow

Figure 14 on page 240 shows the sequence of events. It starts with a Lotus Notes user composing a document that includes a request for information from another system. It concludes with the add-in task updating the original Lotus Notes document with the requested information.

Lotus Notes and MQSeries
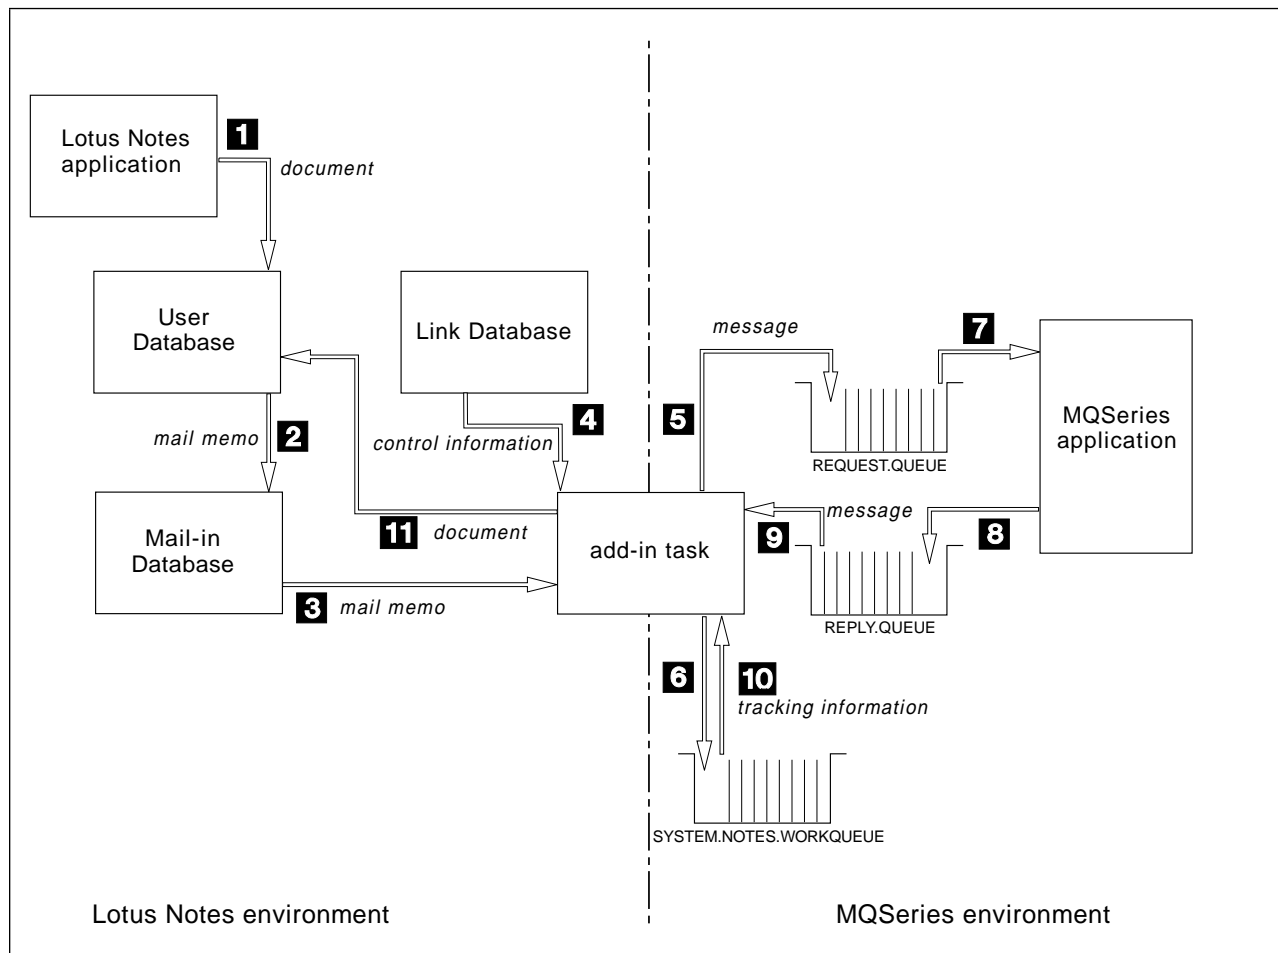
Figure 14. Lotus Notes document and MQSeries message flow

The steps are:

1. The user of the Lotus Notes application composes a document in the User database.

2. A user action, most likely the pressing of a button within the application, activates a macro that includes the @MailSend function. This creates a new mail memo, based on the information supplied in the arguments list, in the add-in task's Mail-In database.

3. The add-in task periodically checks its own Mail-In Database, and removes the mail memo as soon as it is ready to process it.

4. The add-in task, using the name of the link database entry in the mail memo, refers to the information held in the link database to determine what action is required and what queues are to be used.

5. The add-in task, from the information it has found in the link database, constructs an MQSeries message from the mail memo and places the message on the queue named in the link database entry.

6. If a reply is expected (this is only when there is a reply queue name in the link database entry), the add-in task puts a message, containing tracking and control information, on its own work queue (it uses this information to process the reply). Note that the add-in task must have sole use of this work queue.

7. The MQSeries application removes the message from the queue and processes it.

   Note: The following steps occur only if the Lotus Notes application has requested a reply.

8. The MQSeries application places a reply message on the reply queue.

9. The add-in task, after checking its Mail-In database, also checks those queues on which it is expecting reply messages (the names of which are held in the work queue).

10. The control information stored earlier on the work queue is removed and used.

11. The reply message data is used to update the original document on the Lotus Notes User database.

Design considerations

When designing your application to incorporate Lotus Notes, there are several configuration factors to consider:

- The add-in task can be installed in the MQSeries client environment or the MQI server environment. Your MQSeries application can be on any MQSeries platform that supports the environment in which the add-in task is running.

- Messages put by the add-in task are persistent.

- The add-in task is not designed to be triggered.

Lotus Notes application

Your Lotus Notes application must include a macro or formula containing the @MailSend function to create a mail memo and send it to the add-in task's Mail-In database.

This mail memo is typically created with a button in the form design, but you can use computed fields to create the mail memo.

The add-in task requires you to provide parameters within the @MailSend function. These include the name of the link database entry, the data you are sending, and any information you are expecting back from MQSeries. For example, in the sample program provided, the form design is:

```
Enter target Mail-In database name:   Mail-In
Enter data to be sent:   SData
Data received:   RData
Time MQ processed message:   Time
SendData
```

Figure 15. Lotus Notes user input sample, form named Sample

In Figure 15 on page 241, the field names are shown; these names are used when you create a link database entry. **SendData** is a button that is selected by a user after the `data to be sent` has been entered. The actions of this button are held in a macro. The macro behind the button for this sample is:

```
Button text:          SendData
Word wrapwidth:
Formula:
@If(@IsNewDoc; @Command([FileSave]); @True);

@SetField("RData"; " ");

@SetField("Time"; " ");

@MailSend(Mail_In; ""; ""; ""; "MQENTRY";
      @Text(@DocumentUniqueID) : @Char(255)   : "Msg" :
      @Char(255) : SData; [Sign] : [DeliveryReportConfirmed]:
[ReturnReceipt]);

@Command([FileCloseWindow])
```

Figure 16. Macro behind the SendData button

If the button is pressed while composing or editing a document in the User database, the macro that issues the `@MailSend` executes.

The *sendto* parameter of the @MailSend function, the first parameter, specifies the target Mail-In database name associated with the add-in task.

The *remarks* parameter of the @MailSend function, the fifth parameter, specifies the link database entry associated with the form.

The *bodyfield1* must contain @Text(@DocumentUniqueID).

The *bodyfield2* must contain @Char(255). This indicates the start of the request you are sending.

The *bodyfield3* must contain the first piece of data you want sent in the message to your MQSeries application. It can be the name of a field in your notes document, or a literal string. In this example, it is the literal `Msg`.

The *bodyfield4* must contain @Char(255). This indicates the end of the first field in the message and the start of the second field of the request you are sending.

The bodyfield5 must contain the second field you want sent in the message to your MQSeries application. In this example, it is the data that has been input from the screen by the user, and known to the application as `SData`.

In this example, there are no more data fields to go in the message to the MQSeries application, so the bodyfields are followed by the flags. These are explained in the Lotus Notes documentation.

For `@Text(@DocumentUniqueID)` to work correctly, the document you are composing must be saved on the user database before the mail memo is generated by the `@MailSend` command.

Notes:

1. In the event of an error, add the field name `error_field_msg` to the form if you want the error message displayed. You must set the field datatype to `Text` and select the **computed for display** option.

2. The Lotus script call, `chr$(255)`, creates a platform-specific character, whereas the macro `CHAR(255)`, creates a character in code page 850. For details of code page translation between diverse platforms, see the description of the Chr function and the @Char Function in the appropriate Lotus Notes documentation.

Link database

This must have an entry for each request type that a Lotus Notes application can send to an MQSeries application.

The add-in task extracts the name of the link database entry it is to use from the mail memo. The add-in task uses the information in the link database entry to:

- Manage the mail memos it gets from its own Mail-In database
- Construct an MQSeries message from the mail memo
- Update the original Lotus Notes document with information in the MQSeries reply message
- Handle all error conditions that result from interaction with the MQSeries application
- Notify the Lotus Notes application of any errors on completion of its task
- Notify the Lotus Notes user of any error messages by sending a note

Link database format

Figure 17 shows a sample definition in the link database:

```
Entry:              MQENTRY
Database Information
    User Database Name  :  AMQSAMPL
    Form to update      :  Sample
Request Offsets
Syntax:
'FieldName <space> Start  <space> End  <space>  CHAR | S390-BINARY | INTEL-BINARY'
              function    0 2   CHAR
              SData       3 104 CHAR
Reply Offsets
Syntax:
'FieldName <space> Start <space> End  <space>  CHAR | S390-BINARY | INTEL-BINARY'
              Rdata        0 101 CHAR
              Time       102 127 CHAR
Message Queuing Parameters
  Request Queue Name   :  SYSTEM.SAMPLE.NOTES.INQUEUE
  Reply Queue Name     :  SYSTEM.SAMPLE.NOTES.OUTQUEUE
  Message Format Field :  MQSTR
Error Handling
Syntax:
'Error Condition <space> start-offset <space> end-offset <space> CHAR'
client 0 5 CHAR
In the event of an error you may provide additional information:
Syntax:
'Fieldname <space> start-offset <space> end-offset<space>  CHAR'
RData 0 101 CHAR
```

Figure 17. Example of a link database entry

Each document must follow the structure given in this section. The most important
aspects to get right are the start and end positions of the fields you are interested
in. Remember that the first character is in position zero.

1. Entry

   The Entry Name given in the *remark* field of the @MailSend request enables
   MQSeries to identify the correct document to use for translation of a particular
   message format from, or to, Lotus Notes format. For example:

   Entry: MQENTRY

2. `Database Information`

   This information identifies the Lotus Notes database associated with this document.

   `User Database Name`: The name of the database that contains the document to be updated

   `Form to update`: The name of the form that was used to compose the original document

   For example:

   ```
   User Database Name      :  AMQSAMPL
   Form to update          :  Sample
   ```

3. `Request Offsets`

   This describes each field you are sending to an MQSeries application. The name you give to each field is of no significance, however you are recommended to use the field names used on your form, for clarity and ease of maintenance. The order that you list the fields *must* match the order in which they appear in the @MailSend parameter list. Each field must be described using the following format:

   `Syntax:`

   `FieldName<space>Start<space>End<space>CHAR | S390-BINARY | INTEL-BINARY`

   This is the layout of the MQSeries request message.

   For example:

   ```
   function 0 2 CHAR
   Msg      3 104 CHAR
   ```

4. `Reply Offsets`

   This describes each field to be updated with the data received from an MQSeries application. The name you give to each field *must* be the same as the field name used in your form.

   `Syntax:`

   `FieldName<space>Start<space>End<space>CHAR | S390-BINARY | INTEL-BINARY`

   This is the layout of the MQSeries reply message to the add-in task.

   For example:

   ```
   RData      0 101 CHAR
   Time     102 127 CHAR
   ```

5. `Message Queuing Parameters`

These parameters contain the following information: `Request Queue Name`:
Outgoing Message queue `Reply Queue Name`: Incoming Message queue
(Optional) `Message Format Field`: Format of the MQSeries request   message
(Optional)

If you do not require MQSeries to send a reply message, leave the Reply
Queue Name blank. If you do include the name of a reply queue, it must be
the name of a local queue.

For example:

```
Request Queue Name   :  SYSTEM.SAMPLE.NOTES.INQUEUE
Reply Queue Name     :  SYSTEM.SAMPLE.NOTES.OUTQUEUE
Message Format Field :  MQSTR
```

6. `Error Handling`

This allows you to check a field in the MQSeries reply message before the
Lotus Notes document is updated. In the event that the error condition is met,
no update occurs and a Lotus Notes mail message is sent to the Lotus Notes
user containing the error message, with any additional information that has
been requested.

Note: The datatype must always be CHAR.

- To set the error condition that you want to check against:

  Syntax:
  `'Error Condition Value<space>start-offset<space>end-offset<space>CHAR`

  Enter the condition that you want to check for. A Lotus Notes mail
  message is generated if the portion of the reply message you specify in the
  error handling condition matches the character string value.

  For example:

  ```
  client 0 5 CHAR
  ```

  In this example, an error is detected if the first six characters of the data
  sent from MQSeries are `client`, in which case a Lotus Notes mail message
  is generated.

  If you want to check for a 'not equal' value, prefix the character string with
  an **!** (exclamation mark).

  Note: The comparison *is* case sensitive.

- To add more information to the Lotus Notes mail message:

  Syntax: `'Fieldname <space>start-offset<space>end-offset<space> CHAR`

  For example:

      Rdata 0 101 CHAR

  If you want information to be copied from the MQSeries message into the Lotus Notes mail message when the error is detected, you add the field descriptions here. The name you give to each field is of no significance.

  This example adds the data in RData (102 bytes long) to the Lotus Notes mail message, but only when the first 6 bytes of the reply data match the characters `client`.

MQSeries application

Your MQSeries application is a standard messaging application that gets messages from specific queues. It responds by putting messages to other named queues if a reply is required.

The *MsgId* of the request message must be copied to the *CorrelId* of the reply message, to link the reply to the correct Lotus Notes document.

Your application should examine the *MsgType* field in the message descriptor of the request message, to determine whether a reply message is required. If the type is MQMT_REQUEST, your application must send a reply message to the Queue Manager and Queue name given in the *ReplyToQMgr* and *ReplyToQ* fields from the message descriptor. If the type is MQMT_DATAGRAM, your application must not send a reply message.

Error handling

All MQSeries activity is recorded in the Lotus Notes Servers Note log. If there is an error with the @MailSend execution, Lotus Notes reports the error. If the @MailSend is successful, MQSeries attempts to process the request.

An error detected by the add-in task causes a message to be generated, a copy of which is placed on the log files used by both Lotus Notes and MQSeries. The Lotus Notes user receives a mail message with details of the error, unless there is a problem communicating with the Lotus Notes user, in which case the message is reported to the log files only.

If no error is encountered during processing of 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 Notes electronic mail.

Note: The add-in task also attempts to update an optional field, error_field_msg, on the form where updates are to take place. If error_field_msg is required on the form, it should be entered as data type 'TEXT' and flagged as 'Computed for Display'. If an error occurs, MQSeries updates this field with the text:

    An error was encountered by the add-in task.
    Check the mail for details.

Programming considerations

There are a few aspects that are specific to using the add-in task provided:

- The add-in task has no knowledge of the applications that are using it. It is totally dependent on the control information held in the link database to tell it what it has to do with a document or message it has retrieved. The add-in task looks for the separating characters @char(255) on the @MailSend command to identify each piece of data it has to reformat.

- If your form includes a field that you have given the data type of number, it must be converted to a text format using the @text function within the @MailSend function. If you do not do this, the message cannot be converted by the add-in task.

- If data conversion is needed by the Lotus Notes document, you need to write a data-conversion exit and add it to the MQSeries environment where both the add-in task and MQSeries application are running. Enter the name of the exit against the 'message field format' prompt in the link database of the 'document entry name' used by the document.

- The add-in task always uses the MQGMO_CONVERT option to get the reply message. The conversion is done as part of the MQGET.

- S390 BINARY fields must be 2 bytes in length.

Data conversion

If your MQSeries application and your Lotus Notes Server are running on platforms that have different character sets or different integer encodings, the data from your Lotus Notes application has to be converted. This can be performed by your application or by MQSeries on behalf of your application.

These are the steps that take place:

1. All field data is passed from the Lotus Notes application to the add-in task as ASCII character strings, irrespective of the representation of that data in the Lotus Notes database itself. These strings are concatenated into a single string by the @MailSend command.

2. The add-in task builds the MQSeries request message using the information in Request Offsets from the link database entry.

   Entries in the Request Offsets fields in the link database are used to control the type of conversion when the add-in task builds the MQSeries message. The format options are:

CHAR	The field is copied across unchanged and so appears in the code page used by the Notes Server.
INTEL-BINARY	The field is encoded as a 2-byte Intel-format integer before it is inserted into the MQSeries message.
S390-BINARY	The field is encoded as a 2-byte Intel-format integer. The order of the bytes is reversed before the field is inserted into the MQSeries message.

3. When the add-in task puts the message on the request queue, it sets the message descriptor format field to the value given in the message format field of the link database entry. The possible values of the message format are:

   - Blank

     Use this option when your application performs any necessary data conversion. Your application must check against the *Encoding* and *CodedCharSetId* fields in the message descriptor.

   - MQSTR

     Use this option when the fields in the message are all of type CHAR. This indicates to MQSeries that the application data is all characters and MQSeries does any necessary data conversion. This is not recommended if you are converting between single-byte and double-byte character sets, as individual fields may expand or contract.

   - User-defined format

     Use this option when your application sends a mixture of character and binary data. You must provide a data conversion exit and associated format name if MQSeries is to do the data conversion.

4. If the channel between your Lotus Notes server and the MQSeries queue manager has an attribute of *convert* set to 'yes', and the format field supplied in the link database was non-blank, MQSeries is converts the data to the receiving platform's character set and encoding before transmitting the data. Therefore, if you define your own formats, you must install your data-conversion exits on the Lotus Notes server machine.

5. If the channel has an attribute of *convert* set to 'no', the data is transmitted unconverted. The format name is copied to the message descriptor of the message. The receiving MQSeries application has to retrieve the information using the MQGMO_CONVERT option. If you are not using a format supplied by MQSeries, you must install your own exit in this environment.

6. If the MQSeries application sends a reply message it must include a *format* field in the message descriptor of the reply.

   The options are :

   - Blank

     No data conversion is to be performed by MQSeries. The application must put the data into the message in the precise formats (ASCII character, S390 binary, Intel binary) specified in the link database entry. You need to know the character set and codepage to be used to represent the ASCII data.

   - MQSTR

     Use this when the fields in the message are all of type CHAR. This indicates to MQSeries that the application data is all characters, so that MQSeries can convert it to the appropriate codepage for the Lotus Notes server environment. This is not recommended if you are converting between single-byte and double-byte character sets, as individual fields may expand or contract.

- User-defined format

  If your application sends a mixture of character and binary data, you must provide a data-conversion exit and associated format name if MQSeries is to do the data conversion.

7. The add-in task retrieves the reply with an option of MQGMO_CONVERT. Therefore, the data-conversion user exit must be installed if a user-defined format has been used.

Part 3. Building an MQSeries application

Chapter 17. Building your application on AIX

The AIX publications describe how to build executable applications from the programs you write. This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for AIX applications to run under AIX. C, C++, and COBOL are supported. For information about preparing your C++ programs, see the *MQSeries Using C++* book.

The tasks you must perform to create an executable application using MQSeries for AIX vary with the programming language your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the MQSeries for AIX include files for the language you are using. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

Preparing C programs

Precompiled C programs are supplied in the /**mqmtop**/samp/bin directory. Use the ANSI compiler and run the following command:

```
$ cc -o <amqsput> <amqsput>.c -lmqm
```

where `amqsput` is a sample program.

If you want to use the programs on a machine which has only the MQSeries client for AIX installed, recompile the programs to link them with the client library (`-lmqic`) instead.

Linking libraries

The following lists the libraries you will need:

• If your application is running in a DCE client environment you will need to copy the DCE library, libxdsom.a, on to your machine.

• You need to link your programs with the appropriate library provided by MQSeries.

In a non-threaded environment you must link to one of the following libraries:

Library file	Program/exit type
libmqm.a	server for C
libmqic.a	client for C

In a threaded environment, you must link to one of the following libraries:

Library file	Program/exit type
libmqm_r.a	server for C
libmqic_r.a	client for C

Notes:

1. If you are writing an installable service (see the *MQSeries Programmable System Management* for further information), you need to link to the libmqmzf.a library in a non-threaded application and to the libmqmzf_r.a library in a threaded application.

2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or

Novell Tuxedo, you need to link to the libmqmxa.a library in a non-threaded application and to the libmqmxa_r.a library in a threaded application.

3. You need to link trusted applications to the threaded MQSeries libraries. However, only one thread in an MQSeries on UNIX systems trusted application can be connected at a time.

4. To run the sample Encina program, link against the following libraries:

   – libmqmxa_r.a
   – libmqm_r.a

   Also, link to the Encina and DCE libraries:

   – libEncServer.a
   – libEncina.a
   – libdce.a

   The sample must be compiled and linked using xlc_r4.

5. You must link MQSeries libraries before any other product libraries (in this case, DCE and Encina). For example:

   ```
   cc -o put put.c -lmqm_r -ldce
   ```

   This ensures that any operating system functions that have been redefined by DCE are also used by MQSeries.

Preparing COBOL programs

You need to link your program with one of the following:

libmqmcb.a   server for COBOL
libmqicb.a   client for COBOL
libmqmcb_r.a server for COBOL (in a threaded application)

You can use the IBM COBOL Set compiler or Micro Focus COBOL compiler depending on the program:

- Programs beginning amqi are suitable for the IBM COBOL Set compiler,

- programs beginning amqm are suitable for the Micro Focus COBOL compiler,

and

- programs beginning amq0 are suitable for either compiler.

Preparing COBOL programs using IBM COBOL SET for AIX

Sample COBOL programs are supplied with MQSeries. To compile such a program, enter:

```
cob2 -o amq0put0 amq0put0.cbl
-L/usr/lpp/mqm/lib
-lmqmcb -qLIB
-I/usr/lpp/mqm/inc
```

Note: For threaded applications, cob2 is used with the libmqmcb_r.a library.

Preparing COBOL programs using Micro Focus COBOL

Set environment variables before compiling your program as follows:

```
export COBCPY=/usr/lpp/mqm/inc
export LIB=/usr/lpp/mqm/lib;$LIB
```

To compile a COBOL program using Micro Focus COBOL, enter:

```
  cob -xvP amq0put0.cbl -lmqmcb
```

See the Micro Focus COBOL documentation for a description of the environment variables that need to be set up.

Preparing PL/I programs

Sample PL/I programs are supplied with MQSeries. PL/I include files are also provided so that the C entry points in the MQSeries libraries can be invoked directly.

To prepare a PL/I program:

1. Link your program with one of the libraries listed in "Linking libraries" on page 253.

2. Compile your program:

   ```
   pli amqpput0.pli -I/usr/lpp/mqm/inc /usr/lpp/mqm/lib/libmqm.a
   ```

Preparing CICS programs

XA switch modules are provided to enable you to link CICS with MQSeries:

Table 11. Essential Code for CICS applications (AIX)

Description	C (source)	C (exec) - add to your XAD.Stanza
XA initialization routine	amqzscix.c	amqzsc21 - CICS for AIX

Always link your transactions with the thread safe MQSeries library libmqm_r.a.

Note: On AIX Version 4, the libmqm_r.a library works with both native and DCE libraries.

You can find more information about supporting CICS transactions in the *MQSeries System Administration* book.

CICS on Open Systems support

MQSeries on UNIX systems support CICS on Open Systems via the XA interface.

You must ensure that CICS COBOL applications are linked to the threaded version of the library. CICS on Open Systems MQSeries transactions must link with libmqm_r, except on Sun Solaris, where you must link with lmqmcs_d.

You can run CICS programs using IBM COBOL SET for AIX or Micro Focus COBOL. The following sections describe the difference between these.

Preparing CICS COBOL programs using IBM COBOL SET for AIX

To use IBM COBOL, follow these steps:

1. Export the following environment variable:

```
export LDFLAGS="-qLIB -bI:/usr/lpp/cics/lib/cicsprIBMCOB.exp  \
                -I/usr/lpp/mqm/inc -I/usr/lpp/cics/include  \
                -e _iwz_cobol_main  \
                -L/usr/lib/dce -lmqmcb_r -ldcelibc_r -ldcepthreads"
```

   where `LIB` is a compiler directive.

2. Translate, compile, and link the program by typing:

```
cicstcl -l IBMCOB <yourprog>.ccp
```

Preparing CICS COBOL programs using Micro Focus COBOL

To use Micro Focus COBOL, follow these steps:

1. Add the MQSeries COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -L/usr/lpp/mqm/lib  \
             /usr/lpp/mqm/lib/ libmqmcbrt.0 -lmqm_r
```

   This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call MQSeries on UNIX systems.

   Note: `cicsmkcobol` must be run only when one of the following is installed:

   new version or release of Micro Focus COBOL
   new version or release of CICS for AIX
   new version or release of any supported database product (for COBOL transactions only)
   CICS for AIX

2. Export the following environment variables:

```
COBCPY=/usr/lpp/mqm/inc  export COBCPY
LDFLAGS="-L/usr/lpp/mqm/lib -lmqm_r"  export LDFLAGS
```

3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e <yourprog>.ccp
```

Preparing CICS C programs

You build CICS C programs using the standard CICS facilities:

1. Export *one* of the following environment variables:

   - LDFLAGS = "-L/usr/lpp/mqm/lib -lmqm_r" export LDFLAGS
   - USERLIB = "-L/usr/lpp/mqm/lib -lmqm_r" export USERLIB

2. Translate, compile, and link the program by typing:

```
cicstcl -l C amqscic0.ccs
```

CICS C sample transaction: Sample C source for a CICS MQSeries transaction is provided by AMQSCIC0.CCS.The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them onto the local queue whose name is contained in the transmission header of the message. Any failures will be sent to the queue SYSTEM.SAMPLE.CICS.DLQ. The sample MQSC script AMQSCIC0.TST may be used to create these queues and sample input queues.

Chapter 18. Building your application on AS/400

The AS/400 publications describe how to build executable applications from the programs you write. This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for AS/400 applications to run on AS/400 systems. COBOL, C, C++, and RPG programming languages are supported.

The tasks you must perform to create an executable MQSeries for AS/400 application depend on the programming language the source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the MQSeries for AS/400 data definition files for the language you are using. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

Preparing C programs

To compile a C program, you can use the OS/400 command, CRTCMOD. Make sure that the library containing the data definition files (QMQM) is in the library list when you perform the compilation.

You must then bind the output of the compiler with the service program AMQZSTUB using the CRTPGM command. An example of the command is:

```
CRTPGM PGM(pgmname) MODULE(pgmname)
BNDSRVPGM(QMQM/AMQZSTUB)
```

where *pgmname* is the name of your program.

Preparing C++ programs

If you are using MQSeries for AS/400 V4R2 you can write your applications in C++. To see how to link your C++ applications and for full details of all aspects of using C++ see the book *MQSeries Using C++*.

Preparing COBOL programs

To compile a COBOL program in the OPM environment, use the OS/400 command, CRTCBLPGM. To compile an ILE COBOL program, use the OS/400 command, CRTBNDCBL, or the two commands CRTCBLMOD, and CRTPGM.

See the *ILE COBOL/400 Programmer's Guide* and *ILE COBOL/400 Reference* for further information.

To make the compiler accept the single quotation mark character (') as the string delimiter in the COBOL copy files, use the compiler option *APOST. Make sure that the library containing the data definition files (QMQM) is in the library list when you perform the compilation.

Note: For more information on compiling a CICS application, see the *CICS for AS/400 Application Programming Guide*, SC33-1386.

When you run a compiled COBOL MQSeries for AS/400 program, make sure the library QMQM is in the library list.

Preparing RPG programs

If you are using MQSeries for AS/400 you can write your applications in RPG. For more information see "Coding in RPG" on page 82, and refer to the *MQSeries for AS/400 Application Programming Reference (RPG)* manual.

AS/400 programming considerations

In MQSeries for AS/400 V4R2 the system behaviour for activation group termination has been altered. In previous releases, calling an entry point (by function pointer or bound call) in an activation group which had been terminated would normally cause a MCH3402 error. In MQSeries for AS/400 V4R2 the system now terminates the process, unless the activation group was explicitly added to the program stack.

If you issue an ENDMQM *IMMED whilst a program is running which is bound to the AMQZSTUB and/or AMQVSTUB service programs, this change causes the following behaviour:

- If the program has been called from a QCMD prompt your program is terminated and you are logged off the system.

- If the program has been called from within an MQSeries panel, for example WRKMQMCHL, the program is terminated but you are not logged off the system.

QMQM activation group

When creating your program on the AS/400 the QMQM activation group should not be used. The QMQM activation group is for the use of MQSeries only.

Chapter 19. Building your application on AT&T GIS UNIX

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for AT&T GIS UNIX applications to run under AT&T GIS UNIX [5]. C and C++ programming languages are supported.

In addition to coding the MQI calls in your source code, you must add the appropriate include files. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

> **Note**
>
> Throughout this chapter, the name **mqmtop** has been used to represent the name of the installation directory. For MQSeries for AT&T GIS UNIX the name of the actual directory is /opt/mqm.

Preparing C programs

Precompiled C programs are found in the /**mqmtop**/samp/bin directory. To build a sample from source code, use the C compiler in /bin/cc, for example:

```
$ /bin/cc -o <yourprog> <yourprog>.c -lmqm -lmqmcs -lmqmzse   \
 -lnet -lnsl -lsocket -ldl
```

Note: The backslash (\) represents the continuation of the line.

C compiler flags

The order of the libraries specified is important. The following is an example of how to build the sample program amqsput0:

```
/bin/cc -o <amqsput0> <amqsput0>.c -lmqic -lmqmcs -lmqmzse \
-lnet -lnsl -lsocket -ldl
```

This links with the client library -lmqic, so allows you to use the programs on a machine which has only the MQSeries client for GIS installed.

If you use the other version of the compiler (/usr/ucb/cc), your application may compile and link successfully. However when you run it, it will fail when it attempts to connect to the queue manager.

---

[5] This platform has become NCR UNIX SVR4 MP-RAS, R3.0.

Linking libraries

You need to link your programs with the appropriate library provided by MQSeries.

You must link to one or more of the following libraries:

Library file	Program/exit type
libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C

Notes:

1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), you need to link to the libmqmzf.so library.

2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo, you need to link to the libmqmxa.a library.

Chapter 20. Building your application on Digital OpenVMS

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for Digital OpenVMS applications to run under Digital OpenVMS. C and COBOL are supported.

In addition to coding the MQI calls in your source code, you must add the appropriate include files. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

Preparing C programs

This section explains the compiler and libraries you need to prepare your C programs.

C compiler version

You must use the DEC C compiler. To invoke the compiler, enter:

```
$ CC/DECC
```

This is the default.

C compiler flags

The include files for MQSeries for Digital OpenVMS are located in the MQS_INCLUDE directory. The following is an example of how to build the sample program AMQSPUT0:

```
$ CC/INCLUDE_DIRECTORY=MQS_INCLUDE  AMQSPUT0
$ LINK AMQSPUT0.OBJ,SYS$INPUT/OPTIONS
SYS$SHARE: MQM/SHAREABLE
Ctrl + Z
```

Linking libraries

You need to link your programs with the appropriate library provided by MQSeries. The libraries are found in SYS$SHARE.

You must link to one or more of the following libraries:

Library file	Program/exit type
mqm.exe	server for C
mqic.exe	client for C
mqmzf.exe	installable service exits for C

Preparing COBOL programs

This section explains the compiler and libraries you need to prepare your COBOL programs.

COBOL compiler flags

You must compile the programs in ANSI mode using the /ANSI switch to the DEC
COBOL compiler. The following is an example of how to build the sample program
AMQ0PUT0:

```
$ COBOL/ANSI  AMQ0PUT0.COB
$ LINK AMQ0PUT0.OBJ,SYS$INPUT/OPTIONS
SYS$SHARE: MQMCB/SHAREABLE
Ctrl + Z
```

Linking libraries

You need to link your program with one of the following:

MQMCB.EXE   COBOL
MQICB.EXE   COBOL MQSeries client

Chapter 21. Building your application on HP-UX

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for HP-UX applications to run under HP-UX. C, C++, and COBOL are supported. For information about preparing your C++ programs, see the *MQSeries Using C++* book.

The tasks you must perform to create an executable application using MQSeries for HP-UX vary with the programming language your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the MQSeries for HP-UX include files for the language you are using. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

> **Note**
>
> Throughout this chapter, the name **mqmtop** has been used to represent the name of the installation directory. For MQSeries for HP-UX the name of the actual directory is /opt/mqm

Preparing C programs

Work in your normal environment. Precompiled C programs are supplied in the /**mqmtop**/samp/bin directory. The following is an example of how to build the sample program amqsput0:

```
cc -Aa -D_HPUX_SOURCE -o amqsput0 amqsput0.c -lmqm
```

If you want to use the programs on a machine which has only the MQSeries client for HP-UX installed, recompile the programs to link them with the client library instead:

```
cc -Aa -D_HPUX_SOURCE -o amqsput0 amqsput0.c -lmqic
```

Linking libraries

The following lists the libraries you will need.

- You need to link your programs with the appropriate library provided by MQSeries.

  In a non-threaded environment, you must link to one or more of the following libraries:

Library file	Program/exit type
libmqm.sl	server for C
libmqic.sl	client for C

  In a threaded environment, you must link to one or more of the following:

Library file	Program/exit type
libmqm_r.sl	server for C
libmqic_r.sl	client for C

Notes:

1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), you need to link to the libmqmzf.sl library.

2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo, you need to link to the libmqmxa.a library in a non-threaded application and to the libmqmxa_r.a library in a threaded application.

3. You must link MQSeries libraries before any other product libraries (in this case, DCE and Encina). This ensures that any operating system functions that have been redefined by DCE are also used by MQSeries.

Preparing COBOL programs

Compile the programs using the Micro Focus compiler. The copy files which declare the structures are in /**mqmtop**/inc:

```
$ export LIB=/usr/mqmtop/lib;$LIB
$ export COBCPY="/mqmtop/inc"
$ cob -vxP <amqsput>.cbl -lmqmcb
```

where amqsput is a sample program.

You must ensure that you have specified adequate runtime stack sizes; 16K is the recommended minimum.

You need to link your program with one of the following:

libmqmcb.sl   server for COBOL
libmqicb.sl   client for COBOL

Programs to run in the MQSeries client environment

If you are using LU 6.2 to connect your MQI client to a server, you must link your application to libsna.a, part of the SNAplusAPI product. Use the −lV3 and −lstr options on your compile and link command.

- The −lV3 option gives your program access to the AT&T signaling library (the SNAPlusAPI uses AT&T signals)
- The −lstr option links your program to the streams component

Note: The −lstr option is not needed on HP–UX Version 10.

If you are not using LU 6.2, consider linking to libsnastubs.a (in /**mqmtop**/lib) to fully resolve function names. The need to link to this library varies with how you are using the -B flag during the linking stage.

Preparing CICS programs

To build the sample CICS transaction, amqscic0.ccs, run the following command:

```
$ export USERLIB="-lmqm_r"
$ cicstcl -l C amqscic0.ccs
```

An XA switch module is provided to enable you to link CICS with MQSeries:

Table 12. Essential Code for CICS applications (HP-UX)		
Description	**C (source)**	**C (exec)**
XA initialization routine	amqzscix.c	amqzsc

You can find more information about supporting CICS transactions in the *MQSeries System Administration* book.

CICS on Open Systems support

MQSeries on UNIX systems supports CICS on Open Systems via the XA interface.

It is very important to ensure that CICS COBOL applications are linked to the threaded version of the library. CICS on Open Systems MQSeries transactions must link with libmqm_r, except on Sun Solaris, where you must link with lmqmcs_d.

CICS C sample transaction: Sample C source for a CICS MQSeries transaction is provided by AMQSCIC0.CCS. The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them onto the local queue whose name is contained in the transmission header of the message. Any failures will be sent to the queue SYSTEM.SAMPLE.CICS.DLQ. The sample MQSC script AMQSCIC0.TST may be used to create these queues and sample input queues.

Preparing CICS programs

Chapter 22. Building your application on MVS/ESA

The CICS, IMS, and MVS publications describe how to build applications that run in these environments. This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for MVS/ESA applications for these environments. COBOL, C, Assembler, and PL/I programming languages are supported.

The tasks you must perform to create an executable MQSeries for MVS/ESA application depend on both the programming language the program is written in, and the environment in which the application will run.

In addition to coding the MQI calls in your program, you must add the appropriate language statements to include the MQSeries for MVS/ESA data definitions file for the language you are using. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

> **Note**
>
> The name **thlqual** is the high-level qualifier of the installation library on MVS/ESA.

Preparing your program to run

After you have written the program for your MQSeries application, to create an executable application you have to compile or assemble it, then link-edit the resulting object code with the stub program that MQSeries for MVS/ESA supplies for each environment it supports. How you prepare your program depends on both the environment (batch, CICS, or IMS(BMP or MPP)) in which the application will run, and the structure of the data sets on your MVS installation. The details are described in the following sections.

"Dynamically calling the MQSeries stub" on page 270 describes an alternative method of making MQI calls in your programs so that you do not need to link-edit an MQSeries stub. This method is not available for all languages and environments.

Building MVS batch applications

To build an MQSeries for MVS/ESA application that runs under MVS batch, create job control language (JCL) that performs these tasks:

1. Compile (or assemble) the program to produce object code. The JCL for your compilation must include SYSLIB statements that make the product data definition files available to the compiler. The data definitions are supplied in the following MQSeries for MVS/ESA libraries:

   > For COBOL, **thlqual**.SCSQCOBC
   > For assembler language, **thlqual**.SCSQMACS
   > For C, **thlqual**.SCSQC370
   > For PL/I, **thlqual**.SCSQPLIC

2. For a C application, prelink the object code created in step 1.

Preparing your programs.

3. Link-edit the object code created in step 1 on page 267 (or step 2 on page 267 for a C/370 application) to produce a load module. When you do this, you must include the MQSeries for MVS/ESA batch stub program (CSQBSTUB). Figure 18 shows fragments of JCL to do this. The stub is language independent and is supplied in library **thlqual**.SCSQLOAD.

4. Store the load module in an application load library.

```
    ⋮
//*
//* MQSERIES FOR MVS/ESA LIBRARY CONTAINING BATCH STUB
//*
//CSQSTUB   DD  DSN=++HLQ.MQM100++.SCSQLOAD,DISP=SHR
//*
    ⋮
//SYSIN     DD *
   INCLUDE CSQSTUB(CSQBSTUB)
    ⋮
/*
```

Figure 18. Fragments of JCL to link-edit the object module in the batch environment

Building CICS applications

To build an MQSeries for MVS/ESA application that runs under CICS, you must:

- Translate the CICS commands in your program into the language in which the rest of your program is written

- Compile or assemble the output from the translator to produce object code

- Link-edit the object code to create a load module

CICS provides a procedure to execute these steps in sequence for each of the programming languages it supports.

- For CICS for MVS/ESA, the *CICS for MVS/ESA System Definition Guide* describes how to use these procedures and the *CICS/ESA Application Programming Guide* gives more information on the translation process.

- For CICS/MVS, the *CICS/MVS Installation Guide* describes how to use these procedures and the *CICS/MVS Application Programmer's Reference* manual gives more information on the translation process.

You must include:

- In the SYSLIB statement of the compilation (or assembly) stage, statements that make the product data definition files available to the compiler. The data definitions are supplied in the following MQSeries for MVS/ESA libraries:

   For COBOL, **thlqual**.SCSQCOBC
   For assembler language, **thlqual**.SCSQMACS
   For C, **thlqual**.SCSQC370
   For PL/I, **thlqual**.SCSQPLIC

- In your link-edit JCL, the MQSeries for MVS/ESA CICS stub program (CSQCSTUB). Figure 19 shows fragments of JCL code to do this. The stub is language independent and is supplied in library **thlqual**.SCSQLOAD.

When you have completed these steps, store the load module in an application load library and define the program to CICS in the usual way.

```
   ⋮
//*
//* MQSERIES FOR MVS/ESA LIBRARY CONTAINING CICS STUB
//*
//CSQSTUB   DD  DSN=++HLQ.MQM100++.SCSQLOAD,DISP=SHR
//*
   ⋮
//LKED.SYSIN DD *
   INCLUDE CSQSTUB(CSQCSTUB)
   ⋮
/*
```

Figure 19. Fragments of JCL to link-edit the object module in the CICS environment

Building IMS (BMP or MPP) applications

If you are building batch DL/I programs, see "Building MVS batch applications" on page 267. To build other applications that run under IMS (either as a BMP or an MPP), create JCL that performs these tasks:

1. Compile (or assemble) the program to produce object code. The JCL for your compilation must include SYSLIB statements that make the product data definition files available to the compiler. The data definitions are supplied in the following MQSeries for MVS/ESA libraries:

    For COBOL, **thlqual**.SCSQCOBC
    For assembler language, **thlqual**.SCSQMACS
    For C, **thlqual**.SCSQC370
    For PL/I, **thlqual**.SCSQPLIC

2. For a C application, prelink the object module created in step 1.

3. Link-edit the object code created in step 1 (or step 2 for a C/370 application) to produce a load module:

    a. Include the IMS language interface module. The name of this module depends on the version of IMS you are using; for more information, see the *MQSeries for MVS/ESA System Management Guide*.

    b. Include the MQSeries for MVS/ESA IMS stub program (CSQQSTUB). Figure 20 on page 270 shows fragments of JCL to do this. The stub is language independent and is supplied in library **thlqual**.SCSQLOAD. Do not link-edit a higher level of CSQQSTUB than that of the version of MQSeries for MVS/ESA on which your program is running. For example, a program running on Version 1.1.3 must not be link-edited with CSQQSTUB supplied with Version 1.1.4.

    Note: If you are using COBOL, you should select the NODYNAM compiler option to enable the linkage editor to resolve references to CSQQSTUB

unless you intend to use dynamic linking as described in "Dynamically calling the MQSeries stub" on page 270.

4. Store the load module in an application load library.

```
   ⋮
//*
//* MQSERIES FOR MVS/ESA LIBRARY CONTAINING IMS STUB
//*
//CSQSTUB   DD  DSN=++HLQ.MQM100++.SCSQLOAD,DISP=SHR
//*
   ⋮
//LKED.SYSIN DD *
  INCLUDE CSQSTUB(CSQQSTUB)
   ⋮
/*
```

Figure 20. Fragments of JCL to link-edit the object module in the IMS environment

Dynamically calling the MQSeries stub

Instead of link-editing the MQSeries stub program with your object code, you can dynamically call the stub from within your program. You can do this in the batch, IMS, and CICS environments. This facility is not supported by programs using PL/I in the CICS environment.

However, this method:

- Increases the complexity of your programs
- Increases the storage required by your programs at execution time
- Reduces the performance of your programs
- Means that you cannot use the same programs in other environments

If you call the stub dynamically, the appropriate stub program and its aliases must be available at execution time. To ensure this, include the MQSeries for MVS/ESA data set SCSQLOAD:

For batch and IMS   In the STEPLIB concatenation of the JCL
For CICS   In the CICS DFHRPL concatenation

For IMS, you must ensure that the library containing the dynamic stub (built as described in the information about installing the IMS adapter in the *MQSeries for MVS/ESA System Management Guide*) is ahead of the data set SCSQLOAD in the STEPLIB concatenation of the region JCL.

Use the names shown in Table 13 when you call the stub dynamically. In PL/I, only declare the call names used in your program.

Table 13. Call names for dynamic linking			
MQI call	**Dynamic call name**		
	Batch	**CICS**	**IMS**
MQBACK	CSQBBACK	not supported	not supported
MQCMIT	CSQBCOMM	not supported	not supported
MQCLOSE	CSQBCLOS	CSQCCLOS	MQCLOSE
MQCONN	CSQBCONN	CSQCCONN	MQCONN
MQDISC	CSQBDISC	CSQCDISC	MQDISC
MQGET	CSQBGET	CSQCGET	MQGET
MQINQ	CSQBINQ	CSQCINQ	MQINQ
MQOPEN	CSQBOPEN	CSQCOPEN	MQOPEN
MQPUT	CSQBPUT	CSQCPUT	MQPUT
MQPUT1	CSQBPUT1	CSQCPUT1	MQPUT1
MQSET	CSQBSET	CSQCSET	MQSET

For examples of how to use this technique, see the following figures:

```
      ⋮
      WORKING-STORAGE SECTION.
      ⋮
         05 WS-MQOPEN                      PIC X(8) VALUE 'CSQBOPEN'.
      ⋮
      PROCEDURE DIVISION.
      ⋮
            CALL WS-MQOPEN WS-HCONN
                        MQOD
                        WS-OPTIONS
                        WS-HOBJ
                        WS-COMPCODE
                        WS-REASON.
      ⋮
```

Figure 21. Dynamic linking using COBOL in the batch environment

```
    ⋮
        WORKING-STORAGE SECTION.
    ⋮
           05 WS-MQOPEN                     PIC X(8) VALUE 'CSQCOPEN'.
    ⋮
        PROCEDURE DIVISION.
    ⋮
             CALL WS-MQOPEN WS-HCONN
                            MQOD
                            WS-OPTIONS
                            WS-HOBJ
                            WS-COMPCODE
                            WS-REASON.
    ⋮
```

Figure 22. Dynamic linking using COBOL in the CICS environment

```
    ⋮
        WORKING-STORAGE SECTION.
    ⋮
           05 WS-MQOPEN                     PIC X(8) VALUE 'MQOPEN'.
    ⋮
        PROCEDURE DIVISION.
    ⋮
             CALL WS-MQOPEN WS-HCONN
                            MQOD
                            WS-OPTIONS
                            WS-HOBJ
                            WS-COMPCODE
                            WS-REASON.
    ⋮
        * -------------------------------------------------------------- *
        *
        *    If the compile option 'DYNAM' is specified
        *    then you may code the MQ calls as follows
        *
        * -------------------------------------------------------------- *
    ⋮
             CALL 'MQOPEN'  WS-HCONN
                            MQOD
                            WS-OPTIONS
                            WS-HOBJ
                            WS-COMPCODE
                            WS-REASON.
    ⋮
```

Figure 23. Dynamic linking using COBOL in the IMS environment

```
   ⋮
        LOAD    EP=CSQBOPEN
   ⋮
        CALL (15),(HCONN,MQOD,OPTIONS,HOBJ,COMPCODE,REASON),VL
   ⋮
        DELETE EP=CSQBOPEN
   ⋮
```

Figure 24. Dynamic linking using assembler language in the batch environment

```
   ⋮
        EXEC CICS LOAD PROGRAM('CSQCOPEN') ENTRY(R15)
   ⋮
        CALL (15),(HCONN,MQOD,OPTIONS,HOBJ,COMPCODE,REASON),VL
   ⋮
        EXEC CICS RELEASE PROGRAM('CSQCOPEN')
   ⋮
```

Figure 25. Dynamic linking using assembler language in the CICS environment

```
   ⋮
        LOAD    EP=MQOPEN
   ⋮
        CALL (15),(HCONN,MQOD,OPTIONS,HOBJ,COMPCODE,REASON),VL
   ⋮
        DELETE EP=MQOPEN
   ⋮
```

Figure 26. Dynamic linking using assembler language in the IMS environment

```
   ⋮
typedef void CALL_ME();
#pragma linkage(CALL_ME, OS)
   ⋮
main()
{
CALL_ME * csqbopen;
   ⋮
csqbopen = (CALL_ME *) fetch("CSQBOPEN");
(*csqbopen)(Hconn,&ObjDesc,Options,&Hobj,&CompCode,&Reason);
   ⋮
```

Figure 27. Dynamic linking using C language in the batch environment

```
   ⋮
typedef void CALL_ME();
#pragma linkage(CALL_ME, OS)
   ⋮
main()
{
CALL_ME * csqcopen;
   ⋮
  EXEC CICS LOAD PROGRAM("CSQCOPEN") ENTRY(csqcopen);
(*csqcopen)(Hconn,&ObjDesc,Options,&Hobj,&CompCode,&Reason);
   ⋮
```

Figure 28. Dynamic linking using C language in the CICS environment

```
   ⋮
typedef void CALL_ME();
#pragma linkage(CALL_ME, OS)
   ⋮
main()
{
CALL_ME * mqopen;
   ⋮
mqopen = (CALL_ME *) fetch("MQOPEN");
(*mqopen)(Hconn,&ObjDesc,Options,&Hobj,&CompCode,&Reason);
   ⋮
```

Figure 29. Dynamic linking using C language in the IMS environment

```
   ⋮
    DCL CSQBOPEN ENTRY EXT OPTIONS(ASSEMBLER INTER);
   ⋮
    FETCH CSQBOPEN;

    CALL CSQBOPEN(HQM,
                  MQOD,
                  OPTIONS,
                  HOBJ,
                  COMPCODE,
                  REASON);

    RELEASE CSQBOPEN;
```

Figure 30. Dynamic linking using PL/I in the batch environment

```
   ⋮
    DCL MQOPEN    ENTRY EXT OPTIONS(ASSEMBLER INTER);
   ⋮
    FETCH MQOPEN;

    CALL    MQOPEN(HQM,
                   MQOD,
                   OPTIONS,
                   HOBJ,
                   COMPCODE,
                   REASON);

    RELEASE   MQOPEN;
```

Figure 31. Dynamic linking using PL/I in the IMS environment

Debugging your programs

The main aids to debugging MQSeries for MVS/ESA application programs are the reason codes returned by each API call. See the *MQSeries Application Programming Reference* manual for a list of these and for more information, including suggestions for corrective action.

This chapter also suggests other debugging tools that you may want to use in particular environments.

Debugging CICS programs

You can use the CICS Execution Diagnostic Facility (EDF) to test your CICS programs interactively without having to modify the program or program-preparation procedure. For more information about EDF, see the *CICS/ESA Application Programming Guide*, or the *CICS/MVS Application Programmer's Reference* manual.

CICS trace

You will probably also find it helpful to use the CICS Trace Control transaction (CETR) to control CICS trace activity. For more information about CETR, see the *CICS for MVS/ESA CICS-Supplied Transactions* manual.

To determine whether CICS trace is active, display connection status using the CKQC panel. This panel also shows the trace number.

To interpret CICS trace entries, see Table 14 on page 276.

The CICS trace entry for these values is AP0*xxx* (where *xxx* is the trace number specified when the CICS adapter was enabled). All trace entries except CSQCTEST are issued by CSQCTRUE. CSQCTEST is issued by CSQCRST and CSQCDSP.

Note: If you are using CICS/MVS, only the first 8 bytes of trace data are produced; trace entries longer than this are truncated.

Debugging programs

Table 14 (Page 1 of 2). CICS adapter trace entries			
Name	**Description**	**Trace sequence**	**Trace data**
CSQCABNT	Abnormal termination	Before issuing END_THREAD ABNORMAL to MQSeries. This is due to the end of the task and therefore an implicit backout could be performed by the application. A ROLLBACK request is included in the END_THREAD call in this case.	Unit of work information. You can use this information when finding out about the status of work. (For example, it can be verified against the output produced by the DISPLAY THREAD command, or the MQSeries for MVS/ESA log print utility.)
CSQCBACK	Syncpoint backout	Before issuing BACKOUT to MQSeries for MVS/ESA. This is due to an explicit backout request from the application.	Unit of work information.
CSQCCCRC	Completion code and reason code	After unsuccessful return from API call.	Completion code and reason code.
CSQCCOMM	Syncpoint commit	Before issuing COMMIT to MQSeries for MVS/ESA. This can be due to a single-phase commit request or the second phase of a two-phase commit request. The request is due to a explicit syncpoint request from the application.	Unit of work information.
CSQCEXER	Execute resolve	Before issuing EXECUTE_RESOLVE to MQSeries for MVS/ESA.	The unit of work information of the unit of work issuing the EXECUTE_RESOLVE. This is the last indoubt unit of work in the resynchronization process.
CSQCGETW	GET wait	Before issuing CICS wait.	Address of the ECB to be waited on.
CSQCGMGD	GET message data	After successful return from MQGET.	Up to 40 bytes of the message data.
CSQCGMGH	GET message handle	Before issuing MQGET to MQSeries for MVS/ESA.	Object handle.
CSQCGMGI	Get message ID	After successful return from MQGET.	Message ID and correlation ID of the message.
CSQCINDL	Indoubt list	After successful return from the second INQUIRE_INDOUBT.	The indoubt units of work list.
CSQCINDO	IBM use only		
CSQCINDS	Indoubt list size	After successful return from the first INQUIRE_INDOUBT and the indoubt list is not empty.	Length of the list. Divided by 64 gives the number of indoubt units of work.
CSQCINQH	INQ handle	Before issuing MQINQ to MQSeries for MVS/ESA.	Object handle.
CSQCLOSH	CLOSE handle	Before issuing MQCLOSE to MQSeries for MVS/ESA.	Object handle.
CSQCLOST	Disposition lost	During the resynchronization process, CICS informs the adapter that it has been cold started so no disposition information regarding the unit of work being resynchronized is available.	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCNIND	Disposition not indoubt	During the resynchronization process, CICS informs the adapter that the unit of work being resynchronized should not have been indoubt (that is, perhaps it is still running).	Unit of work ID known to CICS for the unit of work being resynchronized.

Table 14 (Page 2 of 2). CICS adapter trace entries

Name	Description	Trace sequence	Trace data
CSQCNORT	Normal termination	Before issuing END_THREAD NORMAL to MQSeries for MVS/ESA. This is due to the end of the task and therefore an implicit syncpoint commit may be performed by the application. A COMMIT request is included in the END_THREAD call in this case.	Unit of work information.
CSQCOPNH	OPEN handle	After successful return from MQOPEN.	Object handle.
CSQCOPNO	OPEN object	Before issuing MQOPEN to MQSeries for MVS/ESA.	Object name.
CSQCPMGD	PUT message data	Before issuing MQPUT to MQSeries for MVS/ESA.	Up to 40 bytes of the message data.
CSQCPMGH	PUT message handle	Before issuing MQPUT to MQSeries for MVS/ESA.	Object handle.
CSQCPMGI	PUT message ID	After successful MQPUT from MQSeries for MVS/ESA.	Message ID and Correlation ID of the message.
CSQCPREP	Syncpoint prepare	Before issuing PREPARE to MQSeries for MVS/ESA in the first phase of two-phase commit processing. This call can also be issued from the distributed queuing component as an API call.	Unit of work information.
CSQCP1MD	PUTONE message data	Before issuing MQPUT1 to MQSeries for MVS/ESA.	Up to 40 bytes of data of the message.
CSQCP1MI	PUTONE message ID	After successful return from MQPUT1.	Message ID and correlation ID of the message.
CSQCP1ON	PUTONE object name	Before issuing MQPUT1 to MQSeries for MVS/ESA.	Object name.
CSQCRBAK	Resolved backout	Before issuing RESOLVE_ROLLBACK to MQSeries for MVS/ESA.	Unit of work information.
CSQCRCMT	Resolved commit	Before issuing RESOLVE_COMMIT to MQSeries for MVS/ESA.	Unit of work information.
CSQCRMIR	RMI response	Before returning to the CICS RMI (resource manager interface) from a specific invocation.	Architected RMI response value. Its meaning depends of the type of the invocation. These values are documented in the *CICS for MVS/ESA Customization Guide*. To determine the type of invocation, look at previous trace entries produced by the CICS RMI component.
CSQCRSYN	Resynchronization	Before the resynchronization process starts for the task.	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCSETH	SET handle	Before issuing MQSET to MQSeries for MVS/ESA.	Object handle.
CSQCTASE	IBM use only		
CSQCTEST	Trace test	Used in EXEC CICS ENTER TRACE call to verify the trace number supplied by the user or the trace status of the connection.	No data.
CSQCDCFF	IBM use only		

Debugging TSO programs

The following interactive debugging tools are available for TSO programs:

- TEST tool
- VS COBOL II interactive debugging tool
- INSPECT interactive debugging tool for C and PL/I programs

Chapter 23. Building your application on OS/2 Warp

The OS/2 publications describe how to build executable applications from the programs you write. This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for OS/2 Warp applications to run under OS/2 Warp. C, C++, and COBOL programming languages are supported. For information about preparing your C++ programs, see the *MQSeries Using C++* book.

The tasks you must perform to create an executable application using MQSeries for OS/2 Warp vary with the programming language your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the MQSeries for OS/2 Warp include files for the language you are using. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

Preparing C programs

> **For DOS and Windows 3.1 only**
>
> Applications must be built using the large memory model.

Work in your normal environment; MQSeries for OS/2 Warp requires nothing special.

- You need to link your programs with the appropriate libraries provided by MQSeries:

Library file	Program/exit type
MQM.LIB	server for 32-bit C
MQIC.LIB	client for C

  The following command gives an example of compiling the sample program amqsget0:

  ```
  icc amqsget0.c /Gm /Gd /B "/pmtype:vio" /Fe"amqsget0.exe" mqm.lib
  ```

  Notes:

  1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), link to the MQMZF.LIB library.

  2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo, use the MQRMIXASwitch structure and link to the MQMXA.LIB library.

  3. If you are writing a CICS exit for use with CICS for OS/2 Version 2.0.1, link to the MQMCICS.LIB library. If you are writing a CICS exit for use with Transaction Server for OS/2 Warp, Version 4, link to the MQMCICS3.LIB library.

- **For DOS only:** Your application must also be linked with two of the following libraries, one for each protocol, indicating whether you do or do not require it.

If you require TCP/IP you must also link to SOCKETL from the DOS TCP/IP product.

Library file	Protocol
MQICN.LIB	NetBIOS required
MQICDN.LIB	NetBIOS not required
MQICT.LIB	TCP/IP required
MQICDT.LIB	TCP/IP not required

- You must ensure that you have specified adequate runtime stack and heap sizes:

  – You must link a trusted application with more stack than a normal application. Therefore, a stack size of 200 KB is the recommended minimum.

  – A heap size of 8 KB is the recommended minimum.

- The DLLs must be in the library path (LIBPATH) you have specified.

- If you use lowercase characters whenever possible, you can move from MQSeries for OS/2 Warp to MQSeries on UNIX systems, where use of lowercase is necessary.

Preparing CICS and Transaction Server programs

Sample C source for a CICS MQSeries transaction is provided by AMQSCIC0.CCS. You build it using the standard CICS facilities.

For CICS for OS/2 Version 2:

1. Add the following lines to the CICSENV.CMD file:

```
UserWrk = 'c:\mqm\dll'
UserInclude = 'c:\mqm\tools\c\include;c:\mqm\tools\c\samples'
```

If necessary replace c:\mqm with the path on which you installed the sample code.

2. Compile using the command:

```
CICS32TC AMQSCIC0.CCS LIBS(MQM)
```

This is described in the *CICS for OS/2 V2.0.1 Application Programming* Guide.

For Transaction Server for OS/2 Warp, Version 4:

1. Add the following lines to the CICSENV.CMD file:

```
UserWrk = 'c:\mqm\dll'
UserInclude = 'c:\mqm\tools\c\include;c:\mqm\tools\c\samples'
```

If necessary replace c:\mqm with the path on which you installed the sample code.

2. Compile using the command:

```
CICSCTCL AMQSCIC0.CCS LIBS(MQM)
```

This is described in the *Transaction Server for OS/2 Warp, V4 Application Programming* Guide.

You can find more information about supporting CICS transactions in the *MQSeries System Administration* Guide.

Preparing COBOL programs

To prepare COBOL programs on OS/2, link your programs with one of the following libraries provided by MQSeries:

Library file	Program/exit type
MQMCB16	server for 16-bit Micro Focus COBOL
MQICCB16	client for 16-bit Micro Focus COBOL
MQMCBB	server for 32-bit IBM VisualAge for COBOL
MQMCB32	server for 32-bit Micro Focus COBOL
MQICCBB	client for 32-bit IBM VisualAge for COBOL
MQICCB32	client for 32-bit Micro Focus COBOL

To compile, for example, the sample program amq0put0, using IBM VisualAge for COBOL:

1. Set the SYSLIB environment variable to include the path to the MQSeries VisualAge COBOL copybooks:

   ```
   set SYSLIB=<drive>:\mqm\tools\cobol\copybook\VAcobol;%SYSLIB%
   ```

2. Compile and link the program:

   ```
   cob2 amq0put0.cbl -qlib <drive>:\mqm\tools\lib\mqmcbb.lib
   ```

   (for use on the MQSeries server)

   ```
   cob2 amq0put0.cbl -qlib <drive>:\mqm\tools\lib\mqiccbb.lib
   ```

   (for use on the MQSeries client)

   Note: Although the compiler option CALLINT(SYSTEM) must be used, this is the default for `cob2`.

To prepare Micro Focus COBOL programs, follow these steps:

1. Compile your applications with the LITLINK directive.

2. Specify adequate runtime stack sizes. You must link a trusted application with more stack than a normal application, so a stack size of 200 KB is the recommended minimum. To do this, use:

   ```
   set cobsw=xxxx
   ```

3. Link the object file to the runtime system.

   Set the LIB environment variable to point to the compiler COBOL libraries.

   Link the object file for use on the MQSeries server:

   ```
   cbllink amq0put0.obj mqmcb32.lib
   ```

   or

   Link the object file for use on the MQSeries client:

   ```
   cbllink amq0put0.obj mqiccb32.lib
   ```

4. Add the MQSeries copybook directory (\mqm\tools\cobol\copybook) to the `cobcpy` environment variable.

   ```
   set cobcpy=c:\mqm\tools\cobol\copybook;%COBCPY%
   ```

Preparing Transaction Server programs

To prepare Transaction Server for OS/2 Warp, V4 programs using IBM VisualAge for COBOL:

1. Add the following lines to the CICSENV.CMD file:

```
UserWrk='c:\mqm\dll'
UserCobol='IBM'
UserCobcopy='c:\mqm\tools\cobol\copybook'
UserCobWork='c:\mq-cics\wrk'
```

Where \mq-cics\wrk is the name of a work directory for output from CICSTRAN and CICSCOMP commands (see steps 2 and 3).

2. Translate your program:

```
CICSTRAN MYPROG.CPP
```

This translates your program to a .CBL program.

3. Compile your program:

```
CICSCOMP MYPROG.CBL
```

4. Link your program:

```
CICSLINK MYPROG.OBJ LIBS(MQMCBB)
```

For further information about this, see the *CICS for OS/2 Customization V3.0, SC33-1581-00* and the *Transaction Server for OS/2 Warp, V4 Application Programming* Guide.

Preparing PL/I programs

Sample PL/I programs are supplied with MQSeries. PL/I include files are also provided so that the C entry points in the MQSeries libraries can be invoked directly.

To prepare a PL/I program:

1. Link your program with one of the libraries listed in "Preparing C programs" on page 279.

2. Ensure that \mqm\tools\pli\include is in your INCLUDE environment variable.

3. Compile your program:

```
pli amqpput0.pli
ilink amqpput0.obj mqm.lib
```

Chapter 24. Building your application on SINIX or DC/OSx

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for SINIX and DC/OSx applications to run under SINIX or DC/OSx. COBOL and C programming languages are supported.

In addition to coding the MQI calls in your source code, you must add the appropriate include files. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

> **Note**
>
> Throughout this chapter, the name **mqmtop** has been used to represent the name of the installation directory. For MQSeries for SINIX and DC/OSx the name of the actual directory is /opt/mqm.

Preparing C programs

You need to link your programs with the appropriate library provided by MQSeries.

If you are **not** working in a DCE-threaded environment or using CICS, you must link to one of the following libraries:

Library file	Program/exit type
libmqm.so	server for C
libmqic.so	client for C

If you **are** working in a DCE-threaded environment or using CICS, you must link to the C library, libmqm_r.so.

Notes:

1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), you need to link to the libmqmzf.so library. Installable services must not use DCE.

2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo, link to the libmqmxa.so library in a non-DCE threaded environment and to the libmqmxa_r.so library in a DCE threaded environment.

C compiler flags

When you compile dynamic libraries, or shared objects, for use with MQSeries for SINIX and DC/OSx, you **must** use the cc command in the final step that creates the library or object, and not merely the ld command. This is because the `cc` command automatically links various initialization data that is needed for proper dynamic linking and loading.

The order of the libraries specified is important. The following is an example of how to build the sample program amqsput0 for SINIX:

```
cc -o amqsput0 -lmqm -lmqmcs -lmqmzse -lnsl   \
-lsocket -ldl -lmproc -lext amqsput0.c
```

For DC/OSx Version cd087, include `-liconv -lresolv` on the above command, as shown below:

```
cc -o amqsput0 -lmqm -lmqmcs -lmqmzse -lnsl   \
-lsocket -ldl -liconv -lresolv  -lmproc -lext amqsput0.c
```

In the same way, for versions preceding cd087 of DC/OSx, include `-liconv`.

Note: If you are using an additional product such as ENCINA, you need to find the appropriate header files. You can do this in two ways:

1. Use the `-I` option to scan the extra include directory, for example:

   ```
   cc -c -I/opt/encina/include amqsxaex.c
   ```

2. Symbolically link the header files into `/usr/include`, for example:

   ```
   ln -s /opt/encina/include/* /usr/include
   ```

Preparing COBOL programs

You must compile your COBOL programs using the MicroFocus Cobol compiler for SINIX with the LITLINK directive.

You must ensure that you have specified adequate runtime stack sizes; 16K is the recommended minimum.

You need to link your program with one of the following:

libmqmcbrt.o MQSeries COBOL runtime
libmqmcb.so server for COBOL
libmqicb.so client for COBOL

Export the following variables:

```
COBDIR=/usr/opt/lib/cobol    export COBDIR
COBLIB=$COBDIR/coblib    export COBLIB
COBCPY=/mqmtop/inc    export COBCPY
LD_LIBRARY_PATH=/opt/lib/cobol/coblib   export LD_LIBRARY_PATH
```

Compiling COBOL programs

To compile a COBOL program like amq0gbr0.cbl on MQSeries for SINIX and DC/OSx, enter:

```
$ cob -xU -C warning=2 amq0gbr0.cbl -lmqmcb -lmqm   \
-lmqmcs -lmqmzse -lmproc
```

For DC/OSx Version cd087, include `-liconv -lresolv` on the above command. In the same way, for versions preceding cd087 of DC/OSx, include `-liconv`.

If you want to use the programs on a machine which only has MQSeries client connections, recompile the programs and link them with the mqicb library instead of the mqmcb library.

Note: The mqicb and mqmcb libraries *must* come before the mqm library on the above command line.

If you have DCE, you can link your COBOL batch programs with either DCE threaded libraries or non-DCE threaded libraries.

Notes:

1. A single program cannot contain both DCE threaded and non-DCE threaded modules.

2. Programs running under CICS must always be DCE threaded.

3. DCE threaded libraries are referred to as *reentrant*.

If you do not choose to use DCE threaded libraries, remove /opt/dcelocal/bin from your PATH environment variable before calling the COBOL compiler.

If you do choose to use DCE threaded libraries, export the following:

```
$ export COBLIBLIST="/opt/lib/cobol/coblib/liblist_r"
```

Note: `COBLIBLIST` is used only in the Micro Focus Compiler for SINIX. It is the same as the Micro Focus Compiler expression `LIBLIST`.

Preparing CICS programs

An XA switch module is provided to enable you to link CICS with MQSeries:

Table 15. Essential Code for CICS applications (SINIX)		
Description	**C (source)**	**C (exec) - add one of the following to your XAD.Stanza**
XA initialization routine	amqzscix.c	amqzsc - CICS for Siemens Nixdorf SINIX V2.2

Always link your transactions with the thread safe MQSeries library libmqm_r.so.

You can find more information about supporting CICS transactions in the *MQSeries for SINIX and DC/OSx System Management Guide*.

CICS on Open Systems support

MQSeries for SINIX supports CICS on Open Systems via the XA interface.

Note: MQSeries for DC/OSx does not support CICS.

In order to enable the CICS runtime COBOL library to call MQSeries on UNIX systems, you must add the MQSeries COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol libmqmcbrt.o -lmqm_r
```

It is important to ensure that the COBOL runtime library and CICS are linked to the same (DCE) version of the library. All CICS on Open Systems MQSeries transactions **must** link with libmqm_r.

CICS sample transaction

Sample C source for a CICS MQSeries transaction is provided by
AMQSCIC0.CCS. You build it using the standard CICS facilities. Compile it using
the following commands.

Export the following environment variables:

> export CCFLAGS="-I/**mqmtop**/inc -I/**mqmtop**/samp"
> export USERLIB="-L/**mqmtop**/lib -L/opt/cics/lib -lmqm_r -lmqmcs_r"

Unset the lib path.

Then, use the command:

```
cicstcl -l C amqscic0.ccs
```

The transaction reads messages from the transmission queue
SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places
them onto the local queue whose name is contained in the transmission header of
the message. Any failures will be sent to the queue SYSTEM.SAMPLE.CICS.DLQ.
The sample MQSC script AMQSCIC0.TST may be used to create these queues
and sample input queues.

Linking libraries

You need to link your programs with the appropriate library provided by MQSeries.

You must link to one or more of the following libraries:

Library file	Program/exit type
libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C
libmqmzf.so	installable service exits for C
libmqmxa.a	XA interface for C

If you are using an additional product such as ENCINA, you need to find the
runtime libraries. There are three ways (the first two are preferred, especially if the
module is an exit or trigger monitor):

1. Link the libraries into /usr/lib/, for example:

   > ln -s /opt/encina/lib/*.so /usr/lib

   Note: You need to check these symbolic links when you install a newer
   version of ENCINA.

2. Set LD_LIBRARY_PATH to include the ENCINA library directory (this is in the
   environment when you run the programs), for example:

   > LD_LIBRARY_PATH=/opt/encina/lib export LD_LIBRARY_PATH

3. Set LD_RUN_PATH to include /opt/encina/lib when you compile the
   programs.

To compile an ENCINA program on SINIX which uses the MQI:

```
LD_RUN_PATH=/opt/encina/lib  export LD_RUN_PATH
cc -o amqsxaex -I/opt/encina/include amqsxaex.c -lmqm -lmqmcs   \
-lmqmcs -lmqmzse -lnsl -lsocket -lencina -ldl -lmproc -lext
```

For DC/OSx, include `-liconv` on the above command line.

Linking libraries

Chapter 25. Building your application on SunOS

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for SunOS applications to run under SunOS. COBOL and C programming languages are supported.

In addition to coding the MQI calls in your source code, you must add the appropriate include files. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

> **Note**
>
> Throughout this chapter, the name **mqmtop** has been used to represent the name of the installation directory. For MQSeries for SunOS the name of the actual directory is /usr/mqm.

Preparing C programs

Precompiled C programs are supplied in the /usr/mqm/samp/bin directory. To build a sample from source code, use the SPARCompiler C compiler.

Ensure the environment is set to use the correct versions of the compiler software and man pages, by exporting the environment variables PATH and MANPATH:

```
$ export PATH=/usr/lang/SC3.0.1/bin:/usr/lang/SW3.0.1/bin:$PATH
$ export MANPATH=/usr/lang.SC3.0.1/man:/usr/man:$MANPATH
```

Linking libraries

You need to link your programs with the appropriate library provided by MQSeries.

You must link to one or more of the following libraries:

Library file	Program/exit type
libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C

Notes:

1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), you need to link to the libmqmzf.so library.

2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo, link to the libmqmxa.a library.

C compiler flags

The order of the libraries specified is important. The following is an example of how to build the sample program amqsput0.

```
$ acc -sys5 -o <amqsput0> <amqsput0>.c -L/usr/mqm/lib -lmqm
```

Note: Specify the -sys5 flag so that the compiler passes the correct libraries (/usr/5lib/libc.so.*) to the linker. Failure to do so may cause the compiled program to fail.

If you already have object modules and need to use them, enter the following:

```
$ acc -sys5 -o <yourprog1> <yourprog2>.o (<yourprog2>.o <yourprog3>.o ...)
```

If you want to use the programs on a machine which has only the MQSeries client for SunOS installed, recompile the programs to link them with the client library (mqic) instead of mqm.

For example:

- If the client has SNA:

  ```
  acc -sys5 -o <yourprog> yourprog.c -L/usr/mqm/lib -lmqm  \
  -L/usr/sunlink/snap2p/p2p_lib -lp2p
  ```

- If the client has not got SNA:

  ```
  acc -sys5 -o <yourprog> yourprog.c -L/usr/mqm/lib -lmqm -lsnastubs
  ```

Preparing COBOL programs

Before you start, you need to set up the COBOL compiler. To do this, ask your system administrator to follow these steps:

1. Login as root

2. Change directory to $COBDIR/coblib, for example:

   ```
   cd $COBDIR/coblib
   ```

3. Edit the liblist file by changing the following lines:

   ```
   F-dc
   F-dp
   F-e
   Fstart
   F-X
   F-L/usr/5lib
   ```

   to

   ```
   F-dc
   F-dp
   F-e
   Fstart
   F-X
   F-Bstatic
   F-L/usr/5lib
   F-lc
   F-Bdynamic
   F-L/usr/lib
   F-ldl
   ```

The above tasks have to be carried out only once. The COBOL compiler is now set up for you to compile COBOL programs.

You must ensure that you have specified adequate runtime stack sizes; 16K is the recommended minimum.

You need to link your program with one of the following:

libmqmcb.a   server for COBOL
libmqicb.a   client for COBOL

Compile the programs using the Micro Focus compiler. Ensure that the environment is set correctly by exporting the COBCPY environment variable:

```
$ export COBCPY="/usr/mqm/inc"
$ cob -xF <yourprog>.cbl -L/usr/mqm/lib -lmqmcb
```

On machines which have only the MQSeries client for SunOS installed, link the programs with the client library, libmqicb.a, instead of libmqmcb.a:

```
$ export COBCPY="/usr/mqm/inc"
$ cob -xF <yourprog>.cbl -L/usr/mqm/lib -lmqmicb
```

Preparing COBOL programs

Chapter 26. Building your application on Sun Solaris

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for Sun Solaris applications to run under Sun Solaris. COBOL, C, and C++ programming languages are supported. For information about preparing your C++ programs, see the *MQSeries Using C++* book.

In addition to coding the MQI calls in your source code, you must add the appropriate include files. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

> **Note**
>
> Throughout this chapter, the name **mqmtop** has been used to represent the name of the installation directory. For MQSeries for Sun Solaris the name of the actual directory is /opt/mqm

Preparing C programs

Precompiled C programs are supplied in the /**mqmtop**/samp/bin directory. To build a sample from source code, use the C compiler in /opt/SUNWspro/bin/cc.

To compile, for example, the sample program amqsput0:

1. `$ export LIB=/opt/mqm/lib;$LIB`

2. Ensure the environment is set to use the correct versions of the compiler software and man pages:

   ```
   $ export PATH=/opt/SUNWspro/bin:$PATH
   $ export MANPATH=/opt/SUNWspro/man:/usr/man:$MANPATH
   $ export LD_LIBRARY_PATH=  \
   /opt/SUNWspro/lib:/$OPENWINHOME/lib:$LD_LIBRARY_PATH
   ```

3. Compile the program (the order of the libraries specified is important):

   ```
   $ cc -o <amqsput0> <amqsput0>.c -mt -lmqm -lmqmcs_d -lmqmzse  \
   -lsocket -lnsl -ldl
   ```

If you wish to compile a DCE application, use the following:

```
$ cc -o <amqsput0> <amqsput0>.c -mt -lmqm -lmqmcs -lmqmzse   \
-ldce -lthread-lsocket -lnsl -ldl
```

If you use the other version of the compiler (/usr/ucb/cc), your application may compile and link successfully. However when you run it, it will fail when it attempts to connect to the queue manager.

If you want to use the programs on a machine which has only the MQSeries client for Sun Solaris V2.2 installed, recompile the programs to link them with the client library instead:

```
$ cc -o <amqsput0> <amqsput0>.c -lmqic -lmqmcs -lsocket
```

Linking libraries

You need to link your programs with the appropriate library provided by MQSeries.

You must link to one or more of the following libraries:

Library file	Program/exit type
libmqm.so	server for C
libmqmzse.so	for C
libmqic.so	client for C
libmqmcs.so	client for C

Notes:

1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), link to the libmqmzf.so library.

2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM CICS, Transarc Encina, or Novell Tuxedo, link to the libmqmxa.a library.

3. To run the sample Encina program, link against the following libraries in addition to the libraries listed above.

   • libmqmxa.a

   Also, link against libmqmcs_d.so instead of libmqmcs.so, in addition to the Encina and DCE libraries:

   • libEncServer.so
   • libEncina.so
   • libdce.so

Preparing COBOL programs

Before preparing your COBOL programs, you should check with your system administrator that the COBOL compiler is set up to link with the correct C libraries. By default, the COBOL compiler Version 3.2 links to 3.0 SPARCompiler C libraries. For example, to update the compiler to link with SPARCompiler Version 4.0, ensure that your system administrator has completed the following:

1. Change directory to `$COBDIR/coblib`:

   `$ cd $COBDIR/coblib`

   Note: By default, COBDIR is /opt/lib/cobol.

2. Make a backup copy of liblist:

   `$ cp liblist liblist.saved`

3. Edit the liblist file using a standard UNIX editor like vi:

   `$ vi liblist`

4. Change all references from `SC3.0` to `SC4.0`.

The COBOL compiler is now set up for you to compile COBOL programs.

Precompiled COBOL programs are supplied in the /**mqmtop**/samp/bin directory. Use the Micro Focus compiler from the directory /opt/bin to build a sample from source code.

To compile, for example, the sample program amq0put0:

1. Ensure that the environment is set:

   ```
   $ export COBDIR=/opt/lib/cobol
   $ export PATH=/opt/bin:$PATH
   $ export LD_LIBRARY_PATH=$COBDIR/coblib:$LD_LIBRARY_PATH
   ```

   Note: The above assumes that COBOL is installed in the default directories.

2. Define the location of the copybooks which declare the MQI structures:

   ```
   export COBCPY="/opt/mqm/inc"
   ```

3. Link your program with one of the following libraries when building the application:

   libmqmcb.so server for COBOL
   libmqicb.so client for COBOL

4. Compile the program:

   ```
   cob -vxP amq0put0.cbl -lmqmcb -lmqm -lmqmcs -lmqmzse
   ```

Preparing CICS programs

XA switch modules are provided to enable you to link CICS with MQSeries:

Table 16. Essential Code for CICS applications (Sun Solaris)		
Description	**C (source)**	**C (exec) - add one of the following to your XAD.Stanza**
XA initialization routine	amqzscix.c	amqzsc - CICS for Sun Solaris

Always link your transactions with the thread safe MQSeries library libmqm_so.

You can find more information about supporting CICS transactions in the *MQSeries System Administration* book.

CICS on Open Systems support

MQSeries on UNIX systems supports CICS on Open Systems via the XA interface.

You must ensure that CICS COBOL applications are linked to the threaded version of the library. CICS on Open Systems MQSeries transactions must link with libmqm_r, except on Sun Solaris, where you must link with lmqmcs_d.

Preparing CICS COBOL programs using Micro Focus COBOL
To use Micro Focus COBOL, follow these steps:

1. Add the MQSeries COBOL runtime library module to the runtime library using the following command:

   ```
   cicsmkcobol /mqmtop/lib/libmqmcbrt.o -lmqmcs_d
   ```

   This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call MQSeries on UNIX systems.

   Note: cicsmkcobol must be run only when one of the following is installed:

   New version or release of Micro Focus COBOL
   New version or release of CIS for Sun Solaris

> New version or release of any supported database product (for COBOL transactions only)
> CICS for Sun Solaris

2. Export the following environment variables:

```
COBCPY=/mqmtop/inc  export COBCPY
LDFLAGS="-L/usr/lpp/mqm/lib -lmqmcs_d"  export LDFLAGS
```

3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e <yourprog>.ccp
```

Preparing CICS C programs

You build CICS C programs using the standard CICS facilities:

1. Export *one* of the following environment variables:

   • LDFLAGS = "-L/**mqmtop**/lib -L/opt/cics/lib -lmqmcs_d -lmqm -lmqmzse -lsocket -lnsl -ldl"

2. Translate, compile, and link the program by typing:

```
cicstcl -l C amqscic0.ccs
```

CICS C sample transaction: Sample C source for a CICS MQSeries transaction is provided by AMQSCIC0.CCS. The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them onto the local queue whose name is contained in the transmission header of the message. Any failures are sent to the queue SYSTEM.SAMPLE.CICS.DLQ. The sample MQSC script AMQSCIC0.TST may be used to create these queues and sample input queues.

Chapter 27. Building your application on Tandem NSK

The sample programs and the sample compilation and binding scripts, provided in subvolume ZMQSSMPL, illustrate the main features of the MQI in MQSeries for Tandem NSK, and demonstrate how to compile and bind an application.

This chapter describes some minor differences between the standard Version 2 MQI interface, as documented in the *MQSeries Application Programming Reference* book, and the MQI interface for MQSeries for Tandem NSK.

Unit of work (transaction) management

Transaction management is performed under the control of Tandem's TM/MP product, rather than by MQSeries itself. See "Syncpoints in MQSeries for Tandem NSK applications" on page 194 for details.

General design considerations

Please note that:

* The MQI library (bound into the application process) does not open $RECEIVE and does not open $TMP (TM/MP transaction pseudo-file) itself, so you may code your application to use these features.

* The MQI library uses a SERVERCLASS_SEND_() call in initial communication with the Queue Manager. While connected, it maintains two process file opens (with the LINKMON process and a Local Queue Manager Agent) and a small number of disk file opens (fewer than 10).

MQGMO_BROWSE_* with MQGMO_LOCK

As a consequence of the use of TM/MP, MQGMO_BROWSE_* with MQGMO_LOCK is not supported.

Triggered applications

Triggered MQSeries applications in the Tandem NSK environment receive user data through environment variables set up in the TACL process that is running. This is because there is a limit to the length of the argument list that can be passed to a Tandem C process.

297

Building applications on Tandem NSK

In order to access this information, triggered applications should contain code similar to the following (see sample amqsinqa for more details):

```
MQTMC2  *trig;                              /* trigger message structure    */
MQTMC2  trigdata;                  /* trigger message structure    */
char    *applId;
char    *envData;
char    *usrData;
char    *qmName;

/*******************************************************************/
/*                                                                 */
/*   Set the program argument into the trigger message            */
/*                                                                 */
/*******************************************************************/
trig = (MQTMC2*)argv[1];        /* -> trigger message   */

/* get the environment variables and load the rest of the trigger */
memcpy(&trigdata, trig, sizeof(trigdata));

memset(trigdata.ApplId,   ' ', sizeof(trigdata.ApplId));
memset(trigdata.EnvData,  ' ', sizeof(trigdata.EnvData));
memset(trigdata.UserData, ' ', sizeof(trigdata.UserData));
memset(trigdata.QMgrName, ' ', sizeof(trigdata.QMgrName));


if( (applId = getenv("TRIGAPPLID")) != 0)
{
  strncpy( trigdata.ApplId  ,applId, strlen(applId) );
}

if ( (envData = getenv("TRIGENVDATA")) != 0)
{
  strncpy( trigdata.EnvData , envData, strlen(envData) );
}

if ( (usrData = getenv("TRIGUSERDATA")) != 0)
{
  strncpy( trigdata.UserData, usrData, strlen(usrData) );
}


if ( (qmName = getenv("TRIGQMGRNAME")) != 0)
{
  strncpy( trigdata.QMgrName, qmName, strlen(qmName) );
}

trig = &trigdata;
```

Compiling and binding applications

The MQSeries for Tandem NSK V2.2 MQI is implemented using the Tandem wide memory model (the `int` datatype is 4 bytes) and the Common Runtime Environment (CRE). Applications must be compatible with this environment in order to work correctly. Refer to the sample build files for the correct options for each compiler in order to ensure compatibility.

In particular, TAL and COBOL applications must follow the rules that are required for compatibility with the CRE, documented in the Tandem manuals relating to the CRE.

Four versions of the MQI library are delivered with MQSeries for Tandem NSK V2.2, contained in ZMQSLIB. You must ensure that you use the correct library, as follows:

mqmlibc	for C, nonnative
mqmlibt	for TAL or COBOL, nonnative
mqmlibnc	for native C
mqmlibnt	for native TAL or COBOL

Running applications

In order to be able to connect to a queue manager, the environment of an application program must be correctly defined:

- The PARAM MQDEFAULTPREFIX is mandatory in the environment of all applications.

- If you have chosen an alternative (nondefault) location for your MQSINI file, an application will not be able to connect to the queue manager if the PARAM MQMACHINIFILE is not set correctly.

- TAL and COBOL applications must have the PARAM SAVE-ENVIRONMENT ON defined in their environment, or they will not be able to connect to the queue manager.

An application may run as either low-pin or high-pin. MQSeries executables themselves are configured to run as high-pin.

MQSeries applications are supported in the NSK environment only. No support for OSS applications is provided.

An MQSeries application may run under PATHWAY, from TACL, or as a child process of another process. Applications can even be added to the queue manager PATHWAY configuration itself, provided they behave correctly on queue manager shutdown.

Running applications

Chapter 28. Building your application on Windows

This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for Windows applications to run under Windows. C and Visual Basic programming languages are supported.

The tasks you must perform to create an executable application using MQSeries for Windows depend on the language in which your source code is written. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the MQSeries for Windows data definition files for the language you are using. You should be aware of the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

Linking libraries

You need to link your programs with the appropriate libraries provided by MQSeries:

Library file	Program
MQM.LIB	server for 32-bit C
MQM16.LIB	server for 16-bit C
MQM16.LIB	server for 16-bit Visual Basic
MQMSTD.LIB	server for 32-bit Visual Basic

Linking libraries

Chapter 29. Building your application on Windows NT

The Windows NT publications describe how to build executable applications from the programs you write. This chapter describes the additional tasks, and the changes to the standard tasks, you must perform when building MQSeries for Windows NT applications to run under Windows NT. COBOL, C, and C++ programming languages are supported. For information about preparing your C++ programs, see the *MQSeries Using C++* book.

The tasks you must perform to create an executable application using MQSeries for Windows NT vary with the programming language your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the MQSeries for Windows NT include files for the language you are using. You should make yourself familiar with the contents of these files. See Appendix G, "MQSeries data definition files" on page 509 for a full description.

Preparing C programs

> ──── **For DOS and Windows 3.1 only** ────────────────────────────────
>
> Applications must be built using the large memory model.

Work in your normal environment; MQSeries for Windows NT requires nothing special.

- You need to link your programs with the appropriate libraries provided by MQSeries:

Library file	Program/exit type
MQM.LIB	server for 32-bit C
MQIC.LIB	client for 16-bit C
MQIC32.LIB	client for 32-bit C

  The following command gives an example of compiling the sample program amqsget0 (using the Microsoft Visual C++ compiler):

  ```
  cl amqsget0.c /link mqm.lib
  ```

  Notes:

  1. If you are writing an installable service (see the *MQSeries Programmable System Management* book for further information), you need to link to the MQMZF.LIB library.

  2. If you are producing an XA switch load file for external coordination by an XA-compliant transaction manager such as IBM Transaction Server, Transarc Encina, or Novell Tuxedo, use the MQRMIXASwitch structure and link to the MQMXA.LIB library.

  3. If you are producing an XA switch load file using the MQRMIXASwitchDynamic structure, link to the Encina MQMENC.LIB library.

  4. To build the Encina sample, link against the following libraries:

     – MQM.LIB

– MQMENC.LIB

Also, link against the Encina and DCE libraries:

– libEncServer.lib
– libEncina.lib
– libdce.lib

5. If you are writing a CICS exit, link to the MQMCICS.LIB library.

- **For DOS only:** Your application must also be linked with two of the following libraries, one for each protocol, indicating whether you do or do not require it. If you require TCP/IP you must also link to SOCKETL from the DOS TCP/IP product.

Library file	Protocol
MQICN.LIB	NetBIOS required
MQICDN.LIB	NetBIOS not required
MQICT.LIB	TCP/IP required
MQICDT.LIB	TCP/IP not required

- You must ensure that you have specified adequate runtime heap and stack sizes. A heap size of 8K and stack size of 16K are the recommended minimum size.

- The DLLs must be in the path (PATH) you have specified.

- If you use lowercase characters whenever possible, you can move from MQSeries for Windows NT to MQSeries on UNIX systems, where use of lowercase is necessary.

Preparing CICS and Transaction Server programs

Sample C source for a CICS MQSeries transaction is provided by AMQSCIC0.CCS. You build it using the standard CICS facilities:

For CICS for Windows NT V2:

1. Add the following lines to the CICSENV.CMD file:

```
UserWork = 'c:\mqm\dll'
UserIncl = 'c:\mqm\tools\c\include;c:\mqm\tools\c\samples'
```

If necessary replace c:\mqm with the path on which you installed the sample code.

2. Edit the CICSCCL.CMD file (found in <drive>:\CNT200\UTIL) and add the library mqm.lib to the set of libraries.

3. To the LIB environment variable add:

```
<drive>:\MQM\TOOLS\LIB
```

4. To the INCLUDE environment variable add:

```
<drive>:\MQM\TOOLS\C\INCLUDE
<drive>:\CNT200\INCLUDE
```

5. Compile using the command:

```
CICSCTCL AMQSCIC0
```

This is described in the *CICS for Windows NT V2.0 Application Programming Guide*.

For Transaction Server for Windows NT, V4:

1. Set the environment variable:

   ```
   set CICS_IBMC_FLAGS=-IC:\MQM\TOOLS\C\INCLUDE; %CICS_IBMC_FLAGS%
   ```

2. Set the USERLIB environment variable:

   ```
   set USERLIB=MQM.LIB;%USERLIB%
   ```

3. Translate, compile, and link the sample program:

   ```
   cicstcl -l IBMC amqscic0.ccs
   ```

This is described in the *Transaction Server for Windows NT Application Programming Guide (CICS) V4*.

You can find more information about supporting CICS transactions in the *MQSeries System Administration* Guide.

Preparing COBOL programs

To prepare COBOL programs on Windows NT, link your program to one of the following libraries provided by MQSeries:

Library file	Program/exit type
MQMCBB	server for 32-bit IBM COBOL
MQMCB32	server for 32-bit Micro Focus COBOL
MQICCBB	client for 32-bit IBM COBOL
MQICCB32	client for 32-bit Micro Focus COBOL
MQMCB16	server for 16-bit Micro Focus COBOL
MQICCB16	client for 16-bit Micro Focus COBOL

When you are running a program in the MQI client environment, ensure the DOSCALLS library appears before any COBOL or MQSeries library.

Micro Focus

You must relink any existing MQSeries Micro Focus COBOL programs using either mqmcb3.lib or mqiccb32.lib rather than the mqmcbb and mqiccbb libraries.

To compile, for example, the sample program amq0put0, using IBM VisualAge for COBOL:

1. Set the SYSLIB environment variable to include the path to the MQSeries VisualAge COBOL copybooks:

   ```
   set SYSLIB=<drive>:\mqm\tools\cobol\copybook\VAcobol;%SYSLIB%
   ```

2. Compile and link the program:

   ```
   cob2 amq0put0.cbl -qlib <drive>:\mqm\tools\lib\mqmcbb.lib
   ```

   (for use on the MQSeries server)

   ```
   cob2 amq0put0.cbl -qlib <drive>:\mqm\tools\lib\mqiccbb.lib
   ```

   (for use on the MQSeries client)

   Note: Although the compiler option CALLINT(SYSTEM) must be used, this is the default for `cob2`.

Preparing COBOL programs

To compile, for example, the sample program amq0put0, using Micro Focus COBOL:

1. Set the COBCPY environment variable to point to the MQSeries COBOL copybooks:

   ```
   set COBCPY=<drive>\mqm\tools\cobol\copybook
   ```

2. Compile the program to give you an object file:

   ```
   cobol amq0put0 LITLINK
   ```

3. Link the object file to the runtime system.

   Set the LIB environment variable to point to the compiler COBOL libraries.

   Link the object file for use on the MQSeries server:

   ```
   cbllink amq0put0.obj mqmcb32.lib
   ```

   or

   Link the object file for use on the MQSeries client:

   ```
   cbllink amq0put0.obj mqiccb32.lib
   ```

Preparing CICS and Transaction Server programs

To compile and link a Transaction Server for Windows NT, V4 program using IBM VisualAge for COBOL:

1. Set the environment variable:

   ```
   set CICS_IBMCOB_FLAGS=c:\mqm\tools\cobol\copybook\VAcobol;%CICS_IBMCOB_FLAGS%
   ```

2. Set the USERLIB environment variable:

   ```
   set USERLIB=MQMCBB.LIB
   ```

3. Translate, compile, and link your program:

   ```
   cicstcl -l IBMCOB myprog.ccp
   ```

This is described in the *Transaction Server for Windows NT, V4 Application Programming* Guide.

To compile and link a CICS for Windows NT V2 program using Micro Focus COBOL:

• Edit the CICSLINK.CMD file and add the library mqmcbb.lib to the set of libraries.(This file is called by the CICSTCL.CMD utility.)

• Set the COBCPY environment variable:

   ```
   set
   cobcpy=<drive>:\mqm\tools\cobol\copybook;<drive>:\cnt200\copybook
   ```

• To the LIB environment variable add:

   ```
   <drive.>:\mqm\tools\lib
   ```

   ```
   <drive.>:\cobol32\lib
   ```

• Edit the CICSCOMP.CMD file, change LITLINK(2) to LITLINK to enable link-time, not run-time resolution of the MQI calls.

• Compile using the command:

   ```
   CICSTCL MQMXADC
   ```

Where MQMXADC.CCP (not actually provided as a sample program) is the name of the program. This creates a MQMXADC.DLL.

This is described in the *CICS for Windows NT V2.0 Application Programming Guide.*

Preparing PL/I programs

Sample PL/I programs are supplied with MQSeries. PL/I include files are also provided so that the C entry points in the MQSeries libraries can be invoked directly.

To prepare a PL/I program:

1. Link your program with one of the libraries listed in "Preparing C programs" on page 303.

2. Ensure that \mqm\tools\pli\include is in your INCLUDE environment variable.

3. Compile your program:

```
pli amqpput0.pli
ilink amqpput0.obj mqm.lib
```

Preparing PL/I programs

Part 4. Sample MQSeries programs

Chapter 30. Sample programs (all platforms except MVS/ESA)

| This chapter describes the sample programs delivered with MQSeries, written in C,
| COBOL, PL/I, and TAL The samples demonstrate typical uses of the Message
Queue Interface (MQI).

The samples are not intended to demonstrate general programming techniques, so
some error checking that you may want to include in a production program has
been omitted. However, these samples are suitable for use as a base for your own
message queuing programs.

The source code for all the samples is provided with the product; this source
includes comments that explain the message queuing techniques demonstrated in
the programs.

| **DCE sample exits:** For information on compiling and linking DCE sample exits
| (amqsdsc0.c) see the *MQSeries Intercommunication* book.

C++ sample programs: See the *MQSeries Using C++* book for a description of the
sample programs available in C++.

| **RPG sample programs:** See the *MQSeries for AS/400 Application Programming
| Reference (RPG)* book for a description of the sample programs available in RPG.

Note for MQSeries on UNIX systems

Throughout this chapter and the remainder of the book, the name **mqmtop** has
been used to represent the name of the installation directory.

For MQSeries for AIX the name of the actual directory is /usr/lpp/mqm.

For other UNIX platforms the name of the actual directory is /opt/mqm.

| The names of the samples start with the prefix amq, except for TAL programs that
| start zmq. The fourth character indicates the programming language, and the
compiler where necessary.

s	C language
0	COBOL language on both IBM and Micro Focus compilers
i	COBOL language on IBM compilers only
m	COBOL language on Micro Focus compilers only
v	COBOL language on DEC COBOL V2.3 and subsequent releases
p	PL/I language

Features demonstrated in the sample programs

| The following tables show the techniques demonstrated by the MQSeries sample
| programs on all systems except MVS/ESA (see Chapter 31, "Sample programs for
| MQSeries for MVS/ESA" on page 375). All the samples open and close queues
| using the MQOPEN and MQCLOSE calls, so these techniques are not listed

| separately in the tables. See the heading that includes the platform you are
| interested in:

Samples for Digital OpenVMS and UNIX systems

Table 17 shows the techniques demonstrated by the sample programs for MQSeries on UNIX systems and MQSeries for Digital OpenVMS.

Table 17 (Page 1 of 2). MQSeries on UNIX and Digital OpenVMS sample programs demonstrating use of the MQI				
Technique	**C (source) (1)**	**COBOL (source) (2)**	**C (executable)**	**Client (3) (executable)**
Putting messages using the MQPUT call	amqsput0, amqsputw	amq0put0	amqsput	amqsputc, amqsputw
Putting a single message using the MQPUT1 call	amqsinqa, amqsecha	amqminqx, amqmechx, amqiinqx, amqiechx, amqvinqx, amqviechx	amqsinq, amqsech	no sample
Putting messages to a distribution list (4)	amqsptl0	no sample	amqsptl	amqsptlc
Replying to a request message	amqsinq0	amqminqx, amqiinqx, amqvinqx	amqsinq	no sample
Getting messages (no wait)	amqsgbr0	amq0gbr0	amqsgbr	no sample
Getting messages (wait with a time limit)	amqsget0, amqsgetw	amq0get0	amqsget	amqsgetc, amqsgetw
Getting messages (unlimited wait)	amqstrg0	no sample	amqstrg	amqstrgc
Getting messages (with data conversion)	amqsecha	no sample	amqsech	no sample
Putting reference messages to a queue (4)	amqsprma	no sample	amqsprm	amqsprmc
Getting reference messages from a queue (4)	amqrgrma	no sample	amqsgrm	amqsgrmc
Reference message channel exit (4)	amqsxrma	no sample	amqsxrm	no sample
Browsing first 20 characters of a message	amqsgbr0	amq0gbr0	amqsgbr	no sample
Browsing complete messages	amqsbcg0	no sample	amqsbcg	no sample
Using a shared input queue	amqsinq0	amqminqx, amqiinqx, amqvinqx	amqsinq	no sample
Using an exclusive input queue	amqstrg0	amq0req0	amqstrg	amqstrgc
Using the MQINQ call	amqsinqa	amqminqx, amqiinqx, amqvinqx	amqsinq	no sample
Using the MQSET call	amqsseta	amqmsetx, amqisetx, amqvsetx	amqsset	no sample

Table 17 (Page 2 of 2). MQSeries on UNIX and Digital OpenVMS sample programs demonstrating use of the MQI

Technique	C (source) (1)	COBOL (source) (2)	C (executable)	Client (3) (executable)
Using a reply-to queue	amqsreq0	amq0req0	amqsreq	no sample
Requesting message exceptions	amqsreq0	amq0req0	amqsreq	no sample
Accepting a truncated message	amqsgbr0	amq0gbr0	amqsgbr	no sample
Using a resolved queue name	amqsgbr0	amq0gbr0	amqsgbr	no sample
Triggering a process	amqstrg0	no sample	amqstrg	amqstrgc
Using data conversion	amqsvfcx	no sample	no sample	no sample
MQSeries (coordinating XA-compliant database managers) accessing a single database using SQL (4)	amqsxas0.sqc	amq0xas0.sqb	no sample	no sample
MQSeries (coordinating XA-compliant database managers) accessing two databases using SQL (4)	amqsxag0.c, amqsxab0.sqc, amqsxaf0.sqc	amq0xag0.cbl, amq0xab0.sqb, amq0xaf0.sqb	no sample	no sample
CICS transaction (5)	amqscic0.ccs	no sample	amqscic0	no sample
Encina transaction (4)	amqsxae0	no sample	amqsxae0	no sample
TUXEDO transaction to put messages (6)	amqstxpx	no sample	no sample	no sample
TUXEDO transaction to get messages (6)	amqstxgx	no sample	no sample	no sample
Server for TUXEDO (6)	amqstxsx	no sample	no sample	no sample
Dead-letter queue handler	(7)	no sample	amqsdlq	no sample
From an MQI client, putting a message	amqsputw	no sample	no sample	amqsputc, amqsputw
Lotus Notes and MQSeries link	amqslnk0	no sample	no sample	no sample
From an MQI client, getting a message	amqsgetw	no sample	no sample	amqsgetc, amqsgetw

Notes:

1. The executable version of the MQSeries client samples share the same source as the samples that run in a server environment.

2. COBOL is not supported by MQSeries for AT&T GIS UNIX. Compile programs beginning 'amqm' with the Micro Focus COBOL, beginning 'amqi' with the IBM COBOL compiler, and beginning 'amq0' with either.

3. The executable versions of the MQSeries client samples are not available on MQSeries for HP-UX or MQSeries for Digital OpenVMS.

4. Supported on MQSeries for AIX, HP-UX, and Sun Solaris only.

5. CICS is supported by MQSeries for AIX and MQSeries for HP-UX only.

6. TUXEDO is not supported by MQSeries for Digital OpenVMS, AS/400, and Windows.

7. The source for the dead-letter queue handler is made up of several files and provided in a separate directory.

Samples for OS/2 Warp and Windows NT

Table 18 shows the techniques demonstrated by the sample programs for MQSeries for OS/2 Warp and Windows NT.

Table 18 (Page 1 of 2). MQSeries for OS/2 Warp and Windows NT sample programs demonstrating use of the MQI				
Technique	**C (source)**	**COBOL (source)**	**C (executable)**	**Client (executable)**
Putting messages using the MQPUT call	amqsput0	amq0put0	amqsput	amqsputc, amqsputw
Putting a single message using the MQPUT1 call	amqsinqa, amqsecha	amqminq2, amqmech2, amqiinq2, amqiech2	amqsinq, amqsech	amqsinqc, amqsechc
Putting messages to a distribution list	amqsptl0	no sample	amqsptl	amqsptlc
Replying to a request message	amqsinq0	amqminq2, amqiinq2	amqsinq	amqsinqc
Getting messages (no wait)	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Getting messages (wait with a time limit)	amqsget0, amqsgetw	amq0get0	amqsget	amqsgetc, amqsgetw
Getting messages (unlimited wait)	amqstrg0	no sample	amqstrg	amqstrgc
Getting messages (with data conversion)	amqsecha	no sample	amqsech	amqsechc
Putting reference messages to a queue	amqsprma	no sample	amqsprm	amqsprmc
Getting reference messages from a queue	amqsgrma	no sample	amqsgrm	amqsgrmc
Reference message channel exit	amqsxrma	no sample	amqsxrm	no sample
Browsing first 20 characters of a message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Browsing complete messages	amqsbcg0	no sample	amqsbcg	amqsbcgc
Using a shared input queue	amqsinq0	amqminq2, amqiinq2	amqsinq	amqsinqc
Using an exclusive input queue	amqstrg0	amq0req0	amqstrg	amqstrgc
Using the MQINQ call	amqsinqa	amqminq2, amqiinq2	amqsinq	amqsinqc
Using the MQSET call	amqsseta	amqmset2, amqiset2	amqsset	amqssetc
Using a reply-to queue	amqsreq0	amq0req0	amqsreq	amqsreqc
Requesting message exceptions	amqsreq0	amq0req0	amqsreq	amqsreqc
Accepting a truncated message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Using a resolved queue name	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Triggering a process	amqstrg0	no sample	amqstrg	amqstrgc
Using data conversion	amqsvfc2 amqsvfcn	no sample	no sample	no sample
CICS transaction	amqscic0.ccs	no sample	amqscic0 (1)	no sample
TUXEDO transaction to put messages (Windows NT only)	amqstxpx	no sample	no sample	no sample

Table 18 (Page 2 of 2). MQSeries for OS/2 Warp and Windows NT sample programs demonstrating use of the MQI				
Technique	**C (source)**	**COBOL (source)**	**C (executable)**	**Client (executable)**
TUXEDO transaction to get messages (Windows NT only)	amqstxgx	no sample	no sample	no sample
Server for TUXEDO (Windows NT only)	amqstxsx	no sample	no sample	no sample
Encina transaction	amqsxae0	no sample	amqsxae0	no sample
Dead-letter queue handler	(2)	no sample	amqsdlq	no sample
Linking with Lotus Notes	amqslnk0	no sample	amqslnk0	no sample
From an MQSeries client, putting a message	amqsputw	no sample	no sample	amqsputc, amqsputw
From an MQSeries client, getting a message	amqsgetw	no sample	no sample	amqsgetc, amqsgetw
Notes:				
1. The executable version on OS/2 is for Transaction Server for OS/2 Warp, Version 4; the version on Windows NT is for Transaction Server for Windows NT, Version 4.				
2. The source for the dead-letter queue handler is made up of several files and provided in a separate directory.				

The following list shows the techniques demonstrated by the MQSeries for Windows sample programs:

Sample program	**Technique**
AMQSPUTW	Putting a message on a specified queue
AMQSGETW	Getting a message from a specified queue
AMQSBCGW	Browsing a message and its header

For further information about these sample programs, see the following:

- *MQSeries for Windows V2.0 User's Guide*.
- *MQSeries for Windows V2.1 User's Guide*.

PL/I samples for AIX, OS/2 Warp, and Windows NT

Table 19 shows the techniques demonstrated by the MQSeries for AIX, MQSeries for OS/2 Warp, and MQSeries for Windows NT sample programs.

Table 19. MQSeries for AIX, OS/2 Warp, and Windows NT sample programs demonstrating use of the MQI		
Technique	**PL/I (source)**	**PL/I (executable)**
Putting messages using the MQPUT call	amqpput0	no sample
Getting messages (wait with a time limit)	amqpget0	no sample

Samples for AS/400

Table 20 shows the techniques demonstrated by the MQSeries for AS/400 sample programs. Some techniques occur in more than one sample program, but only one program is listed in the table.

Table 20. MQSeries for AS/400 sample programs demonstrating use of the MQI

Technique	C (source)	COBOL (source)
Using the MQCONN and MQDISC calls	AMQSECHA, AMQSINQA	AMQ0ECHA, AMQ0INQA
Implicitly connecting and disconnecting	AMQSPUT4	AMQ0PUT4
Putting messages using the MQPUT call	AMQSPUT4	AMQ0PUT4
Putting a single message using the MQPUT1 call	AMQSINQA, AMQSECHA	AMQ0INQA, AMQ0ECHA
Putting messages to a distribution list	AMQSPTL0	no sample
Replying to a request message	AMQSINQA	AMQ0INQA
Getting messages (no wait)	AMQSGBR4	AMQ0GBR4
Getting messages (wait with a time limit)	AMQSGET4	AMQ0GET4
Getting messages (unlimited wait)	AMQSTRG4	no sample
Getting messages (with data conversion)	AMQSECHA	AMQ0ECHA
Putting reference messages to a queue	AMQSPRMA	no sample
Getting reference messages from a queue	AMQSGRMA	no sample
Reference message channel exit	AMQSXRM4	no sample
Browsing first 20 characters of a message	AMQSGBR4	AMQ0GBR4
Browsing complete messages	AMQSBCG0	no sample
Using a shared input queue	AMQSINQA	AMQ0INQA
Using an exclusive input queue	AMQSREQ4	AMQ0REQ4
Using the MQINQ call	AMQSINQA	AMQ0INQA
Using the MQSET call	AMQSSETA	AMQ0SETA
Using a reply-to queue	AMQSREQ4	AMQ0REQ4
Requesting message exceptions	AMQSREQ4	AMQ0REQ4
Accepting a truncated message	AMQSGBR4	AMQ0GBR4
Using a resolved queue name	AMQSGBR4	AMQ0GBR4
Triggering a process	AMQSTRG4	no sample
Trigger server	AMQSERV4	no sample
Using a trigger server (including CICS transactions)	AMQSVFC4	no sample
Using data conversion	AMQSVFC4	no sample

In addition to these, the MQSeries for AS/400 sample option includes a sample data file which can be used as input to the sample programs, AMQSDATA and sample CL programs that demonstrate administration tasks. The CL samples are described in the *MQSeries for AS/400 Administration Guide*. You could use the sample CL program amqsamp4 to create queues to use with the sample programs described in this chapter.

| Samples for Tandem NSK

| The following C and COBOL sample programs are supplied with MQSeries for
| Tandem NSK, Version 2.2:

| Table 21. MQSeries for Tandem NSK C and COBOL sample programs demonstrating use of the MQI

Description	C (source)	C (executable)	COBOL85 (source)	COBOL85 (executable)
Putting messages using the MQPUT call	amqsput0	amqsput	amq0put0	amq0put
Putting a single message using the MQPUT1 call	amqsinqa	amqsinq	No sample	No sample
Getting messages (no wait)	amqsgbr0	amqsgbr	amq0gbr0	amq0gbr
Getting messages (wait with a time limit)	amqsget0	amqsget	amq0get0	amq0get
Getting messages (unlimited wait)	amqstrg0	amqstrg	No sample	No sample
Getting messages (with data conversion)	amqsecha	amqsech	amq0ech0	amq0ech
Browsing complete messages	amqsbcg0	amqsbcg	No sample	No sample
Use a shared input queue	No sample	No sample	amq0inq0	amq0inq
Using the MQSET call	amqsseta	amqsset	amq0set0	amq0set
Using a reply-to queue	amqsreq0	amqsreq	amq0req0	amq0req
Using data conversion	amqsvfcn	No sample	No sample	No sample
Sample skeleton for channel exit	amqsvchn	No sample	No sample	No sample

| The following TAL sample programs are supplied with MQSeries for Tandem NSK
| Version 2.2:

| Table 22. MQSeries for Tandem NSK TAL sample programs demonstrating use of the MQI

Description	TAL (source)	TAL (executable)
Read _n_ messages from a queue	zmqreadt	zmqread
Write _n_ messages of _n_ length to a queue	zmqwritt	zmqwrit

Preparing and running the sample programs

The following sections help you find the samples that you need to run on the different platforms.

AS/400

The MQSeries for AS/400 sample programs are provided in library QMQMSAMP as members of QCSRC, QLBLSRC, QCBLLESRC. Before you can run the samples, you must compile them as you would any other MQSeries for AS/400 applications.

UNIX systems

Table 23. Where to find the samples for MQSeries on UNIX systems

Content	Directory
source files	/**mqmtop**/samp
C source file for Windows 3.1 sample	/**mqmtop**/win_client/samp
dead-letter queue handler source files	/**mqmtop**/samp/dlq
executable files	/**mqmtop**/samp/bin
Other MQSeries client executable files	/**mqmtop**/dos_client/samp/bin /**mqmtop**/os2_client/samp/bin /**mqmtop**/win_client/samp/bin

The MQSeries on UNIX systems sample files will be in the directories listed in Table 23 if the defaults were used at installation time. To run the samples, either use the executable versions supplied or compile the source versions as you would any other applications, using an ANSI compiler. For information on how to do this, see "Running the sample programs" on page 323.

Digital OpenVMS

Table 24. Where to find the samples for MQSeries for Digital OpenVMS

Content	Directory
source files	MQS_EXAMPLES
C source file for Windows 3.1 sample	[.WIN_CLIENT.SAMP] under MQS_EXAMPLES
dead-letter queue handler source files	[.DLQ] under MQS_EXAMPLES
executable files	[.BIN] under MQS_EXAMPLES
Other MQSeries client executable files	[.DOS_CLIENT.SAMP.BIN] under MQS_EXAMPLES [OS2_CLIENT.SAMP.BIN] under MQS_EXAMPLES [WIN_CLIENT.SAMP.BIN] under MQS_EXAMPLES

The MQSeries for Digital OpenVMS sample files are in the directories listed in Table 24 if the defaults were used at installation time. To run the samples, either use the executable versions supplied or compile the source versions as you would any other applications, using an ANSI compiler. For information on how to do this, see "Running the sample programs" on page 323.

OS/2 and Windows NT

Table 25. Where to find the samples for MQSeries for OS/2 Warp and MQSeries for Windows NT

Content	Directory
C source code	<drive:directory>\MQM\TOOLS\C\SAMPLES
Source code for dead-letter handler sample	<drive:directory>\MQM\TOOLS\C\SAMPLES\DLQ
C source code for Windows 3.1 sample	<drive:directory>\MQM\WIN
COBOL source code	<drive:directory>\MQM\TOOLS\COBOL\SAMPLES
C executable files	<drive:directory>\MQM\TOOLS\C\SAMPLES\BIN
Other MQSeries client executable files	<drive:directory>\MQM\DOS <drive:directory>\MQM\AIX <drive:directory>\MQM\WIN
Sample MQSC files	<drive:directory>\MQM\TOOLS\MQSC\SAMPLES
PL/I source code	<drive:directory>\MQM\TOOLS\PLI\SAMPLES

The MQSeries for OS/2 Warp sample files will be in the directories listed in Table 25 if the defaults were used at installation time, the <drive:directory> will default to <c:>. To run the samples, either use the executable versions supplied or compile the source versions as you would any other MQSeries for OS/2 Warp applications. For information on how to do this, see "Running the sample programs" on page 323.

Tandem NSK

See the section for the programming language you are using.

Building C sample programs
The subvolume ZMQSSMPL contains the following TACL macro files to be used for building sample C applications:

CSAMP   Usage: CSAMP *source-code-file-name*

This is a basic macro for compiling a source file using the include files contained in subvolume ZMQSLIB. For example, to compile the sample AMQSBCG0, use CSAMP AMQSBCG0. If the compilation is successful, the macro produces an object file with the last character of the file name replaced by the letter O; for example, AMQSBCGO.

BSAMP   Usage: BSAMP *exe-file-name*

This is a basic macro used to bind an object file with the user libraries in ZMQSLIB. For example, to bind the compiled sample AMQSBCG0, use BSAMP AMQSBCG. The macro produces an executable file called *exe-filename*E; for example, AMQSBCGE.

	COMPALL	Usage: `COMPALL`
		This TACL macro compiles each of the sample source code files using the CSAMP macro.
	BINDALL	Usage: `BINDALL`
		This TACL macro binds each of the sample object files into executables using the BSAMP macro.
	BUILDC	Usage: `BUILDC`
		This TACL macro compiles and binds all of the C sample files using the macros COMPALL and BINDALL.

Building COBOL sample programs

The subvolume ZMQSSMPL contains the following files to be used for building sample COBOL applications.

	CCBSMPLS	Usage: `CCBSMPLS`
		This TACL macro compiles each of the COBOL sample source code files.
	BCBSMPLS	Usage: `BIND /IN BCBSMPLS/`
		This bind input file binds each of the COBOL sample object files into executables.
	BUILDCOB	Usage: `BUILDCOB`
		This TACL macro compiles and binds all of the COBOL sample files using the macros CCBSMPLS and BCBSMPLS.

Building TAL sample programs

The subvolume ZMQSSMPL contains the following files to be used for building sample TAL programs.

	CTLSMPLS	Usage: `CTLSMPLS`
		This TACL macro compiles each of the TAL sample source code files.
	BTLSMPLS	Usage: `BIND /IN BTLSMPLS/`
		This bind input file binds each of the TAL sample object files into executables.
	BUILDCOB	Usage: `BUILDCOB`
		This TACL macro compiles and binds all of the TAL sample files using the macros CTLSMPLS and BTLSMPLS.

Windows

For information about MQSeries for Windows, see the following:

- *MQSeries for Windows V2.0 User's Guide*.
- *MQSeries for Windows V2.1 User's Guide*.

Running the sample programs

Before you can run any of the sample programs, a queue manager must be created and the default definitions set up. This is explained in the *MQSeries System Administration* guide for MQSeries for AIX, HP-UX, OS/2, Sun Solaris, and Windows NT; for other platforms, see the appropriate *System Management Guide*.

On all platforms except AS/400

The samples need a set of queues to work with. Either use your own queues or run the sample MQSC file amqscos0.tst to create a set

To do this on UNIX systems and Digital OpenVMS, enter:

```
runmqsc QManagerName <amqscos0.tst >/tmp/sampobj.out
```

Check the `sampobj.out` file to ensure that there are no errors.

To do this on OS/2 and Windows NT enter:

```
runmqsc QManagerName  <amqscos0.tst > sampobj.out
```

Check the `sampobj.out` file to ensure that there are no errors. This file will be found in your current directory.

To do this on Tandem NSK enter:

```
runmqsc -i $SYSTEM.ZMQSSMPL.AMQSCOMA
```

Check the `sampobj.out` file to ensure that there are no errors. This file will be found in your current directory.

The sample applications can now be run. Enter the name of the sample application followed by any parameters, for example:

```
amqsput myqueue qmanagername
```

o where `myqueue` is the name of the queue on which the messages are going to be put, and `qmanagername` is the queue manager that owns `myqueue`.

See the description of the individual samples for information on the parameters that each of them expects.

On AS/400

You can use your own queues when you run the samples, or you can run the sample program amqsamp4 to create some sample queues. The source for this program is shipped in file QCLSRC in library QMQMSAMP. It can be compiled using the CRTCLPGM command.

To call one of the sample programs using data from member PUT in file AMQSDATA of library QMQMSAMP, use a command like:

```
CALL PGM(QMQMSAMP/AMQSPUT4) PARM('QMQMSAMP/AMQSDATA(PUT)')
```

The sample data only applies to the C/400 sample programs.

Length of queue name

For the COBOL sample programs, when you pass queue names as parameters, you must provide 48 characters, padding with blank characters if necessary. Anything other than 48 characters causes the program to fail with reason code 2085.

Inquire, Set, and Echo examples

For the Inquire, Set, and Echo examples, the sample definitions cause the C versions of these samples to be triggered. If you want the COBOL versions you must change the process definitions:

> SYSTEM.SAMPLE.INQPROCESS

> SYSTEM.SAMPLE.SETPROCESS

> SYSTEM.SAMPLE.ECHOPROCESS

On OS/2, Windows NT, and UNIX do this by editing the amqscos0.tst file and changing the C executable file names to the COBOL executable file names before using the **runmqsc** command above.

On AS/400, you can use the CHGMQMPRC command (described in the *MQSeries for AS/400 Administration Guide*), or edit and run AMQSAMP4 with the alternative definition.

The Put sample programs

The Put sample programs put messages on a queue using the MQPUT call. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

Running the amqsput and amqsputc samples

These programs each take 2 parameters:

1. The name of the target queue (required)
2. The name of the queue manager (optional)

If a queue manager is not specified, amqsput connects to the default queue manager and amqsputc connects to the queue manager identified by an environment variable or the client channel definition file. To run these programs, enter one of the following:

    amqsput myqueue qmanagername

    amqsputc myqueue qmanagername

where `myqueue` is the name of the queue on which the messages are going to be put, and `qmanagername` is the queue manager that owns `myqueue`.

Running the amqsputw sample

This program has no visible interface; all messages are put in the output file.

This program takes 4 parameters:

1. The name of the output file (required)
2. The name of the input file (required)
3. The name of the queue manager (required)
4. The name of the target queue (optional)

To run amqsputw from the Windows program manager:

1. Select File and click on Run...

2. On the run dialog, enter into the command line entry field the program name followed by the parameters.

For example:

```
amqsputw outfile.out infile.in qmanagername myqueue
```

where:

`outfile.out` is used to hold the messages generated when the program runs.

`infile.in` contains the data to be put onto the target queue. Each line of data is put as a message. This must be an ASCII file.

`qmanagername` is the queue manager that owns `myqueue`.

`myqueue` is the name of the target queue on which the messages are going to be put. If you don't enter a queue name, the default queue for the queue manager is used.

Here is an example of what you would see in the output file if you supplied a target queue name:

```
Sample AMQSPUTW start
Output file "OUTFILE.OUT" opened
Input file "INFILE.IN" opened
Queue Manager name "QMANAGERNAME" will be used
target queue is MYQUEUE
MQPUT OK - message contents: <AMQSPUTW: Windows Client Test Message 1>
MQPUT OK - message contents: <AMQSPUTW: Windows Client Test Message 2>
MQPUT OK - message contents: <AMQSPUTW: Windows Client Test Message 3>
Sample AMQSPUTW end
```

Here is an example of what you would see in the output file if you did not enter a target queue name (for example, `amqsputw outfil2.out c: \infil2.in qmanagernam2`):

```
Sample AMQSPUTW start
Output file "OUTFIL2.OUT" opened
Input file "C:\INFIL2.IN" opened
Queue Manager name "QMANAGERNAM2" will be used
No parameter for Queue Name. Default Queue Name will be used
target queue is QDEF2.Q
MQPUT OK - message contents: <AMQSPUTW: Windows Client Test Message 1>
MQPUT OK - message contents: <AMQSPUTW: Windows Client Test Message 2>
MQPUT OK - message contents: <AMQSPUTW: Windows Client Test Message 3>
Sample AMQSPUTW end
```

where QDEF2.Q is the name of the default queue for the queue manager.

It is important *always* to look in the output file to see what has happened as there is no visible indication of success or failure when you run this program.

Running the amq0put sample

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the target queue
```

It takes input from StdIn and adds each line of input to the target queue. A blank line indicates there is no more data.

Running the AMQSPUT4 C sample

The C program creates messages by reading data from a member of a source file. You must specify the name of the file as a parameter when you start the program. The structure of the file must be:

```
queue name
text of message 1
text of message 2

    ⋮
text of message n
blank line
```

A sample of input for the put samples is supplied in library QMQMSAMP file AMQSDATA member PUT.

Note: Remember that queue names are case sensitive. All the queues created by the sample file create program AMQSAMP4 have names created in uppercase characters.

The C program puts messages on the queue named in the first line of the file—you could use the supplied queue SYSTEM.SAMPLE.LOCAL. The program puts the text of each of the following lines of the file into separate datagram messages, and stops when it reads a blank line at the end of the file.

Using the example data file the command is:

```
CALL PGM(QMQMSAMP/AMQSPUT4) PARM('QMQMSAMP/AMQSDATA(PUT)')
```

Running the AMQ0PUT4 COBOL sample

The COBOL program creates messages by accepting data from the keyboard. To start the program, call the program and give the name of your target queue as a program parameter. The program accepts input from the keyboard into a buffer and creates a datagram message for each line of text. The program stops when you enter a blank line at the keyboard.

Design of the Put sample program

The program uses the MQOPEN call with the MQOO_OUTPUT option to open the target queue for putting messages. If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call. To keep the program simple, on this and on subsequent MQI calls, the program uses default values for many of the options.

For each line of input, the program reads the text into a buffer and uses the MQPUT call to create a datagram message containing the text of that line. The program continues until either it reaches the end of the input or the MQPUT call fails. If the program reaches the end of the input, it closes the queue using the MQCLOSE call.

The Distribution List sample program

The Distribution List sample amqsptl0 gives an example of putting a message on several message queues. It is based on the MQPUT sample, amqsput0.

Running the Distribution List sample amqsptl0

The distribution list sample runs in a similar way to the Put samples. It takes the following parameters:

- The names of the queues
- The names of the queue managers

These values are entered as pairs. For example:

```
amqsptl0 queue1 qmanagername1 queue2 qmanagername2
```

The queues are opened using MQOPEN and messages are put to the queues using MQPUT. Reason codes are returned if any of the queue or queue manager names are not recognized.

Design of the Distribution List sample

Put Message Records (MQPMRs) specify message attributes on a per destination basis. The sample chooses to provide values for *MsgId* and *CorrelId*, and these override the values specified in the MQMD structure. The *PutMsgRecFields* field in the MQPMO structure indicates which fields are present in the MQPMRs:

```
MQLONG PutMsgRecFields=MQPMRF_MSG_ID + MQPMRF_CORREL_ID;
```

Next, the sample allocates the response records and object records. The object records (MQORs) require at least one pair of names and an even number of names, that is, *ObjectName* and *ObjectQMgrName*.

The next stage involves connecting to the queue managers using MQCONN. The sample attempts to connect to the queue manager associated with the first queue in the MQOR; if this fails, it goes through the object records in turn. You are informed if it is not possible to connect to any queue manager and the program exits.

The target queues are opened using MQOPEN and the message is put to these queues using MQPUT. Any problems and failures are reported in the response records (MQRRs).

Finally, the target queues are closed using MQCLOSE and the program disconnects from the queue manager using MQDISC. The same response records are used for each call stating the *CompCode* and *Reason*.

The Browse sample programs

The Browse sample programs browse messages on a queue using the MQGET call. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

OS/2, UNIX systems, Digital OpenVMS, and Windows NT

The C version of the program takes 2 parameters

1. The name of the source queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it will connect to the default one. For example, enter one of the following:

```
amqsgbr myqueue qmanagername

amqsgbrc myqueue qmanagername

amq0gbr0 myqueue
```

where `myqueue` is the name of the queue that the messages will be viewed from, and `qmanagername` is the queue manager that owns `myqueue`.

If you omit the `qmanagername`, when running the C sample, it will assume that the default queue manager owns the queue.

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the target queue
```

Each program retrieves copies of all the messages on the queue you specify when you call the program; the messages remain on the queue. You could use the supplied queue SYSTEM.SAMPLE.LOCAL; run the Put sample program first to put some messages on the queue. You could use the queue SYSTEM.SAMPLE.ALIAS, which is an alias name for the same local queue. The program continues until it reaches the end of the queue or an MQI call fails.

Only the first 20 characters of each message are displayed, followed by `- - - truncated` when this is the case.

AS/400

Each program retrieves copies of all the messages on the queue you specify when you call the program; the messages remain on the queue. You could use the supplied queue SYSTEM.SAMPLE.LOCAL; run the Put sample program first to put some messages on the queue. You could use the queue SYSTEM.SAMPLE.ALIAS, which is an alias name for the same local queue. The program continues until it reaches the end of the queue or an MQI call fails.

An example of a command to call the C program is:

```
CALL PGM(QMQMSAMP/AMQSGBR4) PARM('SYSTEM.SAMPLE.LOCAL')
```

Design of the Browse sample program

The program opens the target queue using the MQOPEN call with the MQOO_BROWSE option. If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call.

For each message on the queue, the program uses the MQGET call to copy the message from the queue, then displays the data contained in the message. The MQGET call uses these options:

MQGMO_BROWSE_NEXT

> After the MQOPEN call, the browse cursor is positioned logically before the first message in the queue, so this option causes the **first** message to be returned when the call is first made.

MQGMO_NO_WAIT

> The program does not wait if there are no messages on the queue.

MQGMO_ACCEPT_TRUNCATED_MSG

> The MQGET call specifies a buffer of fixed size. If a message is longer than this buffer, the program displays the truncated message, together with a warning that the message has been truncated.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call, because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The program continues to the end of the queue; at this point the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code and the program displays a warning message. If the MQGET call fails, the program displays an error message that contains the reason code.

The program then closes the queue using the MQCLOSE call.

The Browser sample program

This sample is not available on AS/400.

The Browser sample program is written as a utility not just to demonstrate a technique. It reads and outputs both the message descriptor and the message content fields of all the messages on a queue. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

This program takes 2 parameters:

1. The name of the source queue
2. The name of the queue manager

Both input parameters for this program are mandatory. For example, enter one of the following:

    amqsbcg myqueue qmanagername

    amqsbcgc myqueue qmanagername

where `myqueue` is the name of the queue on which the messages are going to be browsed, and `qmanagername` is the queue manager that owns `myqueue`.

It reads each message from the queue and outputs the following to the stdout:

Formatted message descriptor fields

Message data (dumped in hex and, where possible, character format)

The program is restricted to printing the first 32767 characters of the message, and will fail with the reason 'truncated msg' if a longer message is read.

See the *System Management Guide* for your platform, for examples of the output from this utility.

The Get sample programs

The Get sample programs get messages from a queue using the MQGET call.
See "Features demonstrated in the sample programs" on page 313 for the names
of these programs.

Running the amqsget and amqsgetc samples

These programs each take two parameters:

1. The name of the source queue (required)
2. The name of the queue manager (optional)

If a queue manager is not specified, amqsget connects to the default queue
manager, and amqsgetc connects to the queue manager identified by an
environment variable or the client channel definition file.

To run these programs, enter, one of the following:

```
amqsget myqueue qmanagername
```

```
amqsgetc myqueue qmanagername
```

where `myqueue` is the name of the queue on which the messages are going to be
put, and `qmanagername` is the queue manager that owns `myqueue`.

If you omit the `qmanagername`, the programs assume the default, or, in the case of
the MQI client, the queue manager identified by an environment variable or the
client channel definition file.

Running the amqsgetw sample

This program has no visible interface, all messages are put in the output file, not to
stdout.

This program takes 3 parameters:

1. The name of the output file (required)
2. The name of the queue manager (required)
3. The name of the target queue (optional)

To run amqsgetw from the Windows 3.1 program manager:

1. Select File and click on Run...

2. On the run dialog, enter into the command line entry field the program name
   followed by the parameters.

For example:

```
amqsgetw outfile.out qmanagername myqueue
```

where:

outfile.out is used to hold the messages generated when the program runs.

qmanagername is the queue manager that owns myqueue.

myqueue is the name of the target queue on from which the program will get messages. If you don't enter a queue name, the default queue for the queue manager is used.

Here is an example of the contents of the output file:

```
Sample AMQSGETW start
Output file "OUTFILE.OUT" opened
Queue Manager name "QMANAGERNAME" will be used
Queue Name "MYQUEUE" will be used
MQGET OK - message contents: <AMQSPUTW: Windows Client Test Message 1>
MQGET OK - message contents: <AMQSPUTW: Windows Client Test Message 2>
MQGET OK - message contents: <AMQSPUTW: Windows Client Test Message 3>
no more messages
Sample AMQSGETW end
```

It is important **always** to look in the output file to see what has happened as there is no visible indication of success or failure when you run this program.

Running the amq0get sample

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the source queue
```

Each program removes messages from the queue you specify when you call the program. You could use the supplied queue SYSTEM.SAMPLE.LOCAL; run the Put sample program first to put some messages on the queue. You could use the queue SYSTEM.SAMPLE.ALIAS, which is an alias name for the same local queue. The program continues until the queue is empty or an MQI call fails.

Running the AMQSGET4 and the AMQ0GET4 samples

The Get sample programs get messages from a queue using the MQGET call. The programs are named:

C language   AMQSGET4
COBOL language   AMQ0GET4

Each program removes messages from the queue you specify when you call the program. You could use the supplied queue SYSTEM.SAMPLE.LOCAL; run the Put sample program first to put some messages on the queue. You could use the queue SYSTEM.SAMPLE.ALIAS, which is an alias name for the same local queue. The program continues until the queue is empty or an MQI call fails.

An example of a command to call the C program is:

```
CALL PGM(QMQMSAMP/AMQSGET4) PARM('SYSTEM.SAMPLE.LOCAL')
```

Design of the Get sample program

The program opens the target queue using the MQOPEN call with the MQOO_INPUT_AS_Q_DEF option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

For each message on the queue, the program uses the MQGET call to remove the message from the queue, then displays the data contained in the message. The MQGET call uses the MQGMO_WAIT option, specifying a *WaitInterval* of 15 seconds, so that the program waits for this period if there is no message on the queue. If no message arrives before this interval expires, the call fails and returns the MQRC_NO_MSG_AVAILABLE reason code.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The MQGET call specifies a buffer of fixed size. If a message is longer than this buffer, the call fails and the program stops.

The program continues until either the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code or the MQGET call fails. If the call fails, the program displays an error message that contains the reason code.

The program then closes the queue using the MQCLOSE call.

The Reference Message sample programs

The reference message samples allow a large object to be transferred from one node to another (usually on different systems) without the need for the object to be stored on MQSeries queues at either the source or the destination nodes.

A set of sample programs is provided to demonstrate how reference messages can be 1) put to a queue, 2) received by message exits, and 3) taken from a queue. The sample programs use reference messages to move files. If you want to move other objects such as databases, or if you want to perform security checks, you must define your own exit, based on our sample, amqsxrm. The following sections describe the Reference Message sample programs.

Notes for AS/400 users

To receive a reference message using the sample message exit, specify a file in the root file system of IFS or any sub-directory so that a stream file can be created. The sample message exit on AS/400 creates the file, converts the data to EBCDIC, and sets the code page to your system code page. You then have the option of copying this file to the QSYS.LIB file system using the CPYFRMSTMF command. For example:

```
| CPYFRMSTMF FROMSTMF('JANEP/TEST.TXT')
|             TOMBR('qsys.lib.janep.lib/test.fie/test.mbr') MBROPT(*REPLACE)
|             CVTDTA(*NONE)
```

| Note that the CPYFRMSTMF command does not create the file, it must be created
| before running this command.

| If you send a file from QSYS.LIB no changes are required to the samples. For any
| other file system ensure that the CCSID specified in the CodedCharSetId field in
| the MQRMH structure matches the bulk data you are sending.

| When using the integrated file system, create program modules with the
| SYSIFCOPT(*IFSIO) option set. If you want to move database or fixed-length
| record files, define your own exit based on the supplied sample AMQSXRM4.

Running the Reference Message samples

The reference message samples run as follows:
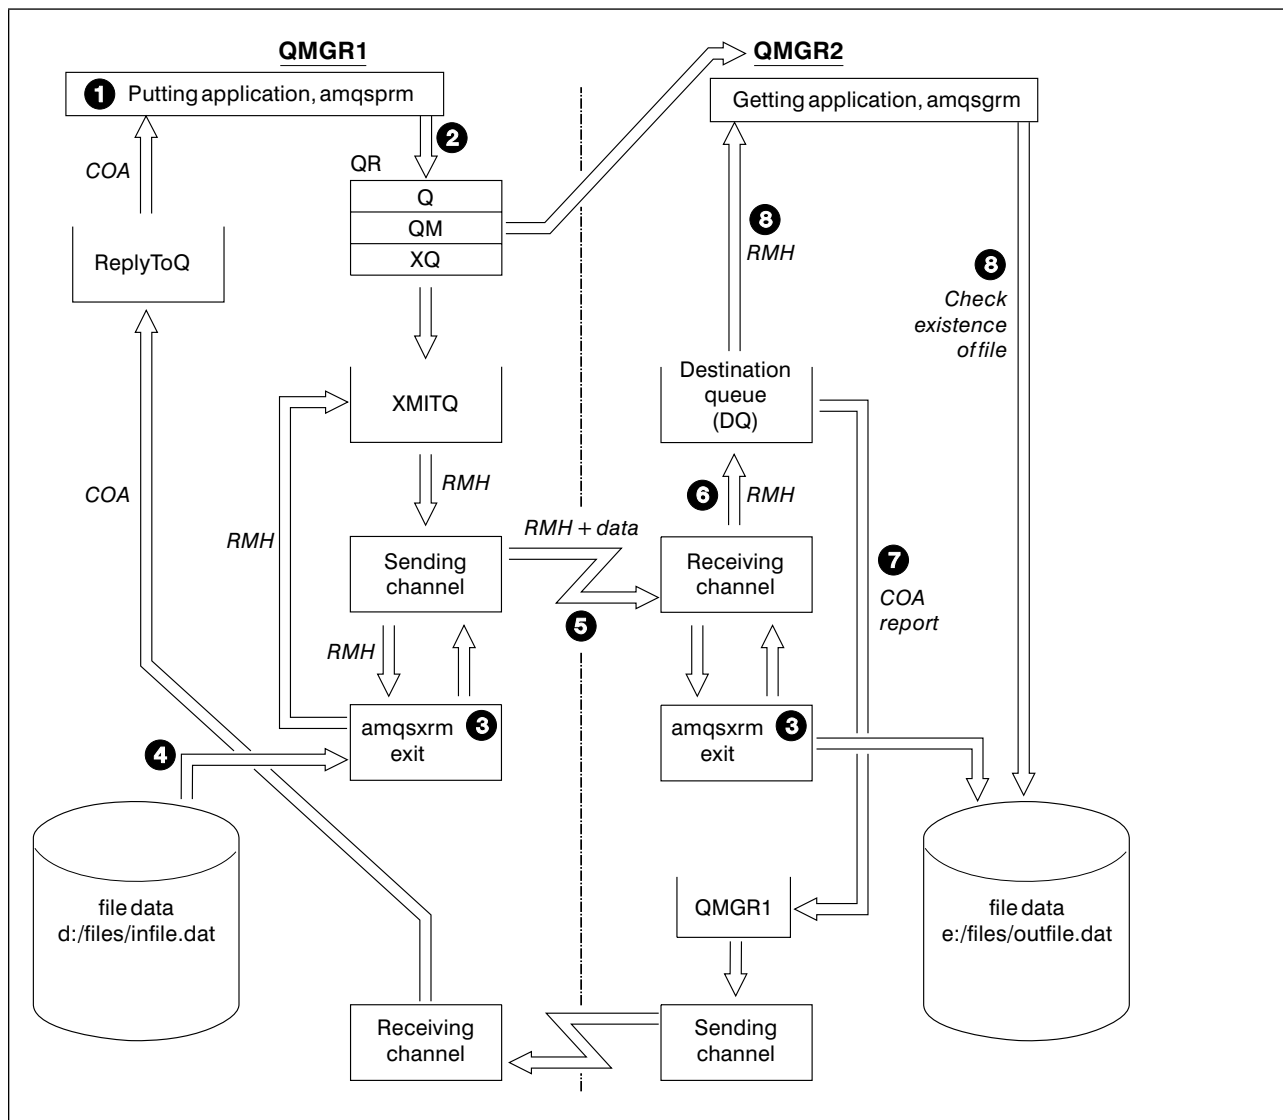
Figure 32. Running the reference message samples

1. Set up the environment to start the listeners, channels, and trigger monitors, and define your channels and queues.

   For the purposes of describing how to set-up the reference message example this refers to the sending machine as MACHINE1 with a queue manager called QMGR1 and the receiving machine as MACHINE2 with a queue manager called QMGR2.

   For AS/400 users: On the AS/400 the sample programs must be created, this can be done by issuing the following commands. You may wish to put the modules and programs into a different library but for this example they are kept in QMQMSAMP. It is important to note that the activation group QMQM must not be used for any exit programs - this activition group is for the use of MQSeries only.

   ```
   CRTCMOD MODULE(QMQMSAMP/AMQSPRMA) SRCFILE(QMQMSAMP/QCSRC)
   CRTPGM PGM(QMQMSAMP/AMQSPRMA) BNDSRVPGM(QMQM/AMQZSTUB)

   CRTCMOD MODULE(QMQMSAMP/AMQSGRMA) SRCFILE(QMQMSAMP/QCSRC)
   CRTPGM PGM(QMQMSAMP/AMQSGRMA) BNDSRVPGM(QMQM/AMQZSTUB)

   CRTCMOD MODULE(QMQMSAMP/AMQSXRM4) SRCFILE(QMQMSAMP/QCSRC)
   CRTPGM PGM(QMQMSAMP/AMQSXRM4) BNDSRVPGM(QMQM/AMQZSTUB QMQM/AMQVSTUB)
   ```

   For all users: The following definitions allow a reference message to be built to send a file with an object type of FLATFILE from queue manager QMGR1 to QMGR2 and to recreate the file as defined in the call to AMQSPRM (or AMQSPRMA on AS/400). The reference message (including the file data) is sent using channel CHL1 and transmission queue XMITQ and placed on queue DQ. Exception and COA reports are sent back to QMGR1 using the channel REPORT and transmission queue QMGR1.

   The application that receives the reference message (AMQSGRM or AMQSGRMA on the AS/400) is triggered using the initiation queue INITQ and process PROC. You need to ensure the CONNAME fields are set correctly and the MSGEXIT field reflects your directory structure, depending on machine type and where the MQSeries product is installed.

   The MQSC definitions have used an OS/2 style for defining the exits, if you are using MQSC on the AS/400 you will need to modify these accordingly. It is important to note that the message data FLATFILE is case sensitive and the sample will not work unless it is in uppercase.

   On machine MACHINE1, queue manager QMGR1

   MQSC syntax:

   ```
   define chl(chl1) chltype(sdr) trptype(tcp) conname('machine2') xmitq(xmitq)
   msgdata(FLATFILE) msgexit('d:\mqm\tools\c\samples\bin\amqsxrm.dll(MsgExit)')

   define ql(xmitq) usage(xmitq)

   define chl(report) chltype(rcvr) trptype(tcp) replace

   define qr(qr) rname(dq) rqmname(qmgr2) xmitq(xmitq) replace
   ```

   AS/400 command syntax:

Reference message samples

```
CRTMQMCHL CHLNAME(CHL1) CHLTYPE(*SDR) CONNAME(MACHINE2) TMQNAME(XMITQ)
      MSGEXIT(QMQMSAMP/AMQSXRM4) MSGUSRDATA(FLATFILE)

CRTMQMQ QNAME(XMITQ) QTYPE(*LCL) USAGE(*TMQ)

CRTMQMCHL CHLNAME(REPORT) CHLTYPE(*RCVR)

CRTMQMQ QNAME(QR) QTYPE(*RMT) RMTQNAME(DQ) RMTMQMNAME(QMGR2) TMQNAME(XMITQ)
```

On machine MACHINE2, queue manager QMGR2

MQSC syntax:

```
define chl(chl1) chltype(rcvr) trptype(tcp)
msgexit('d:\mqm\tools\c\samples\bin\amqsxrm.dll(MsgExit)')
      msgdata(flatfile)

define chl(report) chltype(sdr) trptype(tcp) conname('MACHINE1')
      xmitq(qmgr1)

define ql(initq)

define ql(qmgr1) usage(xmitq)

define pro(proc) applicid('d:\mqm\tools\c\samples\bin\amqsgrm')

define ql(dq) initq(initq) process(proc) trigger trigtype(first)
```

AS/400 command syntax:

```
CRTMQMCHL CHLNAME(CHL1) CHLTYPE(*RCVR) MSGEXIT(QMQMSAMP/AMQSXRM4)
  MSGUSRDATA(FLATFILE)

CRTMQMCHL CHLNAME(REPORT) CHLTYPE(*SDR) CONNAME(MQCHINE1) TMQNAME(QMGR1)

CRTMQMQ QNAME(INITQ) QTYPE(*LCL) USAGE(*NORMAL)

CRTMQMQ QNAME(QMGR1) QTYPE(*LCL) USAGE(*TMQ)

CRTMQMPRC PRCNAME(PROC) APPID('QMQMSAMP/AMQSGRMA')

CRTMQMQ QNAME(DQ) QTYPE(*LCL) PRCNAME(PROC) TRGENBL(*YES)
      INITQNAME(INITQ)
```

2. Once the above MQ objects have been created:

   a. Where applicable to the platform, start the listener for the sending and receiving queue managers

   b. Start the channels CHL1 and REPORT

   c. On the receiving queue manager start the trigger monitor for the initiation queue INITQ

3. Invoke the put reference message sample AMQSPRM (AMQSPRMA on the AS/400) from the command line using the following parameters:

-m Name of the local queue manager, this defaults to the default queue manager.
-i Name and location of source file
-o Name and location of destination file.
-q Name of queue
-g Name of queue manager where the queue, defined in the -q parameter exists. This defaults to the queue manager specified in the -m parameter.
-t Object type
-w Wait interval, that is, the waiting time for exception and COA reports from the receiving queue manager.

For example, to use the sample with the objects defined above you would use the following parameters:

```
-mQMGR1 -iInput File -oOutput File -qQR -tFLATFILE -w120
```

Increasing the waiting time will allow time for a large file to be sent across a network before the putting program times out.

```
amqsprm -q QR -m QMGR1 -i d:\x\file.in -o d:\y\file.out -t FLATFILE
```

AS/400 users: On the AS/400 use the following command:

```
CALL PGM(QMQMSAMP/AMQSPRMA) PARM('-mQMGR1' '-iLIBRARY/FILEIN'
'-oLIBRARY/FILEOUT' '-qDQ' '-tFLATFILE'
```

To use the IFS, use the following commands:

```
CRTCMOD MODULE(QMQMSAMP/AMQSXRM4) SRCFILE(QMQMSAMP/QCSRC) SYSIFCOPT(*IFSIO)
CRTCMOD MODULE(QMQMSAMP/AMQSGRMA) SRCFILE(QMQMSAMP/QCSRC) SYSIFCOPT(*IFSIO)
```

You may use the root directory, but it is recommended you create one using the CRTDIR command.

When calling the putting program the output file name will need to reflect the IFS naming convention, for instance /TEST/FILENAME will create a file called FILENAME in the directory TEST.

Note: You can use either a forward slash (/) or a dash (-) when specifying parameters.

For example:

```
  amqsprm /i d:\files\infile.dat /o e:\files\outfile.dat /q QR
  /m QMGR1 /w 30 /t FLATFILE
```

Note: For UNIX platforms, you must use two slashes (\\) instead of one to denote the destination file directory. Therefore, the above command looks like this:

```
  amqsprm -i /files/infile.dat -o e:\\files\\outfile.dat -q QR
  -m QMGR1 -w 30 -t FLATFILE
```

This demonstrates the following:

- The reference message will be put to queue QR on queue manager QMGR1.

- The source file and path is `d:\files\infile.dat` and exists on the system where the example command is issued.

- If the queue QR is a remote queue, the reference message is sent to another queue manager, on a different system, where a file is created with the name and path `e:\files\outfile.dat`. The contents of this file are the same as the source file.

- amqsprm waits for 30 seconds for a COA report from the destination queue manager.

- The object type is `flatfile`, so the channel used to move messages from the queue QR must specify this in the *MsgData* field.

4. When you define your channels, select the message exit at both the sending and receiving ends to be amqsxrm. This is defined on MQSeries for OS/2 Warp, and Windows NT as follows:

   ```
   msgexit('<pathname>\amqsxrm.dll(MsgExit)')
   ```

   This is defined on MQSeries for AIX, HP-UX, and Sun Solaris as follows:

   ```
   msgexit('<pathname>/amqsxrm(MsgExit)')
   ```

   If a pathname is specified, the complete name must be specified (with `.dll` on OS/2). If a pathname is not specified, it is assumed that the program is in the path specified in the qm.ini file. This is explained fully in the *MQSeries Intercommunication* book.

5. The channel exit reads the reference message header and finds the file that it is referring to.

6. It can then choose to segment the file before sending it down the channel along with the header. On MQSeries for AIX, HP-UX, and Sun Solaris, you must change the group owner of the target directory to 'mqm' so that the sample message exit can create the file in that directory. Also, change the permissions of the target directory to allow mqm group members to write to it. The file data is not stored on the MQSeries queues.

7. When the last segment of the file is processed by the receiving message exit, the reference message is put to the destination queue specified by amqsprm. If this queue is triggered (that is, the definition specifies *Trigger*, *InitQ*, and *Process* queue attributes), the program specified by the PROC parameter of the destination queue is triggered. The program to be triggered must be defined in the *ApplId* field of the *Process* attribute.

8. When the reference message reaches the destination queue (DQ), a COA report is sent back to the putting application (amqsprm).

9. The get reference message sample, amqsgrm, gets messages from the queue specified in the input trigger message and checks the existence of the file.

Design of the Put Reference Message sample (amqsprma.c)

This sample creates a reference message that refers to a file and puts it on a specified queue:

1. The sample connects to a local queue manager using MQCONN.

2. It then opens (MQOPEN) a model queue which is used to receive report messages.

3. The sample builds a reference message containing the values required to move the file, for example, the source and destination file names and the object type.

   As an example, the sample shipped with MQSeries builds a reference message to send the file `d:\x\file.in` from QMGR1 to QMGR2 and to recreate the file as `d:\y\file.out` using the following parameters:

   ```
   amqsprm -q QR -m QMGR1 -i d:\x\file.in -o d:\y\file.out -t FLATFILE
   ```

   Where QR is a remote queue definition that refers to a target queue on QMGR2.

> **Note:** For UNIX platforms, you must use two slashes (\\) instead of one to denote the destination file directory. Therefore, the above command looks like this:
>
> ```
> amqsprm -q QR -m QMGR1 -i /x/file.in -o d:\\y\\file.out -t FLATFILE
> ```

4. The reference message is put (without any file data) to the queue specified by the /q parameter. If this is a remote queue, the message is put to the corresponding transmission queue.

5. The sample waits, for the duration of time specified in the /w parameter (which defaults to 15 seconds), for COA reports, which, along with exception reports, are sent back to the dynamic queue created on the local queue manager (QMGR1).

Design of the Reference Message Exit sample (amqsxrma.c, AMQSXRM4)

This sample recognizes reference messages with an object type that matches the object type in the message exit user data field of the channel definition. For these messages, the following happens:

- At the sender or server channel, the specified length of data is copied from the specified offset of the specified file into the space remaining in the agent buffer after the reference message. If the end of the file is not reached, the reference message is put back on the transmission queue after updating the $DataLogicalOffset$ field.

- At the requester or receiver channel, if the $DataLogicalOffset$ field is zero and the specified file does not exist, it is created. The data following the reference message is added to the end of the specified file. If the reference message is not the last one for the specified file, it is discarded. Otherwise, it is returned to the channel exit, without the appended data, to be put on the target queue.

For sender and server channels, if the $DataLogicalLength$ field in the input reference message is zero, the remaining part of the file, from $DataLogicalOffset$ to the end of the file, is to be sent along the channel. If it is not zero, only the length specified is sent.

If an error occurs (for example, if the sample is unable to open a file), MQCXP.$ExitResponse$ is set to MQXCC_SUPPRESS_FUNCTION so that the message being processed is put to the dead-letter queue instead of continuing to the destination queue. A feedback code is returned in MQCXP.$Feedback$ and returned to the application that put the message in the $Feedback$ field of the message descriptor of a report message. This is because the putting application requested exception reports by setting MQRO_EXCEPTION in the $Report$ field of the MQMD.

If the encoding or $CodedCharacterSetId$ (CCSID) of the reference message is different from that of the queue manager, the reference message is converted to the local encoding and CCSID. In our sample, amqsprm, the format of the object is MQFMT_STRING, so amqsxrm converts the object data to the local CCSID at the receiving end before the data is written to the file.

Compiling the Reference Message Exit sample

To compile amqsxrma, use the following commands:

On AIX:

```
$cc -d -I/mqmtop/inc amqsxrma.c
$ ld -o amqsxrm amqsxrma.o -bE:amqsxrm.exp -H512 -T512 \
-e MQStart -bM:SRE -lc -ls -lmqm
```

On HP-UX:

```
$ cc -c -Aa +z -I/mqmtop/inc amqsxrma.c
$ ld -b -o amqsxrm amqsxrma.o -z +b : -lmqm -lc
```

On Sun Solaris:

```
$ cc -c -KPIC -I/mqmtop/inc amqsxrma.c
$ ld -G -o amqsxrm amqsxrma.o -dy -lmqm -lc -lnsl -ldl
```

Design of the Get Reference Message sample (amqsgrma.c)

The program logic is as follows:

1. The sample is triggered and extracts the queue and queue manager names from the input trigger message.

2. It then connects to the specified queue manager using MQCONN and opens the specified queue using MQOPEN.

3. The sample issues MQGET with a wait interval of 15 seconds within a loop to get messages from the queue.

4. If a message is a reference message, the sample checks the existence of the file that has been transferred.

5. It then closes the queue and disconnects from the queue manager.

The Request sample programs

The Request sample programs demonstrate client/server processing. The samples are the clients that put request messages on a target server queue that is processed by a server program. They wait for the server program to put a reply message on a reply-to queue.

The Request samples put a series of request messages on target server queue using the MQPUT call. These messages specify the local queue, SYSTEM.SAMPLE.REPLY as the reply-to queue, which can be a local or remote queue. The programs wait for reply messages, then display them. Replies are sent only if the target server queue is being processed by a server application, or if an application is triggered for that purpose (the Inquire, Set, and Echo sample programs are designed to be triggered). The C sample waits 1 minute (the COBOL sample waits 5 minutes), for the first reply to arrive (to allow time for a server application to be triggered), and 15 seconds for subsequent replies, but both samples can end without getting any replies. See "Features demonstrated in the sample programs" on page 313 for the names of the Request sample programs.

Running the amqsreq0.c, amqsreq, and amqsreqc samples

The C version of the program takes 2 parameters:

1. The name of the target server queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it will connect to the default one. For example, enter one of the following:

```
amqsreq myqueue qmanagername

amqsreqc myqueue qmanagername

amq0req0 myqueue
```

where `myqueue` is the name of the target server queue, and `qmanagername` is the queue manager that owns `myqueue`.

If you omit the `qmanagername`, when running the C sample, it will assume that the default queue manager owns the queue.

Running the amq0req0.cbl sample

The COBOL version does have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the target server queue
```

The program takes its input from StdIn and adds each line to the target server queue, taking each line of text as the content of a request message. The program ends when a null line is read.

Running the AMQSREQ4 sample

The C program creates messages by reading data from a member of a source file. You must specify the name of the file as a parameter when you start the program. The structure of the file must be:

```
queue name
text of message 1
text of message 2

⋮

text of message n
blank line
```

Samples of input for the request samples are supplied in library QMQMSAMP file AMQSDATA members ECHO, INQ and SET.

Note: Remember that queue names are case sensitive. All the queues created by the sample file create program AMQSAMP4 have names created in uppercase characters.

The C program puts messages on the queue named in the first line of the file—you could use the supplied queue SYSTEM.SAMPLE.TRIGGER. The program puts the text of each of the following lines of the file into separate request messages, and stops when it reads a blank line at the end of the file.

Running the AMQ0REQ4 sample

The COBOL program creates messages by accepting data from the keyboard. To start the program, call the program and specify the name of your target queue as a parameter. The program accepts input from the keyboard into a buffer and creates a request message for each line of text. The program stops when you enter a blank line at the keyboard.

Running the Request sample using triggering

If the sample is used with triggering and one of the Inquire, Set, or Echo sample programs, the line of input must be the queue name of the queue that you want the triggered program to access.

OS/2, UNIX systems, and Windows NT

To run the samples using triggering:

1. Start the trigger monitor program RUNMQTRM in one session (the initiation queue SYSTEM.SAMPLE.TRIGGER is available for you to use).

2. Start the amqsreq program in another session.

3. Make sure you have defined a target server queue.

   The sample queues available to you to use as the target server queue for the request sample to put messages are:

   - SYSTEM.SAMPLE.INQ - for the Inquire sample program
   - SYSTEM.SAMPLE.SET - for the Set sample program
   - SYSTEM.SAMPLE.ECHO - for the Echo sample program

   These queues have a trigger type of FIRST, so if there are already messages on the queues before you run the Request sample, server applications are not triggered by the messages you send.

4. Make sure you have defined a queue for the Inquire, Set or Echo sample program to use.

This means that the trigger monitor is ready when the request sample sends a message.

Note: The sample process definitions created using RUNMQSC and the amqscos0.tst file cause the C samples to be triggered. Change the process definitions in amqscos0.tst and use RUNMQSC with this updated file if the COBOL versions are required.
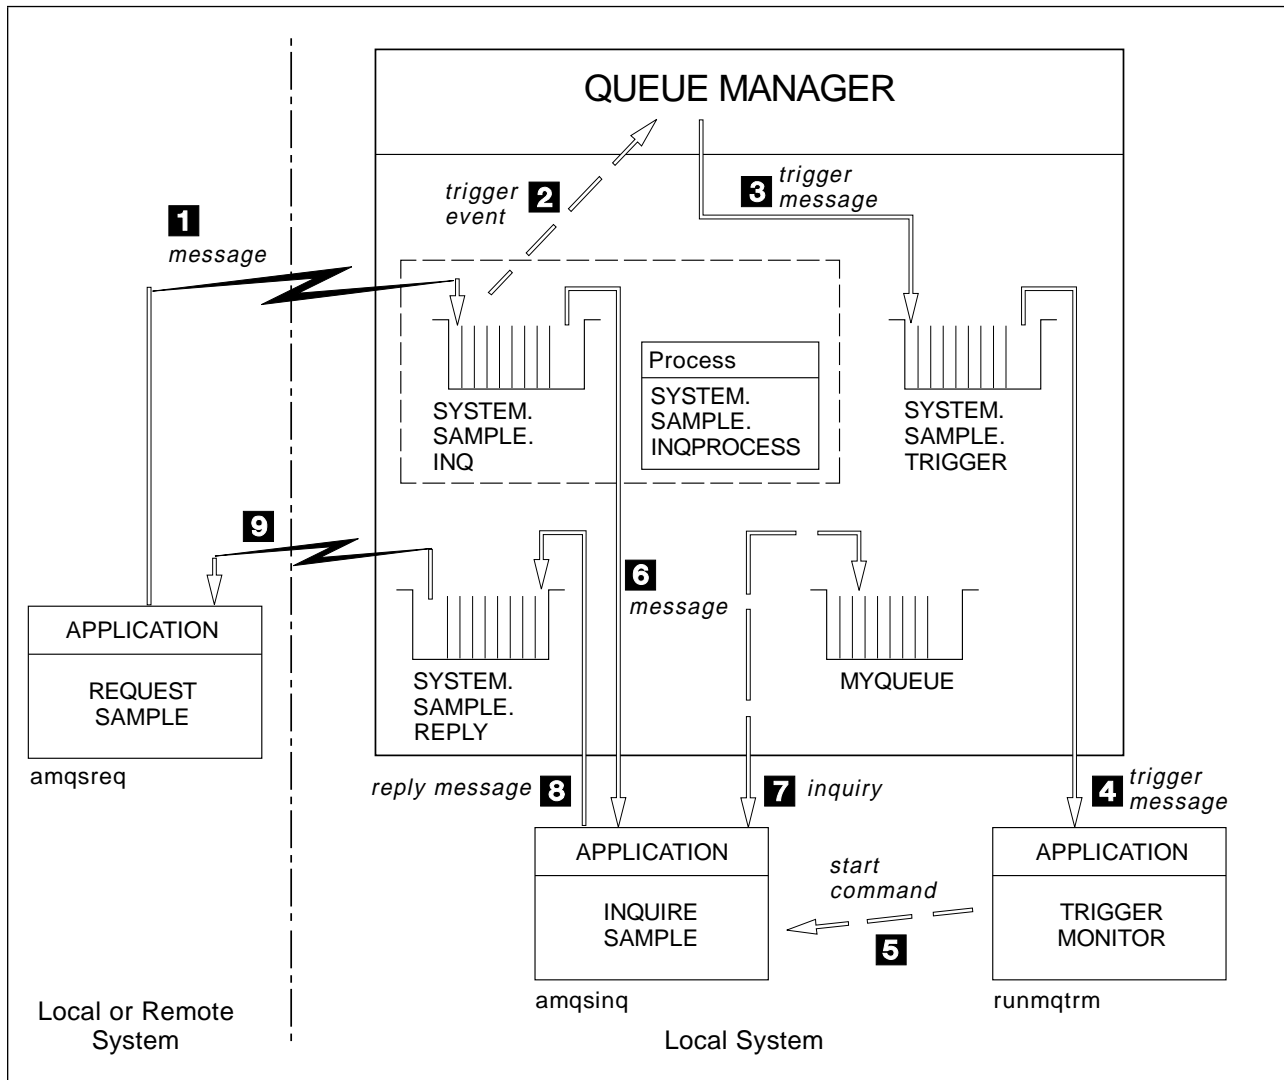
Figure 33. Request and Inquire samples using triggering

In Figure 33 the Request sample puts messages on to the target server queue, SYSTEM.SAMPLE.INQ, and the Inquire sample queries the queue, MYQUEUE. Alternatively, you can use one of the sample queues defined when you ran amqscos0.tst, or any other queue you have defined, for the Inquire sample.

Note: The numbers in Figure 33 show the sequence of events.

To run the Request and Inquire samples, using triggering:

1. Check that the queues you want to use are defined. Run amqscos0.tst, to define the sample queues, and define a queue MYQUEUE.

2. Run the trigger monitor command RUNMQTRM:

   RUNMQTRM -m qmanagername -q SYSTEM.SAMPLE.TRIGGER

3. Run the request sample

   amqsreq SYSTEM.SAMPLE.INQ

   Note: The process object defines what is to be triggered. If the client and server are not running on the same platform, any processes started by the trigger monitor must define *ApplType*, otherwise the server takes its

default definitions (that is, the type of application that is normally associated with the server machine) and causes a failure.

For example, if the trigger monitor is running on a Windows NT client and wants to send a request to an OS/2 server, MQAT_WINDOWS_NT must be defined otherwise OS/2 uses its default definitions (that is, MQAT_OS2) and the process fails.

For a list of application types, see the *MQSeries Application Programming Reference* manual.

4. Enter the name of the queue you want the Inquire sample to use:

   MYQUEUE

5. Enter a blank line (to end the Request program).

6. The request sample will then display a message, containing the data the Inquire program obtained from MYQUEUE.

If you wish, you can use more than one queue. In this case, you enter the names of the other queues at step 4.

For more information on triggering see Chapter 14, "Starting MQSeries applications using triggers" on page 197.

AS/400

To try the samples using triggering on AS/400, start the sample trigger server, AMQSERV4, in one job, then start AMQSREQ4 in another. This means that the trigger server is ready when the Request sample program sends a message.

Notes:

1. The sample definitions created by AMQSAMP4 cause the C versions of the samples to be triggered. If you want to trigger the COBOL versions, you must change the process definitions SYSTEM.SAMPLE.ECHOPROCESS, SYSTEM.SAMPLE.INQPROCESS and SYSTEM.SAMPLE.SETPROCESS; you can use the CHGMQMPRC command (described in the *MQSeries for AS/400 Administration Guide*) to do this, or edit and run your own version of AMQSAMP4.

2. Source code for AMQSERV4 is supplied for the C language only. However, a compiled version (that you can use with the COBOL samples) is supplied in library QMQM.

You could put your request messages on these sample server queues:

- SYSTEM.SAMPLE.ECHO (for the Echo sample programs)
- SYSTEM.SAMPLE.INQ (for the Inquire sample programs)
- SYSTEM.SAMPLE.SET (for the Set sample programs)

A flow chart for the SYSTEM.SAMPLE.ECHO program is shown in Figure 34 on page 346. Using the example data file the command to issue the C program request to this server is:

    CALL PGM(QMQMSAMP/AMQSREQ4) PARM('QMQMSAMP/AMQSDATA(ECHO)')

Note: This sample queue has a trigger type of FIRST, so if there are already messages on the queue before you run the Request sample, server applications are not triggered by the messages you send.

If you want to attempt further examples, you can try the following variations:

- Use AMQSTRG4 instead of AMQSERV4 to submit the job instead, but potential job submission delays could make it less easy to follow what is happening.

- Run the SYSTEM.SAMPLE.INQUIRE and SYSTEM.SAMPLE.SET sample programs. Using the example data file the commands to issue the C program requests to these servers are, respectively:

```
CALL PGM(QMQMSAMP/AMQSREQ4) PARM('QMQMSAMP/AMQSDATA(INQ)')
CALL PGM(QMQMSAMP/AMQSREQ4) PARM('QMQMSAMP/AMQSDATA(SET)')
```

These sample queues also have a trigger type of FIRST.

Design of the Request sample program

The program opens the target server queue so that it can put messages. It uses the MQOPEN call with the MQOO_OUTPUT option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

The program then opens the reply-to queue called SYSTEM.SAMPLE.REPLY so that it can get reply messages. For this, the program uses the MQOPEN call with the MQOO_INPUT_EXCLUSIVE option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

For each line of input, the program then reads the text into a buffer and uses the MQPUT call to create a request message containing the text of that line. On this call the program uses the MQRO_EXCEPTION_WITH_DATA report option to request that any report messages sent about the request message will include the first 100 bytes of the message data. The program continues until either it reaches the end of the input or the MQPUT call fails.

The program then uses the MQGET call to remove reply messages from the queue, and displays the data contained in the replies. The MQGET call uses the MQGMO_WAIT, MQGMO_CONVERT, and MQGMO_ACCEPT_TRUNCATED options. The *WaitInterval* is 5 minutes in the COBOL version, and 1 minute in the C version, for the first reply (to allow time for a server application to be triggered), and 15 seconds for subsequent replies. The program waits for these periods if there is no message on the queue. If no message arrives before this interval expires, the call fails and returns the MQRC_NO_MSG_AVAILABLE reason code. The call also uses the MQGMO_ACCEPT_TRUNCATED_MSG option, so messages longer than the declared buffer size are truncated.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The program continues until either the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code or the MQGET call fails. If the call fails, the program displays an error message that contains the reason code.

The program then closes both the target server queue and the reply-to queue using the MQCLOSE call. Figure 26 shows the changes to the Echo sample program that are necessary to run the Inquire and Set sample programs on AS/400.

Note: The details for the Echo sample program are included as a reference.

Table 26. Client/server sample program details			
Program name	**AMQSAMP/ AMQSDATA data file**	**SYSTEM/SAMPLE queue**	**Program started**
Echo	ECHO	ECHO	AMQSECHA
Inquire	INQ	INQ	AMQSINQA
Set	SET	SET	AMQSSETA
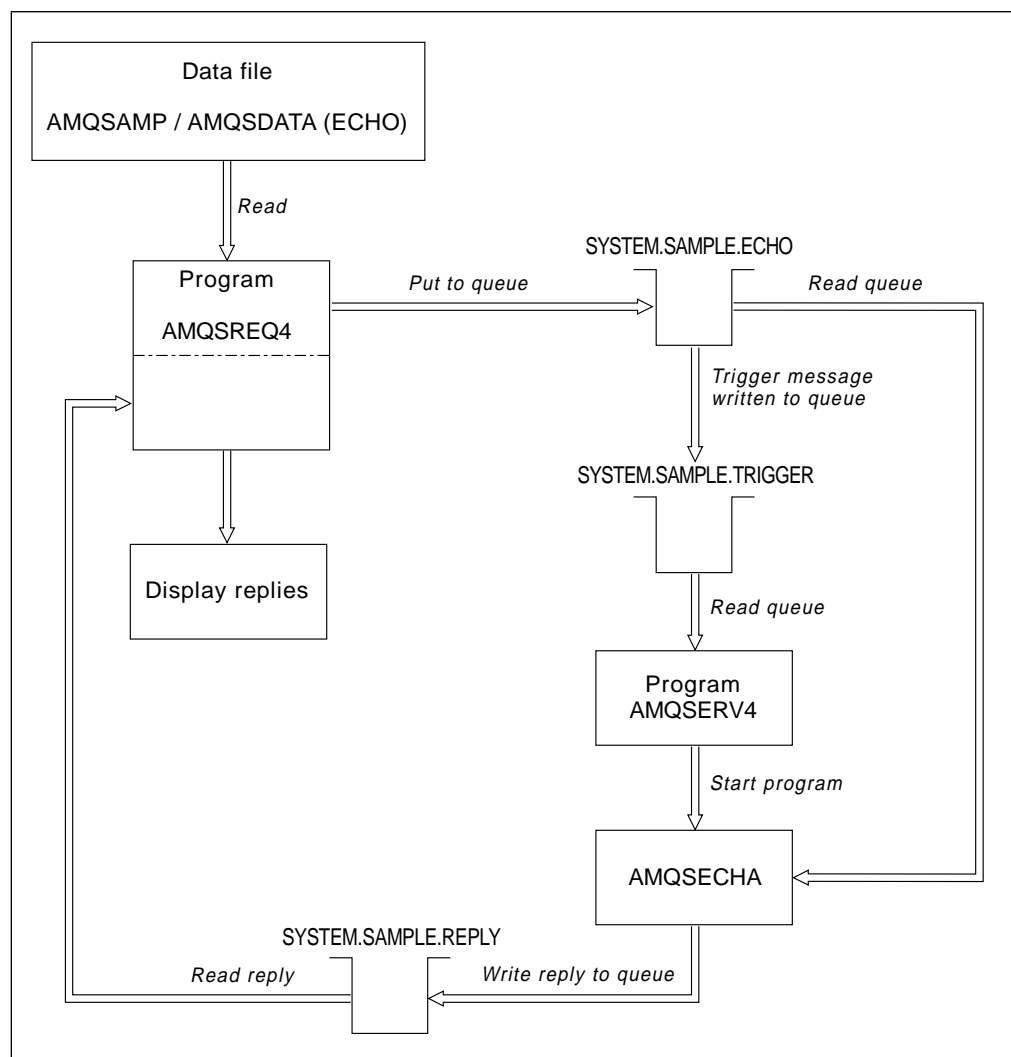
Figure 34. Sample Client/Server (Echo) program flowchart

The Inquire sample programs

The Inquire sample programs inquire about some of the attributes of a queue using the MQINQ call. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

These programs are intended to run as triggered programs, so their only input is an MQTMC2 (trigger message) structure for OS/2, Windows NT, Digital OpenVMS, and UNIX, and an MQTMC structure for AS/400. These structures contain the name of a target queue whose attributes are to be inquired. The C version also uses the queue manager name. The COBOL version uses the default queue manager.

For the triggering process to work, you must ensure that the Inquire sample program you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.INQ. To do this, specify the name of the Inquire sample program you want to use in the *ApplicId* field of the process definition SYSTEM.SAMPLE.INQPROCESS. For AS/400, you can use the CHGMQMPRC command described in the *MQSeries for AS/400 Administration Guide* for this. The sample queue has a trigger type of FIRST; if there are already messages on the queue before you run the request sample, the inquire sample is not triggered by the messages you send.

When you have set the definition correctly:

- For OS/2, UNIX systems, Digital OpenVMS, and Windows NT, start the **runmqtrm** program in one session, then start the amqsreq program in another.

- For AS/400, start the AMQSERV4 program in one session, then start the AMQSREQ4 program in another. You could use AMQSTRG4 instead of AMQSERV4, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send request messages, each containing just a queue name, to queue SYSTEM.SAMPLE.INQ. For each request message, the Inquire sample programs send a reply message containing information about the queue specified in the request message. The replies are sent to the reply-to queue specified in the request message.

On AS/400, if the sample input file member QMQMSAMP.AMQSDATA(INQ) is used, the last queue named does not exist, so the sample returns a report message with a reason code for the failure.

Design of the Inquire sample program

The program opens the queue named in the trigger message structure it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each request message removed from the request queue, the program reads the name of the queue (which we will call the *target queue*) contained in the data and opens that queue using the MQOPEN call with the MQOO_INQ option. The program then uses the MQINQ call to inquire about the values of the *InhibitGet*, *CurrentQDepth*, and *OpenInputCount* attributes of the target queue.

If the MQINQ call is successful, the program uses the MQPUT1 call to put a reply message on the reply-to queue. This message contains the values of the 3 attributes.

If the MQOPEN or MQINQ call is unsuccessful, the program uses the MQPUT1 call to put a report message on the reply-to queue. In the *Feedback* field of the message descriptor of this report message is the reason code returned by either the MQOPEN or MQINQ call, depending on which one failed.

After the MQINQ call, the program closes the target queue using the MQCLOSE call.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The Set sample programs

The Set sample programs inhibit put operations on a queue by using the MQSET call to change the queue's *InhibitPut* attribute. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

The programs are intended to run as triggered programs, so their only input is an MQTMC2 (trigger message) structure that contains the name of a target queue whose attributes are to be inquired. The C version also uses the queue manager name. The COBOL version uses the default queue manager.

For the triggering process to work, you must ensure that the Set sample program you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.SET. To do this, specify the name of the Set sample program you want to use in the *ApplicId* field of the process definition SYSTEM.SAMPLE.SETPROCESS. The sample queue has a trigger type of FIRST; if there are already messages on the queue before you run the Request sample, the Set sample is not triggered by the messages you send.

When you have set the definition correctly:

- For OS/2, UNIX systems, Digital OpenVMS, and Windows NT, start the **runmqtrm** program in one session, then start the amqsreq program in another.

- For AS/400, start the AMQSERV4 program in one session, then start the AMQSREQ4 program in another. You could use AMQSTRG4 instead of AMQSERV4, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send request messages, each containing just a queue name, to queue SYSTEM.SAMPLE.SET. For each request message, the Set sample programs send a reply message containing a confirmation that put operations have been inhibited on the specified queue. The replies are sent to the reply-to queue specified in the request message.

On AS/400, if the sample input file member QMQMSAMP.AMQSDATA(SET) is used, one queue, SYSTEM.SAMPLE.LOCAL has put inhibited.

Design of the Set sample program

The program opens the queue named in the trigger message structure it was passed when it started. (For clarity, we will call this the request queue.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each request message removed from the request queue, the program reads the name of the queue (which we will call the target queue) contained in the data and opens that queue using the MQOPEN call with the MQOO_SET option. The program then uses the MQSET call to set the value of the *InhibitPut* attribute of the target queue to MQQA_PUT_INHIBITED.

If the MQSET call is successful, the program uses the MQPUT1 call to put a reply message on the reply-to queue. This message contains the string `PUT inhibited`.

If the MQOPEN or MQSET call is unsuccessful, the program uses the MQPUT1 call to put a `report` message on the reply-to queue. In the *Feedback* field of the message descriptor of this report message is the reason code returned by either the MQOPEN or MQSET call, depending on which one failed.

After the MQSET call, the program closes the target queue using the MQCLOSE call.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The Echo sample programs

The Echo sample programs echo a message from a message queue to the reply queue. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

The programs are intended to run as triggered programs.

On OS/2, UNIX systems, and Windows NT, their only input is an MQTMC2 (trigger message) structure that contains the name of a target queue and the queue manager. The COBOL version uses the default queue manager.

On AS/400, for the triggering process to work, you must ensure that the Echo sample program you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.ECHO. To do this, specify the name of the Echo sample program you want to use in the *ApplId* field of the process definition SYSTEM.SAMPLE.ECHOPROCESS. (For this, you can use the CHGMQMPRC command, described in the *MQSeries for AS/400 Administration Guide*.) The sample queue has a trigger type of FIRST, so if there are already messages on the queue before you run the Request sample, the Echo sample is not triggered by the messages you send.

When you have set the definition correctly, first start AMQSERV4 in one job, then start AMQSREQ4 in another. You could use AMQSTRG4 instead of AMQSERV4, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send messages to queue SYSTEM.SAMPLE.ECHO. The Echo sample programs send a reply message containing the data in the request message to the reply-to queue specified in the request message.

Design of the Echo sample programs

The program opens the queue named in the trigger message structure it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG, MQGMO_CONVERT, and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each line of input, the program then reads the text into a buffer and uses the MQPUT1 call to put a request message, containing the text of that line, on to the reply-to queue.

If the MQGET call fails, the program puts a report message on the reply-to queue, setting the *Feedback* field of the message descriptor to the reason code returned by the MQGET.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

On AS/400, the program can also respond to messages sent to the queue from platforms other than MQSeries for AS/400, although no sample is supplied for this situation. To make the ECHO program work, you:

- Write a program, correctly specifying the *Format*, *Encoding*, and *CCSID* parameters, to send text request messages.

  The ECHO program requests the queue manager to perform message data conversion, if this is needed.

- Specify CONVERT(*YES) on the MQSeries for AS/400 sending channel, if the program you have written does not provide similar conversion for the reply.

The Data-Conversion sample program

The data-conversion sample program is a skeleton of a data conversion exit routine. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

Design of the data-conversion sample

Each data-conversion exit routine converts a single named message format. This skeleton is intended as a wrapper for code fragments generated by the data-conversion exit generation utility program.

The utility produces one code fragment for each data structure; several such structures make up a format, so several code fragments are added to this skeleton to produce a routine to do data conversion of the entire format.

The program then checks whether the conversion is a success or failure, and returns the values required to the caller.

The Triggering sample programs

The function provided in the triggering sample is a subset of that provided in the trigger monitor in the **runmqtrm** program. See "Features demonstrated in the sample programs" on page 313 for the names of these programs.

Running the amqstrg0.c, amqstrg, and amqstrgc samples

The program takes 2 parameters:

1. The name of the initiation queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it will connect to the default one. A sample initiation queue will have been defined when you ran amqscos0.tst. the name of that queue is SYSTEM.SAMPLE.TRIGGER, and you can use it when you run this program.

Note: The function in this sample is a subset of the full triggering function that is supplied in the **runmqtrm** program.

Running the AMQSTRG4 sample

This is a trigger monitor for the AS/400 environment. It submits an AS/400 job for the application to be started, but this means there is a processing overhead associated with each trigger message.

AMQSTRG4 takes one parameter: the name of the initiation queue it is to serve. AMQSAMP4 defines a sample initiation queue, SYSTEM.SAMPLE.TRIGGER, that you can use when you try the sample programs.

Using the example trigger queue the command to issue is:

```
CALL PGM(QMQM/AMQSTRG4) PARM('SYSTEM.SAMPLE.TRIGGER')
```

Design of the triggering sample

The triggering sample program opens the initiation queue using the MQOPEN call with the MQOO_INPUT_AS_Q_DEF option. It gets messages from the initiation queue using the MQGET call with the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, specifying an unlimited wait interval. The program clears the *MsgId* and *CorrelId* fields before each MQGET call to get messages in sequence.

When it has retrieved a message from the initiation queue, the program tests the message:

- It checks the size of the message to make sure it is the same size as an MQTM structure.

- It checks the *ApplType* field to make sure it contains the value MQAT_UNIX.

If either of these tests fail, the program displays a warning.

For valid trigger messages, the triggering sample copies data from these fields: *ApplicId*, *EnvrData*, *Version*, and *ApplType*. The last two of these fields are numeric, so the program creates character replacements to use in an MQTMC2 structure for OS/2, UNIX, and Windows NT, and in an MQTMC structure for AS/400.

The triggering sample issues a start command to the application specified in the *ApplicId* field of the trigger message, and passes an MQTMC2 or MQTMC (a character version of the trigger message) structure. In OS/2, UNIX systems, and Windows NT, the *EnvData* field is used as an extension to the invoking command string. In AS/400, it is used as job submission parameters, for example, the job priority.

Finally, the program closes the initiation queue.

Running the AMQSERV4 sample

This is a trigger server for the AS/400 environment. For each trigger message, this server runs the start command in its own job to start the specified application. The trigger server can call CICS transactions.

AMQSERV4 takes one parameter: the name of the initiation queue it is to serve. AMQSAMP4 defines a sample initiation queue, SYSTEM.SAMPLE.TRIGGER, that you can use when you try the sample programs.

Using the example trigger queue the command to issue is:

```
CALL PGM(QMQM/AMQSERV4) PARM('SYSTEM.SAMPLE.TRIGGER')
```

Design of the trigger server

The design of the trigger server is similar to that of the trigger monitor, except the trigger server:

- Allows MQAT_CICS as well as MQAT_OS400 applications
- Calls AS/400 applications in its own job (or uses STRCICSUSR to start CICS applications) rather than submitting an AS/400 job
- For CICS applications, substitutes the *EnvData*, for example, to specify the CICS region, from the trigger message in the STRCICSUSR command
- Opens the initiation queue for shared input, so many trigger servers can run at the same time

Note: Programs started by AMQSERV4 must not use the MQDISC call because this will stop the trigger server. If programs started by AMQSERV4 use the MQCONN call, they will get the MQRC_ALREADY_CONNECTED reason code.

Ending the triggering sample programs on AS/400

A trigger monitor program can be ended by the sysrequest option 2 (ENDRQS) or by inhibiting gets from the trigger queue. If the sample trigger queue is used the command is:

```
CHGMQMQ QNAME('SYSTEM.SAMPLE.TRIGGER') GETENBL(*NO)
```

Note: To start triggering again on this queue, you ***must*** enter the command:

```
CHGMQMQ QNAME('SYSTEM.SAMPLE.TRIGGER') GETENBL(*YES)
```

Running the samples using remote queues

You can demonstrate remote queuing by running the samples on connected queue managers.

Program amqscos0.tst provides a local definition of a remote queue (SYSTEM.SAMPLE.REMOTE) that uses a remote queue manager named OTHER. To use this sample definition, change OTHER to the name of the second queue manager you want to use. You must also set up a message channel between your two queue managers; for information on how to do this, see the MQSeries Intercommunication book.

The Request sample programs put their own local queue manager name in the *ReplyToQMgr* field of messages they send. The Inquire and Set samples send reply messages to the queue and message queue manager named in the *ReplyToQ* and *ReplyToQMgr* fields of the request messages they process.

Database coordination samples

Two samples are provided which demonstrate how MQSeries can coordinate both MQSeries updates and database updates within the same unit of work:

1. AMQSXAS0 (in C) or AMQ0XAS0 (in COBOL), which updates a single database within an MQSeries unit of work.

2. AMQSXAG0 (in C) or AMQ0XAG0 (in COBOL), AMQSXAB0 (in C) or AMQ0XAB0 (in COBOL), and AMQSXAF0 (in C) or AMQ0XAF0 (in COBOL), which together update two databases within an MQSeries unit of work, showing how multiple databases can be accessed.

   These samples are provided to show the use of the MQBEGIN call, mixed SQL and MQSeries calls, and where and when to connect to a database.

Figure 35 on page 355 shows how the samples provided are used to update databases:
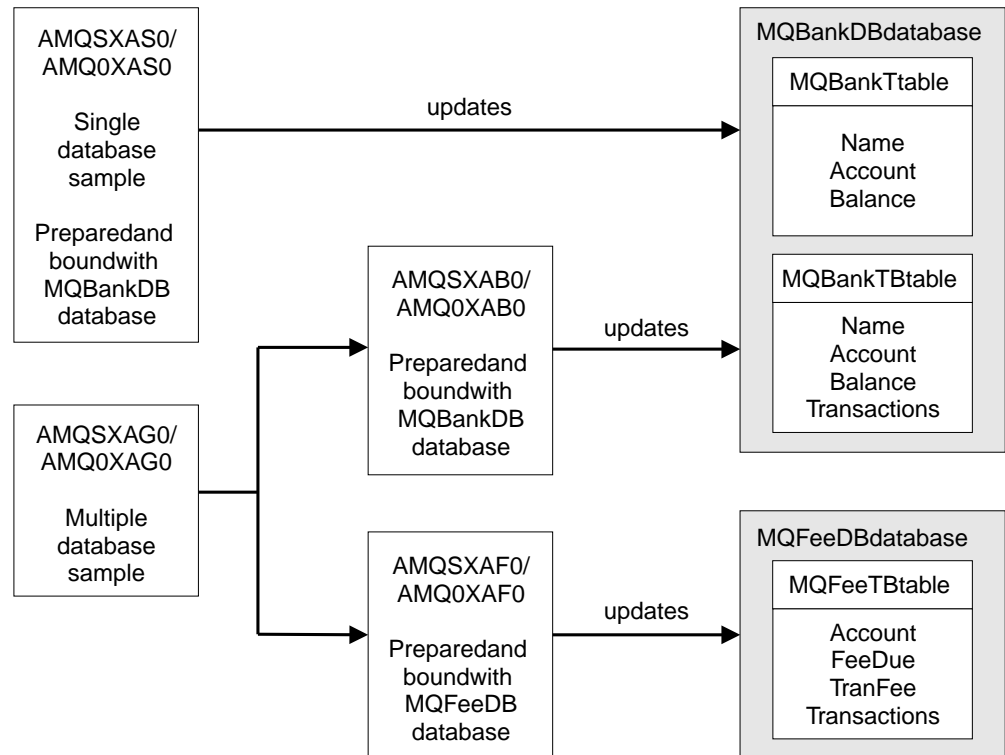
Figure 35. The database coordination samples

The programs read a message from a queue (under syncpoint), then, using the information in the message, obtain the relevant information from the database and update it. The new status of the database is then printed.

The program logic is as follows:

1. Use name of input queue from program argument

2. Connect to default queue manager (or optionally supplied name in C) using MQCONN

3. Open queue (using MQOPEN) for input while no failures

4. Start a unit of work using MQBEGIN

5. Get next message (using MQGET) from queue under syncpoint

6. Get information from databases

7. Update information from databases

8. Commit changes using MQCMIT

9. Print updated information (no message available counts as failure, and loop ends)

10. Close queue using MQCLOSE

11. Disconnect from queue using MQDISC

SQL cursors are used in the samples, so that reads from the databases (that is, multiple instances) are locked whilst a message is being processed, thus multiple instances of these programs can be run simultaneously. The cursors are explicitly opened, but implicitly closed by the MQCMIT call.

Database coordination samples

The single database sample (AMQSXAS0 or AMQ0XAS0) has no SQL CONNECT statements and the connection to the database is implicitly made by MQSeries with the MQBEGIN call. The multiple database sample (AMQSXAG0 or AMQ0XAG0, AMQSXAB0 or AMQ0XAB0, and AMQSXAF0 or AMQ0XAF0) has SQL CONNECT statements, as some database products allow only one active connection. If this is not the case for your database product, or if you are accessing a single database in multiple database products, the SQL CONNECT statements can be removed.

The samples are prepared with IBM's DB2 database product, so they may need some modification to work with other database products.

The SQL error checking uses routines in UTIL.C and CHECKERR.CBL supplied by DB2. These must be compiled or replaced before compiling and linking.

Note: If you are using the Micro Focus COBOL source CHECKERR.MFC for SQL error checking, you must change the program ID to uppercase, that is CHECKERR, for AMQ0XAS0 to link correctly.

Creating the databases and tables

The databases and tables must be created before the samples can be compiled. To create the databases, use the normal method for your database product, for example:

```
DB2 CREATE DB MQBankDB
DB2 CREATE DB MQFeeDB
```

Create the tables using SQL statements as follows:

In C:

```
EXEC SQL CREATE TABLE MQBankT(Name         VARCHAR(40) NOT NULL,
                              Account      INTEGER     NOT NULL,
                              Balance      INTEGER     NOT NULL,
                              PRIMARY KEY (Account));

EXEC SQL CREATE TABLE MQBankTB(Name         VARCHAR(40) NOT NULL,
                               Account      INTEGER     NOT NULL,
                               Balance      INTEGER     NOT NULL,
                               Transactions INTEGER,
                               PRIMARY KEY (Account));

EXEC SQL CREATE TABLE MQFeeTB(Account      INTEGER     NOT NULL,
                              FeeDue       INTEGER     NOT NULL,
                              TranFee      INTEGER     NOT NULL,
                              Transactions INTEGER,
                              PRIMARY KEY (Account));
```

In COBOL:

```
 EXEC SQL CREATE TABLE
   MQBankT(Name       VARCHAR(40) NOT NULL,
           Account    INTEGER     NOT NULL,
           Balance    INTEGER     NOT NULL,
           PRIMARY KEY (Account))
   END-EXEC.

EXEC SQL CREATE TABLE
  MQBankTB(Name       VARCHAR(40) NOT NULL,
           Account    INTEGER     NOT NULL,
           Balance    INTEGER     NOT NULL,
           Transactions INTEGER,
           PRIMARY KEY (Account))
  END-EXEC.

EXEC SQL CREATE TABLE
  MQFeeTB(Account     INTEGER     NOT NULL,
          FeeDue      INTEGER     NOT NULL,
          TranFee     INTEGER     NOT NULL,
          Transactions INTEGER,
          PRIMARY KEY (Account))
  END-EXEC.
```

Fill in the tables using SQL statements as follows:

```
EXEC SQL INSERT INTO MQBankT VALUES ('Mr Fred Bloggs',1,0);
EXEC SQL INSERT INTO MQBankT VALUES ('Mrs S Smith',2,0);
EXEC SQL INSERT INTO MQBankT VALUES ('Ms Mary Brown',3,0);
⋮
EXEC SQL INSERT INTO MQBankTB VALUES ('Mr Fred Bloggs',1,0,0);
EXEC SQL INSERT INTO MQBankTB VALUES ('Mrs S Smith',2,0,0);
EXEC SQL INSERT INTO MQBankTB VALUES ('Ms Mary Brown',3,0,0);
⋮
EXEC SQL INSERT INTO MQFeeTB VALUES (1,0,50,0);
EXEC SQL INSERT INTO MQFeeTB VALUES (2,0,50,0);
EXEC SQL INSERT INTO MQFeeTB VALUES (3,0,50,0);
⋮
```

Note: For COBOL, use the same SQL statements but add END_EXEC at the end of each line.

Preparing, compiling, and linking the samples

The .SQC files (in C) and .SQB files (in COBOL) must be prepared and bound against the appropriate database to produce the .C or .CBL files. To do this, use the normal method for your database product, for example:

In C:

```
SQLPREP AMQSXAS0.SQC MQBankDB
SQLPREP AMQSXAB0.SQC MQBankDB
SQLPREP AMQSXAF0.SQC MQFeeDB
```

In COBOL:

```
db2 connect to MQBankDB
db2 prep AMQ0XAS0.SQB bindfile target ibmcob
db2 bind AMQ0XAS0.BND
db2 connect reset

db2 connect to MQBankDB
db2 prep AMQ0XAB0.SQB bindfile target ibmcob
db2 bind AMQ0XAB0.BND
db2 connect reset

db2 connect to MQFeeDB
db2 prep AMQ0XAF0.SQB bindfile target ibmcob
db2 bind AMQ0XAF0.BND
db2 connect reset
```

Compile the files in the normal way to produce the object files. The samples must be linked against your database libraries as well as the appropriate MQSeries libraries. For more information about compiling and linking, seePart 3, "Building an MQSeries application" on page 251 and your database documentation.

The following lists extra points to consider when compiling and linking:

- If you are working on HP-UX and wish to access DB2, add the following link flags:

```
-E -u allow_unaligned_data_access
```

- If you are working on AIX and wish to access Oracle, use the xlc_r compiler and link to libmqm_r.a.

Running the samples

Before the samples can be run, the queue manager must be configured with the database product you are using. For information about how to do this, see the *MQSeries System Administration* book.

C samples

Messages must be in the following format to be read from a queue:

```
UPDATE Balance change=nnn WHERE Account=nnn
```

AMQSPUT can be used to put the messages on the queue.

The database coordination samples take two parameters:

1. Queue name (required)
2. Queue manager name (optional)

Assuming that you have created and configured a queue manager for the single database sample called singDBQM, with a queue called singDBQ, you increment Mr Fred Bloggs' account by 50 as follows:

```
AMQSPUT singDBQ singDBQM
```

Then key in the following message:

```
UPDATE Balance change=50 WHERE Account=1
```

You can put multiple messages on the queue.

```
AMQSXAS0 singDBQ singDBQM
```

The updated status of Mr Fred Bloggs's account is then printed.

Assuming that you have created and configured a queue manager for the multiple-database sample called multDBQM, with a queue called multDBQ, you decrement Ms Mary Brown's account by 75 as follows:

```
AMQSPUT multDBQ multDBQM
```

Then key in the following message:

```
UPDATE Balance change=-75 WHERE Account=3
```

You can put multiple messages on the queue.

```
AMQSXAG0 multDBQ multDBQM
```

The updated status of Ms Mary Brown's account is then printed.

COBOL samples

Messages must be in the following format to be read from a queue:

```
UPDATE Balance change=snnnnnnnn WHERE Account=nnnnnnnn
```

For simplicity, the `Balance change` must be a signed eight-character number and the `Account` must be an eight-character number.

The sample AMQSPUT can be used to put the messages on the queue.

The samples take no parameters and use the default queue manager. It can be configured to run only one of the samples at any time. Assuming that you have configured the default queue manager for the single database sample, with a queue called singDBQ, you increment Mr Fred Bloggs's account by 50 as follows:

```
AMQSPUT singDBQ
```

Then key in the following message:

```
UPDATE Balance change=+00000050 WHERE Account=00000001
```

You can put multiple messages on the queue.

```
AMQ0XAS0
```

Type in the name of the queue:

```
singDBQ
```

The updated status of Mr Fred Bloggs's account is then printed.

Assuming that you have configured the default queue manager for the multiple database sample, with a queue called multDBQ, you decrement Ms Mary Brown's account by 75 as follows:

```
AMQSPUT multDBQ
```

Then key in the following message:

```
UPDATE Balance change=-00000075 WHERE Account=00000003
```

You can put multiple messages on the queue.

```
AMQ0XAG0
```

Type in the name of the queue:

```
multDBQ
```

The updated status of Ms Mary Brown's account is then printed.

The CICS transaction sample

A sample CICS transaction program is provided, named amqscic0.ccs for source code and amqscic0 for the executable version. Transactions may be built using the standard CICS facilities. See Part 3, "Building an MQSeries application" on page 251 for details on the commands needed for your platform.

The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them on to the local queue, the name of which is contained in the transmission header of the message. Any failures will be sent to the queue SYSTEM.SAMPLE.CICS.DLQ.

Note: A sample MQSC script amqscic0.tst may be used to create these queues and sample input queues.

TUXEDO** samples

Before running these samples, you must build the server environment.

Building the server environment

It is assumed that you have a working TUXEDO environment.

To build the server environment for MQSeries for AIX:

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.

2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

   ```
   $ export CFLAGS="-I mqmtop/inc -I /<APPDIR> -L /mqmtop/lib"
   $ export LDOPTS="-lmqm -lmqmzse -lnet -insl -lsocket -lc -ldl"
   $ export FIELDTBLS=/mqmtop/samp/amqstxvx.flds
   $ export VIEWFILES=/<APPDIR>/amqstxvx.V
   $ export LIBPATH_PATH=$(TUXDIR)/lib:mqmtop/lib:/lib
   ```

3. Add the following to the TUXEDO file udataobj/RM

   ```
   MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
   /mqmtop/lib/libmqmxa.a /mqmtop/lib/libmqm.a
   ```

4. Run the commands:

```
$ mkfldhdr    /mqmtop/samp/amqstxvx.flds
$ viewc       /mqmtop/samp/amqstxvx.v
$ buildtms    -o MQXA -r MQSeries_XA_RMI
$ buildserver -o MQSERV1 -f /mqmtop/samp/amqstxsx.c \
        -f /mqmtop/lib/libmqm.a \
        -r MQSeries_XA_RMI -s MPUT1:MPUT \
        -s MGET1:MGET \
        -v -bshm
$ buildserver -o MQSERV2 -f /mqmtop/samp/amqstxsx.c \
        -f /mqmtop/lib/libmqm.a \
        -r MQSeries_XA_RMI -s MPUT2:MPUT
        -s MGET2:MGET \
        -v -bshm
$ buildclient -o doputs  -f /mqmtop/samp/amqstxpx.c \
        -f /mqmtop/lib/libmqm.a
$ buildclient -o dogets  -f /mqmtop/samp/amqstxgx.c\
        -f /mqmtop/lib/libmqm.a
```

5. Edit ubbstxcx.cfg (see Figure 36 on page 367), and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf    -y mqmtop/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt then appears. At this prompt, enter:

```
> crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ strmqm
```

8. Start Tuxedo:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

To build the server environment for MQSeries for AT&T GIS UNIX and MQSeries for Sun Solaris:

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.

2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-I /<APPDIR>"
$ export FIELDTBLS=amqstxvx.flds
$ export VIEWFILES=amqstxvx.V
$ export SHLIB_PATH=$TUXDIR/lib:mqmtop/lib:lib
$ export LD_LIBRARY_PATH=$(TUXDIR)/lib:mqmtop/lib:/lib
```

3. Add the following to the TUXEDO file udataobj/RM (RM must include /opt/mqm/lib/libmqmcs and /opt/mqm/lib/libmqmzse).

Note: The \ characters should not be entered into the file; they are line continuations.

```
                    MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
                    /mqmtop/lib/libmqmxa.a /mqmtop/lib/libmqm.so \
                    /opt/tuxedo/lib/libtux.so /mqmtop/lib/libmqm.so  \
                    /mqmtop/lib/libmqmzse.so
```

4. Run the commands:

```
        $ mkfldhdr    amqstxvx.flds
        $ viewc       amqstxvx.v
        $ buildtms    -o MQXA -r MQSeries_XA_RMI
        $ buildserver -o MQSERV1 -f amqstxsx.c \
                -f /mqmtop/lib/libmqm.so \
                -r MQSeries_XA_RMI -s MPUT1:MPUT \
                -s MGET1:MGET \
                -v -bshm
                -l -ldl
        $ buildserver -o MQSERV2 -f amqstxsx.c \
                -f /mqmtop/lib/libmqm.so \
                -r MQSeries_XA_RMI -s MPUT2:MPUT \
                -s MGET2:MGET \
                -v -bshm
                -l -ldl
        $ buildclient -o doputs  -f amqstxpx.c \
                -f /mqmtop/lib/libmqm.so  \
                -f /mqmtop/lib/mqmzse.so  \
                -f /mqmtop/lib/mqmcs.so
        $ buildclient -o dogets  -f amqstxgx.c \
                -f /mqmtop/lib/libmqm.so
                -f /mqmtop/lib/mqmzse.so  \
                -f /mqmtop/lib/mqmcs.so
```

5. Edit ubbstxcx.cfg (see Figure 36 on page 367), and add details of the machine name, working directories, and Queue Manager as necessary:

```
        $ tmloadcf    -y ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
        $tmadmin -c
```

A prompt then appears. At this prompt, enter:

```
        > crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
        $ strmqm
```

8. Start Tuxedo:

```
        $ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

To build the server environment for MQSeries for HP-UX:

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.

2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-Aa -D_HPUX_SOURCE"
$ export LDOPTS="-lmqm"
$ export FIELDTBLS=/mqmtop/samp/amqstxvx.flds
$ export VIEWFILES=<APPDIR>/amqstxvx.V
$ export SHLIB_PATH=$SHLIB/lib:mqmtop/lib:lib
$ export LPATH=$TUXDIR/lib:mqmtop/lib:lib
```

3. Add the following to the TUXEDO file udataobj/RM

```
     MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
/mqmtop/lib/libmqmxa.a /mqmtop/lib/libmqm.sl  \
/opt/tuxedo/lib/libtux.sl
```

4. Run the commands:

```
$ mkfldhdr    /mqmtop/samp/amqstxvx.flds
$ viewc       /mqmtop/samp/amqstxvx.v
$ buildtms    -o MQXA -r MQSeries_XA_RMI
$ buildserver -o MQSERV1 -f /mqmtop/samp/amqstxsx.c \
        -f /mqmtop/lib/libmqm.sl \
        -r MQSeries_XA_RMI -s MPUT1:MPUT \
        -s MGET1:MGET \
        -v -bshm
$ buildserver -o MQSERV2 -f /mqmtop/samp/amqstxsx.c \
        -f /mqmtop/lib/libmqm.sl \
        -r MQSeries_XA_RMI -s MPUT2:MPUT \
        -s MGET2:MGET \
        -v -bshm
$ buildclient -o doputs  -f /mqmtop/samp/amqstxpx.c \
        -f /mqmtop/lib/libmqm.sl
$ buildclient -o dogets  -f /mqmtop/samp/amqstxgx.c \
        -f /mqmtop/lib/libmqm.sl
```

5. Edit ubbstxcx.cfg (see Figure 36 on page 367), and add details of the machine name, working directories, and Queue Manager as necessary:

```
$ tmloadcf    -y mqmtop/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt then appears. At this prompt, enter:

```
> crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ strmqm
```

8. Start Tuxedo:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

To build the server environment for MQSeries for SunOS

Note: Use the sys5 libraries and headers during compilation (for example, sys5 flag can be passed to the acc compiler).

1. Export the following environment variables where TUXDIR is the root directory for TUXEDO:

   ```
   $ export FIELDTBLS=amqstxvx.flds
   $ export VIEWFILES=amqstxvx.V
   $ export LD_LIBRARY_PATH=$(TUXDIR)/lib:usr/mqm/lib:lib
   ```

2. Add the following to the TUXEDO file udataobj/RM

   Note: The \ characters should not be entered into the file; they are line continuations.

   ```
   MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
   -L/usr/mqm/lib -lmqmxa -lmqm -L/usr/5lib -lc
   ```

3. Run the commands:

   ```
   $ mkfldhdr    amqstxvx.flds
   $ viewc       amqstxvx.v
   $ buildtms    -o MQXA -r MQSeries_XA_RMI
   $ buildserver -o MQSERV1 -f amqstxsx.c \
           -f "-L/usr/mqm/lib" -f "-lmqm" \
           -r MQSeries_XA_RMI -s MPUT1:MPUT \
           -s MGET1:MGET \
           -v -bshm
   $ buildserver -o MQSERV2 -f amqstxsx.c \
           -f "/usr/mqm/lib" -f "-lmqm" \
           -r MQSeries_XA_RMI -s MPUT2:MPUT \
           -s MGET2:MGET \
           -v -bshm
   $ buildclient -o doputs  -f amqstxpx.c \
           -f "-L/usr/mqm/lib" -f "-lmqm"
   $ buildclient -o dogets  -f amqstxgx.c \
           -f "-L/usr/mqm/lib" -f "-lmqm"
   ```

4. Edit ubbstxcx.cfg (see Figure 36 on page 367), and add details of the machine name, working directories, and Queue Manager as necessary:

   ```
   $ tmloadcf    -y ubbstxcx.cfg
   ```

5. Start the queue manager:

   ```
   $ strmqm
   ```

6. Start Tuxedo:

   ```
   $ tmboot -y
   ```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

To build the server environment for MQSeries for SINIX and DC/OSx

1. Export the following environment variables where TUXDIR is the root directory for TUXEDO:

   ```
   $ export CFLAGS="-lmqm -lmqmcs -lmqmzse -lmqmxa \
   -lnsl -lsocket -ldl -lmproc -lext"
   ```

   Note: For DC/OSx, add "-liconv" to the above.

   ```
   $ export FIELDTBLS=amqstxvx.flds
   $ export VIEWFILES=amqstxvx.V
   $ export VIEWDIR=The path to the directory where the views
   are held
   $ export TUXDIR=The path to the directory where TUXEDO
   is installed (/opt/tuxedo).
   $ export CFLAGS="-lmqm -lmqmcs -lmqmzse -lmqmxa \
   -lnsl -lsocket -ldl -lmproc -lext"
   ```

2. Add the following to the TUXEDO file udataobj/RM

   ```
       MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
     /mqmtop/lib/libmqmxa.so
     /mqmtop/lib/libmqm.so /mqmtop/lib/libmqmcs.s
   ```

3. Ensure that your LD_LIBRARY_PATH contains the path to the Tuxedo libraries (/opt/tuxedo/lib), and that it is exported.

4. Ensure that your PATH contains the path to the Tuxedo bin directory (/opt/tuxedo/bin), and that it is exported.

5. Run the commands:

   ```
   $ mkfldhdr    amqstxvx.flds
   $ viewc       amqstxvx.v
   ```

6. Alter the value of the CFLAGS variable:

   ```
   export CFLAGS="$CFLAGS -LDuMQRMIXASwitchDynamic -lmqmxa"
   ```

7. Run the commands:

   ```
   $ buildtms    -o MQXA -r MQSeries_XA_RMI
   $ buildserver -o MQSERV1 -f amqstxsx.c \
           -f /mqmtop/lib/libmqm.so i \
           -r MQSeries_XA_RMI -s MPUT1:MPUT \
           -s MGET1:MGET \
           -v -bshm
   $ buildserver -o MQSERV2 -f amqstxsx.c \
           -f /mqmtop/lib/libmqm.so \
           -r MQSeries_XA_RMI -s MPUT2:MPUT \
           -s MGET2:MGET \
           -v -bshm
   $ buildclient -o doputs  -f amqstxpx.c \
           -f /mqmtop/lib/libmqm.so
   $ buildclient -o dogets  -f amqstxgx.c \
           -f /mqmtop/lib/libmqm.so
   ```

8. Ensure that your NLS_PATH contains the path to the Tuxedo messages (/opt/tuxedo/locale/C/%N), and that it is exported.

9. Edit ubbstxcx.cfg (see Figure 36 on page 367), and add details of the machine name, working directories, and Queue Manager as necessary.

10. Set the environment variable TUXCONFIG to the value specified in the MACHINES section of the ubbstxcx.cfg file.

11. If you are using the Tuxedo main machine, run the following commands:

    ```
    tmadmin -c
    ```

    At the prompt (>), enter:

    ```
    crdl -z filename
    ```

    where *filename* is the path to the Tuxedo TLOG file.

12. Run the following command:

    ```
    $ tmloadcf    -y ubbstxcx.cfg
    ```

13. Start the queue manager:

    ```
    $ strmqm
    ```

14. Start Tuxedo:

    ```
    $ tmboot -y
    ```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

For further information on building the TUXEDO server environment, see the ReadMe file in the MQSeries sample directory, **mqmtop**.

```
*RESOURCES
IPCKEY           <IPCKey>

#Example:
#IPCKEY          123456

MASTER           <MachineName>
MAXACCESSERS     20
MAXSERVERS       20
MAXSERVICES      50
MODEL            SHM
LDBAL            N

*MACHINES
DEFAULT:
                 APPDIR="<WorkDirectory>"
                 TUXCONFIG="<WorkDirectory>/tuxconfig"
                 ROOTDIR="<RootDirectory>"

<MachineName>    LMID=<MachineName>
                 TLOGDEVICE="<WorkDirectory>/TLOG1"
                 TLOGNAME=TLOG

*GROUPS
GROUP1
       LMID=<MachineName>  GRPNO=1
       TMSNAME=MQXA
       OPENINFO="MQSeries_XA_RMI:MYQUEUEMANAGER"

*SERVERS
DEFAULT:
                 CLOPT="-A -- -m MYQUEUEMANAGER

MQSERV1          SRVGRP=GROUP1 SRVID=1
MQSERV2          SRVGRP=GROUP1 SRVID=2

*SERVICES
MPUT1
MGET1
MPUT2
MGET2
```

Figure 36. Example of ubbstxcx.cfg file for UNIX systems

Note: Other information that you need to add is identified by <> characters. In this file, the queue manager name has been changed to MYQUEUEMANAGER:

To build the server environment for MQSeries for Windows NT:
Note: Change the fields identified by <> in the following, to the directory paths:

<MQMDIR>   the directory path specified when MQSeries was installed, for example g:\mqm

<TUXDIR>   the directory path specified when TUXEDO was installed, for example f:\tuxedo

        <APPDIR>   the directory path to be used for the sample application, for example f:\tuxedo\apps\mqapp

To build the server environment and samples:

1. Create an application directory in which to build the sample application, for example:

   ```
   f:\tuxedo\apps\mqapp
   ```

2. Copy the following sample files from the MQSeries sample directory to the application directory:

   ```
   amqstxmn.mak
   amqstxen.env
   ubbstxcn.cfg
   ```

3. Edit each of these files to set the directory names and directory paths used on your installation.

4. Edit ubbstxcn.cfg (see Figure 37 on page 369) to add details of the machine name and the Queue Manager that you wish to connect to.

5. Add the following line to the TUXEDO file <TUXDIR>udataobj\rm

   ```
   MQSeries_XA_RMI;MQRMIXASwitchDynamic;
       <MQMDIR>\tools\lib\mqmtux.lib <MQMDIR>\tools\lib\mqm.lib
   ```

   where <MQMDIR> is replaced as above. Although shown here as two lines, the new entry must be one line in the file.

6. Set the following environment variables:

   ```
   TUXDIR=<TUXDIR>
   TUXCONFIG=<APPDIR>\tuxconfig
   FIELDTBLS=<MQMDIR>\tools\c\samples\amqstxvx.fld
   LANG=C
   ```

7. Create a TLOG device for TUXEDO. To do this, invoke `tmadmin -c`, and enter the command:

   ```
   crdl -z <APPDIR>\TLOG
   ```

   where *<APPDIR>* is replaced as above.

8. Set the current directory to <APPDIR>, and invoke the sample makefile (amqstxmn.mak) as an external project makefile. For example, with Microsoft Visual C++ Version 2.0, issue the command:

   ```
   msvc amqstxmn.mak
   ```

   Select **build** to build all the sample programs.

```
*RESOURCES
IPCKEY        99999
UID           0
GID           0
MAXACCESSERS  20
MAXSERVERS    20
MAXSERVICES   50
MASTER        SITE1
MODEL         SHM
LDBAL         N

*MACHINES
<MachineName> LMID=SITE1
              TUXDIR=    "f:\tuxedo"
              APPDIR=    "f:\tuxedo\apps\mqapp;g:\mqm\bin"
              ENVFILE=   "f:\tuxedo\apps\mqapp\amqstxen.env"
              TUXCONFIG= "f:\tuxedo\apps\mqapp\tuxconfig"
              ULOGPFX=   "f:\tuxedo\apps\mqapp\ULOG"
              TLOGDEVICE="f:\tuxedo\apps\mqapp\TLOG"
              TLOGNAME=TLOG
              TYPE="i386NT"
              UID=0
              GID=0

*GROUPS
GROUP1
        LMID=SITE1  GRPNO=1
        TMSNAME=MQXA
        OPENINFO="MQSeries_XA_RMI:MYQUEUEMANAGER"

*SERVERS
DEFAULT: CLOPT="-A -- -m MYQUEUEMANAGER"

MQSERV1     SRVGRP=GROUP1 SRVID=1
MQSERV2     SRVGRP=GROUP1 SRVID=2

*SERVICES
MPUT1
MGET1
MPUT2
MGET2
```

Figure 37. Example of ubbstxcn.cfg file for Windows NT

Note: The directory names and directory paths must be changed to match your installation. The queue manager name MYQUEUEMANAGER should also be changed to the name of the queue manager you wish to connect to. Other information that you need to add is identified by <> characters.

The sample ubbconfig file for MQSeries for Windows NT is listed in Figure 37 It is supplied as ubbstxcn.cfg in the MQSeries samples directory.

The sample makefile (see Figure 38 on page 370) supplied for MQSeries for Windows NT is called ubbstxmn.mak, and is held in the MQSeries samples directory.

```
TUXDIR  = f:\tuxedo
MQMDIR  = g:\mqm
APPDIR  = f:\tuxedo\apps\mqapp
MQMLIB  = $(MQMDIR)\tools\lib
MQMINC  = $(MQMDIR)\tools\c\include
MQMSAMP = $(MQMDIR)\tools\c\samples
INC = -f "-I$(MQMINC) -I$(APPDIR)"
DBG = -f "/Zi"

amqstx.exe:
 $(TUXDIR)\bin\mkfldhdr    -d$(APPDIR) $(MQMSAMP)\amqstxvx.fld
 $(TUXDIR)\bin\viewc       -d$(APPDIR) $(MQMSAMP)\amqstxvx.v
 $(TUXDIR)\bin\buildtms    -o MQXA -r MQSeries_XA_RMI
 $(TUXDIR)\bin\buildserver -o MQSERV1 -f $(MQMSAMP)\amqstxsx.c      \
                           -f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)    \
                           -r MQSeries_XA_RMI                       \
                           -s MPUT1:MPUT -s MGET1:MGET
 $(TUXDIR)\bin\buildserver -o MQSERV2 -f $(MQMSAMP)\amqstxsx.c        \
                           -f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)    \
                           -r MQSeries_XA_RMI                       \
                           -s MPUT2:MPUT -s MGET2:MGET
 $(TUXDIR)\bin\buildclient -o doputs -f $(MQMSAMP)\amqstxpx.c       \
                           -f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)
 $(TUXDIR)\bin\buildclient -o dogets -f $(MQMSAMP)\amqstxgx.c       \
                           -f $(MQMLIB)\mqm.lib $(INC) -v $(DBG)
 $(TUXDIR)\bin\tmloadcf    -y $(APPDIR)\ubbstxcn.cfg
```

Figure 38. Sample TUXEDO makefile for MQSeries for Windows NT

Server sample program for TUXEDO

This program is designed to run with the Put (amqstxpx.c) and the Get (amqstxgx.c) sample programs. The sample server program runs automatically when TUXEDO is started.

Note: You must start your queue manager *before* you start TUXEDO.

The sample server provides two TUXEDO services, MPUT1 and MGET1.

The MPUT1 service is driven by the PUT sample and uses MQPUT1 in syncpoint to put a message in a unit of work controlled by TUXEDO. It takes the parameters QName and Message Text, which are supplied by the PUT sample.

The MGET1 service opens and closes the queue each time it gets a message. It takes the parameters QName and Message Text, which are supplied by the GET sample.

Any error messages,reason codes, and status messages are written to the TUXEDO log file.
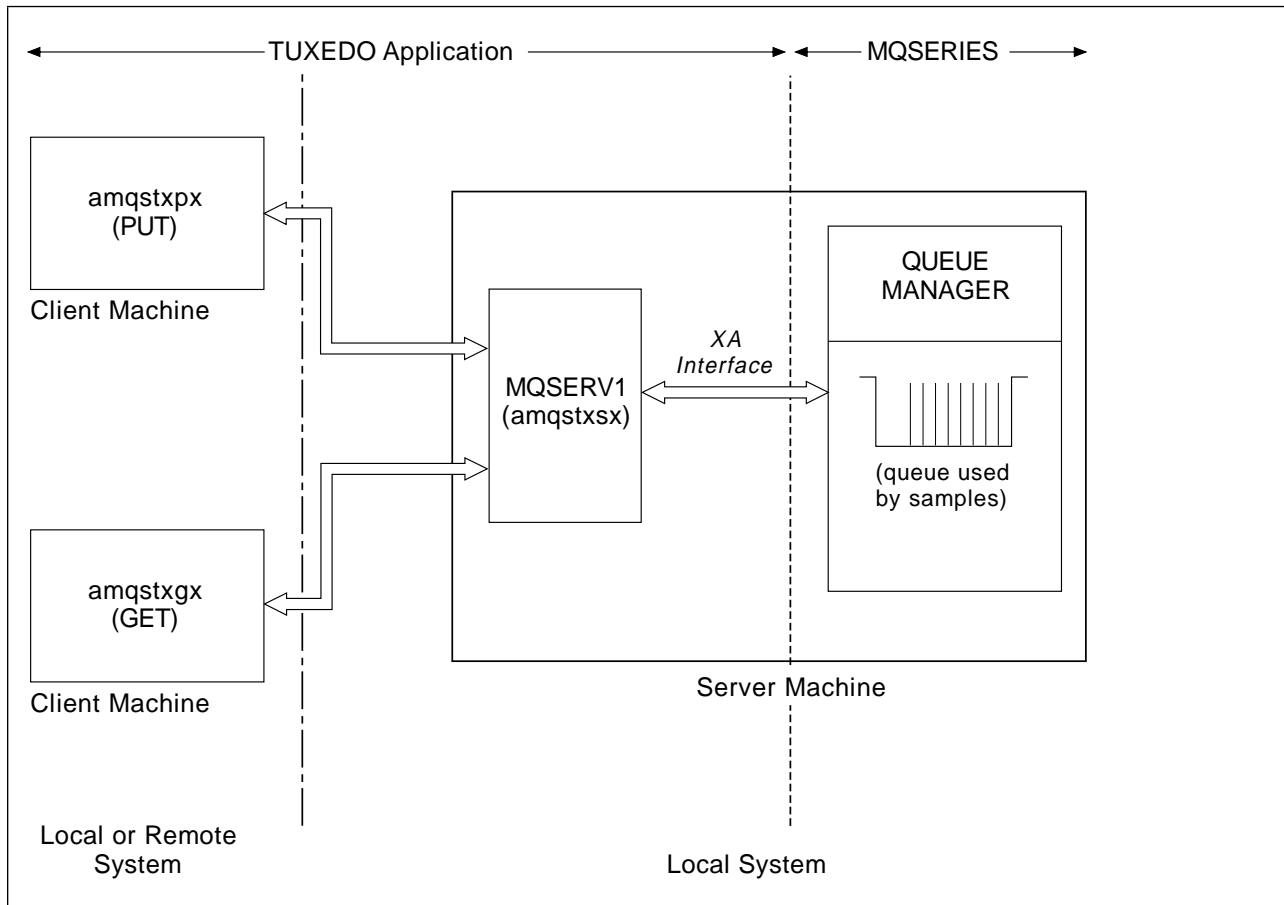
Figure 39. How TUXEDO samples work together

Put sample program for TUXEDO

This sample allows you to put a message on a queue multiple times, in batches, demonstrating syncpointing using TUXEDO as the resource manager. The sample server program amqstxsx must be running for the put sample to succeed - the server sample program makes the connection to the queue manager and uses the XA interface. To run the sample enter:

- doputs –n queuename –b batchsize –c trancount –t message

For example:

- doputs -n myqueue -b 5 -c 6 -t "Hello World"

This puts 30 messages on to the queue named myqueue, in 6 batches each with 5 messages in them. If there are any problems it will back a batch of messages out, otherwise it will commit them.

Any error messages are written to the TUXEDO log file and to stderr. Any reason codes are written to stderr.

Get sample for TUXEDO

This sample allows you to get messages from a queue in batches. The sample server program `amqstxsx` must be running for the put sample to succeed - the server sample program makes the connection to the queue manager and uses the XA interface. To run the sample enter:

- `dogets –n queuename –b batchsize –c trancount`

For example:

- `dogets -n myqueue -b 6 -c 4`

This takes 24 messages off the queue named `myqueue`, in 6 batches each with 4 messages in them. If you ran this after the put example, which put 30 messages on `myqueue`, you would now have only 6 messages on `myqueue`. Note that the number of batches and the batch size can vary between the putting of messages and the getting of them.

Any error messages are written to the TUXEDO log file and to stderr. Any reason codes are written to stderr.

Encina sample program

This program puts 10 messages to the queue, backing out the odd numbered messages and committing the even numbered messages. The message is a 4-byte number.

The queue used by this sample is the SYSTEM.DEFAULT.MODEL.QUEUE, so a temporary dynamic queue is created each time the program is run. You will need to run trace to see what happens when the program runs.

Building the AMQSXAE0.C sample

When compiling for a UNIX or OS/2 platform, ensure that the symbolic constant, `WIN32` is not defined. This constant is used in the preprocessor statements for processing specific to Windows NT:

```
#if defined(WIN32)
```

Compiling and linking on Windows NT

When compiling, specify the following options (in addition to those usually specified for an MQSeries application) to the C compiler:

```
-MD -DWIN32 -DDEC_DCE -Gz
```

The sample contains references to the Encina header files:

```
#include <tc/tc.h>
#include <tmxa/tmxa_status.h>
#include <tmxa/tmxa.h>
```

At compile time, also include the parent directory path name containing these files, using the compiler `-I` option with a value which names the directory. For example:

```
-Ic:\opt\encina\include
```

At link time, the directory path names containing the Encina and DCE library files must also be specified to the linker, by setting the LIB environment variable. For example:

```
SET LIB=C:\OPT\ENCINA\LIB;C:\OPT\DCE\LIB;%LIB%
```

When linking, specify the following library files:

- mqm.lib
- mqmenc.lib
- libEncServer.lib
- libEncina.lib
- msvcrt.lib
- pthreads.lib
- libdce.lib

Compiling and linking on Sun Solaris

Use the following invocation:

```
cc -I/opt/encina/include -c amqsxae0.c && cc -mt -o amqsxae0 amqsxae0.o \
   -L/opt/encina/lib -L/opt/mqm/lib -lmqm -lmqmcs_d -lmqmzse -lmqmxa     \
   -lsocket -lnsl -ldce -lthread -lEncServer -lEncina -lc -lm
```

Dead-letter queue handler sample

A sample dead-letter queue handler is provided, the name of the executable version is amqsdlq. If you want a dead-letter queue handler that is different to RUNMQDLQ, the source of the sample is available for you to use your base.

The sample is similar to the dead-letter handler provided within the product but trace and error reporting are different. There are two environment variables available to you:

ODQ_TRACE set to YES or yes to switch tracing on

ODQ_MSG   set to the name of the file containing error and information messages. The file provided is called amqsdlq.msg.

These need to be made known to your environment using either the **export** or **set** commands, depending on your platform; trace is turned off using the **unset** command.

You can modify the error message file, amqsdlq.msg, to suit your own requirements. The sample puts messages out to stdout, *not* to the MQSeries error log file.

The *System Management Guide* for your platform explains how the dead-letter handler works, and how you run it.

Lotus Notes and MQSeries link sample

The sample program is named amqslnk0. This program:

- Enables you to check that Lotus Notes and MQSeries have been set up correctly to allow the applications to communicate
- Provides an example of a macro and a link database entry

Lotus Notes link sample

To run the program, follow these steps:

1. Check that database and queues are available for use:

   - Start the queue manager.

   - Check the queues are available. If this is the first time the sample is to be run, or the queues no longer exist, run the supplied MQSC file, amqslnk0.tst.

   - Create a Mail-In database in the Lotus Notes workspace.

   - Where appropriate, check channels are active.

   - Ensure the sample link database is in the directory, C:NOTES.

2. Enter:

   ```
   start amqslnk0
   ```

3. Go to the Lotus Notes workspace.

4. Open the MQSeries sample database.

5. Compose the MQSeries sample.

   There is a description at the top of the form, with further information available under the help menu.

6. Enter the name of the Mail-In database, if you are not using the default.

7. Enter the data to be sent to MQSeries. This can be anything you like.

8. Click on the **Send** button.

9. Re-open the MQSeries sample database to see the reply data in the form - data received and time.

For details of the link database entry and the macro behind the button, see Chapter 16, "Lotus Notes" on page 239. For details on installation of the add-in task and configuration of your environment, see the *MQSeries System Administration* Guide for MQSeries for AIX, HP-UX, OS/2, Sun Solaris, and Windows NT; for other platforms, see the appropriate *System Management Guide*.

Chapter 31. Sample programs for MQSeries for MVS/ESA

This chapter describes the sample applications that are delivered with MQSeries for MVS/ESA. These samples demonstrate typical uses of the Message Queue Interface (MQI).

MQSeries for MVS/ESA also provides a sample API-crossing exit program, described in the "The API-crossing exit for MVS/ESA" on page 223, and sample data-conversion exits, described in Chapter 11, "Data-conversion exits" on page 157.

The sample applications are supplied in source form only. The source modules include pseudocode that describes the program logic. If you want to run the samples, follow the instructions given in:

- "Preparing and running sample applications for the batch environment" on page 380
- "Preparing the sample applications for the CICS for MVS/ESA environment" on page 383
- "Preparing the sample application for the IMS environment" on page 387

Notes:

1. The sample programs for CICS use features that are not supported by CICS/MVS. However, they can be used as a base when designing applications for this version of CICS.

2. Although some of the sample applications have basic panel-driven interfaces, they do not aim to demonstrate how to design the "look and feel" of your applications. For more information on how to design panel-driven interfaces for nonprogrammable terminals, see the *SAA Common User Access: Basic Interface Design Guide* (SC26-4583) and its addendum (GG22-9508). These provide guidelines to help you design applications that are consistent both within the application and across other applications.

Features demonstrated in the sample applications

This section summarizes the MQI features demonstrated in each of the sample applications, shows the programming languages that each sample is written in, and the environment in which each sample runs. Descriptions of the design of each sample are given in:

- "The Put samples" on page 388
- "The Get samples" on page 391
- "The Browse sample" on page 394
- "The Print Message sample" on page 396
- "The Queue Attributes sample" on page 400
- "The Mail Manager sample" on page 401
- "The Credit Check sample" on page 409
- "The Message Handler sample" on page 423

Put samples

The Put samples demonstrate how to put messages on a queue using the MQPUT call.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQPUT
- MQCLOSE
- MQDISC

The program is delivered in COBOL and C, and runs in the batch and CICS for MVS/ESA environment. See Table 29 on page 380 for the batch application and Table 34 on page 384 for the CICS application.

Get samples

The Get samples demonstrate how to get messages from a queue using the MQGET call.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQGET
- MQCLOSE
- MQDISC

The program is delivered in COBOL and C, and runs in the batch and CICS for MVS/ESA environment. See Table 29 on page 380 for the batch application and Table 34 on page 384 for the CICS application.

Browse sample

The Browse sample demonstrates how to browse a message, print it, then step through the messages on a queue.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQGET for browsing messages
- MQCLOSE
- MQDISC

The program is delivered in the COBOL, assembler, PL/I, and C languages. The application runs in the batch environment. See Table 30 on page 381 for the batch application.

Print Message sample

The Print Message sample demonstrates how to remove a message from a queue and print the data in the message, together with all the fields of its message descriptor. By removing comment characters from two lines in the source module, you can change the program so that it browses, rather than removes, the messages on a queue. This program can usefully be used for diagnosing problems with an application that is putting messages on a queue.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQGET for removing messages from a queue (with an option to browse)
- MQCLOSE
- MQDISC

The program is delivered in the C language. The application runs in the batch environment. See Table 31 on page 381 for the batch application.

Queue Attributes sample

The Queue Attributes sample demonstrates how to inquire about and set the values of MQSeries for MVS/ESA object attributes.

The application uses these MQI calls:

- MQOPEN
- MQINQ
- MQSET
- MQCLOSE

The program is delivered in the COBOL, assembler, and C languages. The application runs in the CICS for MVS/ESA environment. See Table 35 on page 385 for the CICS application.

Mail Manager sample

The Mail Manager sample demonstrates these techniques:

- Using alias queues
- Using a model queue to create a temporary dynamic queue
- Using reply-to queues
- Using syncpoints in the CICS for MVS/ESA and batch environments
- Sending commands to the system-command input queue
- Testing return codes
- Sending messages to remote queue managers, both by using a local definition of a remote queue and by putting messages directly on a named queue at a remote queue manager

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQPUT1
- MQGET
- MQINQ
- MQCMIT

- MQCLOSE
- MQDISC

Three versions of the application are provided:

- A CICS for MVS/ESA application written in COBOL
- A TSO application written in COBOL
- A TSO application written in C

The TSO applications use the MQSeries for MVS/ESA batch adapter and include some ISPF panels.

See Table 32 on page 382 for the TSO application, and Table 36 on page 385 for the CICS application.

Credit Check sample

The Credit Check sample is a suite of programs that demonstrates these techniques:

- Developing an application that runs in more than one environment
- Using a model queue to create a temporary dynamic queue
- Using a correlation identifier
- The setting and passing of context information
- Using message priority and persistence
- Starting programs by using triggering
- Using reply-to queues
- Using alias queues
- Using a dead-letter queue
- Using a namelist
- Testing return codes

The application uses these MQI calls:

- MQOPEN

- MQPUT

- MQPUT1

- MQGET for browsing and getting messages, using the wait and signal options, and for getting a specific message

- MQINQ

- MQSET

- MQCLOSE

The sample can run as a stand-alone CICS for MVS/ESA application. However, to demonstrate how to design a message queuing application that uses the facilities provided by both the CICS for MVS/ESA and IMS environments, one module is also supplied as an IMS batch message processing program.

The CICS for MVS/ESA programs are delivered in C and COBOL. The single IMS program is delivered in C.

See Table 37 on page 386 for the CICS application, and Table 38 on page 388 for the IMS application.

The Message Handler sample

The Message Handler sample allows you to browse, forward, and delete messages on a queue.

The application uses these MQI calls:

- MQCONN
- MQOPEN
- MQINQ
- MQPUT1
- MQCMIT
- MQBACK
- MQGET
- MQCLOSE
- MQDISC

The program is delivered in C and COBOL programming languages. The application runs under TSO. See Table 33 on page 383 for the TSO application.

Distributed queuing exit samples

The names of the source programs of the distributed queuing exit samples are listed in the following table:

Table 27. Source for the distributed queuing exit samples			
Member name	**For language**	**Description**	**Supplied in library**
CSQ4BAX0	Assembler	Source program	SCSQASMS
CSQ4BCX1	C	Source program	SCSQC37S
CSQ4BCX2	C	Source program	SCSQC37S

Note: The source programs are link-edited with CSQXSTUB.

See the *MQSeries Intercommunication* book for a description of the distributed queuing exit samples.

Data-conversion exit samples

A skeleton is provided for a data-conversion exit routine, and a sample is shipped with MQSeries illustrating the MQXCNVC call. The names of the source programs of the data-conversion exit samples are listed in the following table:

Table 28. Source for the data conversion exit samples (Assembler language only)		
Member name	**Description**	**Supplied in library**
CSQ4BAX8	Source program	SCSQASMS
CSQ4BAX9	Source program	SCSQASMS
CSQ4CAX9	Source program	SCSQASMS

Note: The source programs are link-edited with CSQASTUB.

See Chapter 11, "Data-conversion exits" on page 157 for more information.

Preparing and running sample applications for the batch environment

To prepare a sample application that runs in the batch environment, first perform the same steps that you would when building any batch MQSeries for MVS/ESA application. These steps are listed in "Building MVS batch applications" on page 267.

Note: The assembler language version of the Browse sample uses data control blocks (DCBs), so you must link-edit it using `RMODE(24)`.

The library members that you will use are listed in Table 29, Table 30, and Table 31 on page 381.

You must edit the run JCL supplied for the samples that you want to use (see Table 29, Table 30, and Table 31 on page 381).

The PARM statement in the supplied JCL contains a number of parameters that you need to modify. To run the C sample programs, separate the parameters by spaces; to run the Assembler, COBOL, and PL/I sample programs, separate them by commas. For example, if the name of your queue manager is CSQ1 and you want to run the application with a queue named LOCALQ1, in the COBOL, PL/I, and assembler-language JCL, your PARM statement should look like this:

```
PARM=(CSQ1,LOCALQ1)
```

In the C language JCL, your PARM statement should look like this:

```
PARM=('CSQ1 LOCALQ1')
```

You are now ready to submit the jobs.

Names of the sample batch applications

The names of the source programs and JCL that are supplied for each of the sample batch applications are listed in the following tables:

Put and Get samples	Table 29
Browse sample	Table 30 on page 381
Print message sample	Table 31 on page 381

Table 29. Source and JCL for the Put and Get samples

Member name	For language	Description	Supplied in library
CSQ4BCJ1	C	Source program	SCSQC37S
CSQ4BCK1	C	Source program	SCSQC37S
CSQ4BVJ1	COBOL	Source program	SCSQCOBS
CSQ4BVK1	COBOL	Source program	SCSQCOBS
CSQ4BCJR	C	Sample run JCL	SCSQPROC
CSQ4BVJR	COBOL	Sample run JCL	SCSQPROC

Table 30. Source and JCL for the Browse sample

Member name	For language	Description	Supplied in library
CSQ4BVA1	COBOL	Source program	SCSQCOBS
CSQ4BVAR	COBOL	Sample run JCL	SCSQPROC
CSQ4BAA1	Assembler	Source program	SCSQASMS
CSQ4BAAR	Assembler	Sample run JCL	SCSQPROC
CSQ4BCA1	C	Source program	SCSQC37S
CSQ4BCAR	C	Sample run JCL	SCSQPROC
CSQ4BPA1	PL/I	Source program	SCSQPLIS
CSQ4BPAR	PL/I	Sample run JCL	SCSQPROC

Table 31. Source for the Print Message sample (C language only)

Member name	Description	Supplied in library
CSQ4BCG1	Source program	SCSQC37S
CSQ4BCGR	Sample run JCL	SCSQPROC

Preparing sample applications for the TSO environment

To prepare a sample application that runs in the TSO environment, first perform the same steps that you would when building any batch MQSeries for MVS/ESA application—these steps are listed in "Building MVS batch applications" on page 267. The library members you will use are listed in Table 32 on page 382.

Names of the sample TSO applications

The names of the source programs that are supplied for each of the sample TSO applications are listed in the following tables:

Mail manager sample   Table 32 on page 382
Message handler sample   Table 33 on page 383

Preparing and running samples

Table 32. Source and JCL for the Mail Manager (TSO) sample			
Member name	**For language**	**Description**	**Supplied in library**
CSQ4CVD	independent	MQSeries for MVS/ESA object definitions	SCSQPROC
CSQ40	independent	ISPF messages	SCSQMSGE
CSQ4RVD1	COBOL	CLIST to initiate CSQ4TVD1	SCSQCLSS
CSQ4TVD1	COBOL	Source program for Menu program	SCSQCOBS
CSQ4TVD2	COBOL	Source program for Get Mail program	SCSQCOBS
CSQ4TVD4	COBOL	Source program for Send Mail program	SCSQCOBS
CSQ4TVD5	COBOL	Source program for Nickname program	SCSQCOBS
CSQ4VDP1-6	COBOL	Panel definitions	SCSQPNLA
CSQ4VD0	COBOL	Data definition	SCSQCOBC
CSQ4VD1	COBOL	Data definition	SCSQCOBC
CSQ4VD2	COBOL	Data definition	SCSQCOBC
CSQ4VD4	COBOL	Data definition	SCSQCOBC
CSQ4RCD1	C	CLIST to initiate CSQ4TCD1 and CSQ4TVH1	SCSQCLSS
CSQ4TCD1	C	Source program for Menu program	SCSQC37S
CSQ4TCD2	C	Source program for Get Mail program	SCSQC37S
CSQ4TCD4	C	Source program for Send Mail program	SCSQC37S
CSQ4TCD5	C	Source program for Nickname program	SCSQC37S
CSQ4CDP1-6	C	Panel definitions	SCSQPNLA
CSQ4TC0	C	Include file	SCSQC370

Table 33. Source for the Message Handler sample			
Member name	**For language**	**Description**	**Supplied in library**
CSQ4TCH0	C	Data definition	SCSQC370
CSQ4TCH1	C	Source program	SCSQC37S
CSQ4TCH2	C	Source program	SCSQC37S
CSQ4TCH3	C	Source program	SCSQC37S
CSQ4RCH1	C and COBOL	CLIST to initiate CSQ4TCH1 or CSQ4TVH1	SCSQCLST
CSQ4CHP1	C and COBOL	Panel definition	SCSQPNLA
CSQ4CHP2	C and COBOL	Panel definition	SCSQPNLA
CSQ4CHP3	C and COBOL	Panel definition	SCSQPNLA
CSQ4CHP9	C and COBOL	Panel definition	SCSQPNLA
CSQ4TVH0	COBOL	Data definition	SCSQCOBC
CSQ4TVH1	COBOL	Source program	SCSQCOBS
CSQ4TVH2	COBOL	Source program	SCSQCOBS
CSQ4TVH3	COBOL	Source program	SCSQCOBS

Preparing the sample applications for the CICS for MVS/ESA environment

Before you run the CICS sample programs, you must log on to CICS using a LOGMODE of 32702. This is because the sample programs have been written to use a 3270 mode 2 screen.

To prepare a sample application that runs in the CICS for MVS/ESA environment, perform the following steps:

1. Create the symbolic description map and the physical screen map for the sample by assembling the BMS screen definition source (supplied in library **thlqual**.SCSQMAPS, where **thlqual** is the high-level qualifier used by your installation). When you name the maps, use the name of the BMS screen definition source (not available for Put and Get sample programs), but omit the last character of that name.

2. Perform the same steps that you would when building any CICS MQSeries for MVS/ESA application—these steps are listed in "Building CICS applications" on page 268. The library members that you will use are listed in Table 34 on page 384, Table 35 on page 385, Table 36 on page 385, and Table 37 on page 386.

3. Identify the map set, programs, and transaction to CICS by updating the CICS for MVS/ESA system definition (CSD) data set. The definitions you require are in the member **thlqual**.SCSQPROC(CSQ4S100). For guidance on how to do this, see the *MQSeries for MVS/ESA System Management Guide*.

   Note: For the Credit Check sample application, you will get an error message at this stage if you have not already created the VSAM data set that the sample uses.

4. For the Credit Check and Mail Manager sample applications, ensure that the queues they use are available on your system. For the Credit Check sample, they are defined in the member **thlqual**.SCSQPROC(CSQ4CVB) for COBOL, and **thlqual**.SCSQPROC(CSQ4CCB) for C. For the Mail Manager sample, they are defined in the member **thlqual**.SCSQPROC(CSQ4CVD). To ensure that these queues are always available, you could add these members to your CSQINP2 initialization input data set, or use the CSQUTIL program to load these queue definitions.

For the Queue Attributes sample application, you could use one or more of the queues that are supplied for the other sample applications. Alternatively, you could use your own queues. However, note that in the form that it is supplied, this sample works only with queues that have the characters `CSQ4SAMP` in the first eight bytes of their name.

QLOP abend

When the CICS sample applications supplied with MQSeries for MVS/ESA use MQI calls, they do not test for the return codes that indicate that the queue manager is not available. If the queue manager is not available when you attempt to run one of the CICS samples, the sample abends with the CICS abend code QLOP. If this happens, you must connect your queue manager to your CICS system before you attempt to start the sample application again. For information about starting a connection, see the *MQSeries for MVS/ESA System Management Guide*.

Names of the sample CICS applications

The source and JCL files that are supplied for each of the sample CICS applications are listed in the following tables:

Put and Get samples	Table 34
Queue attributes sample	Table 35 on page 385
Mail Manager (CICS) sample	Table 36 on page 385
Credit Check (CICS) sample	Table 37 on page 386

Table 34. Source and JCL for the Put and Get samples

Member name	For language	Description	Supplied in library
CSQ4CCK1	C	Source program	SCSQC37S
CSQ4CCJ1	C	Source program	SCSQC37S
CSQ4CVJ1	COBOL	Source program	SCSQCOBS
CSQ4CVK1	COBOL	Source program	SCSQCOBS
CSQ4S100	independent	CICS system definition data set	SCSQPROC

Table 35. Source for the Queue Attributes sample

Member name	For language	Description	Supplied in library
CSQ4CVC1	COBOL	Source program	SCSQCOBS
CSQ4VMSG	COBOL	Message definition	SCSQCOBC
CSQ4VCMS	COBOL	BMS screen definition	SCSQMAPS
CSQ4CAC1	Assembler	Source program	SCSQASMS
CSQ4AMSG	Assembler	Message definition	SCSQMACS
CSQ4ACMS	Assembler	BMS screen definition	SCSQMAPS
CSQ4CCC1	C	Source program	SCSQC37S
CSQ4CMSG	C	Message definition	SCSQC370
CSQ4CCMS	C	BMS screen definition	SCSQMAPS
CSQ4S100	independent	CICS system definition data set	SCSQPROC

Table 36. Source and JCL for the Mail Manager (CICS) sample (COBOL only)

Member name	Description	Supplied in library
CSQ4CVD	MQSeries for MVS/ESA object definitions	SCSQPROC
CSQ4CVD1	Source for Menu program	SCSQCOBS
CSQ4CVD2	Source for Get Mail program	SCSQCOBS
CSQ4CVD3	Source for Display Message program	SCSQCOBS
CSQ4CVD4	Source for Send Mail program	SCSQCOBS
CSQ4CVD5	Source for Nickname program	SCSQCOBS
CSQ4VDMS	BMS screen definition source	SCSQMAPS
CSQ4S100	CICS system definition data set	SCSQPROC
CSQ4VD0	Data definition	SCSQCOBC
CSQ4VD3	Data definition	SCSQCOBC
CSQ4VD4	Data definition	SCSQCOBC

Preparing and running samples

Table 37. Source and JCL for the Credit Check CICS sample			
Member name	**For language**	**Description**	**Supplied in library**
CSQ4CVB	independent	MQSeries object definitions	SCSQPROC
CSQ4CCB	independent	MQSeries object definitions	SCSQPROC
CSQ4CVB1	COBOL	Source for user-interface program	SCSQCOBS
CSQ4CVB2	COBOL	Source for credit application manager	SCSQCOBS
CSQ4CVB3	COBOL	Source for checking-account program	SCSQCOBS
CSQ4CVB4	COBOL	Source for distribution program	SCSQCOBS
CSQ4CVB5	COBOL	Source for agency-query program	SCSQCOBS
CSQ4CCB1	C	Source for user-interface program	SCSQC37S
CSQ4CCB2	C	Source for credit application manager	SCSQC37S
CSQ4CCB3	C	Source for checking-account program	SCSQC37S
CSQ4CCB4	C	Source for distribution program	SCSQC37S
CSQ4CCB5	C	Source for agency-query program	SCSQC37S
CSQ4CB0	C	Include file	SCSQC370
CSQ4CBMS	C	BMS screen definition source	SCSQMAPS
CSQ4VBMS	COBOL	BMS screen definition source	SCSQMAPS
CSQ4VB0	COBOL	Data definition	SCSQCOBC
CSQ4VB1	COBOL	Data definition	SCSQCOBC
CSQ4VB2	COBOL	Data definition	SCSQCOBC
CSQ4VB3	COBOL	Data definition	SCSQCOBC
CSQ4VB4	COBOL	Data definition	SCSQCOBC
CSQ4VB5	COBOL	Data definition	SCSQCOBC
CSQ4VB6	COBOL	Data definition	SCSQCOBC
CSQ4VB7	COBOL	Data definition	SCSQCOBC
CSQ4VB8	COBOL	Data definition	SCSQCOBC
CSQ4BAQ	independent	Source for VSAM data set	SCSQPROC
CSQ4FILE	independent	JCL to build VSAM data set used by CSQ4CVB3	SCSQPROC
CSQ4S100	independent	CICS system definition data set	SCSQPROC

Preparing the sample application for the IMS environment

Part of the Credit Check sample application can run in the IMS environment. To prepare this part of the application to run with the CICS sample you must first perform the steps described in "Preparing the sample applications for the CICS for MVS/ESA environment" on page 383.

Then perform the following steps:

1. Perform the same steps that you would when building any IMS MQSeries for MVS/ESA application—these steps are listed in "Building IMS (BMP or MPP) applications" on page 269. The library members that you will use are listed in Table 38 on page 388.

2. Identify the application program and database to IMS. Samples are provided with PSBGEN, DBDGEN, ACB definition, IMSGEN, and IMSDALOC statements to enable this.

3. Load the database CSQ4CA by tailoring and running the sample JCL provided for this purpose. This JCL loads the database with data from the file CSQ4BAQ.

4. Start the checking-account program as a batch message processing (BMP) program by tailoring and running the sample JCL provided for this purpose. This JCL starts a batch-oriented BMP program. To run the program as a message-oriented BMP program, remove the comment characters from the line in the JCL that contains the IN= statement.

Names of the sample IMS application

The source and JCL that are supplied for the Credit Check sample IMS application are listed in Table 38.

Table 38. Source and JCL for the Credit Check IMS sample (C only)		
Member name	**Description**	**Supplied in library**
CSQ4CVB	MQSeries object definitions	SCSQPROC
CSQ4ICB3	Source for checking-account program	SCSQC37S
CSQ4ICBL	Source for loading the checking-account database	SCSQC37S
CSQ4CBI	Data definition	SCSQC370
CSQ4PSBL	PSBGEN JCL for database-load program	SCSQPROC
CSQ4PSB3	PSBGEN JCL for checking-account program	SCSQPROC
CSQ4DBDS	DBDGEN JCL for database CSQ4CA	SCSQPROC
CSQ4GIMS	IMSGEN macro definitions for CSQ4IVB3 and CSQ4CA	SCSQPROC
CSQ4ACBG	Application control block (ACB) definition for CSQ4IVB3	SCSQPROC
CSQ4BAQ	Source for database	SCSQPROC
CSQ4ILDB	Sample run JCL for database-load job	SCSQPROC
CSQ4ICBR	Sample run JCL for checking-account program	SCSQPROC
CSQ4DYNA	IMSDALOC macro definitions for database	SCSQPROC

The Put samples

The Put sample programs put messages on a queue using the MQPUT call.

The source programs are supplied in C and COBOL in the batch and CICS environments (see Table 29 on page 380 and Table 34 on page 384).

Design of the Put sample

The flow through the program logic is:

1. Connect to the queue manager using the MQCONN call. If this call fails, print the completion and reason codes and stop processing.

   Note: If you are running the sample in a CICS environment, you do not need to issue an MQCONN call; if you do, it returns DEF_HCONN. You can use the connection handle MQHC_DEF_HCONN for the MQI calls that follow.

2. Open the queue using the MQOPEN call with the MQOO_OUTPUT option. On input to this call, the program uses the connection handle that is returned in step 1. For the object descriptor structure (MQOD), it uses the default values for all fields except the queue name field which is passed as a parameter to the program. If the MQOPEN call fails, print the completion and reason codes and stop processing.

3. Create a loop within the program issuing MQPUT calls until the required number of messages are put on the queue. If an MQPUT call fails, the loop is abandoned early, no further MQPUT calls are attempted, and the completion and reason codes are returned.

4. Close the queue using the MQCLOSE call with the object handle returned in step 2. If this call fails, print the completion and reason codes.

5. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 1. If this call fails, print the completion and reason codes.

   Note: If you are running the sample in a CICS environment, you do not need to issue an MQDISC call.

The Put samples for the batch environment

To run the samples, you must edit and run the sample JCL, as described in "Preparing and running sample applications for the batch environment" on page 380.

The programs take the following parameters in an EXEC PARM, separated by spaces in C and commas in COBOL:

1. The name of the queue manager (4 characters)
2. The name of the target queue (48 characters)
3. The number of messages (up to 4 digits)
4. The padding character to be written in the message (1 character)
5. The number of characters to write in the message (up to 4 digits)
6. The persistence of the message (1 character: 'P' for persistent or 'N' for nonpersistent)

If you enter any of the above parameters wrongly, you will receive appropriate error messages.

Any messages from the samples are written to the SYSPRINT data set.

Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. However, these differences are minimized if the layout of the parameters shown in the sample run JCL, CSQ4BCJR, and CSQ4BVJR, is used. None of the differences relate to the MQI.

- CSQ4BCK1 allows you to enter more than four digits for the number of messages sent and the length of the messages.

- In CSQ4BVK1, entries of numeric parameters are position sensitive. For example, if you enter

  ```
  //  PARM=('CSQ1,A.TEST.QUEUE,-1,*,-1,P')
  ```

  100 messages each of 100 bytes in length are sent.

To send a single message, you must enter 0001, as the value '1' will result in 1000 messages being sent. To send a message of 1 byte in length, you must enter 0001, as the value '1' will result in a message of 1000 bytes in length.

- For both programs, CSQ4BCK1 and CSQ4BVK1, you must enter 'P' in the persistence parameter, ++PER++, if you require the message to be persistent. If you fail to do so, the message will be nonpersistent.

The Put samples for the CICS environment

The transactions take the following parameters separated by commas:

1. The number of messages (up to 4 digits)

2. The padding character to be written in the message (1 character)

3. The number of characters to write in the message (up to 4 digits)

4. The persistence of the message (1 character: 'P' for persistent or 'N' for nonpersistent)

5. The name of the target queue (48 characters)

If you enter any of the above parameters wrongly, you will receive appropriate error messages.

For the COBOL sample, invoke the Put sample in the CICS environment by entering:

```
MVPT,9999,*,9999,P,QUEUE.NAME
```

For the C sample, invoke the Put sample in the CICS environment by entering:

```
MCPT,9999,*,9999,P,QUEUE.NAME
```

Any messages from the samples are displayed on the screen.

Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. None of the differences relate to the MQI.

- If you enter a queue name that is longer than 48 characters, its length is truncated to the maximum of 48 characters but no error message is returned.

- To put a single message, you must enter 0001, as the value '1' will result in 1000 messages being sent. Also, if you enter non-numeric or negative values, you receive an error or some messages are sent.

- To send a message of 1 byte in length in the COBOL program, you must enter 0001, as the value '1' will result in a message of 1000 bytes in length. Also, if you enter non-numeric or negative values, you receive an error or some messages are sent. For example, if you enter '-1', the COBOL program will send a 1 byte message, and the C program will abend with an error from malloc().

- For both programs, CSQ4CCK1 and CSQ4CVK1, you must enter 'P' in the persistence parameter, if you require the message to be persistent. If you fail to do so, the message will be nonpersistent.

- The messages are put in syncpoint because default values are used for all parameters except those set during program invocation.

The Get samples

The Get sample programs get messages from a queue using the MQGET call.

The source programs are supplied in C and COBOL in the batch and CICS environments (see Table 29 on page 380 and Table 34 on page 384).

Design of the Get sample

The flow through the program logic is:

1. Connect to the queue manager using the MQCONN call. If this call fails, print the completion and reason codes and stop processing.

   Note: If you are running the sample in a CICS environment, you do not need to issue an MQCONN call; if you do, it returns DEF_HCONN. You can use the connection handle MQHC_DEF_HCONN for the MQI calls that follow.

2. Open the queue using the MQOPEN call with the MQOO_INPUT_SHARED and MQOO_BROWSE options. On input to this call, the program uses the connection handle that is returned in step 1. For the object descriptor structure (MQOD), it uses the default values for all fields except the queue name field which is passed as a parameter to the program. If the MQOPEN call fails, print the completion and reason codes and stop processing.

3. Create a loop within the program issuing MQGET calls until the required number of messages are retrieved from the queue. If an MQGET call fails, the loop is abandoned early, no further MQGET calls are attempted, and the completion and reason codes are returned.

   The following options are specified on the MQGET call:

   - MQGMO_NO_WAIT
   - MQGMO_ACCEPT_TRUNCATED_MESSAGE
   - MQGMO_SYNCPOINT or MQGMO_NO_SYNCPOINT
   - MQGMO_BROWSE_FIRST and MQGMO_BROWSE_NEXT

   For a description of these options, see the *MQSeries Application Programming Reference* manual.

   For each message, the message number is printed followed by the length of the message and the message data.

4. Close the queue using the MQCLOSE call with the object handle returned in step 2. If this call fails, print the completion and reason codes.

5. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 1. If this call fails, print the completion and reason codes.

   Note: If you are running the sample in a CICS environment, you do not need to issue an MQDISC call.

The Get samples for the batch environment

To run the samples, you must edit and run the sample JCL, as described in "Preparing and running sample applications for the batch environment" on page 380.

The programs take the following parameters in an EXEC PARM, separated by spaces in C and commas in COBOL:

1. The name of the queue manager (4 characters)
2. The name of the target queue (48 characters)
3. The number of messages to get (up to 4 digits)
4. The browse/get message option (1 character: 'B' to browse or 'D' to destructively get the messages)
5. The syncpoint control (1 character: 'S' for syncpoint or 'N' for no syncpoint)

If you enter any of the above parameters wrongly, you will receive appropriate error messages.

Output from the samples is written to the SYSPRINT data set:

```
=====================================
PARAMETERS PASSED :
    QMGR       - VC9
    QNAME      - A.Q
    NUMMSGS    - 000000002
    GET        - D
    SYNCPOINT  - N
=====================================
MQCONN SUCCESSFUL
MQOPEN SUCCESSFUL
000000000 : 000000010 : **********
000000001 : 000000010 : **********
000000002 MESSAGES GOT FROM QUEUE
MQCLOSE SUCCESSFUL
MQDISC SUCCESSFUL
```

Usage notes

- To keep the samples simple, there are some minor functional differences between language versions. However, these differences are minimized if the layout of the parameters shown in the sample run JCL, CSQ4BCJR, and CSQ4BVJR, are used. None of the differences relate to the MQI.

- CSQ4BCJ1 allows you to enter more than four digits for the number of messages retrieved.

- Messages longer than 64 KB are truncated.

- CSQ4BCJ1 can only correctly display character messages as it only displays until the first NULL (\0) character is displayed.

- In CSQ4BVJ1, the entry of the parameter is position sensitive. For example, if you enter

    ```
    //  PARM=('CSQ1,A.TEST.QUEUE,-1,D,S')
    ```

  you will retrieve 100 messages.

  To get a single message, you must enter 0001, as the value '1' will result in 1000 messages being sent.

- For both programs, CSQ4BCJ1 and CSQ4BVJ1, you must enter 'B' in the get parameter, ++GET++, if you want to browse the messages.

- For both programs, CSQ4BCJ1 and CSQ4BVJ1, you must enter 'S' in the syncpoint parameter, ++SYNC++, for messages to be retrieved in syncpoint.

The Get samples for the CICS environment

The transactions take the following parameters in an EXEC PARM, separated by commas:

1. The number of messages to get (up to 4 digits)

2. The browse/get message option (1 character: 'B' to browse or 'D' to destructively get the messages)

3. The syncpoint control (1 character: 'S' for syncpoint or 'N' for no syncpoint)

4. The name of the target queue (48 characters)

If you enter any of the above parameters wrongly, you will receive appropriate error messages.

For the COBOL sample, invoke the Put sample in the CICS environment by entering:

```
MVGT,9999,B,S,QUEUE.NAME
```

For the C sample, invoke the Put sample in the CICS environment by entering:

```
MCGT,9999,B,S,QUEUE.NAME
```

When the messages are retrieved from the queue, they are put on a CICS temporary storage queue with the same name as the CICS transaction (for example, MCGT for the C sample).

Here is example output of the Get samples:

```
*************************  TOP OF QUEUE  ************************
000000000 : 000000010: **********
000000001 : 000000010 :**********
*************************  BOTTOM OF QUEUE  **********************
```

Usage notes
- To keep the samples simple, there are some minor functional differences between language versions. None of the differences relate to the MQI.

- If you enter a queue name that is longer than 48 characters, its length is truncated to the maximum of 48 characters but no error message is returned.

- CSQ4CCJ1 can only correctly display character messages as it only displays until the first NULL (\0) character is displayed.

- To get a single message, you must enter 0001, as the value '1' will result in 1000 messages being sent. Also, if you enter non-numeric or negative values, you receive an error or some messages are retrieved.

- Messages longer than 24K-50 bytes in C and 9999-49 bytes in COBOL are truncated. This is due to the way the CICS temporary storage queues are used.

- For both programs, CSQ4CCK1 and CSQ4CVK1, you must enter 'B' in the get parameter, if you want to browse the messages.

- For both programs, CSQ4CCJ1 and CSQ4CVJ1, you must enter 'S' in the syncpoint parameter for messages to be retrieved in syncpoint.

The Browse sample

The Browse sample is a batch application that demonstrates how to browse messages on a queue using the MQGET call. The application steps through all the messages in a queue, printing the first 80 bytes of each one. You could use this application to look at the messages on a queue without changing them.

Source programs and sample run JCL are supplied in the COBOL, assembler, PL/I, and C languages (see Table 30 on page 381).

To start the application, you must edit and run the sample run JCL, as described in "Preparing and running sample applications for the batch environment" on page 380. You can look at messages on one of your own queues by specifying the name of the queue in the run JCL.

When you run the application (and there are some messages on the queue), the output data set looks this:

```
07/12/93                          SAMPLE QUEUE REPORT          PAGE    1
                         QUEUE MANAGER NAME : VC4
                               QUEUE NAME : CSQ4SAMP.DEAD.QUEUE
          RELATIVE
          MESSAGE   MESSAGE
          NUMBER    LENGTH ------------------ MESSAGE DATA -------------

             1        740 HELLO. PLEASE CALL ME WHEN YOU GET BACK.
             2        429 CSQ4BQRM
             3        429 CSQ4BQRM
             4        429 CSQ4BQRM
             5         22 THIS IS A TEST MESSAGE
             6          8 CSQ4TEST
             7         36 CSQ4MSG - ANOTHER TEST MESSAGE.....!
             8          9 CSQ4STOP
                                   ********** END OF REPORT **********
```

If there are no messages on the queue, the data set contains the headings and the "End of report" message only. If an error occurs with any of the MQI calls, the completion and reason codes are added to the output data set.

Design of the Browse sample

The Browse sample application uses a single program module—one is provided in each of the supported programming languages.

The flow through the program logic is:

1. Open a print data set and print the title line of the report. Check that names of the queue manager and queue have been passed from the run JCL. If both names have been passed, print the lines of the report that contain the names. If they have not, print an error message, close the print data set, and stop processing.

   The way that the program tests the parameters it is passed from the JCL depends on the language in which the program is written—for more information, see "Language-dependent design considerations" on page 396.

2. Connect to the queue manager using the MQCONN call. If this call is not successful, print the completion and reason codes, close the print data set, and stop processing.

3. Open the queue using the MQOPEN call with the MQOO_BROWSE option. On input to this call, the program uses the connection handle returned in step 2. For the object descriptor structure (MQOD), it uses the default values for all the fields except the queue name (which was passed in step 1 on page 394). If this call is not successful, print the completion and reason codes, close the print data set, and stop processing.

4. Browse the first message on the queue, using the MQGET call. On input to this call, the program specifies:

    • The connection and queue handles from steps 2 and 3

    • An MQMD structure with all fields set to their initial values

    • Two options:
        – MQGMO_BROWSE_FIRST
        – MQGMO_ACCEPT_TRUNCATED_MSG

    • A buffer of size 80 bytes to hold the data copied from the message

    The MQGMO_ACCEPT_TRUNCATED_MSG option allows the call to complete even if the message is longer than the 80-byte buffer specified in the call. If the message is longer than the buffer, the message is truncated to fit the buffer, and the completion and reason codes are set to show this. The sample was designed so that messages are truncated to 80 characters simply to make the report easy to read. The buffer size is set by a `DEFINE` statement, so you can easily change it if you want to.

5. Perform the following loop until the MQGET call fails:

    a. Print a line of the report showing:

        • The sequence number of the message (this is a count of the browse operations).
        • The true length of the message (not the truncated length). This value is returned in the *DataLength* field of the MQGET call.
        • The first 80 bytes of the message data.

    b. Reset the *MsqId* and *CorrelId* fields of the MQMD structure to nulls

    c. Browse the next message, using the MQGET call with these two options:

        • MQGMO_BROWSE_NEXT
        • MQGMO_ACCEPT_TRUNCATED_MSG

6. If the MQGET call fails, test the reason code to see if the call has failed because the browse cursor has got to the end of the queue. In this case, print the "End of report" message and go to step 7; otherwise, print the completion and reason codes, close the print data set, and stop processing.

7. Close the queue using the MQCLOSE call with the object handle returned in step 3.

8. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 2.

9. Close the print data set and stop processing.

Language-dependent design considerations

Source modules are provided for the Browse sample in four programming languages. There are two main differences between the source modules:

- When testing the parameters passed from the run JCL, the COBOL, PL/I, and assembler-language modules search for the comma character (,). If the JCL passes `PARM=(,LOCALQ1)`, the application attempts to open queue LOCALQ1 on the default queue manager. If there is no name after the comma (or no comma), the application returns an error. The C module does not search for the comma character. If the JCL passes a single parameter (for example, `PARM=('LOCALQ1')`), the C module uses this as a queue name on the default queue manager.

- To keep the assembler-language module simple, it uses the date format yy/ddd (for example, 93/116) when it creates the print report. The other modules use the calendar date in mm/dd/yy format.

The Print Message sample

The Print Message sample is a simple batch application that demonstrates how to remove all the messages from a queue using the MQGET call. It also prints, for each message, the fields of the message descriptor, followed by the message data. The program prints the data both in hexadecimal and as characters (if they are printable). If a character is not printable, the program replaces it with a period character (.). You can use the program when diagnosing problems with an application that is putting messages on a queue.

You can change the application so that it browses the messages, rather than removing them from the queue. To do this, remove the comment characters from two lines in the code, as indicated in "Design of the sample" on page 398.

The application has a single source program, which is written in the C language. Sample run JCL code is also supplied (see Table 31 on page 381).

To start the application, you must edit and run the sample run JCL, as described in "Preparing and running sample applications for the batch environment" on page 380. When you run the application (and there are some messages on the queue), the output data set looks like that in Figure 40 on page 397.

```
 MQCONN to VC4
 MQOPEN - 'CSQ4SAMP.DEAD.QUEUE'


 MQGET of message number 1
****Message descriptor****
  StrucId  : 'MD '  Version : 1
  Report   : 0  MsgType : 2
  Expiry   : -1  Feedback : 0
  Encoding : 785  CodedCharSetId : 500
  Format : '        '
  Priority : 3  Persistence : 0
  MsgId : X'C3E2D840E5C3F44040404040404040404040A6FE06A95105C620'
  CorrelId : X'C3E2D840E5C3F44040404040404040404040A6FE062950C2F125'
  BackoutCount : 0
  ReplyToQ       : '                                                 '
  ReplyToQMgr    : 'VC4                                        '
  ** Identity Context
  UserIdentifier : 'CICSUSER    '
  Account.Token :
   X'160DD5E3E2D5C5E34BC9C7D7C2F6F1FE060D3B55B60001000000000000000000000'
  ApplIdentData  : '                                '
  ** Origin Context
  PutApplType    : '1'
  PutApplName    : 'VICAUT4 MVB5             '
  PutDate  : '19930203'    PutTime  : '20165982'
  ApplOriginData : '    '
```

Figure 40 (Part 1 of 2). Example of a report from the Print Message sample application

```
 ****    Message       ****
  length - 429 bytes

 00000000:  C3E2 D8F4 C2D8 D9D4 4040 4040 4040 4040 'CSQ4BQRM        '
 00000010:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000020:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000030:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000040:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000050:  4040 4040 4040 40D1 D6C8 D540 D140 4040 '       JOHN J   '
 00000060:  4040 4040 4040 4040 4040 40F1 F2F3 F4F5 '           12345'
 00000070:  F6F7 F8F9 C6C9 D9E2 E340 C7C1 D3C1 C3E3 '6789FIRST GALACT'
 00000080:  C9C3 40C2 C1D5 D240 4040 4040 4040 4040 'IC BANK         '
 00000090:  4040 E2D6 D4C5 E3C8 C9D5 C740 C4C9 C6C6 '  SOMETHING DIFF'
 000000A0:  C5D9 C5D5 E340 4040 4040 4040 4040 4040 'ERENT           '
 000000B0:  F3F5 F0F1 F6F7 F6F2 F1F2 F1F0 F0F0 F0F0 '3501676212100000'
 000000C0:  D985 A297 9695 A285 4086 9996 9440 C3E2 'Response from CS'
 000000D0:  D8F4 E2C1 D4D7 4BC2 F74B D4C5 E2E2 C1C7 'Q4SAMP.B7.MESSAG'
 000000E0:  C5E2 4040 4040 4040 4040 4040 4040 4040 'ES              '
 000000F0:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000100:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000110:  4040 4040 40D3 9681 9540 8194 96A4 95A3 '     Loan amount'
 00000120:  40F1 F0F0 F0F0 F040 8696 9940 D1D6 C8D5 ' 100000 for JOHN'
 00000130:  40D1 4040 4040 4040 4040 4040 4040 4040 ' J              '
 00000140:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000150:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 00000160:  4040 4040 C399 8584 89A3 40A6 9699 A388 '     Credit worth'
 00000170:  8995 85A2 A240 8995 8485 A740 6040 C2C1 'iness index - BA'
 00000180:  C440 4040 4040 4040 4040 4040 4040 4040 'D               '
 00000190:  4040 4040 4040 4040 4040 4040 4040 4040 '                '
 000001A0:  4040 4040 4040 4040 4040 4040 40       '                '


  No more messages
 MQCLOSE
  MQDISC
```

Figure 40 (Part 2 of 2). Example of a report from the Print Message sample application

Design of the sample

The Print message sample application uses a single program written in the C language.

The flow through the program logic is:

1. Check that names of the queue manager and queue have been passed from the run JCL. If they have not, print an error message and stop processing.

2. Connect to the queue manager using the MQCONN call. If this call is not successful, print the completion and reason codes and stop processing; otherwise print the name of the queue manager.

3. Open the queue using the MQOPEN call with the MQOO_INPUT_SHARED option.

Note: If you want the application to browse the messages rather than remove them from the queue, remove the comment characters from the line in the program that adds the MQOO_BROWSE option.

On input to this call, the program uses the connection handle returned in step 2. For the object descriptor structure (MQOD), it uses the default values for all the fields except the queue name (which was passed in step 1). If this call is not successful, print the completion and reason codes and stop processing; otherwise, print the name of the queue.

4. Perform the following loop until the MQGET call fails:

   a. Initialize the buffer to blanks so that the message data does not get corrupted by any data already in the buffer.

   b. Set the *MsgId* and *CorrelId* fields of the MQMD structure to nulls so that the MQGET call selects the first message from the queue.

   c. Get a message from the queue, using the MQGET call. On input to this call, the program specifies:

      • The connection and object handles from steps 2 and 3.

      • An MQMD structure with all fields set to their initial values. (Note that *MsgId* and *CorrelId* are reset to nulls for each MQGET call.)

      • The option MQGMO_NO_WAIT.

        Note: If you want the application to browse the messages rather than remove them from the queue, remove the comment characters from the line in the program that adds the MQOO_BROWSE_NEXT option. When this option is used on a call against a queue for which no browse cursor has previously been used with the current object handle, the browse cursor is positioned logically before the first message.

      • A buffer of size 32 KB to hold the data copied from the message.

   d. Call the printMD subroutine. This prints the name of each field in the message descriptor, followed by its contents.

   e. Print the length of the message, followed by the message data. Each line of message data is in this format:

      • Relative position (in hexadecimal) of this part of the data

      • 16 bytes of hexadecimal data

      • The same 16 bytes of data in character format, if it is printable (nonprintable characters are replaced by periods)

5. If the MQGET call fails, test the reason code to see if the call failed because there are no more messages on the queue. In this case, print the message: "No more messages"; otherwise, print the completion and reason codes. In both cases, go to step 6 on page 400.

Note: The MQGET call fails if it finds a message that has more than 32 KB of data. To change the program to handle larger messages, you could do one of the following:

   • Add the MQGMO_ACCEPT_TRUNCATED_MSG option to the MQGET call, so that the call gets the first 32 KB of data and discards the remainder

   • Make the program leave the message on the queue when it finds one with this amount of data

- Increase the size of the buffer

6. Close the queue using the MQCLOSE call with the object handle returned in step 3 on page 398.

7. Disconnect from the queue manager using the MQDISC call with the connection handle returned in step 2 on page 398.

The Queue Attributes sample

The Queue Attributes sample is a conversational-mode CICS for MVS/ESA application that demonstrates the use of the MQINQ and MQSET calls. It shows how to inquire about the values of the *InhibitPut* and *InhibitGet* attributes of queues, and how to change them so that programs cannot put messages on, or get messages from, a queue. You may want to **lock** a queue in this way when you are testing a program.

To prevent accidental interference with your own queues, this sample works only on a queue object that has the characters CSQ4SAMP in the first eight bytes of its name. However, the source code includes comments to show you how to remove this restriction.

Source programs are supplied in the COBOL, assembler, and C languages (see Table 35 on page 385).

The assembler-language version of the sample uses reenterable code. To do this, you will notice that the code for each MQI call in that version of the sample includes the MF keyword; for example:

```
CALL MQCONN,(NAME,HCONN,COMPCODE,REASON),MF=(E,PARMAREA),VL
```

(The VL keyword means that you can use the CICS Execution Diagnostic Facility (CEDF) supplied transaction for debugging the program.) For more information on writing reenterable programs, see "Writing reenterable programs" on page 82.

To start the application, start your CICS system and use the following CICS transactions:

- For COBOL, MVC1
- For Assembler language, MAC1
- For C, MCC1

You can change the name of any of these transactions by changing the CSD data set mentioned in step 3 on page 383.

Design of the sample

When you start the sample, firstly it displays a screen map that has fields for:

- Name of the queue
- User request (valid actions are: inquire, allow, or inhibit)
- Current status of put operations for the queue
- Current status of get operations for the queue

The first two fields are for user input. The last two fields are filled by the application: they show the word INHIBITED or the word ALLOWED.

The application validates the values you enter in the first two fields. It checks that the queue name starts with the characters CSQ4SAMP and that you entered one of

the three valid requests in the Action field. The application converts all your input to uppercase, so you cannot use any queues with names that contain lowercase characters.

If you enter 'inquire' in the Action field, the flow through the program logic is:

1. Open the queue using the MQOPEN call with the MQOO_INQUIRE option

2. Call MQINQ using the selectors MQIA_INHIBIT_GET and MQIA_INHIBIT_PUT

3. Close the queue using the MQCLOSE call

4. Analyze the attributes that are returned in the *IntAttrs* parameter of the MQINQ call and move the words 'INHIBITED' or 'ALLOWED', as appropriate, to the relevant screen fields

If you enter 'inhibit' in the Action field, the flow through the program logic is:

1. Open the queue using the MQOPEN call with the MQOO_SET option

2. Call MQSET using the selectors MQIA_INHIBIT_GET and MQIA_INHIBIT_PUT, and with the values MQQA_GET_INHIBITED and MQQA_PUT_INHIBITED in the *IntAttrs* parameter

3. Close the queue using the MQCLOSE call

4. Move the word 'INHIBITED' to the relevant screen fields

If you enter 'allow' in the Action field, the application performs similar processing to that for an 'inhibit' request. The only differences are the settings of the attributes and the words displayed on the screen.

When the application opens the queue, it uses the default connection handle to the queue manager. (CICS establishes a connection to the queue manager when you start your CICS system.) The application can trap the following errors at this stage:

• The application is not connected to the queue manager
• The queue does not exist
• The user is not authorized to access the queue
• The application is not authorized to open the queue

For other MQI errors, the application displays the completion and reason codes.

The Mail Manager sample

The Mail Manager sample application is a suite of programs that demonstrates the sending and receiving of messages, both within a single environment and across different environments. The application is a simple electronic mailing system that allows users to exchange messages, even if they use different queue managers.

The application demonstrates how to create queues using the MQOPEN call and by putting MQSeries for MVS/ESA commands on the system-command input queue.

Three versions of the application are provided:

• A CICS for MVS/ESA application written in COBOL
• A TSO application written in COBOL
• A TSO application written in C

Preparing the sample

The Mail Manager is provided in versions that run in two environments. The preparation you must carry out before you run the application depends on the environment you want to use.

A user can access mail queues and nickname queues from both TSO and CICS so long as their sign-on user IDs are the same on each system.

Before you can send messages to another queue manager, you must set up a message channel to that queue manager. To do this, use the channel control function of MQSeries, described in the MQSeries Intercommunication book.

Preparing the sample for the TSO environment

Follow these steps:

1. Prepare the sample as described in "Preparing sample applications for the TSO environment" on page 381.

2. Tailor the CLIST provided for the sample to define:

   - The location of the panels
   - The location of the message file
   - The location of the load modules
   - The name of the queue manager you want to use with the application

   A separate CLIST is provided for each language version of the sample:

   For the COBOL version: CSQ4RVD1
   For the C version: CSQ4RCD1

3. Ensure that the queues used by the application are available on the queue manager. (The queues are defined in CSQ4CVD.)

Note: VS COBOL II does not support multitasking with ISPF. This means that you cannot use the Mail Manager sample application on both sides of a split screen. If you do, the results are unpredictable.

Preparing the sample for the CICS/ESA environment

Perform the steps in "Preparing the sample applications for the CICS for MVS/ESA environment" on page 383.

If you run your CICS. system so that users do not need to sign on, and the default user ID is not CICSUSER, change the user ID in the source code for the sample.

Running the sample

To start the sample in the TSO environment, execute your tailored version of the CLIST from the TSO command processor within ISPF.

To start the sample in the CICS for MVS/ESA environment, run transaction MAIL. If you have not already signed-on to CICS, the application prompts you to enter a user ID to which it can send your mail.

When you start the application, it opens your mail queue. If this queue does not already exist, the application creates one for you. Mail queues have names of the form CSQ4SAMP.MAILMGR.*userid*, where *userid* depends on the environment:

In TSO

The user's TSO ID

In CICS for MVS/ESA

The user's CICS sign-on or the user ID entered by the user when prompted when the Mail Manager started.

All parts of the queue names that the Mail Manager uses must be uppercase.

The application then presents a menu panel that has options for:

- Read incoming mail
- Send mail
- Create nickname

The menu panel also shows you how many messages are waiting on your mail queue. Each of the menu options displays a further panel:

Read incoming mail

The Mail Manager displays a list of the messages that are on your mail queue. (Note that only the first 99 messages on the queue are displayed.) For an example of this panel, see Figure 43 on page 409. When you select a message from this list, the contents of the message are displayed (see Figure 44 on page 409).

Send mail

A panel prompts you to enter:

- The name of the user to whom you want to send a message
- The name of the queue manager that owns their mail queue
- The text of your message

In the user name field you can enter either a user ID or a nickname that you created using the Mail Manager. You can leave the queue manager name field blank if the user's mail queue is owned by the same queue manager that you are using, and you must leave it blank if you entered a nickname in the user name field:

- If you specify only a user name, the program first assumes that the name is a nickname, and sends the message to the object defined by that name. If there is no such nickname, the program attempts to send the message to a local queue of that name.

- If you specify both a user name and a queue manager name, the program sends the message to the mail queue that is defined by those two names.

For example, if you want to send a message to user JONESM on remote queue manager QM12, you could send them a message in either of two ways:

- Use both fields to specify user JONESM at queue manager QM12.

- Define a nickname (for example, MARY) for that user and send them a message by putting MARY in the user name field and nothing in the queue manager name field.

Create nickname

You can define an easy-to-remember name that you can use when you send a message to another user who you contact frequently. You are prompted to enter the user ID of the other user and the name of the queue manager that owns their mail queue.

Nicknames are queues that have names of the form CSQ4SAMP.MAILMGR.*userid.nickname*, where *userid* is your own user ID and *nickname* is the nickname that you want to use. With names structured in this way, users can each have their own set of nicknames.

The type of queue that the program creates depends on how you fill in the fields of the Create Nickname panel:

- If you specify only a user name, or the queue manager name is the same as that of the queue manager to which the Mail Manager is connected, the program creates an alias queue.

- If you specify both a user name and a queue manager name (and the queue manager is not the one to which the Mail Manager is connected), the program creates a local definition of a remote queue. The program does not check the existence of the queue to which this definition resolves, or even that the remote queue manager exists.

For example, if your own user ID is SMITHK and you create a nickname called MARY for user JONESM (who uses the remote queue manager QM12), the nickname program creates a local definition of a remote queue named CSQ4SAMP.MAILMGR.SMITHK.MARY. This definition resolves to Mary's mail queue, which is CSQ4SAMP.MAILMGR.JONESM at queue manager QM12. If you are using queue manager QM12 yourself, the program instead creates an alias queue of the same name (CSQ4SAMP.MAILMGR.SMITHK.MARY).

The C version of the TSO application makes greater use of ISPF's message-handling capabilities than does the COBOL version. You may notice that different error messages are displayed by the C and COBOL versions.

Design of the sample

The following sections describe each of the programs that comprise the Mail Manager sample application. The relationships between the programs and the panels that the application uses is shown in Figure 41 on page 405 for the TSO version, and Figure 42 on page 406 for the CICS for MVS/ESA version.
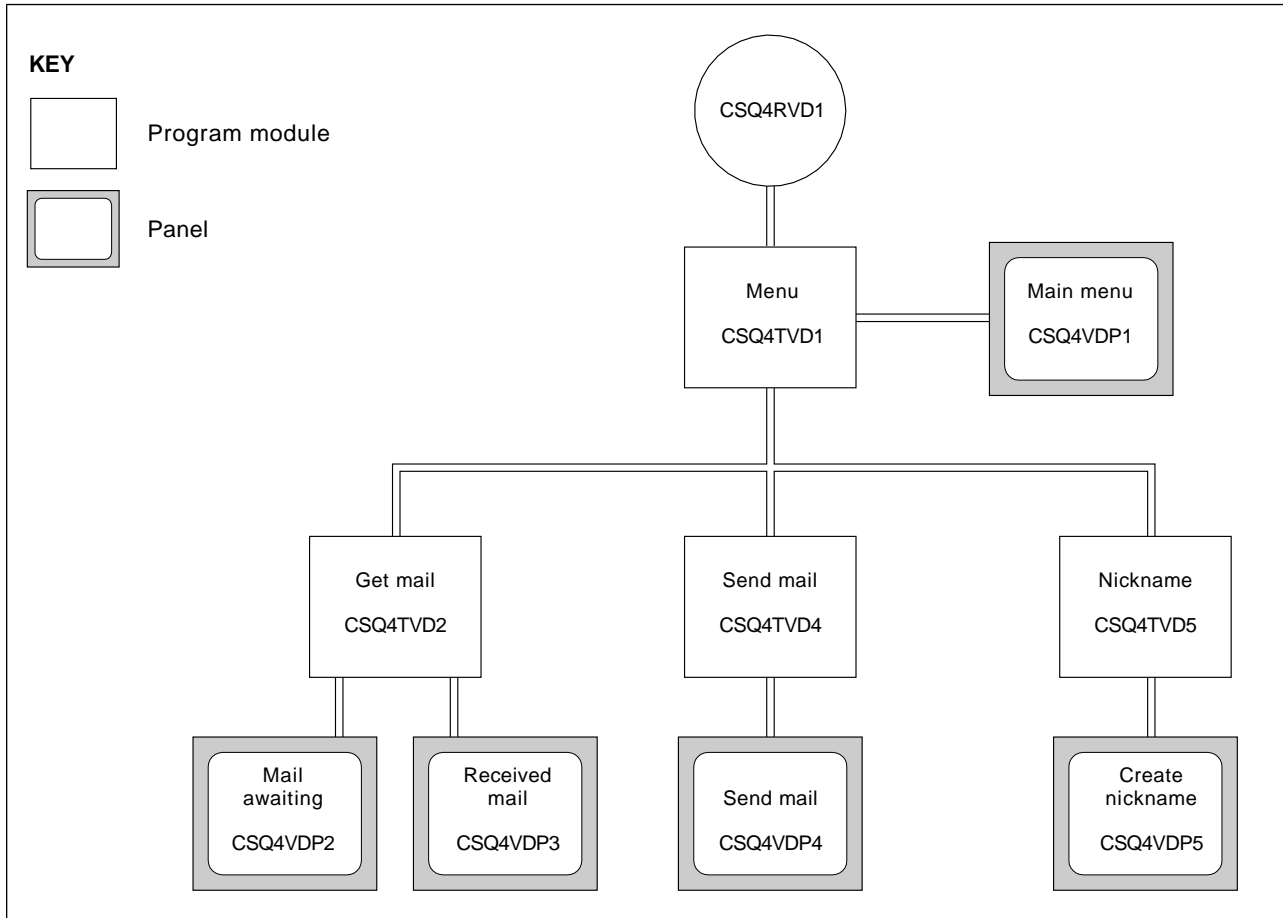
Figure 41. Programs and panels for the TSO versions of the Mail Manager. This figure shows the names for the COBOL version.

Mail manager sample
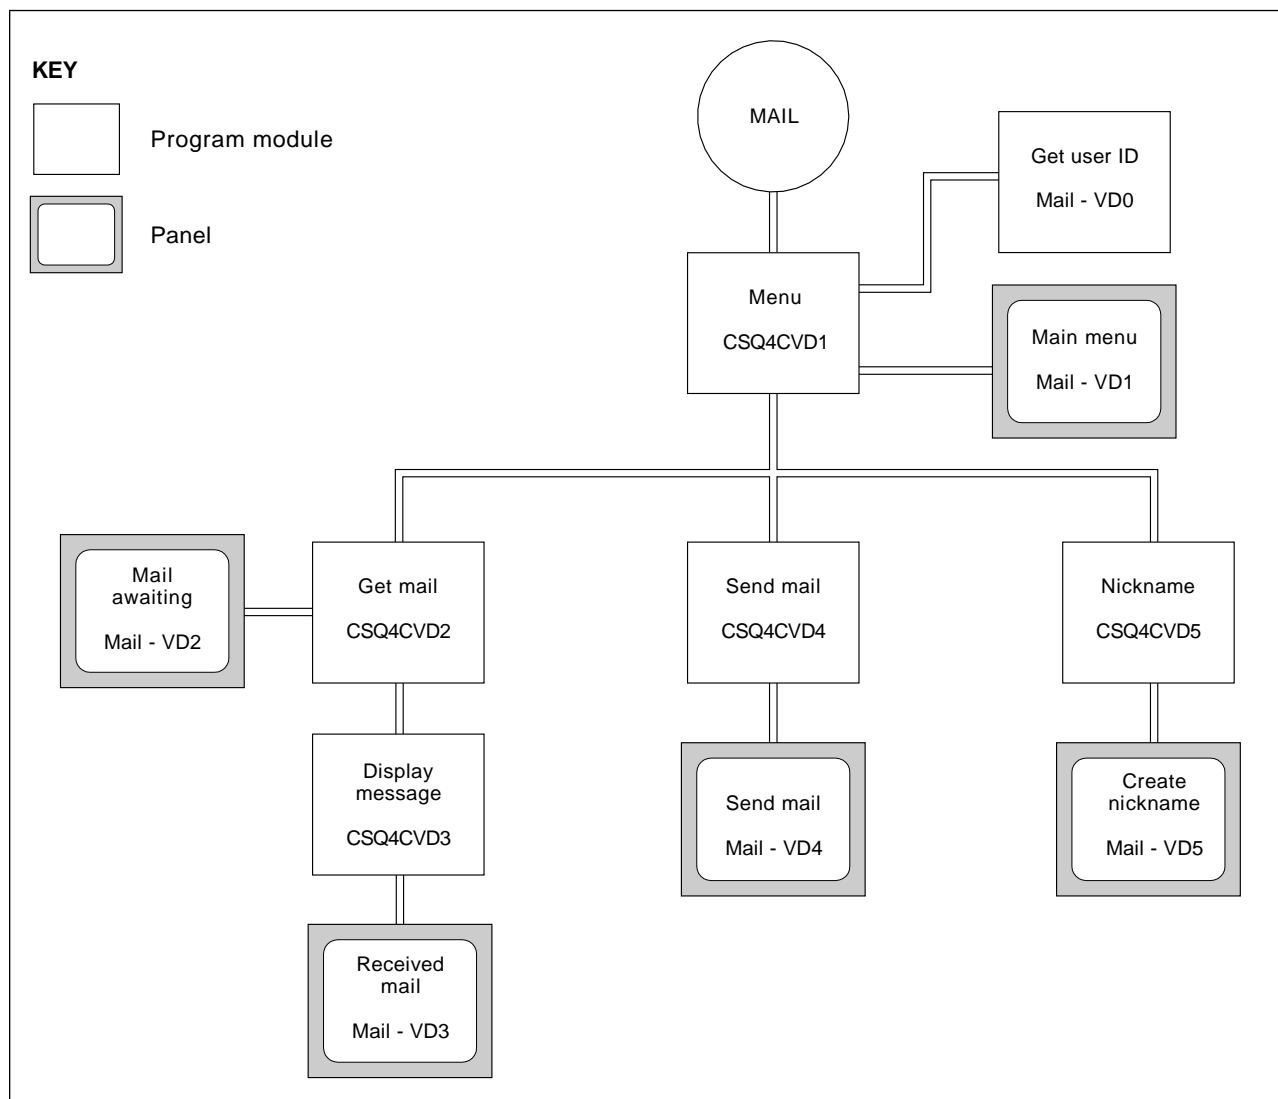
Figure 42. Programs and panels for the CICS/ESA version of the Mail Manager

Menu program

In the TSO environment, the menu program is invoked by the CLIST. In the CICS for MVS/ESA environment, the program is invoked by transaction MAIL.

The menu program is the initial program in the suite. It displays the menu and invokes the other programs when they are selected from the menu.

The program first obtains the user's ID:

- In the CICS for MVS/ESA version of the program, if the user has signed on to CICS, the user ID is obtained by using the CICS command ASSIGN USERID. If the user has not signed on, the program displays the sign-on panel (CSQ4VD0) to prompt the user to enter a user ID. There is no security processing within this program—the user can give *any* user ID.

- In the TSO version, the user's ID is obtained from TSO in the CLIST. It is passed to the menu program as a variable in the ISPF shared pool.

After the program has obtained the user ID, it checks to ensure that the user has a mail queue (CSQ4SAMP.MAILMGR.*userid*). If a mail queue does not exist, the program creates one by putting a message on the system-command input queue. The message contains the MQSeries for MVS/ESA command DEFINE QLOCAL. The object definition that this command uses sets the maximum depth of the queue to 9999 messages.

The program also creates a temporary dynamic queue to handle replies from the system-command input queue. To do this, the program uses the MQOPEN call, specifying the SYSTEM.DEFAULT.MODEL.QUEUE as the template for the dynamic queue. The queue manager creates the temporary dynamic queue with a name that has the prefix CSQ4SAMP; the remainder of the name is generated by the queue manager.

The program then opens the user's mail queue and finds the number of messages on the queue by inquiring about the current depth of the queue. To do this, the program uses the MQINQ call, specifying the MQIA_CURRENT_Q_DEPTH selector.

The program then performs a loop that displays the menu and processes the selection that the user makes. The loop is stopped when the user presses the PF3 key. When a valid selection is made, the appropriate program is started; otherwise an error message is displayed.

Get-mail and display-message programs
In the TSO versions of the application, the get-mail and display-message functions are performed by the same program. In the CICS for MVS/ESA version of the application, these functions are performed by separate programs.

The Mail Awaiting panel (see Figure 43 on page 409 for an example) shows all the messages that are on the user's mail queue. To create this list, the program uses the MQGET call to browse all the messages on the queue, saving information about each one. In addition to the information displayed, the program records the *MsgId* and *CorrelId* of each message.

From the Mail Awaiting panel the user can select one message and display the contents of the message. The program uses the MQGET call to remove this message from the queue, using the *MsgId* and *CorrelId* that the program noted when it browsed all the messages. This MQGET call is performed using the MQGMO_SYNCPOINT option. The program displays the contents of the message, then declares a syncpoint: this commits the MQGET call, so the message now no longer exists.

An obvious extension to the function provided by the Mail Manager is to give the user the option to leave the message on the queue after viewing its contents. To do this, you would have to back out the MQGET call that removes the message from the queue, after displaying the message.

Send-mail program

When the user has completed the Send Mail panel, the send-mail program puts the message on the receiver's mail queue. To do this, the program uses the MQPUT1 call. The destination of the message depends on how the user has filled the fields in the Send Mail panel:

- If the user has specified only a user name, the program first assumes that the name is a nickname, and sends the message to the object defined by that name. If there is no such nickname, the program attempts to send the message to a local queue of that name.

- If the user has specified both a user name and a queue manager name, the program sends the message to the mail queue that is defined by those two names.

The program does not accept blank messages, and it removes leading blanks from each line of the message text.

If the MQPUT1 call is successful, the program displays a message that shows the user name and queue manager name to which the message was put. If the call is unsuccessful, the program checks specifically for the reason codes that indicate the queue or the queue manager do not exist; these are MQRC_UNKNOWN_OBJECT_NAME and MQRC_UNKNOWN_OBJECT_Q_MGR. The program displays its own error message for each of these errors; for other errors, the program displays the completion and reason codes returned by the call.

Nickname program

When the user defines a nickname, the program creates a queue that has the nickname as part of its name. The program does this by putting a message on the system-command input queue. The message contains the MQSeries for MVS/ESA command DEFINE QALIAS or DEFINE QREMOTE. The type of queue that the program creates depends on how the user has filled the fields of the Create Nickname panel:

- If the user has specified only a user name, or the queue manager name is the same as that of the queue manager to which the Mail Manager is connected, the program creates an alias queue.

- If the user has specified both a user name and a queue manager name, (and the queue manager is not the one to which the Mail Manager is connected), the program creates a local definition of a remote queue. The program does not check the existence of the queue to which this definition resolves, or even that the remote queue manager exists.

The program also creates a temporary dynamic queue to handle replies from the system-command input queue.

If the queue manager cannot create the nickname queue for a reason that the program expects (for example, the queue already exists), the program displays its own error message. If the queue manager cannot create the queue for a reason that the program does not expect, the program displays up to two of the error messages that are returned to the program by the command server.

Note: For each nickname, the nickname program creates only an alias queue or a local definition of a remote queue. The local queues to which these queue names resolve are created only when the user ID that is contained in the nickname is used to start the Mail Manager application.

Credit check sample

```
       --------------------MQSeries for MVS/ESA Sample Programs ------- ROW 16 OF 29
        COMMAND ==>                                         Scroll   ===> PAGE
                                                            USERID - NTSFV02
                                     Mail Manager System    QMGR   - VC4
                                        Mail Awaiting

                           Msg     Mail       Date        Time
                           No      From       Sent        Sent
                     16
                           16      Deleted
                           17      JOHNJ      01/06/1993  12:52:02
                           18      JOHNJ      01/06/1993  12:52:02
                           19      JOHNJ      01/06/1993  12:52:03
                           20      JOHNJ      01/06/1993  12:52:03
                           21      JOHNJ      01/06/1993  12:52:03
                           22      JOHNJ      01/06/1993  12:52:04
                           23      JOHNJ      01/06/1993  12:52:04
                           24      JOHNJ      01/06/1993  12:52:04
                           25      JOHNJ      01/06/1993  12:52:05
                           26      JOHNJ      01/06/1993  12:52:05
                           27      JOHNJ      01/06/1993  12:52:05
                           28      JOHNJ      01/06/1993  12:52:06
                           29      JOHNJ      01/06/1993  12:52:06
```

Figure 43. Example of a panel showing a list of waiting messages

```
       --------------------MQSeries for MVS/ESA Sample Programs --------------------
        COMMAND ==>
                                                            USERID - NTSFV02
                                     Mail Manager System    QMGR   - VC4
                                        Received Mail

         Mail sent from JOHNJ    at VC4

         Sent on the 01/06/1993 at 12:52:02
         --------------------------------- Message ------------------------------
        |                                                                        |
        |  HELLO FROM JOHNJ                                                       |
        |                                                                        |
        |                                                                        |
        |                                                                        |
        |                                                                        |
        |                                                                        |
        |                                                                        |
        |                                                                        |
        |                                                                        |
        |_____|
```

Figure 44. Example of a panel showing the contents of a message

The Credit Check sample

The Credit Check sample application is a suite of programs that demonstrates how
to use many of the features provided by MQSeries for MVS/ESA. It shows how the
many component programs of an application can pass messages to each other
using message queuing techniques.

The sample can run as a stand-alone CICS for MVS/ESA application. However, to
demonstrate how to design a message queuing application that uses the facilities

Chapter 31. Sample programs for MQSeries for MVS/ESA   **409**

provided by both the CICS for MVS/ESA and IMS environments, one module is also supplied as an IMS batch message processing program. This extension to the sample is described in "The IMS extension to the Credit Check sample" on page 422.

You can also run the sample on more than one queue manager, and send messages between each instance of the application. To do this, see "The Credit Check sample with multiple queue managers" on page 422.

The CICS programs are delivered in C and COBOL. The single IMS program is delivered only in C. The supplied data sets are shown in Table 37 on page 386 and Table 38 on page 388.

The application demonstrates a method of assessing the risk when bank customers ask for loans. The application shows how a bank could work in two ways to process loan requests:

- When dealing directly with a customer, bank staff want immediate access to account and credit-risk information.

- When dealing with written applications, bank staff can submit a series of requests for account and credit-risk information, and deal with the replies at a later time.

The financial and security details in the application have been kept simple so that the message queuing techniques are clear.

Preparing and running the Credit Check sample

To prepare and run the Credit Check sample, perform the following steps:

1. Create the VSAM data set that holds information about some example accounts. Do this by editing and running the JCL supplied in data set CSQ4FILE.

2. Perform the steps in "Preparing the sample applications for the CICS for MVS/ESA environment" on page 383. (The additional steps you must perform if you want to use the IMS extension to the sample are described in "The IMS extension to the Credit Check sample" on page 422.)

3. Start the CKTI trigger monitor (supplied with MQSeries for MVS/ESA) against queue CSQ4SAMP.INITIATION.QUEUE, using the CICS transaction CKQC.

4. To start the application, start your CICS system and use the transaction MVB1.

5. Select **Immediate** or **Batch** inquiry from the first panel.

   The immediate and batch inquiry panels are similar—Figure 45 on page 411 shows the Immediate Inquiry panel.

```
 CSQ4VB2               MQSeries for MVS/ESA Sample Programs

                      Credit Check - Immediate Inquiry

 Specify details of the request, then press Enter.
       Name . . . . . . . . .  _____
       Social security number  __ __ ___
       Bank account name  . .  _____
       Account number . . . .  _____
       Amount requested . . .  012345
 Response from CHECKING ACCOUNT for name : _____
       Account information not found
       Credit worthiness index - NOT KNOWN
 ..
 ..
 ..
 ..
 ..
 ..
 ..
 ..
 ..
 MESSAGE LINE
 F1=Help  F3=Exit  F5=Make another inquiry
```

Figure 45. Immediate Inquiry panel for the Credit Check sample application

6. Enter an account number and loan amount in the appropriate fields. See "Entering information in the inquiry panels" for guidance on what information you should enter in these fields.

Entering information in the inquiry panels

The Credit Check sample application checks that the data you enter in the 'Amount requested' field of the inquiry panels is in the form of integers.

If you enter one of the following account numbers, the application finds the appropriate account name, average account balance, and credit worthiness index in the VSAM data set CSQ4BAQ:

> 2222222222
> 3111234329
> 3256478962
> 3333333333
> 3501676212
> 3696879656
> 4444444444
> 5555555555
> 6666666666
> 7777777777

You can enter any, or no, information in the other fields. The application retains any information that you do enter and returns the same information in the reports that it generates.

Design of the sample

This section describes the design of each of the programs that comprise the Credit Check sample application. For a discussion of some of the techniques that were considered during the design of the application, see "Design considerations" on page 419.

Figure 46 on page 413 shows the programs that make up the application, and also the queues that these programs serve. In this figure, the prefix CSQ4SAMP has been omitted from all the queue names to make the figure easier to understand.
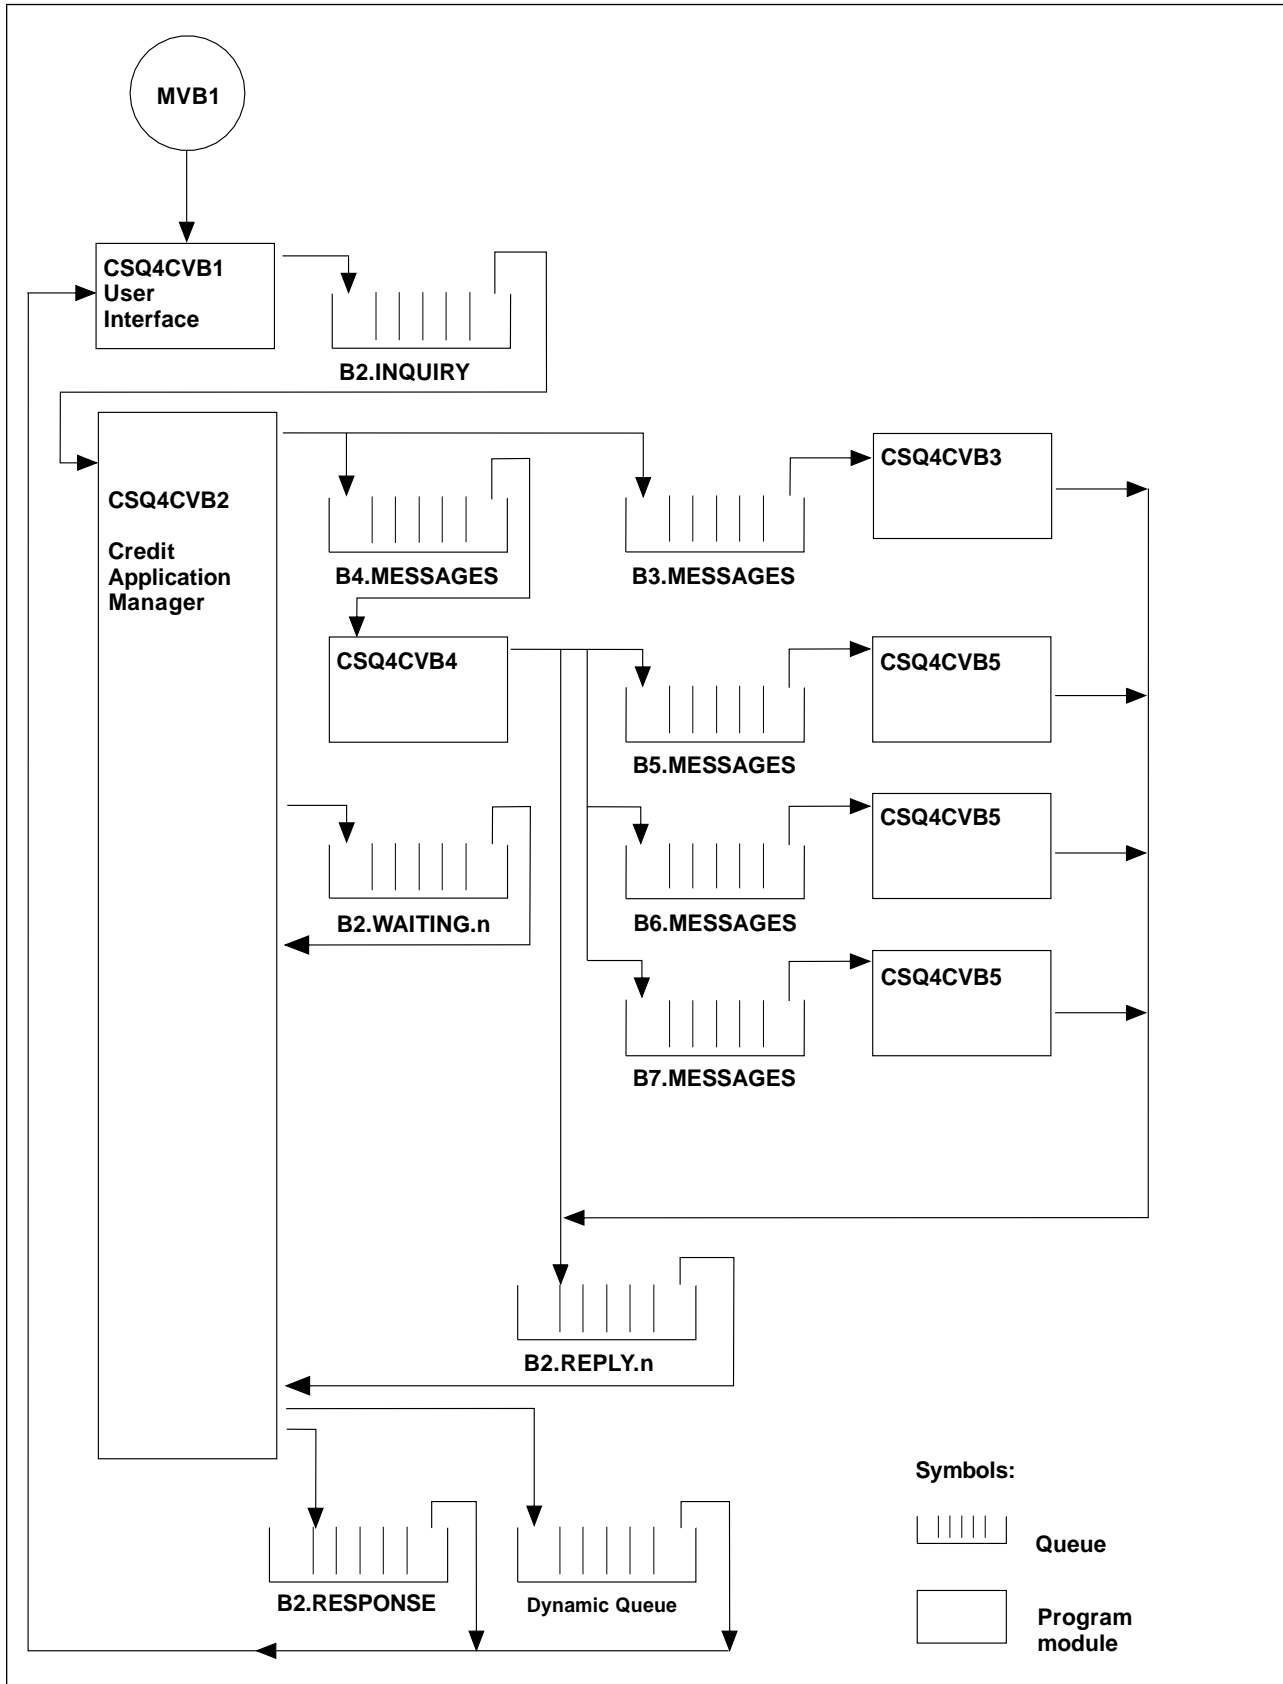
Figure 46. Programs and queues for the Credit Check sample application (COBOL programs only). In the sample application, the queue names shown in this figure have the prefix 'CSQ4SAMP.'

User-interface program (CSQ4CVB1)

When you start the conversational-mode CICS transaction MVB1, this starts the user-interface program for the application. This program puts inquiry messages on queue CSQ4SAMP.B2.INQUIRY and gets replies to those inquiries from a reply-to queue that it specifies when it makes the inquiry. From the user interface you can submit either immediate or batch inquiries:

- For immediate inquiries, the program creates a temporary dynamic queue that it uses as a reply-to queue. This means that each inquiry has its own reply-to queue.

- For batch inquiries, the user-interface program gets replies from the queue CSQ4SAMP.B2.RESPONSE. For simplicity, the program gets replies for all its inquiries from this one reply-to queue. It is easy to see that a bank might want to use a separate reply-to queue for each user of MVB1, so that they could each see replies to only those inquiries they had initiated.

Important differences between the properties of messages used in the application when in batch and immediate mode are:

- For batch working, the messages have a low priority, so they are processed after any loan requests that are entered in immediate mode. Also, the messages are persistent, so they are recovered if the application or the queue manager has to restart.

- For immediate working, the messages have a high priority, so they are processed before any loan requests that are entered in batch mode. Also, messages are not persistent so they are discarded if the application or the queue manager has to restart.

However, in all cases, the properties of loan request messages are propagated throughout the application. So, for example, all messages that result from a high-priority request will also have a high priority.

Credit application manager (CSQ4CVB2)

The Credit Application Manager (CAM) program performs most of the processing for the Credit Check application.

The CAM is started by the CKTI trigger monitor (supplied with MQSeries for MVS/ESA) when a trigger event occurs on either queue CSQ4SAMP.B2.INQUIRY or queue CSQ4SAMP.B2.REPLY.*n*, where *n* is an integer that identifies one of a set of reply queues. The trigger message contains data that includes the name of the queue on which the trigger event occurred.

The CAM uses queues with names of the form CSQ4SAMP.B2.WAITING.n to store information about inquiries it is processing. The queues are named so that they are each paired with a reply-to queue; for example, queue CSQ4SAMP.B2.WAITING.3 contains the input data for a particular inquiry, and queue CSQ4SAMP.B2.REPLY.3 contains a set of reply messages (from programs that query databases) all relating to that same inquiry. To understand the reasons behind this design, see "Separate inquiry and reply queues in the CAM" on page 419.

Start-up logic: If the trigger event occurs on queue CSQ4SAMP.B2.INQUIRY, the CAM opens the queue for shared access. It then tries to open each reply queue until a free one is found. If it cannot find a free reply queue, the CAM logs the fact and terminates normally.

If the trigger event occurs on queue CSQ4SAMP.B2.REPLY.n, the CAM opens the queue for exclusive access. If the return code reports that the object is already in use, the CAM terminates normally. If any other error occurs, the CAM logs the error and terminates. The CAM opens the corresponding waiting queue and the inquiry queue, then starts getting and processing messages. From the waiting queue, the CAM recovers details of partially-completed inquiries.

For the sake of simplicity in this sample, the names of the queues used are held in the program. In a business environment, the queue names would probably be held in a file accessed by the program.

Getting a message: The CAM first attempts to get a message from the inquiry queue using the MQGET call with the MQGMO_SET_SIGNAL option. If a message is available immediately, the message is processed; if no message is available, a signal is set.

The CAM then attempts to get a message from the reply queue, again using the MQGET call with the same option. If a message is available immediately, the message is processed; otherwise a signal is set.

When both signals are set, the program waits until one of the signals is posted. If a signal is posted to indicate a message is available, the message is retrieved and processed. If the signal expires or the queue manager is terminating, the program terminates.

Processing the message retrieved: A message retrieved by the CAM may be one of four types:

- An inquiry message
- A reply message
- A propagation message
- An unexpected or unwanted message

The CAM processes these messages as follows:

Inquiry message

Inquiry messages come from the user-interface program. It creates an inquiry message for each loan request.

For all loan requests, the CAM requests the average balance of the customer's checking account. It does this by putting a request message on alias queue CSQ4SAMP.B2.OUTPUT.ALIAS. This queue name resolves to queue CSQ4SAMP.B3.MESSAGES, which is processed by the checking-account program, CSQ4CVB3. When the CAM puts a message on this alias queue, it specifies the appropriate CSQ4SAMP.B2.REPLY.n queue for the reply-to queue. An alias queue is used here so that program CSQ4CVB3 can easily be replaced by another program that processes a base queue of a different name. To do this, you simply redefine the alias queue so that its name resolves to the new queue. Also, you could assign differing access authorities to the alias queue and to the base queue.

If a user requests a loan that is larger than 10000 units, the CAM initiates checks on other databases as well. It does this by putting a request message on queue CSQ4SAMP.B4.MESSAGES, which is processed by the distribution program, CSQ4CVB4. The process serving this queue propagates the message to queues served by programs that have access to other records such as credit card history, savings accounts, and mortgage payments. The data from these programs is returned to the reply-to queue specified in the put operation. Additionally, a propagation message is sent to the reply-to queue by this program to specify how many propagation messages have been sent.

In a business environment, the distribution program would probably reformat the data provided to match the format required by each of the other types of bank account.

Any of the queues referred to here can be on a remote system.

For each inquiry message, the CAM initiates an entry in the memory-resident Inquiry Record Table (IRT). This record contains:

- The *MsgId* of the inquiry message
- In the ReplyExp field, the number of responses expected (equal to the number of messages sent)
- In the ReplyRec field, the number of replies received (zero at this stage)
- In the PropsOut field, an indication of whether a propagation message is expected

The CAM copies the inquiry message on to the waiting queue with:

- *Priority* set to 3
- *CorrelId* set to the *MsgId* of the inquiry message
- The other message-descriptor fields set to those of the inquiry message

Propagation message
A propagation message contains the number of queues to which the distribution program has forwarded the inquiry. The message is processed as follows:

1. Add to the ReplyExp field of the appropriate record in the IRT the number of messages sent. This information is in the message.

2. Increment by 1 the ReplyRec field of the record in the IRT.

3. Decrement by 1 the PropsOut field of the record in the IRT.

4. Copy the message on to the waiting queue. The CAM sets the *Priority* to 2 and the other fields of the message descriptor to those of the propagation message.

Reply message
A reply message contains the response to one of the requests to the checking-account program or to one of the agency-query programs. Reply messages are processed as follows:

1. Increment by 1 the ReplyRec field of the record in the IRT.

2. Copy the message on to the waiting queue with *Priority* set to 1 and the other fields of the message descriptor set to those of the reply message.

3. If ReplyRec = ReplyExp, and PropsOut = 0, set the MsgComplete flag.

Other messages

The application does not expect other messages. However, the application might receive messages broadcast by the system, or reply messages with unknown *CorrelId*s.

The CAM puts these messages on queue CSQ4SAMP.DEAD.QUEUE, where they can be examined. If this put operation fails, the message is lost and the program continues. For more information on the design of this part of the program, see "How the sample handles unexpected messages" on page 420.

Sending an answer: When the CAM has received all the replies it is expecting for an inquiry, it processes the replies and creates a single response message. It consolidates into one message all the data from all reply messages that have the same *CorrelId*. This response is put on the reply-to queue specified in the original loan request. The response message is put within the same unit of work that contains the retrieval of the final reply message. This is to simplify recovery by ensuring that there is never a completed message on queue CSQ4SAMP.B2.WAITING.n.

Recovery of partially-completed inquiries: The CAM copies on to queue CSQ4SAMP.B2.WAITING.n all the messages that it receives. It sets the fields of the message descriptor like this:

- *Priority* is determined by the type of message:

    – For request messages, priority = 3
    – For datagrams, priority = 2
    – For reply messages, priority = 1

- *CorrelId* is set to the *MsgId* of the loan request message

- Other MQMD fields are copied from those of the received message

When an inquiry has been completed, the messages for a specific inquiry are removed from the waiting queue during answer processing. Therefore, at any time, the waiting queue contains all messages relevant to in-progress inquiries. These messages are used to recover details of in-progress inquiries if the program has to restart. The different priorities are set so that inquiry messages are recovered before propagations or reply messages.

Checking-account program (CSQ4CVB3)

The checking-account program is started by a trigger event on queue CSQ4SAMP.B3.MESSAGES. After it has opened the queue, this program gets a message from the queue using the MQGET call with the wait option, and with the wait interval set to 30 seconds.

The program searches VSAM data set CSQ4BAQ for the account number in the loan request message. It retrieves the corresponding account name, average balance, and credit worthiness index, or notes that the account number is not in the data set.

The program then puts a reply message (using the MQPUT1 call) on the reply-to queue named in the loan request message. For this reply message, the program:

- Copies the *CorrelId* of the loan request message
- Uses the MQPMO_PASS_IDENTITY_CONTEXT option

The program continues to get messages from the queue until the wait interval expires.

Distribution program (CSQ4CVB4)

The distribution program is started by a trigger event on queue CSQ4SAMP.B4.MESSAGES. To simulate the distribution of the loan request to other agencies that have access to records such as credit card history, savings accounts, and mortgage payments, the program puts a copy of the same message on all the queues in the namelist CSQ4SAMP.B4.NAMELIST. There are three of these queues, with names of the form CSQ4SAMP.B*n*.MESSAGES, where *n* is 5, 6, or 7. In a business application, the agencies could be at separate locations, so these queues could be remote queues. If you want to modify the sample application to show this, see "The Credit Check sample with multiple queue managers" on page 422.

The distribution program performs the following steps:

1. From the namelist, gets the names of the queues the program is to use. The program does this by using the MQINQ call to inquire about the attributes of the namelist object.

2. Opens these queues and also CSQ4SAMP.B4.MESSAGES.

3. Performs the following loop until there are no more messages on queue CSQ4SAMP.B4.MESSAGES:

    a. Get a message using the MQGET call with the wait option, and with the wait interval set to 30 seconds.

    b. Put a message on each queue listed in the namelist, specifying the name of the appropriate CSQ4SAMP.B2.REPLY.n queue for the reply-to queue. The program copies the *CorrelId* of the loan request message to these copy messages, and it uses the MQPMO_PASS_IDENTITY_CONTEXT option on the MQPUT call.

    c. Send a datagram message to queue CSQ4SAMP.B2.REPLY.n to show how many messages it has successfully put.

    d. Declare a syncpoint.

Agency-query program (CSQ4CVB5/CSQ4CCB5)

The agency-query program is supplied as both a COBOL program and a C program. Both programs have the same design. This shows that programs of different types can easily coexist within an MQSeries application, and that the program modules that comprise such an application can easily be replaced.

An instance of the program is started by a trigger event on any of these queues:

- For the COBOL program (CSQ4CVB5):

    – CSQ4SAMP.B5.MESSAGES
    – CSQ4SAMP.B6.MESSAGES
    – CSQ4SAMP.B7.MESSAGES

- For the C program (CSQ4CCB5), queue CSQ4SAMP.B8.MESSAGES

Note: If you want to use the C program, you must alter the definition of the namelist CSQ4SAMP.B4.NAMELIST to replace the queue CSQ4SAMP.B7.MESSAGES with CSQ4SAMP.B8.MESSAGES. To do this, you can use any one of:

  - The MQSeries for MVS/ESA operations and control panels
  - The ALTER NAMELIST command (described in the *MQSeries Command Reference* manual)
  - The CSQUTIL utility (described in the *MQSeries for MVS/ESA System Management Guide*)

After it has opened the appropriate queue, this program gets a message from the queue using the MQGET call with the wait option, and with the wait interval set to 30 seconds.

The program simulates the search of an agency's database by searching the VSAM data set CSQ4BAQ for the account number that was passed in the loan request message. It then builds a reply that includes the name of the queue it is serving and a credit-worthiness index. To simplify the processing, the credit-worthiness index is selected at random.

When putting the reply message, the program uses the MQPUT1 call and:

  - Copies the *CorrelId* of the loan request message
  - Uses the MQPMO_PASS_IDENTITY_CONTEXT option

The program sends the reply message to the reply-to queue named in the loan request message. (The name of the queue manager that owns the reply-to queue is also specified in the loan request message.)

Design considerations

This section discusses:

  - Why the CAM uses separate inquiry and reply queues
  - How the sample handles errors
  - How the sample handles unexpected messages
  - How the sample uses syncpoints
  - How the sample uses message context information

Separate inquiry and reply queues in the CAM

The application could use a single queue for both inquiries and replies, but it was designed to use separate queues for the following reasons:

  - When the program is handling the maximum number of inquiries, further inquiries can be left on the queue. If a single queue were being used, these would have to be taken off the queue and stored elsewhere.

  - Other instances of the CAM could be started automatically to service the same inquiry queue if message traffic was high enough to warrant it. But the program must track in-progress inquiries, and to do this it must get back all replies to inquiries it has initiated. If only one queue were used, the program would have to browse the messages to see if they were for this program or for another. This would make the operation much less efficient.

    The application can support multiple CAMs and can recover in-progress inquiries effectively by using paired reply-to and waiting queues.

  - The program can wait on multiple queues effectively by using signaling.

How the sample handles errors

The user-interface program handles errors very simply by reporting them directly to the user. The other programs do not have user interfaces, so they have to handle errors in other ways. Also, in many situations (for example, if an MQGET call fails) these other programs do not know the identity of the user of the application.

The other programs put error messages on a CICS temporary storage queue called CSQ4SAMP. You can browse this queue using the CICS-supplied transaction CEBR. The programs also write error messages to the CICS CSML log.

How the sample handles unexpected messages

When you design a message-queuing application, you must decide how to handle messages that arrive on a queue unexpectedly. The two basic choices are:

- The application must do no more work until it has processed the unexpected message. This probably means that the application must notify an operator, terminate itself, and ensure that it is not restarted automatically (it can do this by setting triggering off). This choice means that all processing for the application can be halted by a single unexpected message, and the intervention of an operator is required to restart the application.

- The application must remove the message from the queue it is serving, put the message in another location, and continue processing. The best place to put this message is on the system dead-letter queue.

If you choose the second option:

- An operator, or another program, should examine the messages that are put on the dead-letter queue to find out where the messages are coming from.

- An unexpected message is lost if it cannot be put on the dead-letter queue.

- An long unexpected message is truncated if it is longer than the limit for messages on the dead-letter queue, or longer than the buffer size in the program.

To ensure that the application smoothly handles all inquiries with minimal impact from outside activities, the Credit Check sample application uses the second option. To allow you to keep the sample separate from other applications that use the same queue manager, the Credit Check sample does not use the system dead-letter queue: instead, it uses its own dead-letter queue. This queue is named CSQ4SAMP.DEAD.QUEUE. The sample truncates any messages that are longer than the buffer area provided for the sample programs. You can use the Browse sample application to browse messages on this queue, or use the Print Message sample application to print the messages together with their message descriptors.

However, if you extend the sample to run across more than one queue manager, unexpected messages, or messages that cannot be delivered, could be put on the system dead-letter queue by the queue manager.

How the sample uses syncpoints

The programs in the Credit Check sample application declare syncpoints to ensure that:

- Only one reply message is sent in response to each expected message

- Multiple copies of unexpected messages are never put on the sample's dead-letter queue

- The CAM can recover the state of all partially-completed inquiries by getting persistent messages from its waiting queue

To achieve this, a single unit of work is used to cover the getting of a message, the processing of that message, and any subsequent put operations.

How the sample uses message context information

When the user-interface program (CSQ4CVB1) sends messages, it uses the MQPMO_DEFAULT_CONTEXT option. This means that the queue manager generates both identity and origin context information. The queue manager gets this information from the transaction that started the program (MVB1) and from the user ID that started the transaction.

When the CAM sends inquiry messages, it uses the MQPMO_PASS_IDENTITY_CONTEXT option. This means that the identity context information of the message being put is copied from the identity context of the original inquiry message. With this option, origin context information is generated by the queue manager.

When the CAM sends reply messages, it uses the MQPMO_ALTERNATE_USER_AUTHORITY option. This causes the queue manager to use an alternate user ID for its security check when the CAM opens a reply-to queue. The CAM uses the user ID of the submitter of the original inquiry message. This means that users are allowed to see replies to only those inquiries they have originated. The alternate user ID is obtained from the identity context information in the message descriptor of the original inquiry message.

When the query programs (CSQ4CVB3/4/5) send reply messages, they use the MQPMO_PASS_IDENTITY_CONTEXT option. This means that the identity context information of the message being put is copied from the identity context of the original inquiry message. With this option, origin context information is generated by the queue manager.

Note: The user ID associated with the MVB3/4/5 transactions requires access to the B2.REPLY.n queues. These user IDs may not be the same as those associated with the request being processed. To get around this possible security exposure, the query programs could use the MQPMO_ALTERNATE_USER_AUTHORITY option when putting their replies. This would mean that each individual user of MVB1 needs authority to open the B2.REPLY.n queues.

Use of message and correlation identifiers in the CAM

The application has to monitor the progress of all the "live" inquiries it is processing at any one time. To do this it uses the unique message identifier of each loan request message to associate all the information it has about each inquiry.

The CAM copies the *MsgId* of the inquiry message into the *CorrelId* of all the request messages it sends for that inquiry. The other programs in the sample (CSQ4CVB3 - 5) copy the *CorrelId* of each message they receive into the *CorrelId* of their reply message.

The Credit Check sample with multiple queue managers

You can use the Credit Check sample application to demonstrate distributed queuing by installing the sample on two queue managers. After you have installed the sample on each queue manager, you need to:

1. Set up the communication link between the two queue managers. For information on how to do this, see *MQSeries Intercommunication* book.

2. On one queue manager, create a local definition for each of the remote queues (on the other queue manager) that you want to use. These queues can be any of CSQ4SAMP.B*n*.MESSAGES, where *n* is 3, 5, 6, or 7. (These are the queues that are served by the checking-account program and the agency-query program.) For information on how to do this, see the description of the DEFINE QREMOTE command in the *MQSeries Command Reference* manual.

3. Change the definition of the namelist (CSQ4SAMP.B4.NAMELIST) so that it contains the names of the remote queues you choose to use. For information on how to do this, see the description of the ALTER NAMELIST command in the *MQSeries Command Reference* manual.

The IMS extension to the Credit Check sample

A version of the checking-account program is supplied as an IMS batch message processing (BMP) program. It is written in the C language.

The program performs the same function as the CICS version, except that to obtain the account information, the program reads an IMS database instead of a VSAM file. If you replace the CICS version of the checking-account program with the IMS version, you see no difference in the method of using the application.

To prepare and run the IMS version you must:

1. Follow the steps in "Preparing and running the Credit Check sample" on page 410.

2. Follow the steps in "Preparing the sample application for the IMS environment" on page 387.

3. Alter the definition of the alias queue CSQ4SAMP.B2.OUTPUT.ALIAS to resolve to queue CSQ4SAMP.B3.IMS.MESSAGES (instead of CSQ4SAMP.B3.MESSAGES). To do this, you can use any one of:

   • The MQSeries for MVS/ESA operations and control panels

   • The ALTER QALIAS command (described in the *MQSeries Command Reference* manual)

Another way of using the IMS checking-account program is to make it serve one of the queues that receives messages from the distribution program. In the delivered form of the Credit Check sample application, there are three of these queues (B5/6/7.MESSAGES), all served by the agency-query program. This program searches a VSAM data set. To compare the use of the VSAM data set and the IMS database, you could make the IMS checking-account program serve one of these queues instead. To do this, you must alter the definition of the namelist CSQ4SAMP.B4.NAMELIST to replace one of the CSQ4SAMP.B*n*.MESSAGES queues with the CSQ4SAMP.B3.IMS.MESSAGES queue. You can use any one of:

   • The MQSeries for MVS/ESA operations and control panels

- The ALTER NAMELIST command (described in the *MQSeries Command Reference* manual)

You can then run the sample from CICS transaction MVB1 as usual. The user sees no difference in operation or response. The IMS BMP stops either after receiving a stop message or after being inactive for five minutes.

Design of the IMS checking-account program (CSQ4ICB3)

This program runs as a BMP. You must start the program using its JCL before any MQSeries messages are sent to it.

The program searches an IMS database for the account number in the loan request messages. It retrieves the corresponding account name, average balance, and credit worthiness index.

The program sends the results of the database search to the reply-to queue named in the MQSeries message being processed. The message returned appends the account type and the results of the search to the message received so that the transaction building the response can confirm that the correct query is being processed. The message is in the form of three 79-character groups, as follows:

```
'Response from CHECKING ACCOUNT for name : JONES J B'
'     Opened 870530, 3-month average balance = 000012.57'
'     Credit worthiness index - BBB'
```

When running as a message-oriented BMP, the program drains the IMS message queue, then reads messages from the MQSeries for MVS/ESA queue and processes them. No information is received from the IMS message queue. The program reconnects to the queue manager after each checkpoint because the handles have been closed.

When running in a batch-oriented BMP, the program continues to be connected to the queue manager after each checkpoint because the handles are not closed.

The Message Handler sample

The Message Handler sample TSO application allows you to browse, forward, and delete messages on a queue. The sample is available in C and COBOL.

Preparing and running the sample

Follow these steps:

1. Prepare the sample as described in "Preparing sample applications for the TSO environment" on page 381.

2. Tailor the CLIST (CSQ4RCH1) provided for the sample to define:
   - The location of the panels
   - The location of the message file
   - The location of the load modules

CLIST CSQ4RCH1 may be used to run both the C and the COBOL version of the sample. The supplied version of CSQ4RCH1 runs the C version, and contains instructions on the tailoring necessary for the COBOL version.

Notes:

1. There are no sample queue definitions provided with the sample.

2. VS COBOL II does not support multitasking with ISPF, so you should not use the Message Handler sample application on both sides of a split screen. If you do, the results are unpredictable.

Using the sample

Having installed the sample and invoked it from the tailored CLIST CSQ4RCH1, the screen shown in Figure 47 on page 425 is displayed.

Enter the Queue Manager and Queue name to be viewed (case sensitive) and the message list screen is displayed (see Figure 48 on page 425). This screen shows the first 99 messages on the queue and, for each, shows the following fields:

Msg No Message number

Put Date MM/DD/YY Date the message was put on the queue (GMT)

Put Time HH:MM:SS Time the message was put on the queue (GMT)

Format Name MQMD.Format field

User Identifier MQMD.UserIdentifier field

Put Application Type MQMD.PutApplType field

Put Application Name MQMD.PutApplName field

The total number of messages on the queue is also displayed.

From this screen a message can be chosen, by number not by cursor position, and then displayed, see Figure 49 on page 426.

Once the message has been displayed it can be deleted, left on the queue, or forwarded to another queue. The `Forward to Q Mgr` and `Forward to Queue` fields are initialised with values from the MQMD, these can be changed prior to forwarding the message.

The sample design will only allow messages with unique `MsgId` / `CorrelId` combinations to be selected and displayed, this is because the message is retrieved using the `MsgId` and `CorrelId` as the key. If the key is not unique the sample cannot retrieve the chosen message with certainty.

```
       ---------------------- MQSeries for MVS/ESA - Samples -----------------------
        COMMAND ===>
                                                               User Id : JOHNJ



         Enter information. Press ENTER :

          Queue Manager Name    : _____ :

          Queue Name            : _____ :






          F1=HELP      F2=SPLIT     F3=END      F4=RETURN    F5=RFIND    F6=RCHANGE
          F7=UP        F8=DOWN      F9=SWAP     F10=LEFT     F11=RIGHT   F12=RETRIEVE
```

Figure 47. Initial screen for Message Handler sample

```
       ---------------------- MQSeries for MVS/ESA - Samples ------- Row 1 to 4 of 4
        COMMAND ==>

         Queue Manager   : VM03                                      :
         Queue           : MQEI.IMS.BRIDGE.QUEUE                     :

         Message number    01 of 04

         Msg  Put Date  Put Time  Format    User       Put Application
         No   MM/DD/YY  HH:MM:SS  Name      Identifier  Type     Name
         01   10/16/97  13:51:19  MQIMS     NTSFV02     00000002 NTSFV02A
         02   10/16/97  13:55:45  MQIMS     JOHNJ       00000011 EDIT\CLASSES\BIN\PROGTS
         03   10/16/97  13:54:01  MQIMS     NTSFV02     00000002 NTSFV02B
         04   10/16/97  13:57:22  MQIMS     johnj       00000011 EDIT\CLASSES\BIN\PROGTS
         ****************************** Bottom of data ********************************
```

Figure 48. Message list screen for Message Handler sample

```
--------------------- MQSeries for MVS/ESA - Samples ----- Row 1 to 35 of 35
COMMAND ==>

Queue Manager   : VM03                                             :
Queue           : MQEI.IMS.BRIDGE.QUEUE                            :
Forward to Q Mgr : VM03                                            :
Forward to Queue : QL.TEST.ISCRES1                                 :

Action : _ :   (D)elete  (F)orward

Message Content :
-------------------------------------------------------------------------------
Message Descriptor
  StrucId        :  MD
  Version        : 000000001
  Report         : 000000000
  MsgType        : 000000001
  Expiry         : -00000001
  Feedback       : 000000000
  Encoding       : 000000785
  CodedCharSetId : 000000500
  Format         :  MQIMS
  Priority       : 000000000
  Persistence    : 000000001
  MsgId          :  C3E2D840E5D4F0F34040404040404040AF6B30F0A89B7605 X
  CorrelId       :  0000000000000000000000000000000000000000000000000 X
  BackoutCount   : 000000000
  ReplyToQ       :  QL.TEST.ISCRES1
  ReplyToQMgr    :  VM03
  UserIdentifier :  NTSFV02
  AccountingToken :
          06F2F5F5F3F0F1000000000000000000000000000000000000000000000000000 X
  ApplIdentityData :
  PutApplType    : 000000002
  PutApplName    :  NTSFV02A
  PutDate        :  19971016
  PutTime        :  13511903
  ApplOriginData :

Message Buffer :   108 byte(s)
00000000 :  C9C9 C840 0000 0001 0000 0054 0000 0311    IIH ............
00000010 :  0000 0000 4040 4040 4040 4040 0000 0000    ....        ....
00000020 :  4040 4040 4040 4040 4040 4040 4040 4040
00000030 :  4040 4040 4040 4040 4040 4040 4040 4040
00000040 :  0000 0000 0000 0000 0000 0000 0000 0000    ................
00000050 :  40F1 C300 0018 0000 C9C1 D7D4 C4C9 F2F8     1C.....IAPMDI28
00000060 :  40C8 C5D3 D3D6 40E6 D6D9 D3C4              HELLO WORLD
****************************** Bottom of data ******************************
```

Figure 49. Chosen message is displayed

Design of the sample

This section describes the design of each of the programs that comprise the Message Handler sample application.

Object validation program

This requests a valid queue and queue manager name. If you do not specify a queue manager name, the default queue manager is used, if available. Only local queues can be used; an MQINQ is issued to check the queue type and an error is reported if the queue is not local. If the queue is not opened successfully, or the MQGET call is inhibited on the queue, error messages are returned indicating the CompCode and Reason return code.

Message list program

This displays a list of messages on a queue with information about them such as the putdate, puttime and the message format. The maximum number of messages stored in the list is 99. If there are more messages on the queue than this, the current queue depth is also displayed. To choose a message for display, type the message number into the entry field (the default is 01). If your entry is invalid, you will receive an appropriate error message.

Message content program

This displays message content. The content is formatted and split into two parts:

1. the message descriptor
2. the message buffer

The message descriptor shows the contents of each field on a separate line.

The message buffer is formatted depending on its contents. If the buffer holds a dead letter header (MQDLH) or a transmission queue header (MQXQH), these are formatted and displayed before the buffer itself.

Before the buffer data is formatted, a title line shows the buffer length of the message in bytes. The maximum buffer size is 32768 bytes, and any message longer than this is truncated. The full size of the buffer is displayed along with a message indicating that only the first 32768 bytes of the message are displayed.

The buffer data is formatted in two ways:

1. After the offset into the buffer is printed, the buffer data is displayed in HEX.

2. The buffer data is then displayed again as EBCDIC values. If any EBCDIC value cannot be printed, it prints a '.' instead.

You may enter 'D' for delete, or 'F' for forward into the action field. If you choose to forward the message, the *forward-to queue* and *queue manager name* must be filled in appropriately. The defaults for these fields are read from the message descriptor ReplyToQ and ReplyToQMgr fields.

If you forward a message, any header block stored in the buffer is stripped. If the message is forwarded successfully, it is removed from the original queue. If you enter invalid actions, error messages are displayed.

An example help panel is also available called CSQ4CHP9.

Message handler sample

Part 5. Appendixes

Appendix A. Language compilers and assemblers

Table 39 lists the language compilers and assemblers supported.

Table 39 (Page 1 of 3). Language compilers and assemblers		
Platform	**Language**	**Compiler/Assembler**
MQSeries for AIX	COBOL	Micro Focus COBOL V3.1 and V4.0 for UNIX systems IBM COBOL Set for AIX V1.0
	C	IBM C for AIX V3.1.4 IBM C Set++ for AIX V3.1 (C bindings only)
	C++	IBM C Set++ for AIX V3.1
	PL/I	IBM PL/I Set for AIX V1.1
MQSeries for AS/400	COBOL	IBM ILE COBOL for AS/400, V4R2 (5769-CB1)
	C	IBM ILE C for AS/400, V4R2 (5769-CX2)
	C++	IBM VisualAge for C++ for AS/400, V4R2 (5769-CX4)
	RPG	IBM SAA AD/Cycle RPG/400 compiler IBM ILE RPG for AS/400, V4R2 (5769-RG1)
MQSeries for AT&T GIS UNIX	C	AT&T GIS High Performance C V1.0b
	C++	AT&T C++ language system for AT&T GIS UNIX
MQSeries for Digital OpenVMS	COBOL	DEC COBOL, V5.0 (VAX) V2.2 (AXP)
	C	DEC C V5.0
	C++	DEC C++ V5.0 (VAX) V5.2 (AXP)
MQSeries for HP-UX	COBOL	COBOL Softbench V4.0 Micro Focus COBOL V4.0 for UNIX systems
	C	C Softbench V5.0 HP-UX ANSI C HP C++ V3.1
	C++	HP C++ V3.1
MQSeries for MVS/ESA	COBOL	VS COBOL II IBM SAA AD/Cycle COBOL/370
	C	C/370 Release 2.1.0 IBM SAA AD/Cycle C/370
	Assembler	Assembler H assembler IBM High Level Assembler/MVS assembler
	PL/I	OS PL/I Optimizing compiler IBM SAA AD/Cycle PL/I Compiler
MQSeries for OS/2 Warp	COBOL	Micro Focus COBOL, V4.0 IBM VisualAge for COBOL for OS/2 V1.1
	C	IBM C Set++ V2.1 (C bindings only) IBM VisualAge for C++ for OS/2, V3.0 (C bindings only) Borland C++ V2 (C bindings only)
	C++	IBM VisualAge for C++ for OS/2, V3.0
	PL/I	IBM VisualAge for PL/I for OS/2 IBM PL/I for OS/2 V1.2

433

Compilers and assemblers

Table 39 (Page 2 of 3). Language compilers and assemblers		
Platform	**Language**	**Compiler/Assembler**
MQSeries for SINIX and DC/OSx	COBOL	Micro Focus COBOL V3.2
	C	SINIX: C compiler (C-DS, MIPS) V1.1 DC/OSx: C4.0 compiler V4.0.1
MQSeries for SunOS	COBOL	Micro Focus COBOL V3.0
	C	SPARCompiler C V3.0.1
MQSeries for Sun Solaris	COBOL	Micro Focus COBOL for UNIX systems, V4.0
	C	SPARCompiler C V4.0 and V4.2
	C++	SPARCompiler C++ V4.1 and V4.2
MQSeries for Tandem NSK	COBOL	D30 or later
	C	D30 or later using the WIDE memory model (32-bit integers)
	TAL	D30 or later
MQSeries for Windows NT	COBOL	Micro Focus Object COBOL V3.3 or V4.0 for Windows NT Micro Focus Object COBOL V3.1J for Windows NT IBM VisualAge for COBOL for Windows NT V2.1
	C	IBM VisualAge for C++ for Windows V3.5 Microsoft Visual C++ V4.0 for Windows NT
	C++	IBM VisualAge for C++ for Windows V3.5 Microsoft Visual C++ V4.0 for Windows NT
	PL/I	IBM VisualAge for PL/I for Windows IBM PL/I for Windows V1.2
MQSeries for Windows V2.0	16-bit C	Microsoft Visual C++ V1.5
	32-bit C	Microsoft Visual C++ V2.0
	16-bit Basic	Microsoft Visual Basic V3.0 or V4.0
	32-bit Basic	Microsoft Visual Basic V4.0
MQSeries for Windows V2.1	C	Microsoft Visual C++ V4.0 Borland C
	Basic	Microsoft Visual Basic V4.0
DOS clients	COBOL	Micro Focus COBOL V3.3
	C	Microsoft C/C++ V7 Microsoft Visual C++ for Windows V4.0
Windows 3.1 clients	COBOL	Micro Focus Visual COBOL for Windows V3.3
	C	Microsoft C/C++ V7 Microsoft Visual C++ V2.0
	C++	Microsoft Visual C++ V1.5
Windows 95 clients	COBOL	Micro Focus COBOL Workbench V4.0
	C	Microsoft Visual C++ V2.0
	C++	IBM VisualAge for C++ V3.5 Microsoft Visual C++ V2.0

Table 39 (Page 3 of 3). Language compilers and assemblers		
Platform	**Language**	**Compiler/Assembler**
VM/ESA clients	COBOL	IBM VS COBOL II
	C	IBM C for VM Release 3.1
	PL/I	IBM OS/PL/I Release 2.3
	Assembler	IBM Assembler
	REXX	IBM VM/ESA REXX/VM
Note: RPG bindings are shown for the IBM SAA AD/Cycle RPG/400 compiler.		

Compilers and assemblers

Appendix B. MQI names in RPG

> ── **Note for RPG programmers** ──────────────────
>
> The names of MQI calls, call parameters, data types, fields of structures, and
> constants (applicable to RPG) are shown in this book with long names. When
> writing RPG programs, use the equivalent short RPG names given in the tables
> in this appendix. The RPG names are used in the *MQSeries for AS/400
> Application Programming Reference (RPG)* manual.

Names of calls

Table 40. Names of calls	
Long name	**RPG name**
MQCLOSE	CID=MQCLOS
MQCONN	CID=MQCONN
MQDISC	CID=MQDISC
MQGET	CID=MQGET
MQINQ	CID=MQINQ
MQOPEN	CID=MQOPEN
MQPUT	CID=MQPUT
MQPUT1	CID=MQPUT1
MQSET	CID=MQSET

Names of call parameters

Table 41 (Page 1 of 2). Names of call parameters	
Long name	**RPG name**
Buffer	*BUFFER*
BufferLength	*BUFLEN*
CharAttrs	*CHRATR*
CharAttrLength	*CALEN*
CompCode	*CMPCOD*
DataLength	*DATLEN*
GetMsgOpts	*GMO*
Hconn	*HCONN*
Hobj	*HOBJ*
IntAttrs	*INTATR*
IntAttrCount	*IACNT*
MsgDesc	*MSGDSC*

Table 41 (Page 2 of 2). Names of call parameters	
Long name	**RPG name**
QMgrName	*QMNAME*
ObjDesc	*OBJDSC*
Options	*OPTS*
PutMsgOpts	*PMO*
Reason	*REASON*
Selectors	*SELS*
SelectorCount	*SELCNT*

Data types

Many of the data types have the same RPG name and long name, for example MQDLH, MQGMO, MQMD, and so on. The following table shows only those data types where the RPG representation is different from the long name.

Table 42. Data types	
Long name	**RPG representation**
MQBYTE	1-byte bit string
MQBYTEn	n-byte bit string
MQCHAR	1 character
MQCHARn	n-byte character string
MQHCONN (structure field)	9-digit binary integer
MQHCONN (call parameter)	9-digit decimal integer
MQHOBJ (structure field)	9-digit binary integer
MQHOBJ (call parameter)	9-digit decimal integer
MQLONG (structure field)	9-digit binary integer
MQLONG (call parameter)	9-digit decimal integer
PMQLONG	9-digit binary integer

Names of fields in structures

Table 43 (Page 1 of 9). Names of fields in structures	
Long name	**RPG name**
AccountingToken (MQMD)	*MDACC*
AccountingToken (MQPMR)	*PRACC*
AlternateUserId (MQOD)	*ODAU*
ApplId (MQTMC)	*TCAI*
ApplId (MQTM)	*TMAI*
ApplIdentityData (MQMD)	*MDAID*
ApplOriginData (MQMD)	*MDAOD*
ApplType (MQTMC)	*TCAT*
ApplType (MQTM)	*TMAT*

Table 43 (Page 2 of 9). Names of fields in structures	
Long name	**RPG name**
AppOptions (MQDXP)	*DXAOP*
Authenticator (MQIIH)	*IIAUT*
BackoutCount (MQMD)	*MDBOC*
BatchInterval (MQCD)	*CDBI*
BatchSize (MQCD)	*CDBS*
CapabilityFlags(MQCXP)	*CXCAP*
ChannelName (MQCD)	*CDCHN*
ChannelType (MQCD)	*CDCHT*
CodedCharSetId (MQDH)	*DHCSI*
CodedCharSetId (MQDLH)	*DLCSI*
CodedCharSetId (MQDXP)	*DXCSI*
CodedCharSetId (MQIIH)	*IICSI*
CodedCharSetId (MQMD)	*MDCSI*
CodedCharSetId (MQMDE)	*MECSI*
CodedCharSetId (MQRMH)	*RMCSI*
CodedCharSetId (MQCFSL)	*SLCSI*
CodedCharSetId (MQCFST)	*STCSI*
Command (MQCFH)	*FHCMD*
CommitMode (MQIIH)	*IICMT*
CompCode (MQDXP)	*DXCC*
CompCode (MQCFH)	*FHCMP*
CompCode (MQRR)	*RRCC*
ConnectionName (MQCD)	*CDCON*
Context (MQPMO)	*PMCT*
Control (MQCFH)	*FHCTL*
CorrelId (MQMD)	*MDCID*
CorrelId (MQPMR)	*PRCID*
Count (MQCFIL)	*ILCNT*
Count (MQCFSL)	*SLCNT*
DataConversion (MQCD)	*CDDC*
DataLength (MQ)	*DXLEN*
DataLogicalLength (MQRMH)	*RMDL*
DataLogicalOffset (MQRMH)	*RMDO*
DataLogicalOffset2 (MQRMH)	*RMDO2*
Desc (MQCD)	*CDDES*
DestEnvLength (MQRMH)	*RMDEL*
DestEnvOffset (MQRMH)	*RMDEO*
DestNameLength (MQRMH)	*RMDNL*

MQI names in RPG

Table 43 (Page 3 of 9). Names of fields in structures	
Long name	**RPG name**
DestNameOffset (MQRMH)	*RMDNO*
DestQMgrName (MQDLH)	*DLDM*
DestQName (MQDLH)	*DLDQ*
DiscInterval (MQCD)	*CDDI*
DynamicQName (MQOD)	*ODDN*
Encoding (MQDH)	*DHENC*
Encoding (MQDLH)	*DLENC*
Encoding (MQDXP)	*DXENC*
Encoding (MQIIH)	*IIENC*
Encoding (MQMD)	*MDENC*
Encoding (MQMDE)	*MEENC*
Encoding (MQRMH)	*RMENC*
EnvData (MQTMC)	*TCED*
EnvData (MQTM)	*TMED*
ExitData (MQCXP)	*CXDAT*
ExitDataLength (MQCD)	*CDXDL*
ExitId (MQCXP)	*CXXID*
ExitNameLength (MQCD)	*CDXNL*
ExitNumber (MQCXP)	*CXEXN*
ExitOptions (MQDXP)	*DXXOP*
ExitReason (MQCXP)	*CXREA*
ExitResponse (MQCXP)	*CXRES*
ExitResponse (MQDXP)	*DXRES*
ExitResponse2 (MQCXP)	*CXRE2*
ExitUserArea (MQCXP)	*CXUA*
Expiry (MQMD)	*MDEXP*
Feedback (MQCXP)	*CXFB*
Feedback (MQMD)	*MDFB*
Feedback (MQPMR)	*PRFB*
Flags (MQDH)	*DHFLG*
Flags (MQIIH	*IIFLG*
Flags (MQMDE)	*MEFLG*
Flags (MQRMH)	*RMFLG*
Format (MQDH)	*DHFMT*
Format (MQDLH)	*DLFMT*
Format (MQIIH)	*IIFMT*
Format (MQMD)	*MDFMT*
Format (MQMDE)	*MEFMT*

Table 43 (Page 4 of 9). Names of fields in structures	
Long name	**RPG name**
Format (MQRMH)	*RMFMT*
FAPLevel (MQCXP)	*CXFAP*
GroupId (MQMD)	*MDGID*
GroupId (MQMDE)	*MEGID*
GroupId (MQPMR)	*PRGID*
GroupStatus (MQGMO)	*GMGST*
Hconn (MQDXP)	*DXHCN*
HeaderLength (MQCXP)	*CXHDL*
HeartbeatInterval (MQCD)	*CDHBI*
InvalidDestCount (MQOD)	*ODIDC*
InvalidDestCount (MQPMO)	*PMIDC*
KnownDestCount (MQOD)	*ODKDC*
KnownDestCount (MQPMO)	*PMKDC*
LongRetryCount (MQCD)	*CDLRC*
LongRetryInterval (MQCD)	*CDLRI*
LTermOverride (MQIIH)	*IILTO*
MatchOptions (MQGMO)	*GMMO*
MaxSegmentLength (MQCXP)	*CXMSL*
ModeName (MQCD)	*CDMOD*
MsgDesc (MQXQH)	*XQMD*
MsgExit (MQCD)	*CDMSX*
MsgExitsDefined (MQCD)	*CDMXD*
MsgExitPtr (MQCD)	*CDMXP*
MsgFlags (MQMD)	*MDMFL*
MsgFlags (MQMDE)	*MEMFL*
MsgId (MQMD)	*MDMID*
MsgId (MQPMR)	*PRMID*
MsgRetryCount (MQCD)	*CDMRC*
MsgRetryCount (MQCXP)	*CXMRC*
MsgRetryExit (MQCD)	*CDMRX*
MsgRetryInterval (MQCD)	*CDMRI*
MsgRetryInterval (MQCXP)	*CXMRI*
MsgRetryReason (MQCXP)	*CXMRR*
MsgRetryUserData (MQCD)	*CDMRD*
MsgSeqNumber (MQCFH)	*FHSEQ*
MsgSeqNumber (MQMD)	*MDSEQ*
MsgSeqNumber (MQMDE)	*MESEQ*
MsgType (MQMD)	*MDMT*

Table 43 (Page 5 of 9). Names of fields in structures	
Long name	**RPG name**
MsgUserData (MQCD)	CDMSD
MsgUserDataPtr (MQCD)	CDMUP
MCAName (MQCD)	CDMCA
MCAType (MQCD)	CDCAT
MCAUserIdentifier (MQCD)	CDAUI
MFSMapName (MQIIH)	IIMMN
NonPersistentMsgSpeed (MQCD)	CDNPM
ObjectInstanceId (MQRMH)	RMOII
ObjectName (MQOD)	ODON
ObjectName (MQOR)	ORON
ObjectQMgrName (MQOD)	ODMN
ObjectQMgrName (MQOD)	ODMN
ObjectQMgrName (MQOR)	ORMN
ObjectRecOffset (MQDH)	DHORO
ObjectRecOffset (MQOD)	ODORO
ObjectType (MQOD)	ODOT
ObjectRecPtr (MQOD)	ODORP
ObjectType (MQRMH)	RMOT
Offset (MQMD)	MDOFF
Offset (MQMDE)	MEOFF
Options (MQGMO)	GMOPT
Options (MQPMO)	PMOPT
OriginalLength (MQMD)	MDOLN
OriginalLength (MQMDE)	MEOLN
Parameter (MQCFIL)	ILPRM
Parameter (MQCFIN)	INPRM
Parameter (MQCFSL)	SLPRM
Parameter (MQCFST)	STPRM
ParameterCount (MQCFH)	FHCNT
PartnerName (MQCXP)	CXPNM
Password (MQCD)	CDPW
Persistence (MQMD)	MDPER
Priority (MQMD)	MDPRI
Persistence (MQMD)	MDPER
ProcessName (MQTM)	TMPN
PutApplName (MQDLH)	DLPAN
PutApplName (MQMD)	MDPAN
PutApplType (MQDLH)	DLPAT

Table 43 (Page 6 of 9). Names of fields in structures	
Long name	**RPG name**
PutApplType (MQMD)	*MDPAT*
PutAuthority (MQCD)	*CDPA*
PutDate (MQDLH)	*DLPD*
PutDate (MQMD)	*MDPD*
PutMsgRecFields (MQDH)	*DHPRF*
PutMsgRecFields (MQPMO)	*PMPRF*
PutMsgRecOffset (MQDH)	*DHPRO*
PutMsgRecOffset (MQPMO)	*PMPRO*
PutMsgRecPtr (MQPMO)	*PMPRP*
PutTime (MQDLH)	*DLPT*
PutTime (MQMD)	*MDPT*
QMgrName (MQCD)	*CDQM*
QName (MQTMC)	*TCQN*
QName (MQTM)	*TMQN*
Reason (MQDLH)	*DLREA*
Reason (MQDXP)	*DXREA*
Reason (MQCFH)	*FHREA*
Reason (MQRR)	*RRREA*
ReceiveExit (MQCD)	*CDRCX*
ReceiveExitsDefined (MQCD)	*CDRXD*
ReceiveExitPtr (MQCD)	*CDRXP*
ReceiveUserData (MQCD)	*CDRCD*
ReceiveUserDataPtr (MQCD)	*CDRUP*
RecsPresent (MQDH)	*DHCNT*
RecsPresent (MQOD)	*ODREC*
RecsPresent (MQOD)	*PMREC*
RemotePassword (MQCD)	*CDRPW*
RemoteQMgrName (MQXQH)	*XQRQM*
RemoteQMgrName (MQXQH)	*XQRQM*
RemoteQName (MQXQH)	*XQRQ*
RemoteUserIdentifier (MQCD)	*CDRUI*
ReplyToFormat (MQIIH)	*IIRFM*
ReplyToQ (MQMD)	*MDRQ*
ReplyToQMgr (MQMD)	*MDRM*
Report (MQMD)	*MDREP*
Reserved (MQIIH)	*IIRSV*
Reserved1 (MQGMO)	*GMRE1*
ResolvedQMgrName (MQPMO)	*PMRMN*

Table 43 (Page 7 of 9). Names of fields in structures	
Long name	**RPG name**
ResolvedQName (MQGMO)	*GMRQN*
ResolvedQName (MQPMO)	*PMRQN*
ResponseRecOffset (MQOD)	*ODRRO*
ResponseRecOffset (MQPMO)	*PMRRO*
ResponseRecPtr (MQOD)	*ODRRP*
ResponseRecPtr (MQPMO)	*PMRRP*
SecurityExit (MQCD)	*CDSCX*
SecurityScope (MQIIH)	*IISEC*
SecurityUserData (MQCD)	*CDSCD*
Segmentation (MQGMO)	*GMSEG*
SegmentStatus (MQGMO)	*GMSST*
SendExit (MQCD)	*CDSNX*
SendExitsDefined (MQCD)	*CDSXD*
SendExitPtr (MQCD)	*CDSXP*
SendUserData (MQCD)	*CDSND*
SendUserDataPtr (MQCD)	*CDSUP*
SeqNumberWrap (MQCD)	*CDSNW*
ShortConnectionName (MQCD)	*CDSCN*
ShortRetryCount (MQCD)	*CDSRC*
ShortRetryInterval (MQCD)	*CDSRI*
Signal1 (MQGMO)	*GMSG1*
Signal2 (MQGMO)	*GMSG2*
SrcEnvLength (MQRMH)	*RMSEL*
SrcEnvOffset (MQRMH)	*RMSEO*
SrcNameLength (MQRMH)	*RMSNL*
SrcNameOffset (MQRMH)	*RMSNO*
String (MQCFST)	*STSTR*
Strings (MQCFSL)	*SLSTR*
StringLength (MQCFSL)	*SLSTL*
StringLength (MQCFST)	*STSTL*
StrucId (MQCXP)	*CXSID*
StrucId (MQDH)	*DHSID*
StrucId (MQDLH)	*DLSID*
StrucId (MQDXP)	*DXSID*
StrucId (MQGMO)	*GMSID*
StrucId (MQIIH)	*IISID*
StrucId (MQMD)	*MDSID*
StrucId (MQMDE)	*MESID*

Table 43 (Page 8 of 9). Names of fields in structures	
Long name	**RPG name**
StrucId (MQOD)	*ODSID*
StrucId (MQPMO)	*PMSID*
StrucId (MQRMH)	*RMSID*
StrucId (MQTMC)	*TCSID*
StrucId (MQTM)	*TMSID*
StrucId (MQXQH)	*XQSID*
StrucLength (MQCD)	*CDLEN*
StrucLength (MQDH)	*DHLEN*
StrucLength (MQCFH)	*FHLEN*
StrucLength (MQIIH)	*IILEN*
StrucLength (MQCFIL	*ILLEN*
StrucLength (MQCFIN	*INLEN*
StrucLength (MQMDE)	*MELEN*
StrucLength (MQRMH)	*RMLEN*
StrucLength (MQCFSL	*SLLEN*
StrucLength (MQCFST	*STLEN*
Timeout (MQPMO)	*PMTO*
TpName (MQCD)	*CDTP*
TransportType (MQCD)	*CDTRT*
TranInstanceId (MQIIH)	*IITID*
TranState (MQIIH)	*IITST*
TriggerData (MQTMC)	*TCTD*
TriggerData (MQTM)	*TMTD*
TriggerData (MQTM)	*TMTD*
Type (MQCFH)	*FHTYP*
Type (MQCFIL)	*ILTYP*
Type (MQCFIN)	*INTYP*
Type (MQCFSL)	*SLTYP*
Type (MQCFST)	*STTYP*
UnknownDestCount (MQOD)	*ODUDC*
UnknownDestCount (MQPMO)	*PMUDC*
UserData (MQTMC)	*TCUD*
UserData (MQTM)	*TMUD*
UserIdentifier (MQCD)	*CDUID*
UserIdentifier (MQMD)	*MDUID*
Value (MQCFIN)	*INVAL*
Values (MQCFIL)	*ILVAL*
Version (MQCD)	*CDVER*

Table 43 (Page 9 of 9). Names of fields in structures	
Long name	**RPG name**
Version (MQCXP)	*CXVER*
Version (MQDH)	*DHVER*
Version (MQDLH)	*DLVER*
Version (MQDXP)	*DXVER*
Version (MQCFH)	*FHVER*
Version (MQGMO)	*GMVER*
Version (MQIIH)	*IIVER*
Version (MQMD)	*MDVER*
Version (MQMDE)	*MEVER*
Version (MQOD)	*ODVER*
Version (MQPMO)	*PMVER*
Version (MQRMH)	*RMVER*
Version (MQTMC)	*TCVER*
Version (MQTM)	*TMVER*
Version (MQXQH)	*XQVER*
WaitInterval (MQGMO)	*GMWI*
XmitQName (MQCD)	*CDXQ*

Names of constants

Table 44 (Page 1 of 21). Names of constants	
Long name	**RPG name**
MQ_ABEND_CODE_LENGTH	LNABNC
MQ_ACCOUNTING_TOKEN_LENGTH	LNACCT
MQ_APPL_IDENTITY_DATA_LENGTH	LNAIDD
MQ_APPL_NAME_LENGTH	LNAPPN
MQ_APPL_ORIGIN_DATA_LENGTH	LNAORD
MQ_ATTENTION_ID_LENGTH	LNATID
MQ_AUTHENTICATOR_LENGTH	LNAUTH
MQ_BRIDGE_NAME_LENGTH	LNBRGN
MQ_CANCEL_CODE_LENGTH	LNCNCL
MQ_CHANNEL_DATE_LENGTH	LNCDAT
MQ_CHANNEL_DESC_LENGTH	LNCHD
MQ_CHANNEL_NAME_LENGTH	LNCHN
MQ_CHANNEL_TIME_LENGTH	LNCTIM
MQ_CONN_NAME_LENGTH	LNCONN
MQ_CORREL_ID_LENGTH	LNCID
MQ_CREATION_DATE_LENGTH	LNCRTD
MQ_CREATION_TIME_LENGTH	LNCRTT

Table 44 (Page 2 of 21). Names of constants	
Long name	**RPG name**
MQ_EXIT_DATA_LENGTH	LNEXDA
MQ_EXIT_NAME_LENGTH	LNEXN
MQ_EXIT_USER_AREA_LENGTH	LNEXUA
MQ_FACILITY_LENGTH	LNFAC
MQ_FACILITY_LIKE_LENGTH	LNFACL
MQ_FORMAT_LENGTH	LNFMT
MQ_FUNCTION_LENGTH	LNFUNC
MQ_GROUP_ID_LENGTH	LNGID
MQ_LTERM_OVERRIDE_LENGTH	LNLTOV
MQ_LUWID_LENGTH	LNLUWI
MQ_MCA_JOB_NAME_LENGTH	LNMCAJ
MQ_MCA_NAME_LENGTH	LNMCAN
MQ_MFS_MAP_NAME_LENGTH	LNMFMN
MQ_MODE_NAME_LENGTH	LNMODN
MQ_MSG_HEADER_LENGTH	LNMHD
MQ_MSG_ID_LENGTH	LNMID
MQ_NAMELIST_DESC_LENGTH	LNNLD
MQ_NAMELIST_NAME_LENGTH	LNNLN
MQ_OBJECT_INSTANCE_ID_LENGTH	LNOIID
MQ_PASSWORD_LENGTH	LNPWRD
MQ_PROCESS_APPL_ID_LENGTH	LNPROA
MQ_PROCESS_DESC_LENGTH	LNPROD
MQ_PROCESS_ENV_DATA_LENGTH	LNPROE
MQ_PROCESS_NAME_LENGTH	LNPRON
MQ_PROCESS_USER_DATA_LENGTH	LNPROU
MQ_PUT_APPL_NAME_LENGTH	LNPAN
MQ_PUT_DATE_LENGTH	LNPDAT
MQ_PUT_TIME_LENGTH	LNPTIM
MQ_Q_DESC_LENGTH	LNQD
MQ_Q_MGR_DESC_LENGTH	LNQMD
MQ_Q_MGR_NAME_LENGTH	LNQMN
MQ_Q_NAME_LENGTH	LNQN
MQ_REMOTE_SYS_ID_LENGTH	LNRSID
MQ_SHORT_CONN_NAME_LENGTH	LNSCON
MQ_START_CODE_LENGTH	LNSTCO
MQ_STORAGE_CLASS_LENGTH	LNSTGC
MQ_TOTAL_EXIT_DATA_LENGTH	LNTEXD
MQ_TOTAL_EXIT_NAME_LENGTH	LNTEXN

Table 44 (Page 3 of 21). Names of constants	
Long name	**RPG name**
MQ_TP_NAME_LENGTH	LNTPN
MQ_TRAN_INSTANCE_ID_LENGTH	LNTIID
MQ_TRANSACTION_ID_LENGTH	LNTRID
MQ_TRIGGER_DATA_LENGTH	LNTRGD
MQ_USER_ID_LENGTH	LNUID
MQACT_NONE	ACNONE
MQAT_AIX	ATAIX
MQAT_CICS	ATCICS
MQAT_DEFAULT	ATDEF
MQAT_DOS	ATDOS
MQAT_IMS	ATIMS
MQAT_IMS_BRIDGE	ATIMSB
MQAT_MVS	ATMVS
MQAT_NO_CONTEXT	ATNCON
MQAT_OS2	ATOS2
MQAT_OS400	AT400
MQAT_QMGR	ATQM
MQAT_UNIX	ATUNIX
MQAT_UNKNOWN	ATUNK
MQAT_USER_FIRST	ATUFST
MQAT_USER_LAST	ATULST
MQAT_WINDOWS	ATWIN
MQAT_WINDOWS_NT	ATWINT
MQAT_XCF	ATXCF
MQCA_APPL_ID	CAAPPI
MQCA_BACKOUT_REQ_Q_NAME	CABRQN
MQCA_BASE_Q_NAME	CABASQ
MQCA_CHANNEL_AUTO_DEF_EXIT	CACADX
MQCA_COMMAND_INPUT_Q_NAME	CACMDQ
MQCA_CREATION_DATE	CACRTD
MQCA_CREATION_TIME	CACRTT
MQCA_DEAD_LETTER_Q_NAME	CADLQ
MQCA_DEF_XMIT_Q_NAME	CADXQN
MQCA_ENV_DATA	CAENVD
MQCA_FIRST	CAFRST
MQCA_INITIATION_Q_NAME	CAINIQ
MQCA_LAST	CALAST
MQCA_LAST_USED	CALSTU

Table 44 (Page 4 of 21). Names of constants	
Long name	**RPG name**
MQCA_NAMELIST_DESC	CALSTD
MQCA_NAMELIST_NAME	CALSTN
MQCA_NAMES	CANAMS
MQCA_PROCESS_DESC	CAPROD
MQCA_PROCESS_NAME	CAPRON
MQCA_Q_DESC	CAQD
MQCA_Q_MGR_DESC	CAQMD
MQCA_Q_MGR_NAME	CAQMN
MQCA_Q_NAME	CAQN
MQCA_REMOTE_Q_MGR_NAME	CARQMN
MQCA_REMOTE_Q_NAME	CARQN
MQCA_STORAGE_CLASS	CASTGC
MQCA_TRIGGER_DATA	CATRGD
MQCA_USER_DATA	CAUSRD
MQCA_XMIT_Q_NAME	CAXQN
MQCC_FAILED	CCFAIL
MQCC_OK	CCOK
MQCC_WARNING	CCWARN
MQCCSI_DEFAULT	CSDEF
MQCCSI_EMBEDDED	CSEMBD
MQCCSI_Q_MGR	CSQM
MQCD_CURRENT_VERSION	CDVERC
MQCD_LENGTH_4	CDLEN4
MQCD_VERSION_1	CDVER1
MQCD_VERSION_2	CDVER2
MQCD_VERSION_3	CDVER3
MQCD_VERSION_4	CDVER4
MQCDC_NO_SENDER_CONVERSION	DCNSND
MQCDC_SENDER_CONVERSION	DCSND
MQCF_DIST_LISTS	CFDIST
MQCF_NONE	CFNONE
MQCFC_LAST	CFCLST
MQCFC_NOT_LAST	CFCNOT
MQCFH_CURRENT_VERSION	FHVERC
MQCFH_STRUC_LENGTH	FHLENV
MQCFH_VERSION_1	FHVER1
MQCFIL_STRUC_LENGTH_FIXED	ILLENV
MQCFIN_STRUC_LENGTH	INLENV

Table 44 (Page 5 of 21). Names of constants	
Long name	**RPG name**
MQCFSL_STRUC_LENGTH_FIXED	SLLENV
MQCFST_STRUC_LENGTH_FIXED	STLENV
MQCFT_COMMAND	CFTCMD
MQCFT_EVENT	CFTEVT
MQCFT_INTEGER	CFTINT
MQCFT_INTEGER_LIST	CFTINL
MQCFT_RESPONSE	CFTRSP
MQCFT_STRING	CFTSTR
MQCFT_STRING_LIST	CFTSTL
MQCHAD_DISABLED	CHADDI
MQCHAD_ENABLED	CHADEN
MQCHT_ALL	CTALL
MQCHT_CLNTCONN	CTCLCN
MQCHT_RECEIVER	CTRCVR
MQCHT_REQUESTER	CTREQR
MQCHT_SENDER	CTSEND
MQCHT_SERVER	CTSERV
MQCHT_SVRCONN	CTSVCN
MQCI_NEW_SESSION	CINEWS
MQCI_NONE	CINONE
MQCMD_CHANGE_CHANNEL	CMCHCH
MQCMD_CHANGE_PROCESS	CMCHPR
MQCMD_CHANGE_Q	CMCHQ
MQCMD_CHANGE_Q_MGR	CMCHQM
MQCMD_CHANNEL_EVENT	CMCHEV
MQCMD_CLEAR_Q	CMCLQ
MQCMD_COPY_CHANNEL	CMCPCH
MQCMD_COPY_PROCESS	CMCPPR
MQCMD_COPY_Q	CMCPQ
MQCMD_CREATE_CHANNEL	CMCRCH
MQCMD_CREATE_PROCESS	CMCRPR
MQCMD_CREATE_Q	CMCRQ
MQCMD_DELETE_CHANNEL	CMDLCH
MQCMD_DELETE_PROCESS	CMDLPR
MQCMD_DELETE_Q	CMDLQ
MQCMD_ESCAPE	CMESC
MQCMD_INQUIRE_CHANNEL	CMINCH
MQCMD_INQUIRE_CHANNEL_NAMES	CMINCN

Table 44 (Page 6 of 21). Names of constants	
Long name	**RPG name**
MQCMD_INQUIRE_CHANNEL_STATUS	CMINCS
MQCMD_INQUIRE_PROCESS	CMINPR
MQCMD_INQUIRE_PROCESS_NAMES	CMINPN
MQCMD_INQUIRE_Q	CMINQ
MQCMD_INQUIRE_Q_MGR	CMINQM
MQCMD_INQUIRE_Q_NAMES	CMINQN
MQCMD_PERFM_EVENT	CMPFEV
MQCMD_PING_CHANNEL	CMPNCH
MQCMD_PING_Q_MGR	CMPNQM
MQCMD_Q_MGR_EVENT	CMQMEV
MQCMD_RESET_CHANNEL	CMRSCH
MQCMD_RESET_Q_STATS	CMRSQS
MQCMD_RESOLVE_CHANNEL	CMRVCH
MQCMD_START_CHANNEL	CMSTCH
MQCMD_START_CHANNEL_INIT	CMSTCI
MQCMD_START_CHANNEL_LISTENER	CMSTCL
MQCMD_STOP_CHANNEL	CMSPCH
MQCMDL_LEVEL_1	CMLVL1
MQCMDL_LEVEL_101	CML101
MQCMDL_LEVEL_110	CML101
MQCMDL_LEVEL_114	CML114
MQCMDL_LEVEL_120	CML120
MQCMDL_LEVEL_200	CML200
MQCMDL_LEVEL_201	CML201
MQCMDL_LEVEL_221	CML221
MQCMDL_LEVEL_230	CML230
MQCMDL_LEVEL_320	CML320
MQCMDL_LEVEL_500	CML500
MQCO_DELETE	CODEL
MQCO_DELETE_PURGE	COPURG
MQCO_NONE	CONONE
MQCXP_CURRENT_VERSION	CXVERC
MQCXP_STRUC_ID	CXSIDV
MQCXP_VERSION_1	CXVER1
MQCXP_VERSION_2	CXVER2
MQCXP_VERSION_3	CXVER3
MQDCC_DEFAULT_CONVERSION	DCCDEF
MQDCC_NONE	DCCNON

Table 44 (Page 7 of 21). Names of constants	
Long name	**RPG name**
MQDCC_SOURCE_ENC_FACTOR	DCCSFA
MQDCC_SOURCE_ENC_MASK	DCCSMA
MQDCC_SOURCE_ENC_NATIVE	DCCSNA
MQDCC_SOURCE_ENC_NORMAL	DCCSNO
MQDCC_SOURCE_ENC_REVERSED	DCCSRE
MQDCC_SOURCE_ENC_UNDEFINED	DCCSUN
MQDCC_TARGET_ENC_FACTOR	DCCTFA
MQDCC_TARGET_ENC_MASK	DCCTMA
MQDCC_TARGET_ENC_NATIVE	DCCTNA
MQDCC_TARGET_ENC_NORMAL	DCCTNO
MQDCC_TARGET_ENC_REVERSED	DCCTRE
MQDCC_TARGET_ENC_UNDEFINED	DCCTUN
MQDH_CURRENT_VERSION	DHVERC
MQDH_STRUC_ID	DHSIDV
MQDH_VERSION_1	DHVER1
MQDHF_NEW_MSG_IDS	DHFNEW
MQDHF_NONE	DHFNON
MQDL_NOT_SUPPORTED	DLNSUP
MQDL_SUPPORTED	DLSUPP
MQDLH_CURRENT_VERSION	DLVERC
MQDLH_STRUC_ID	DLSIDV
MQDLH_VERSION_1	DLVER1
MQDXP_CURRENT_VERSION	DXVERC
MQDXP_STRUC_ID	DXSIDV
MQDXP_VERSION_1	DXVER1
MQEI_UNLIMITED	EIULIM
MQENC_DECIMAL_MASK	ENDMSK
MQENC_DECIMAL_NORMAL	ENDNOR
MQENC_DECIMAL_REVERSED	ENDREV
MQENC_DECIMAL_UNDEFINED	ENDUND
MQENC_FLOAT_IEEE_NORMAL	ENFNOR
MQENC_FLOAT_IEEE_REVERSED	ENFREV
MQENC_FLOAT_MASK	ENFMSK
MQENC_FLOAT_S390	ENF390
MQENC_FLOAT_UNDEFINED	ENFUND
MQENC_INTEGER_MASK	ENIMSK
MQENC_INTEGER_NORMAL	ENINOR
MQENC_INTEGER_REVERSED	ENIREV

Table 44 (Page 8 of 21). Names of constants	
Long name	**RPG name**
MQENC_INTEGER_UNDEFINED	ENIUND
MQENC_NATIVE	ENNAT
MQENC_RESERVED_MASK	ENRMSK
MQEVR_DISABLED	EVRDIS
MQEVR_ENABLED	EVRENA
MQFB_	FB
MQFB_APPL_CANNOT_BE_STARTED	FBABEG
MQFB_APPL_FIRST	FBAFST
MQFB_APPL_LAST	FBALST
MQFB_APPL_TYPE_ERROR	FBATYP
MQFB_BUFFER_OVERFLOW	FBBUFO
MQFB_COA	FBCOA
MQFB_COD	FBCOD
MQFB_DATA_LENGTH_NEGATIVE	FBDLN
MQFB_DATA_LENGTH_TOO_BIG	FBDLTB
MQFB_DATA_LENGTH_ZERO	FBDLZ
MQFB_EXPIRATION	FBEXP
MQFB_IIH_ERROR	FBIIH
MQFB_IMS_ERROR	FBIERR
MQFB_IMS_FIRST	FBIFST
MQFB_IMS_LAST	FBILST
MQFB_LENGTH_OFF_BY_ONE	FBLOB1
MQFB_NAN	FBNAN
MQFB_NONE	FBNONE
MQFB_NOT_AUTHORIZED_FOR_IMS	FBNAFI
MQFB_PAN	FBPAN
MQFB_QUIT	FBQUIT
MQFB_STOPPED_BY_MSG_EXIT	FBSBMX
MQFB_SYSTEM_FIRST	FBSFST
MQFB_SYSTEM_LAST	FBSLST
MQFB_TM_ERROR	FBTM
MQFB_XMIT_Q_MSG_ERROR	FBXQME
MQFMT_ADMIN	FMADMN
MQFMT_COMMAND_1	FMCMD1
MQFMT_COMMAND_2	FMCMD2
MQFMT_DEAD_LETTER_HEADER	FMDLH
MQFMT_DIST_HEADER	FMDH
MQFMT_EVENT	FMEVNT

Table 44 (Page 9 of 21). Names of constants	
Long name	**RPG name**
MQFMT_IMS	FMIMS
MQFMT_IMS_VAR_STRING	FMIMVS
MQFMT_MD_EXTENSION	FMMDE
MQFMT_NONE	FMNONE
MQFMT_PCF	FMPCF
MQFMT_REF_MSG_HEADER	FMRMH
MQFMT_STRING	FMSTR
MQFMT_TRIGGER	FMTM
MQFMT_XMIT_Q_HEADER	FMXQH
MQGI_NONE	GINONE
MQGMO_ACCEPT_TRUNCATED_MSG	GMATM
MQGMO_ALL_MSGS_AVAILABLE	GMAMSA
MQGMO_ALL_SEGMENTS_AVAILABLE	GMASGA
MQGMO_BROWSE_FIRST	GMBRWF
MQGMO_BROWSE_MSG_UNDER_CURSOR	GMBRWC
MQGMO_BROWSE_NEXT	GMBRWN
MQGMO_COMPLETE_MSG	GMCMPM
MQGMO_CONVERT	GMCONV
MQGMO_CURRENT_VERSION	GMVERC
MQGMO_FAIL_IF_QUIESCING	GMFIQ
MQGMO_LOCK	GMLK
MQGMO_LOGICAL_ORDER	GMLOGO
MQGMO_MSG_UNDER_CURSOR	GMMUC
MQGMO_NO_SYNCPOINT	GMNSYP
MQGMO_NO_WAIT	GMNWT
MQGMO_NONE	GMNONE
MQGMO_STRUC_ID	GMSIDV
MQGMO_SYNCPOINT	GMSYP
MQGMO_SYNCPOINT_IF_PERSISTENT	GMPSYP
MQGMO_UNLOCK	GMUNLK
MQGMO_VERSION_1	GMVER1
MQGMO_VERSION_2	GMVER2
MQGMO_WAIT	GMWT
MQGS_LAST_MSG_IN_GROUP	GSLMIG
MQGS_MSG_IN_GROUP	GSMIG
MQGS_NOT_IN_GROUP	GSNIG
MQHC_DEF_HCONN	HCDEFH
MQHC_UNUSABLE_HCONN	HCUNUH

Table 44 (Page 10 of 21). Names of constants	
Long name	**RPG name**
MQHO_UNUSABLE_HOBJ	HOUNUH
MQIA_APPL_TYPE	IAAPPT
MQIA_AUTHORITY_EVENT	IAAUTE
MQIA_BACKOUT_THRESHOLD	IABTHR
MQIA_CHANNEL_AUTO_DEF	IACAD
MQIA_CHANNEL_AUTO_DEF_EVENT	IACADE
MQIA_CODED_CHAR_SET_ID	IACCSI
MQIA_COMMAND_LEVEL	IACMDL
MQIA_CURRENT_Q_DEPTH	IACDEP
MQIA_DEF_INPUT_OPEN_OPTION	IADINP
MQIA_DEF_PERSISTENCE	IADPER
MQIA_DEF_PRIORITY	IADPRI
MQIA_DEFINITION_TYPE	IADEFT
MQIA_DIST_LISTS	IADIST
MQIA_FIRST	IAFRST
MQIA_HARDEN_GET_BACKOUT	IAHGB
MQIA_HIGH_Q_DEPTH	IAHQD
MQIA_INDEX_TYPE	IAINDT
MQIA_INHIBIT_EVENT	IAINHE
MQIA_INHIBIT_GET	IAIGET
MQIA_INHIBIT_PUT	IAIPUT
MQIA_LAST	IALAST
MQIA_LAST_USED	IALSTU
MQIA_LOCAL_EVENT	IALCLE
MQIA_MAX_HANDLES	IAMHND
MQIA_MAX_MSG_LENGTH	IAMLEN
MQIA_MAX_PRIORITY	IAMPRI
MQIA_MAX_Q_DEPTH	IAMDEP
MQIA_MAX_UNCOMMITTED_MSGS	IAMUNC
MQIA_MSG_DELIVERY_SEQUENCE	IAMDS
MQIA_MSG_DEQ_COUNT	IAMDC
MQIA_MSG_ENQ_COUNT	IAMEC
MQIA_NAME_COUNT	IANAMC
MQIA_OPEN_INPUT_COUNT	IAOIC
MQIA_OPEN_OUTPUT_COUNT	IAOOC
MQIA_PERFORMANCE_EVENT	IAPFME
MQIA_PLATFORM	IAPLAT
MQIA_Q_DEPTH_HIGH_EVENT	IAQDHE

Table 44 (Page 11 of 21). Names of constants	
Long name	**RPG name**
MQIA_Q_DEPTH_HIGH_LIMIT	IAQDHL
MQIA_Q_DEPTH_LOW_EVENT	IAQDLE
MQIA_Q_DEPTH_LOW_LIMIT	IAQDLL
MQIA_Q_DEPTH_MAX_EVENT	IAQDME
MQIA_Q_SERVICE_INTERVAL	IAQSI
MQIA_Q_SERVICE_INTERVAL_EVENT	IAQSIE
MQIA_Q_TYPE	IAQTYP
MQIA_REMOTE_EVENT	IARMTE
MQIA_RETENTION_INTERVAL	IARINT
MQIA_SCOPE	IASCOP
MQIA_SHAREABILITY	IASHAR
MQIA_START_STOP_EVENT	IASSE
MQIA_SYNCPOINT	IASYNC
MQIA_TIME_SINCE_RESET	IATSR
MQIA_TRIGGER_CONTROL	IATRGC
MQIA_TRIGGER_DEPTH	IATRGD
MQIA_TRIGGER_INTERVAL	IATRGI
MQIA_TRIGGER_MSG_PRIORITY	IATRGP
MQIA_TRIGGER_TYPE	IATRGT
MQIA_USAGE	IAUSAG
MQIAUT_NONE	IAUNON
MQIAV_NOT_APPLICABLE	IAVNA
MQIAV_UNDEFINED	IAVUND
MQICM_COMMIT_THEN_SEND	ICMCTS
MQICM_SEND_THEN_COMMIT	ICMSTC
MQIIH_CURRENT_VERSION	IIVERC
MQIIH_LENGTH_1	IILEN1
MQIIH_NONE	IINONE
MQIIH_STRUC_ID	IISIDV
MQIIH_VERSION_1	IIVER1
MQISS_CHECK	ISSCHK
MQISS_FULL	ISSFUL
MQIT_CORREL_ID	ITCORI
MQIT_MSG_ID	ITMSGI
MQIT_NONE	ITNONE
MQITII_NONE	ITINON
MQITS_IN_CONVERSATION	ITSIC
MQITS_NOT_IN_CONVERSATION	ITSNIC

Table 44 (Page 12 of 21). Names of constants	
Long name	**RPG name**
MQMCAT_PROCESS	MCPROC
MQMCAT_THREAD	MCTHRD
MQMD_CURRENT_VERSION	MDVERC
MQMD_STRUC_ID	MDSIDV
MQMD_VERSION_1	MDVER1
MQMD_VERSION_2	MDVER2
MQMDE_CURRENT_VERSION	MEVERC
MQMDE_LENGTH_2	MELEN2
MQMDE_STRUC_ID	MESIDV
MQMDE_VERSION_2	MEVER2
MQMDEF_NONE	MEFNON
MQMDS_FIFO	MSFIFO
MQMDS_PRIORITY	MSPRIO
MQMF_ACCEPT_UNSUP_IF_XMIT_MASK	MFAUXM
MQMF_ACCEPT_UNSUP_MASK	MFAUM
MQMF_LAST_MSG_IN_GROUP	MFLMIG
MQMF_LAST_SEGMENT	MFLSEG
MQMF_MSG_IN_GROUP	MFMIG
MQMF_NONE	MFNONE
MQMF_REJECT_UNSUP_MASK	MFRUM
MQMF_SEGMENT	MFSEG
MQMF_SEGMENTATION_ALLOWED	MFSEGA
MQMF_SEGMENTATION_INHIBITED	MFSEGI
MQMI_NONE	MINONE
MQMO_MATCH_CORREL_ID	MOCORI
MQMO_MATCH_GROUP_ID	MOGRPI
MQMO_MATCH_MSG_ID	MOMSGI
MQMO_MATCH_MSG_SEQ_NUMBER	MOSEQN
MQMO_MATCH_OFFSET	MOOFFS
MQMO_NONE	MONONE
MQMT_APPL_FIRST	MTAFST
MQMT_APPL_LAST	MTALST
MQMT_DATAGRAM	MTDGRM
MQMT_REPLY	MTRPLY
MQMT_REPORT	MTRPRT
MQMT_REQUEST	MTRQST
MQMT_SYSTEM_FIRST	MTSFST
MQMT_SYSTEM_LAST	MTSLST

Table 44 (Page 13 of 21). Names of constants	
Long name	**RPG name**
MQNPMS_FAST	NPFAST
MQNPMS_NORMAL	NPNORM
MQOD_CURRENT_LENGTH	ODLENC
MQOD_CURRENT_VERSION	ODVERC
MQOD_STRUC_ID	ODSIDV
MQOD_VERSION_1	ODVER1
MQOD_VERSION_2	ODVER2
MQOII_NONE	OIINON
MQOL_UNDEFINED	OLUNDF
MQOO_ALTERNATE_USER_AUTHORITY	OOALTU
MQOO_BROWSE	OOBRW
MQOO_FAIL_IF_QUIESCING	OOFIQ
MQOO_INPUT_AS_Q_DEF	OOINPQ
MQOO_INPUT_EXCLUSIVE	OOINPX
MQOO_INPUT_SHARED	OOINPS
MQOO_INQUIRE	OOINQ
MQOO_OUTPUT	OOOUT
MQOO_PASS_ALL_CONTEXT	OOPASA
MQOO_PASS_IDENTITY_CONTEXT	OOPASI
MQOO_SAVE_ALL_CONTEXT	OOSAVA
MQOO_SET	OOSET
MQOO_SET_ALL_CONTEXT	OOSETA
MQOO_SET_IDENTITY_CONTEXT	OOSETI
MQOT_PROCESS	OTPRO
MQOT_Q	OTQ
MQOT_Q_MGR	OTQM
MQPA_CONTEXT	PACTXT
MQPA_DEFAULT	PADEF
MQPER_NOT_PERSISTENT	PENPER
MQPER_PERSISTENCE_AS_Q_DEF	PEQDEF
MQPER_PERSISTENT	PEPER
MQPL_AIX	PLAIX
MQPL_MVS	PLMVS
MQPL_OS2	PLOS2
MQPL_OS400	PL400
MQPL_UNIX	PLUNIX
MQPL_WINDOWS	PLWIN
MQPL_WINDOWS_NT	PLWINT

Table 44 (Page 14 of 21). Names of constants	
Long name	**RPG name**
MQPMO_ALTERNATE_USER_AUTHORITY	PMALTU
MQPMO_CURRENT_LENGTH	PMLENC
MQPMO_CURRENT_VERSION	PMVERC
MQPMO_DEFAULT_CONTEXT	PMDEFC
MQPMO_FAIL_IF_QUIESCING	PMFIQ
MQPMO_LOGICAL_ORDER	PMLOGO
MQPMO_NEW_CORREL_ID	PMNCID
MQPMO_NEW_MSG_ID	PMNMID
MQPMO_NO_CONTEXT	PMNOC
MQPMO_NO_SYNCPOINT	PMNSYP
MQPMO_NONE	PMNONE
MQPMO_PASS_ALL_CONTEXT	PMPASA
MQPMO_PASS_IDENTITY_CONTEXT	PMPASI
MQPMO_SET_ALL_CONTEXT	PMSETA
MQPMO_SET_IDENTITY_CONTEXT	PMSETI
MQPMO_STRUC_ID	PMSIDV
MQPMO_SYNCPOINT	PMSYP
MQPMO_VERSION_1	PMVER1
MQPMO_VERSION_2	PMVER2
MQPMRF_ACCOUNTING_TOKEN	PFACC
MQPMRF_CORREL_ID	PFCID
MQPMRF_FEEDBACK	PFFB
MQPMRF_GROUP_ID	PFGID
MQPMRF_MSG_ID	PFMID
MQPMRF_NONE	PFNONE
MQPRI_PRIORITY_AS_Q_DEF	PRQDEF
MQQA_BACKOUT_HARDENED	QABH
MQQA_BACKOUT_NOT_HARDENED	QABNH
MQQA_GET_ALLOWED	QAGETA
MQQA_GET_INHIBITED	QAGETI
MQQA_NOT_SHAREABLE	QANSHR
MQQA_PUT_ALLOWED	QAPUTA
MQQA_PUT_INHIBITED	QAPUTI
MQQA_SHAREABLE	QASHR
MQQDT_PERMANENT_DYNAMIC	QDPERM
MQQDT_PREDEFINED	QDPRE
MQQDT_TEMPORARY_DYNAMIC	QDTEMP
MQQSIE_HIGH	QSIEHI

Table 44 (Page 15 of 21). Names of constants	
Long name	**RPG name**
MQQSIE_NONE	QSIENO
MQQSIE_OK	QSIEOK
MQQT_ALIAS	QTALS
MQQT_LOCAL	QTLOC
MQQT_MODEL	QTMOD
MQQT_REMOTE	QTREM
MQRC_ALIAS_BASE_Q_TYPE_ERROR	RC2001
MQRC_ALREADY_CONNECTED	RC2002
MQRC_BRIDGE_STARTED	RC2125
MQRC_BRIDGE_STOPPED	RC2126
MQRC_BUFFER_ERROR	RC2004
MQRC_BUFFER_LENGTH_ERROR	RC2005
MQRC_CALL_IN_PROGRESS	RC2219
MQRC_CFH_ERROR	RC2235
MQRC_CFIL_ERROR	RC2236
MQRC_CFIN_ERROR	RC2237
MQRC_CFSL_ERROR	RC2238
MQRC_CFST_ERROR	RC2239
MQRC_CHANNEL_ACTIVATED	RC2295
MQRC_CHANNEL_AUTO_DEF_ERROR	RC2234
MQRC_CHANNEL_AUTO_DEF_OK	RC2233
MQRC_CHANNEL_CONV_ERROR	RC2284
MQRC_CHANNEL_NOT_ACTIVATED	RC2296
MQRC_CHANNEL_STARTED	RC2282
MQRC_CHANNEL_STOPPED	RC2283
MQRC_CHAR_ATTR_LENGTH_ERROR	RC2006
MQRC_CHAR_ATTRS_ERROR	RC2007
MQRC_CHAR_ATTRS_TOO_SHORT	RC2008
MQRC_CICS_BRIDGE_RESTRICTION	RC2187
MQRC_CONNECTION_BROKEN	RC2009
MQRC_CONTEXT_HANDLE_ERROR	RC2097
MQRC_CONTEXT_NOT_AVAILABLE	RC2098
MQRC_CONVERTED_MSG_TOO_BIG	RC2120
MQRC_DATA_LENGTH_ERROR	RC2010
MQRC_DBCS_ERROR	RC2150
MQRC_DEF_XMIT_Q_TYPE_ERROR	RC2198
MQRC_DEF_XMIT_Q_USAGE_ERROR	RC2199
MQRC_DEST_ENV_ERROR	RC2263

Table 44 (Page 16 of 21). Names of constants	
Long name	**RPG name**
MQRC_DEST_NAME_ERROR	RC2264
MQRC_DH_ERROR	RC2135
MQRC_DLH_ERROR	RC2141
MQRC_DYNAMIC_Q_NAME_ERROR	RC2011
MQRC_EXPIRY_ERROR	RC2013
MQRC_FEEDBACK_ERROR	RC2014
MQRC_FORMAT_ERROR	RC2110
MQRC_GET_INHIBITED	RC2016
MQRC_GMO_ERROR	RC2186
MQRC_GROUP_ID_ERROR	RC2258
MQRC_HANDLE_NOT_AVAILABLE	RC2017
MQRC_HCONN_ERROR	RC2018
MQRC_HEADER_ERROR	RC2142
MQRC_HOBJ_ERROR	RC2019
MQRC_IIH_ERROR	RC2148
MQRC_INCOMPLETE_GROUP	RC2241
MQRC_INCOMPLETE_MSG	RC2242
MQRC_INCONSISTENT_BROWSE	RC2259
MQRC_INCONSISTENT_CCSIDS	RC2243
MQRC_INCONSISTENT_ENCODINGS	RC2244
MQRC_INCONSISTENT_PERSISTENCE	RC2185
MQRC_INCONSISTENT_UOW	RC2245
MQRC_INHIBIT_VALUE_ERROR	RC2020
MQRC_INT_ATTR_COUNT_ERROR	RC2021
MQRC_INT_ATTR_COUNT_TOO_SMALL	RC2022
MQRC_INT_ATTRS_ARRAY_ERROR	RC2023
MQRC_INVALID_MSG_UNDER_CURSOR	RC2246
MQRC_MATCH_OPTIONS_ERROR	RC2247
MQRC_MD_ERROR	RC2026
MQRC_MDE_ERROR	RC2248
MQRC_MISSING_REPLY_TO_Q	RC2027
MQRC_MSG_FLAGS_ERROR	RC2249
MQRC_MSG_SEQ_NUMBER_ERROR	RC2250
MQRC_MSG_TOO_BIG_FOR_CHANNEL	RC2218
MQRC_MSG_TOO_BIG_FOR_Q	RC2030
MQRC_MSG_TOO_BIG_FOR_Q_MGR	RC2031
MQRC_MSG_TYPE_ERROR	RC2029
MQRC_MULTIPLE_REASONS	RC2136

Table 44 (Page 17 of 21). Names of constants

Long name	RPG name
MQRC_NAME_NOT_VALID_FOR_TYPE	RC2194
MQRC_NO_EXTERNAL_PARTICIPANTS	RC2121
MQRC_NO_MSG_AVAILABLE	RC2033
MQRC_NO_MSG_LOCKED	RC2209
MQRC_NO_MSG_UNDER_CURSOR	RC2034
MQRC_NONE	RCNONE
MQRC_NOT_AUTHORIZED	RC2035
MQRC_NOT_CONVERTED	RC2119
MQRC_NOT_OPEN_FOR_BROWSE	RC2036
MQRC_NOT_OPEN_FOR_INPUT	RC2037
MQRC_NOT_OPEN_FOR_INQUIRE	RC2038
MQRC_NOT_OPEN_FOR_OUTPUT	RC2039
MQRC_NOT_OPEN_FOR_PASS_ALL	RC2093
MQRC_NOT_OPEN_FOR_PASS_IDENT	RC2094
MQRC_NOT_OPEN_FOR_SET	RC2040
MQRC_NOT_OPEN_FOR_SET_ALL	RC2095
MQRC_NOT_OPEN_FOR_SET_IDENT	RC2096
MQRC_OBJECT_ALREADY_EXISTS	RC2100
MQRC_OBJECT_CHANGED	RC2041
MQRC_OBJECT_DAMAGED	RC2101
MQRC_OBJECT_IN_USE	RC2042
MQRC_OBJECT_NAME_ERROR	RC2152
MQRC_OBJECT_Q_MGR_NAME_ERROR	RC2153
MQRC_OBJECT_RECORDS_ERROR	RC2155
MQRC_OBJECT_TYPE_ERROR	RC2043
MQRC_OD_ERROR	RC2044
MQRC_OFFSET_ERROR	RC2251
MQRC_OPEN_FAILED	RC2137
MQRC_OPTION_NOT_VALID_FOR_TYPE	RC2045
MQRC_OPTIONS_ERROR	RC2046
MQRC_ORIGINAL_LENGTH_ERROR	RC2252
MQRC_OUTCOME_MIXED	RC2123
MQRC_OUTCOME_PENDING	RC2124
MQRC_PARTICIPANT_NOT_AVAILABLE	RC2122
MQRC_PCF_ERROR	RC2149
MQRC_PERSISTENCE_ERROR	RC2047
MQRC_PERSISTENT_NOT_ALLOWED	RC2048
MQRC_PMO_ERROR	RC2173

Table 44 (Page 18 of 21). Names of constants	
Long name	**RPG name**
MQRC_PMO_RECORD_FLAGS_ERROR	RC2158
MQRC_PRIORITY_ERROR	RC2050
MQRC_PRIORITY_EXCEEDS_MAXIMUM	RC2049
MQRC_PUT_INHIBITED	RC2051
MQRC_PUT_MSG_RECORDS_ERROR	RC2159
MQRC_Q_DELETED	RC2052
MQRC_Q_DEPTH_HIGH	RC2224
MQRC_Q_DEPTH_LOW	RC2225
MQRC_Q_FULL	RC2053
MQRC_Q_MGR_ACTIVE	RC2222
MQRC_Q_MGR_NAME_ERROR	RC2058
MQRC_Q_MGR_NOT_ACTIVE	RC2223
MQRC_Q_MGR_NOT_AVAILABLE	RC2059
MQRC_Q_MGR_QUIESCING	RC2161
MQRC_Q_MGR_STOPPING	RC2162
MQRC_Q_NOT_EMPTY	RC2055
MQRC_Q_SERVICE_INTERVAL_HIGH	RC2226
MQRC_Q_SERVICE_INTERVAL_OK	RC2227
MQRC_Q_SPACE_NOT_AVAILABLE	RC2056
MQRC_Q_TYPE_ERROR	RC2057
MQRC_RECS_PRESENT_ERROR	RC2154
MQRC_REMOTE_Q_NAME_ERROR	RC2184
MQRC_REPORT_OPTIONS_ERROR	RC2061
MQRC_RESOURCE_PROBLEM	RC2102
MQRC_RESPONSE_RECORDS_ERROR	RC2156
MQRC_RMH_ERROR	RC2220
MQRC_SECURITY_ERROR	RC2063
MQRC_SEGMENT_LENGTH_ZERO	RC2253
MQRC_SELECTOR_COUNT_ERROR	RC2065
MQRC_SELECTOR_ERROR	RC2067
MQRC_SELECTOR_LIMIT_EXCEEDED	RC2066
MQRC_SELECTOR_NOT_FOR_TYPE	RC2068
MQRC_SOURCE_BUFFER_ERROR	RC2145
MQRC_SOURCE_CCSID_ERROR	RC2111
MQRC_SOURCE_DECIMAL_ENC_ERROR	RC2113
MQRC_SOURCE_FLOAT_ENC_ERROR	RC2114
MQRC_SOURCE_INTEGER_ENC_ERROR	RC2112
MQRC_SOURCE_LENGTH_ERROR	RC2143

Table 44 (Page 19 of 21). Names of constants	
Long name	**RPG name**
MQRC_SRC_ENV_ERROR	RC2261
MQRC_SRC_NAME_ERROR	RC2262
MQRC_STORAGE_NOT_AVAILABLE	RC2071
MQRC_SYNCPOINT_LIMIT_REACHED	RC2024
MQRC_SYNCPOINT_NOT_AVAILABLE	RC2072
MQRC_TARGET_BUFFER_ERROR	RC2146
MQRC_TARGET_CCSID_ERROR	RC2115
MQRC_TARGET_DECIMAL_ENC_ERROR	RC2117
MQRC_TARGET_FLOAT_ENC_ERROR	RC2118
MQRC_TARGET_INTEGER_ENC_ERROR	RC2116
MQRC_TARGET_LENGTH_ERROR	RC2144
MQRC_TM_ERROR	RC2265
MQRC_TMC_ERROR	RC2191
MQRC_TRIGGER_CONTROL_ERROR	RC2075
MQRC_TRIGGER_DEPTH_ERROR	RC2076
MQRC_TRIGGER_MSG_PRIORITY_ERR	RC2077
MQRC_TRIGGER_TYPE_ERROR	RC2078
MQRC_TRUNCATED_MSG_ACCEPTED	RC2079
MQRC_TRUNCATED_MSG_FAILED	RC2080
MQRC_UNEXPECTED_ERROR	RC2195
MQRC_UNKNOWN_ALIAS_BASE_Q	RC2082
MQRC_UNKNOWN_DEF_XMIT_Q	RC2197
MQRC_UNKNOWN_OBJECT_NAME	RC2085
MQRC_UNKNOWN_OBJECT_Q_MGR	RC2086
MQRC_UNKNOWN_REMOTE_Q_MGR	RC2087
MQRC_UNKNOWN_REPORT_OPTION	RC2104
MQRC_UNKNOWN_XMIT_Q	RC2196
MQRC_UOW_IN_PROGRESS	RC2128
MQRC_UOW_NOT_AVAILABLE	RC2255
MQRC_WAIT_INTERVAL_ERROR	RC2090
MQRC_WRONG_GMO_VERSION	RC2256
MQRC_WRONG_MD_VERSION	RC2257
MQRC_XMIT_Q_TYPE_ERROR	RC2091
MQRC_XMIT_Q_USAGE_ERROR	RC2092
MQRC_XQH_ERROR	RC2260
MQRMH_CURRENT_VERSION	RMVERC
MQRMH_STRUC_ID	RMSIDV
MQRMH_VERSION_1	RMVER1

Table 44 (Page 20 of 21). Names of constants	
Long name	**RPG name**
MQRMHF_LAST	RMLAST
MQRMHF_NOT_LAST	RMNLST
MQRO_ACCEPT_UNSUP_IF_XMIT_MASK	ROAUXM
MQRO_ACCEPT_UNSUP_MASK	ROAUM
MQRO_COA	ROCOA
MQRO_COA_WITH_DATA	ROCOAD
MQRO_COA_WITH_FULL_DATA	ROCOAF
MQRO_COD	ROCOD
MQRO_COD_WITH_DATA	ROCODD
MQRO_COD_WITH_FULL_DATA	ROCODF
MQRO_COPY_MSG_ID_TO_CORREL_ID	ROCMTC
MQRO_DEAD_LETTER_Q	RODLQ
MQRO_DISCARD_MSG	RODISC
MQRO_EXCEPTION	ROEXC
MQRO_EXCEPTION_WITH_DATA	ROEXCD
MQRO_EXCEPTION_WITH_FULL_DATA	ROEXCF
MQRO_EXPIRATION	ROEXP
MQRO_EXPIRATION_WITH_DATA	ROEXPD
MQRO_EXPIRATION_WITH_FULL_DATA	ROEXPF
MQRO_NAN	RONAN
MQRO_NEW_MSG_ID	RONMI
MQRO_NONE	RONONE
MQRO_PAN	ROPAN
MQRO_PASS_CORREL_ID	ROPCI
MQRO_PASS_MSG_ID	ROPMI
MQRO_REJECT_UNSUP_MASK	RORUM
MQSEG_ALLOWED	SEGALW
MQSEG_INHIBITED	SEGIHB
MQSP_AVAILABLE	SPAVL
MQSP_NOT_AVAILABLE	SPNAVL
MQSS_LAST_SEGMENT	SSLSEG
MQSS_NOT_A_SEGMENT	SSNSEG
MQSS_SEGMENT	SSSEG
MQTC_OFF	TCOFF
MQTC_ON	TCON
MQTM_CURRENT_VERSION	TMVERC
MQTM_STRUC_ID	TMSIDV
MQTM_VERSION_1	TMVER1

Table 44 (Page 21 of 21). Names of constants	
Long name	**RPG name**
MQTMC_STRUC_ID	TCSIDV
MQTMC_VERSION_1	TCVER1
MQTT_DEPTH	TTDPTH
MQTT_EVERY	TTEVRY
MQTT_FIRST	TTFRST
MQTT_NONE	TTNONE
MQUS_NORMAL	USNORM
MQUS_TRANSMISSION	USTRAN
MQWI_UNLIMITED	WIULIM
MQXCC_CLOSE_CHANNEL	XCCLSC
MQXCC_OK	XCOK
MQXCC_SEND_AND_REQUEST_SEC_MSG	XCREQS
MQXCC_SEND_SEC_MSG	XCSNDS
MQXCC_SUPPRESS_EXIT	XCSUPE
MQXCC_SUPPRESS_FUNCTION	XCSUPF
MQXDR_CONVERSION_FAILED	XRFAIL
MQXDR_OK	XROK
MQXPT_LU62	XTLU62
MQXPT_NETBIOS	XTNET
MQXPT_SPX	XTSPX
MQXPT_TCP	XTTCP
MQXQH_CURRENT_VERSION	XQVERC
MQXQH_STRUC_ID	XQSIDV
MQXQH_VERSION_1	XQVER1
MQXR2_CONTINUE_CHAIN	X2CONC
MQXR2_DEFAULT_CONTINUATION	X2DEFC
MQXR2_PUT_WITH_DEF_ACTION	X2PDA
MQXR2_PUT_WITH_DEF_USERID	X2PDU
MQXR2_PUT_WITH_MSG_USERID	X2PMU
MQXR2_SUPPRESS_CHAIN	X2SUPC
MQXR2_USE_AGENT_BUFFER	X2UAB
MQXR2_USE_EXIT_BUFFER	X2UXB
MQXT_CHANNEL_MSG_EXIT	XTCMSG
MQXT_CHANNEL_MSG_RETRY_EXIT	XTCMR
MQXT_CHANNEL_RCV_EXIT	XTCRCV
MQXT_CHANNEL_SEC_EXIT	XTCSEC
MQXT_CHANNEL_SEND_EXIT	XTCSND
MQXUA_NONE	XUNONE

Appendix C. C language examples

The extracts in this appendix are mostly taken from the MQSeries for MVS/ESA
sample applications. They are applicable to all platforms, and any exception to this
is noted.

The examples in this appendix demonstrate the following techniques:

467

Connecting to a queue manager

Figure 50 demonstrates how to use the MQCONN call to connect a program to a queue manager in MVS batch. This extract is taken from the Browse sample application (program CSQ4BCA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
#include <cmqc.h>
  ⋮
static char Parm1[MQ_Q_MGR_NAME_LENGTH] ;
  ⋮
int  main(int argc, char *argv[] )
   {
   /*                                                */
   /*      Variables for MQ calls                    */
   /*                                                */
   MQHCONN Hconn;       /* Connection handle         */
   MQLONG  CompCode;    /* Completion code           */
   MQLONG  Reason;      /* Qualifying reason          */
   ⋮
   /* Copy the queue manager name, passed in the     */
   /* parm field, to Parm1                           */
   strncpy(Parm1,argv[1],MQ_Q_MGR_NAME_LENGTH);
   ⋮
   /*                                                */
   /* Connect to the specified queue manager.        */
   /*   Test the output of the connect call.  If the */
   /*   call fails, print an error message showing the */
   /*   completion code and reason code, then leave the */
   /*   program.                                     */
   /*                                                */
   MQCONN(Parm1,
         &Hconn,
         &CompCode,
         &Reason);

   if ((CompCode != MQCC_OK) | (Reason != MQRC_NONE))
      {
      sprintf(pBuff, MESSAGE_4_E,
            ERROR_IN_MQCONN, CompCode, Reason);
      PrintLine(pBuff);
      RetCode = CSQ4_ERROR;
      goto AbnormalExit2;
      }
   ⋮
   }
```

Figure 50. Using the MQCONN call (C language)

Disconnecting from a queue manager

Figure 51 demonstrates how to use the MQDISC call to disconnect a program from a queue manager in MVS batch. This extract is taken from the Browse sample application (program CSQ4BCA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   ⋮
   /*                                                */
   /*    Disconnect from the queue manager.  Test the */
   /*    output of the disconnect call.  If the call  */
   /*    fails, print an error message showing the    */
   /*    completion code and reason code.             */
   /*                                                */
   MQDISC(&Hconn,
         &CompCode,
         &Reason);

   if ((CompCode != MQCC_OK) || (Reason != MQRC_NONE))
      {
      sprintf(pBuff, MESSAGE_4_E,
            ERROR_IN_MQDISC, CompCode, Reason);
      PrintLine(pBuff);
      RetCode = CSQ4_ERROR;
      }
   ⋮
```

Figure 51. Using the MQDISC call (C language). The variables used in this code extract are those that were set in Figure 50.

Creating a dynamic queue

Figure 52 demonstrates how to use the MQOPEN call to create a dynamic queue. This extract is taken from the Mail Manager sample application (program CSQ4TCD1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   :
MQLONG  HCONN = 0;   /* Connection handle    */
MQHOBJ  HOBJ;        /* MailQ Object handle  */
MQHOBJ  HobjTempQ;   /* TempQ Object Handle  */
MQLONG  CompCode;    /* Completion code      */
MQLONG  Reason;      /* Qualifying reason    */
MQOD    ObjDesc = {MQOD_DEFAULT};
                     /* Object descriptor    */
MQLONG  OpenOptions; /* Options control MQOPEN */
   :
  /*----------------------------------------- */
  /* Initialize the Object Descriptor (MQOD)  */
  /* control block.  (The remaining fields    */
  /* are already initialized.)                */
  /*-----------------------------------------*/
  strncpy( ObjDesc.ObjectName,
           SYSTEM_REPLY_MODEL,
           MQ_Q_NAME_LENGTH );

  strncpy( ObjDesc.DynamicQName,
           SYSTEM_REPLY_INITIAL,
           MQ_Q_NAME_LENGTH );

  OpenOptions = MQOO_INPUT_AS_Q_DEF;
  /*-----------------------------------------*/
  /* Open the model queue and, therefore,    */
  /* create and open a temporary dynamic     */
  /* queue                                    */
  /*-----------------------------------------*/
  MQOPEN( HCONN,
          &ObjDesc,
          OpenOptions,
          &HobjTempQ,
          &CompCode,
          &Reason );

  if ( CompCode == MQCC_OK ) {
  :
  }
  else {
     /*--------------------------------------*/
     /* Build an error message to report the */
     /* failure of the opening of the model  */
     /* queue                                */
     /*--------------------------------------*/
     MQMErrorHandling( "OPEN TEMPQ", CompCode,
                        Reason );
     ErrorFound = TRUE;
  }
  return ErrorFound;
}
  :
```

Figure 52. Using the MQOPEN call to create a dynamic queue (C language)

Opening an existing queue

Figure 53 demonstrates how to use the MQOPEN call to open a queue that has already been defined. This extract is taken from the Browse sample application (program CSQ4BCA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
#include <cmqc.h>
   :
static char Parm1[MQ_Q_MGR_NAME_LENGTH];
   :
int  main(int argc, char *argv[] )
   {
   /*
   /*      Variables for MQ calls                 */
   /*
   MQHCONN Hconn ;            /* Connection handle    */
   MQLONG  CompCode;          /* Completion code      */
   MQLONG  Reason;            /* Qualifying reason    */
   MQOD    ObjDesc = { MQOD_DEFAULT };
                             /* Object descriptor    */
   MQLONG  OpenOptions;      /* Options that control */
                             /*  the MQOPEN call     */
   MQHOBJ  Hobj;             /* Object handle        */
   :
   /* Copy the queue name, passed in the parm field,  */
   /* to Parm2 strncpy(Parm2,argv[2],          */
   /* MQ_Q_NAME_LENGTH);                            */
   :
   /*                                             */
   /* Initialize the object descriptor (MQOD) control */
   /* block.  (The initialization default sets StrucId, */
   /* Version, ObjectType, ObjectQMgrName,         */
   /* DynamicQName, and AlternateUserid fields)    */
   /*                                             */
   strncpy(ObjDesc.ObjectName,Parm2,MQ_Q_NAME_LENGTH);
   :
   /* Initialize the other fields required for the open */
   /* call (Hobj is set by the MQCONN call).       */
   /*                                             */
   OpenOptions  = MQOO_BROWSE;
   :
   /*                                             */
   /* Open the queue.                             */
   /*   Test the output of the open call.  If the call */
   /*   fails, print an error message showing the  */
   /*   completion code and reason code, then bypass */
   /*   processing, disconnect and leave the program. */
   /*                                             */
   MQOPEN(Hconn,
          &ObjDesc,
          OpenOptions,
          &Hobj,
          &CompCode,
          &Reason);
```

Figure 53 (Part 1 of 2). Using the MQOPEN call to open an existing queue (C language)

```
   if ((CompCode != MQCC_OK) || (Reason != MQRC_NONE))
      {
      sprintf(pBuff, MESSAGE_4_E,
              ERROR_IN_MQOPEN, CompCode, Reason);
      PrintLine(pBuff);
      RetCode = CSQ4_ERROR;
      goto AbnormalExit1;        /* disconnect processing */
      }
   ⋮
   } /* end of main */
```

Figure 53 (Part 2 of 2). Using the MQOPEN call to open an existing queue (C language)

Closing a queue

Figure 54 demonstrates how to use the MQCLOSE call to close a queue. This extract is taken from the Browse sample application (program CSQ4BCA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   ⋮
   /*                                              */
   /*    Close the queue.                          */
   /*    Test the output of the close call.  If the call  */
   /*    fails, print an error message showing the    */
   /*    completion code and reason code.           */
   /*                                              */
   MQCLOSE(Hconn,
           &Hobj,
           MQCO_NONE,
           &CompCode,
           &Reason);

   if ((CompCode != MQCC_OK) || (Reason != MQRC_NONE))
      {
      sprintf(pBuff, MESSAGE_4_E,
              ERROR_IN_MQCLOSE, CompCode, Reason);
      PrintLine(pBuff);
      RetCode = CSQ4_ERROR;
      }
   ⋮
```

Figure 54. Using the MQCLOSE call (C language)

Putting a message using MQPUT

Figure 55 demonstrates how to use the MQPUT call to put a message on a queue. This extract is not taken from the sample applications supplied with MQSeries. For the names and locations of the sample applications, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313 and Chapter 31, "Sample programs for MQSeries for MVS/ESA" on page 375.

```
   ⋮
qput()
{

    MQMD    MsgDesc;
    MQPMO   PutMsgOpts;

    MQLONG  CompCode;
    MQLONG  Reason;
    MQHCONN Hconn;
    MQHOBJ  Hobj;

    char message_buffer[] = "MY MESSAGE";

   /*------------------------------*/
   /* Set up PMO structure.       */
   /*------------------------------*/
   memset(&PutMsgOpts, '\0', sizeof(PutMsgOpts));
   memcpy(PutMsgOpts.StrucId, MQPMO_STRUC_ID,
          sizeof(PutMsgOpts.StrucId));
    PutMsgOpts.Version = MQPMO_VERSION_1;
    PutMsgOpts.Options = MQPMO_SYNCPOINT;


   /*------------------------------*/
   /* Set up MD structure.        */
   /*------------------------------*/
   memset(&MsgDesc, '\0', sizeof(MsgDesc));
   memcpy(MsgDesc.StrucId, MQMD_STRUC_ID,
          sizeof(MsgDesc.StrucId));
   MsgDesc.Version     = MQMD_VERSION_1;
   MsgDesc.Expiry      = MQEI_UNLIMITED;
   MsgDesc.Report      = MQRO_NONE;
   MsgDesc.MsgType     = MQMT_DATAGRAM;
   MsgDesc.Priority    = 1;
   MsgDesc.Persistence = MQPER_PERSISTENT;
   memset(MsgDesc.ReplyToQ,
          '\0',
          sizeof(MsgDesc.ReplyToQ));
   /*-------------------------------------------------*/
   /* Put the message.                                */
   /*-------------------------------------------------*/
   MQPUT(Hconn, Hobj, &MsgDesc, &PutMsgOpts,
         sizeof(message_buffer), message_buffer,
         &CompCode, &Reason);
```

Figure 55 (Part 1 of 2). Using the MQPUT call (C language)

```
    /*------------------------------------*/
    /* Check completion and reason codes. */
    /*------------------------------------*/
    switch (CompCode)
        {
        case MQCC_OK:
             break;
        case MQCC_FAILED:
             switch (Reason)
                 {
                  case MQRC_Q_FULL:
                  case MQRC_MSG_TOO_BIG_FOR_Q:
                      break;
                  default:
                      break; /* Perform error processing */
                 }
             break;
        default:
             break;              /* Perform error processing */
        }
}
```

Figure 55 (Part 2 of 2). Using the MQPUT call (C language)

Putting a message using MQPUT1

Figure 56 demonstrates how to use the MQPUT1 call to open a queue, put a single message on the queue, then close the queue. This extract is taken from the Credit Check sample application (program CSQ4CCB5) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
    ⋮
MQLONG   Hconn;               /* Connection handle       */
MQHOBJ   Hobj_CheckQ;         /* Object handle           */
MQLONG   CompCode;            /* Completion code         */
MQLONG   Reason;             /* Qualifying reason        */
MQOD     ObjDesc   = {MQOD_DEFAULT};
                              /* Object descriptor       */
MQMD     MsgDesc   = {MQMD_DEFAULT};
                              /* Message descriptor       */
MQLONG   OpenOptions;        /* Control the MQOPEN call  */

MQGMO    GetMsgOpts = {MQGMO_DEFAULT};
                              /* Get Message Options      */
MQLONG   MsgBuffLen;         /* Length of message buffer */
CSQ4BCAQ MsgBuffer;          /* Message structure        */
MQLONG   DataLen;            /* Length of message        */

MQPMO    PutMsgOpts = {MQPMO_DEFAULT};
                              /* Put Message Options      */
CSQ4BQRM PutBuffer;          /* Message structure        */
MQLONG   PutBuffLen = sizeof(PutBuffer);
                              /* Length of message buffer */

    ⋮
```

Figure 56 (Part 1 of 2). Using the MQPUT1 call (C language)

```
void Process_Query(void)
   {
   /*                                             */
   /* Build the reply message                     */
   /*                                             */
    ⋮
   /*                                             */
   /* Set the object descriptor, message descriptor and */
   /* put message options to the values required to     */
   /* create the reply message.                   */
   /*                                             */

   strncpy(ObjDesc.ObjectName, MsgDesc.ReplyToQ,
           MQ_Q_NAME_LENGTH);
   strncpy(ObjDesc.ObjectQMgrName, MsgDesc.ReplyToQMgr,
           MQ_Q_MGR_NAME_LENGTH);

   MsgDesc.MsgType = MQMT_REPLY;
   MsgDesc.Report  = MQRO_NONE;
   memset(MsgDesc.ReplyToQ, ' ', MQ_Q_NAME_LENGTH);
   memset(MsgDesc.ReplyToQMgr, ' ', MQ_Q_MGR_NAME_LENGTH);

   memcpy(MsgDesc.MsgId, MQMI_NONE, sizeof(MsgDesc.MsgId));

   PutMsgOpts.Options = MQPMO_SYNCPOINT +
                        MQPMO_PASS_IDENTITY_CONTEXT;
   PutMsgOpts.Context = Hobj_CheckQ;

   PutBuffLen = sizeof(PutBuffer);

   MQPUT1(Hconn,
         &ObjDesc,
         &MsgDesc,
         &PutMsgOpts,
         PutBuffLen,
         &PutBuffer,
         &CompCode,
         &Reason);

   if (CompCode != MQCC_OK)
      {
      strncpy(TS_Operation, "MQPUT1",
              sizeof(TS_Operation));
      strncpy(TS_ObjName, ObjDesc.ObjectName,
              MQ_Q_NAME_LENGTH);
      Record_Call_Error();
      Forward_Msg_To_DLQ();
      }

   return;
   }
    ⋮
```

Figure 56 (Part 2 of 2). Using the MQPUT1 call (C language)

Getting a message

Figure 57 demonstrates how to use the MQGET call to remove a message from a queue. This extract is taken from the Browse sample application (program CSQ4BCA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
#include "cmqc.h"
  :
#define BUFFERLENGTH 80
  :
int  main(int argc, char *argv[] )
  {
  /*                                            */
  /*      Variables for MQ calls                */
  /*                                            */
  MQHCONN Hconn ;          /* Connection handle */
  MQLONG  CompCode;        /* Completion code    */
  MQLONG  Reason;          /* Qualifying reason  */
  MQHOBJ  Hobj;            /* Object handle      */
  MQMD    MsgDesc = { MQMD_DEFAULT };
                           /* Message descriptor */
  MQLONG  DataLength ;     /* Length of the message */
  MQCHAR  Buffer[BUFFERLENGTH+1];
                           /* Area for message data */
  MQGMO   GetMsgOpts = { MQGMO_DEFAULT };
                           /* Options which control */
                           /*  the MQGET call       */
  MQLONG  BufferLength = BUFFERLENGTH ;
                           /* Length of buffer      */
  :
  /*      No need to change the message descriptor */
  /*      (MQMD) control block because initialization */
  /*      default sets all the fields.               */
  /*                                                 */
  /*      Initialize the get message options (MQGMO) */
  /*      control block (the copy file initializes all */
  /*      the other fields).                          */
  /*                                                 */
  GetMsgOpts.Options = MQGMO_NO_WAIT      +
                       MQGMO_BROWSE_FIRST +
                       MQGMO_ACCEPT_TRUNCATED_MSG;
  /*                                            */
  /* Get the first message.                     */
  /*   Test for the output of the call is carried out */
  /*   in the 'for' loop.                       */
  /*                                            */
  MQGET(Hconn,
        Hobj,
        &MsgDesc,
        &GetMsgOpts,
        BufferLength,
        Buffer,
        &DataLength,
        &CompCode,
        &Reason);
```

Figure 57 (Part 1 of 2). Using the MQGET call (C language)

```
  /*                                            */
  /* Process the message and get the next message, */
  /* until no messages remaining.               */
  :
  /*      If the call fails for any other reason, */
  /*      print an error message showing the completion */
  /*      code and reason code.                  */
  /*                                            */
  if ( (CompCode == MQCC_FAILED) &&
     (Reason   == MQRC_NO_MSG_AVAILABLE) )
    {
  :
    }
  else
    {
    sprintf(pBuff, MESSAGE_4_E,
            ERROR_IN_MQGET, CompCode, Reason);
    PrintLine(pBuff);
    RetCode = CSQ4_ERROR;
    }
  :
  } /* end of main */
```

Figure 57 (Part 2 of 2). Using the MQGET call (C language)

Getting a message using the wait option

Figure 58 demonstrates how to use the wait option of the MQGET call. This code accepts truncated messages. This extract is taken from the Credit Check sample application (program CSQ4CCB5) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
  :
  MQLONG    Hconn;           /* Connection handle    */
  MQHOBJ    Hobj_CheckQ;     /* Object handle        */
  MQLONG    CompCode;        /* Completion code      */
  MQLONG    Reason;          /* Qualifying reason    */
  MQOD      ObjDesc   = {MQOD_DEFAULT};
                             /* Object descriptor    */
  MQMD      MsgDesc   = {MQMD_DEFAULT};
                             /* Message descriptor   */
  MQLONG    OpenOptions;
                             /* Control the MQOPEN call */
  MQGMO     GetMsgOpts = {MQGMO_DEFAULT};
                             /* Get Message Options  */
  MQLONG    MsgBuffLen;      /* Length of message buffer */
  CSQ4BCAQ  MsgBuffer;       /* Message structure    */
  MQLONG    DataLen;         /* Length of message    */
```

Figure 58 (Part 1 of 2). Using the MQGET call with the wait option (C language)

```
     ⋮
void main(void)
  {
     ⋮
  /*                                               */
  /* Initialize options and open the queue for input */
  /*                                               */
     ⋮
    /*                                             */
    /* Get and process messages                    */
    /*                                             */

    GetMsgOpts.Options = MQGMO_WAIT +
                         MQGMO_ACCEPT_TRUNCATED_MSG +
                         MQGMO_SYNCPOINT;
    GetMsgOpts.WaitInterval = WAIT_INTERVAL;
    MsgBuffLen = sizeof(MsgBuffer);

    memcpy(MsgDesc.MsgId, MQMI_NONE,
           sizeof(MsgDesc.MsgId));
    memcpy(MsgDesc.CorrelId, MQCI_NONE,
           sizeof(MsgDesc.CorrelId));
    /*                                             */
    /* Make the first MQGET call outside the loop   */
    /*                                             */

    MQGET(Hconn,
          Hobj_CheckQ,
          &MsgDesc,
          &GetMsgOpts,
          MsgBuffLen,
          &MsgBuffer,
          &DataLen,
          &CompCode,
          &Reason);
     ⋮
    /*                                             */
    /* Test the output of the MQGET call.  If the call */
    /* failed, send an error message showing the    */
    /* completion code and reason code, unless the   */
    /* reason code is NO_MSG AVAILABLE.              */
    /*                                             */

    if (Reason != MQRC_NO_MSG_AVAILABLE)
       {
       strncpy(TS_Operation, "MQGET", sizeof(TS_Operation));
       strncpy(TS_ObjName, ObjDesc.ObjectName,
               MQ_Q_NAME_LENGTH);
       Record_Call_Error();
       }
     ⋮
```

*Figure 58 (Part 2 of 2). Using the MQGET call with
the wait option (C language)*

Getting a message using signaling

*Signaling is available only with MQSeries for
MVS/ESA and MQSeries for Windows V2.1.*

Figure 59 demonstrates how to use the MQGET
call to set a signal so that you are notified when a
suitable message arrives on a queue. This extract
is not taken from the sample applications supplied
with MQSeries.

```
     ⋮
get_set_signal()
{

    MQMD     MsgDesc;
    MQGMO    GetMsgOpts;

    MQLONG   CompCode;
    MQLONG   Reason;
    MQHCONN  Hconn;
    MQHOBJ   Hobj;
    MQLONG   BufferLength;
    MQLONG   DataLength;

    char message_buffer[100];
    long int q_ecb, work_ecb;
    short int signal_sw, endloop;
    long int mask = 255;

    /*--------------------------*/
    /* Set up GMO structure.    */
    /*--------------------------*/
    memset(&GetMsgOpts,'\0',sizeof(GetMsgOpts));
    memcpy(GetMsgOpts.StrucId, MQGMO_STRUC_ID,
           sizeof(GetMsgOpts.StrucId);
    GetMsgOpts.Version       = MQGMO_VERSION_1;
    GetMsgOpts.WaitInterval  = 1000;
    GetMsgOpts.Options       = MQGMO_SET_SIGNAL +
                               MQGMO_BROWSE_FIRST;
    q_ecb               = 0;
    GetMsgOpts.Signal1       = &q_ecb;

    /*--------------------------*/
    /* Set up MD structure.     */
    /*--------------------------*/
    memset(&MsgDesc,'\0',sizeof(MsgDesc));
    memcpy(MsgDesc.StrucId, MQMD_STRUC_ID,
           sizeof(MsgDesc.StrucId);
    MsgDesc.Version  = MQMD_VERSION_1;
    MsgDesc.Report   = MQRO_NONE;
    memcpy(MsgDesc.MsgId,MQMI_NONE,
           sizeof(MsgDesc.MsgId));
    memcpy(MsgDesc.CorrelId,MQCI_NONE,
           sizeof(MsgDesc.CorrelId));
```

*Figure 59 (Part 1 of 3). Using the MQGET call with
signaling (C language)*

```
/*------------------------------------------------*/
/*  Issue the MQGET call.                         */
/*------------------------------------------------*/
 BufferLength = sizeof(message_buffer);
 signal_sw    = 0;

 MQGET(Hconn, Hobj, &MsgDesc, &GetMsgOpts,
       BufferLength, message_buffer, &DataLength,
       &CompCode, &Reason);
/*------------------------------------*/
/*  Check completion and reason codes. */
/*------------------------------------*/
 switch (CompCode)
     {
      case (MQCC_OK):            /* Message retrieved */
          break;
      case (MQCC_WARNING):
          switch (Reason)
          {
           case (MQRC_SIGNAL_REQUEST_ACCEPTED):
              signal_sw = 1;
              break;
           default:
              break;    /* Perform error processing  */
          }
          break;
      case (MQCC_FAILED):
          switch (Reason)
          {
           case (MQRC_Q_MGR_NOT_AVAILABLE):
           case (MQRC_CONNECTION_BROKEN):
           case (MQRC_Q_MGR_STOPPING):
              break;
           default:
              break; /* Perform error processing. */
          }
          break;
      default:
          break;         /* Perform error processing. */
     }
/*------------------------------------------------*/
/* If the SET_SIGNAL was accepted, set up a loop to */
/* check whether a message has arrived at one second */
/* intervals. The loop ends if a message arrives or */
/* the wait interval specified in the MQGMO         */
/* structure has expired.                           */
/*                                                  */
/* If a message arrives on the queue, another MQGET */
/* must be issued to retrieve the message. If other */
/* MQM calls have been made in the intervening      */
/* period, this may necessitate reinitializing the  */
/* MQMD and MQGMO structures.                       */
/* In this code, no intervening calls               */
/* have been made, so the only change required to   */
/* the structures is to specify MQGMO_NO_WAIT,       */
/* since we now know the message is there.          */
/*                                                  */
/* This code uses the EXEC CICS DELAY command to    */
/* suspend the program for a second. A batch program */
/* may achieve the same effect by calling an        */
/* assembler language subroutine which issues an MVS */
/* STIMER macro.                                    */
/*------------------------------------------------*/
```

Figure 59 (Part 2 of 3). Using the MQGET call with signaling (C language)

```
    if (signal_sw == 1)
      {
        endloop = 0;
        do
          {
            EXEC CICS DELAY FOR HOURS(0) MINUTES(0) SECONDS(1);
            work_ecb = q_ecb & mask;
            switch (work_ecb)
              {
                case (MQEC_MSG_ARRIVED):
                    endloop = 1;
                    mqgmo_options = MQGMO_NO_WAIT;
                    MQGET(Hconn, Hobj, &MsgDesc, &GetMsgOpts,
                          BufferLength, message_buffer,
                          &DataLength, &CompCode, &Reason);
                    if (CompCode != MQCC_OK)
                        ;          /* Perform error processing. */
                    break;
                case (MQEC_WAIT_INTERVAL_EXPIRED):
                case (MQEC_WAIT_CANCELED):
                    endloop = 1;
                    break;
                default:
                    break;
              }
          } while (endloop == 0);
      }

    return;
}
```

Figure 59 (Part 3 of 3). Using the MQGET call with signaling (C language)

Inquiring about the attributes of an object

Figure 60 demonstrates how to use the MQINQ call to inquire about the attributes of a queue. This extract is taken from the Queue Attributes sample application (program CSQ4CCC1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
#include <cmqc.h>        /* MQ API header file        */
  ⋮
#define NUMBEROFSELECTORS  2

const MQHCONN Hconn = MQHC_DEF_HCONN;
  ⋮
static void InquireGetAndPut(char   *Message,
                             PMQHOBJ pHobj,
                             char   *Object)
```

Figure 60 (Part 1 of 2). Using the MQINQ call (C language)

```
{
/*      Declare local variables                  */
/*                                               */
MQLONG  SelectorCount = NUMBEROFSELECTORS;
                         /* Number of selectors  */
MQLONG  IntAttrCount  = NUMBEROFSELECTORS;
                         /* Number of int attrs  */
MQLONG  CharAttrLength = 0;
              /* Length of char attribute buffer */
MQCHAR *CharAttrs ;
              /* Character attribute buffer       */
MQLONG  SelectorsTable[NUMBEROFSELECTORS];
                         /* attribute selectors  */
MQLONG  IntAttrsTable[NUMBEROFSELECTORS];
                         /* integer attributes   */
MQLONG  CompCode;         /* Completion code      */
MQLONG  Reason;           /* Qualifying reason    */

/*                                               */
/*    Open the queue.  If successful, do the inquire */
/*    call.                                       */
/*                                               */
  /*                                             */
  /*  Initialize the variables for the inquire   */
  /*  call:                                      */
  /*   - Set SelectorsTable to the attributes whose */
  /*      status is                              */
  /*      required                               */
  /*   - All other variables are already set     */
  /*                                             */
  SelectorsTable[0] = MQIA_INHIBIT_GET;
  SelectorsTable[1] = MQIA_INHIBIT_PUT;
  /*                                             */
  /*  Issue the inquire call                     */
  /*    Test the output of the inquire call. If the */
  /*    call failed, display an error message     */
  /*    showing the completion code and reason code,*/
  /*    otherwise display the status of the       */
  /*    INHIBIT-GET and INHIBIT-PUT attributes     */
  /*                                             */
  MQINQ(Hconn,
        *pHobj,
        SelectorCount,
        SelectorsTable,
        IntAttrCount,
        IntAttrsTable,
        CharAttrLength,
        CharAttrs,
        &CompCode,
        &Reason);

  if (CompCode != MQCC_OK)
    {
    sprintf(Message, MESSAGE_4_E,
          ERROR_IN_MQINQ, CompCode, Reason);
    SetMsg(Message);
    }
  else
    {
    /* Process the changes */
    } /* end if CompCode */
```

Figure 60 (Part 2 of 2). Using the MQINQ call (C language)

Setting the attributes of a queue

Figure 61 demonstrates how to use the MQSET call to change the attributes of a queue. This extract is taken from the Queue Attributes sample application (program CSQ4CCC1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
#include <cmqc.h>      /* MQ API header file       */
  :
#define NUMBEROFSELECTORS  2

const MQHCONN Hconn = MQHC_DEF_HCONN;

static void InhibitGetAndPut(char   *Message,
                             PMQHOBJ pHobj,
                             char   *Object)
  {
  /*                                                 */
  /*      Declare local variables                    */
  /*                                                 */
  MQLONG  SelectorCount = NUMBEROFSELECTORS;
                           /* Number of selectors  */
  MQLONG  IntAttrCount  = NUMBEROFSELECTORS;
                           /* Number of int attrs  */
  MQLONG  CharAttrLength = 0;
                /* Length of char attribute buffer */
  MQCHAR *CharAttrs ;
                /* Character attribute buffer       */
  MQLONG  SelectorsTable[NUMBEROFSELECTORS];
                           /* attribute selectors  */
  MQLONG  IntAttrsTable[NUMBEROFSELECTORS];
                           /* integer attributes   */
  MQLONG  CompCode;         /* Completion code      */
  MQLONG  Reason;           /* Qualifying reason    */
  :
  /*                                               */
  /*    Open the queue.  If successful, do the      */
  /*    inquire call.                               */
  /*                                               */
  :
    /*                                             */
    /*  Initialize the variables for the set call:  */
    /*   - Set SelectorsTable to the attributes to be */
    /*      set                                     */
    /*   - Set IntAttrsTable to the required status  */
    /*   - All other variables are already set       */
    /*                                             */
    SelectorsTable[0] = MQIA_INHIBIT_GET;
    SelectorsTable[1] = MQIA_INHIBIT_PUT;
    IntAttrsTable[0]  = MQQA_GET_INHIBITED;
    IntAttrsTable[1]  = MQQA_PUT_INHIBITED;
  :
```

Figure 61 (Part 1 of 2). Using the MQSET call (C language)

C language examples

```
/*                                          */
/*   Issue the set call.                    */
/*     Test the output of the set call.  If the  */
/*     call fails, display an error message    */
/*     showing the completion code and reason  */
/*     code; otherwise move INHIBITED to the   */
/*     relevant screen map fields             */
/*                                          */
MQSET(Hconn,
      *pHobj,
      SelectorCount,
      SelectorsTable,
      IntAttrCount,
      IntAttrsTable,
      CharAttrLength,
      CharAttrs,
      &CompCode,
      &Reason);

if (CompCode != MQCC_OK)
   {
   sprintf(Message, MESSAGE_4_E,
           ERROR_IN_MQSET, CompCode, Reason);
   SetMsg(Message);
   }
else
   {
      /* Process the changes */
   } /* end if CompCode */
```

Figure 61 (Part 2 of 2). Using the MQSET call (C language)

Appendix D. COBOL examples

The examples in this appendix are taken from the MQSeries for MVS/ESA sample applications. They are applicable to all platforms, and any exception to this is noted.

The examples in this appendix demonstrate the following techniques:

477

Connecting to a queue manager

Figure 62 demonstrates how to use the MQCONN call to connect a program to a queue manager in MVS batch. This extract is taken from the Browse sample application (program CSQ4BVA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*    W02 - Data fields derived from the PARM field
 01  W02-MQM                  PIC X(48) VALUE SPACES.
*    W03 - MQM API fields
 01  W03-HCONN                PIC S9(9) BINARY.
 01  W03-COMPCODE             PIC S9(9) BINARY.
 01  W03-REASON               PIC S9(9) BINARY.
*
*    MQV contains constants (for filling in the control
*    blocks)
*    and return codes (for testing the result of a call)
*
 01  W05-MQM-CONSTANTS.
 COPY CMQV SUPPRESS.
   ⋮
*    Separate into the relevant fields any data passed
*    in the PARM statement
*
     UNSTRING PARM-STRING DELIMITED BY ALL ','
                          INTO W02-MQM
                               W02-OBJECT.
   ⋮
*    Connect to the specified queue manager.
*
     CALL 'MQCONN' USING W02-MQM
                         W03-HCONN
                         W03-COMPCODE
                         W03-REASON.
*
*    Test the output of the connect call.  If the call
*    fails, print an error message showing the
*    completion code and reason code.
*
     IF (W03-COMPCODE NOT = MQCC-OK) THEN
   ⋮
     END-IF.
   ⋮
```

Figure 62. Using the MQCONN call (COBOL)

Disconnecting from a queue manager

Figure 63 demonstrates how to use the MQDISC call to disconnect a program from a queue manager in MVS batch. This extract is taken from the Browse sample application (program CSQ4BVA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   ⋮
*
* Disconnect from the queue manager
*
     CALL 'MQDISC' USING W03-HCONN
                         W03-COMPCODE
                         W03-REASON.
*
*    Test the output of the disconnect call.  If the
*    call fails, print an error message showing the
*    completion code and reason code.
*
     IF (W03-COMPCODE NOT = MQCC-OK) THEN
   ⋮
         END-IF.
   ⋮
```

Figure 63. Using the MQDISC call (COBOL). The variables used in this code extract are those that were set in Figure 62.

Creating a dynamic queue

Figure 64 demonstrates how to use the MQOPEN call to create a dynamic queue. This extract is taken from the Credit Check sample application (program CSQ4CVB1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
     ⋮
* -----------------------------------------------------*
 WORKING-STORAGE SECTION.
* -----------------------------------------------------*
*
*    W02 - Queues processed in this program
*
 01 W02-MODEL-QNAME        PIC X(48) VALUE
    'CSQ4SAMP.B1.MODEL                            '.
 01 W02-NAME-PREFIX        PIC X(48) VALUE
    'CSQ4SAMP.B1.*                                '.
 01 W02-TEMPORARY-Q        PIC X(48).
*
*    W03 - MQM API fields
*
 01 W03-HCONN        PIC S9(9) BINARY VALUE ZERO.
 01 W03-OPTIONS      PIC S9(9) BINARY.
 01 W03-HOBJ         PIC S9(9) BINARY.
 01 W03-COMPCODE     PIC S9(9) BINARY.
 01 W03-REASON       PIC S9(9) BINARY.
*
*    API control blocks
*
 01 MQM-OBJECT-DESCRIPTOR.
    COPY CMQODV.
*
*    CMQV contains constants (for setting or testing
*    field values) and return codes (for testing the
*    result of a call)
*
 01 MQM-CONSTANTS.
 COPY CMQV SUPPRESS.
* -----------------------------------------------------*
 PROCEDURE DIVISION.
* -----------------------------------------------------*
   ⋮
* -----------------------------------------------------*
 OPEN-TEMP-RESPONSE-QUEUE SECTION.
* -----------------------------------------------------*
```

Figure 64 (Part 1 of 2). Using the MQOPEN call to create a dynamic queue (COBOL)

```
*
*  This section creates a temporary dynamic queue
*  using a model queue
*
* -----------------------------------------------------*
*
* Change three fields in the Object Descriptor (MQOD)
* control block.  (MQODV initializes the other fields)
*
     MOVE MQOT-Q           TO MQOD-OBJECTTYPE.
     MOVE W02-MODEL-QNAME  TO MQOD-OBJECTNAME.
     MOVE W02-NAME-PREFIX  TO MQOD-DYNAMICQNAME.
*
     COMPUTE W03-OPTIONS = MQOO-INPUT-EXCLUSIVE.
*
     CALL 'MQOPEN' USING W03-HCONN
                         MQOD
                         W03-OPTIONS
                         W03-HOBJ-MODEL
                         W03-COMPCODE
                         W03-REASON.
*
     IF W03-COMPCODE NOT = MQCC-OK
         MOVE 'MQOPEN'    TO M01-MSG4-OPERATION
         MOVE W03-COMPCODE TO M01-MSG4-COMPCODE
         MOVE W03-REASON   TO M01-MSG4-REASON
         MOVE M01-MESSAGE-4 TO M00-MESSAGE
     ELSE
         MOVE MQOD-OBJECTNAME TO W02-TEMPORARY-Q
     END-IF.
*
 OPEN-TEMP-RESPONSE-QUEUE-EXIT.
*
*    Return to performing section.
*
     EXIT.
     EJECT
*
```

Figure 64 (Part 2 of 2). Using the MQOPEN call to create a dynamic queue (COBOL)

Opening an existing queue

Figure 65 demonstrates how to use the MQOPEN call to open an existing queue. This extract is taken from the Browse sample application (program CSQ4BVA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
  ⋮
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*    W01 - Fields derived from the command area input
*
 01  W01-OBJECT                 PIC X(48).
*
*    W02 - MQM API fields
*
 01  W02-HCONN       PIC S9(9) BINARY VALUE ZERO.
 01  W02-OPTIONS     PIC S9(9) BINARY.
 01  W02-HOBJ        PIC S9(9) BINARY.
 01  W02-COMPCODE    PIC S9(9) BINARY.
 01  W02-REASON      PIC S9(9) BINARY.
*
*    CMQODV defines the object descriptor (MQOD)
*
 01  MQM-OBJECT-DESCRIPTOR.
     COPY CMQODV.
*
* CMQV contains constants (for setting or testing
* field values) and return codes (for testing the
* result of a call)
*
 01  MQM-CONSTANTS.
 COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 E-OPEN-QUEUE SECTION.
* -------------------------------------------------------*
*                                                       *
* This section opens the queue                          *
*
*    Initialize the Object Descriptor (MQOD) control
*    block
*    (The copy file initializes the remaining fields.)
*
     MOVE MQOT-Q       TO MQOD-OBJECTTYPE.
     MOVE W01-OBJECT   TO MQOD-OBJECTNAME.
*
*    Initialize W02-OPTIONS to open the queue for both
*    inquiring about and setting attributes
*
     COMPUTE W02-OPTIONS = MQOO-INQUIRE + MQOO-SET.
```

Figure 65 (Part 1 of 2). Using the MQOPEN call to open an existing queue (COBOL)

```
*
*    Open the queue
*
     CALL 'MQOPEN' USING W02-HCONN
                         MQOD
                         W02-OPTIONS
                         W02-HOBJ
                         W02-COMPCODE
                         W02-REASON.
*
*    Test the output from the open
*
*    If the completion code is not OK, display a
*    separate error message for each of the following
*    errors:
*
*  Q-MGR-NOT-AVAILABLE - MQM is not available
*  CONNECTION-BROKEN   - MQM is no longer connected to CICS
*  UNKNOWN-OBJECT-NAME - The queue does not exist
*  NOT-AUTHORIZED      - The user is not authorized to open
*                        the queue
*
* For any other error, display an error message
* showing the completion and reason codes
*
  IF W02-COMPCODE NOT = MQCC-OK
     EVALUATE TRUE
*
        WHEN W02-REASON = MQRC-Q-MGR-NOT-AVAILABLE
             MOVE M01-MESSAGE-6 TO M00-MESSAGE
*
        WHEN W02-REASON = MQRC-CONNECTION-BROKEN
             MOVE M01-MESSAGE-6 TO M00-MESSAGE
*
        WHEN W02-REASON = MQRC-UNKNOWN-OBJECT-NAME
             MOVE M01-MESSAGE-2 TO M00-MESSAGE
*
        WHEN W02-REASON = MQRC-NOT-AUTHORIZED
             MOVE M01-MESSAGE-3 TO M00-MESSAGE
*
        WHEN OTHER
             MOVE 'MQOPEN'     TO M01-MSG4-OPERATION
             MOVE W02-COMPCODE TO M01-MSG4-COMPCODE
             MOVE W02-REASON   TO M01-MSG4-REASON
             MOVE M01-MESSAGE-4 TO M00-MESSAGE
       END-EVALUATE
     END-IF.
 E-EXIT.
*
*    Return to performing section
*
     EXIT.
     EJECT
```

Figure 65 (Part 2 of 2). Using the MQOPEN call to open an existing queue (COBOL)

Closing a queue

Figure 66 demonstrates how to use the MQCLOSE call. This extract is taken from the Browse sample application (program CSQ4BVA1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   ⋮
*
*     Close the queue
*
      MOVE MQCO-NONE TO W03-OPTIONS.
*
      CALL 'MQCLOSE' USING W03-HCONN
                          W03-HOBJ
                          W03-OPTIONS
                          W03-COMPCODE
                          W03-REASON.
*
*     Test the output of the MQCLOSE call.  If the call
*     fails, print an error message showing the
*     completion code and reason code.
*
      IF (W03-COMPCODE NOT = MQCC-OK) THEN
         MOVE 'CLOSE'      TO W04-MSG4-TYPE
         MOVE W03-COMPCODE  TO W04-MSG4-COMPCODE
         MOVE W03-REASON    TO W04-MSG4-REASON
         MOVE W04-MESSAGE-4 TO W00-PRINT-DATA
         PERFORM PRINT-LINE
         MOVE W06-CSQ4-ERROR TO W00-RETURN-CODE
      END-IF.
*
```

Figure 66. Using the MQCLOSE call (COBOL). The variables used in this code extract are those that were set in Figure 62 on page 478.

Putting a message using MQPUT

Figure 67 demonstrates how to use the MQPUT call using context. This extract is taken from the Credit Check sample application (program CSQ4CVB1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   ⋮
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*     W02 - Queues processed in this program
*
 01  W02-TEMPORARY-Q          PIC X(48).
*
*     W03 - MQM API fields
*
 01  W03-HCONN          PIC S9(9) BINARY VALUE ZERO.
 01  W03-HOBJ-INQUIRY   PIC S9(9) BINARY.
 01  W03-OPTIONS        PIC S9(9) BINARY.
 01  W03-BUFFLEN        PIC S9(9) BINARY.
 01  W03-COMPCODE       PIC S9(9) BINARY.
 01  W03-REASON         PIC S9(9) BINARY.
*
 01  W03-PUT-BUFFER.
*
     05 W03-CSQ4BIIM.
     COPY CSQ4VB1.
*
*     API control blocks
*
 01  MQM-MESSAGE-DESCRIPTOR.
     COPY CMQMDV.
 01  MQM-PUT-MESSAGE-OPTIONS.
     COPY CMQPMOV.
*
*     MQV contains constants (for filling in the
*     control blocks) and return codes (for testing
*     the result of a call).
*
 01  MQM-CONSTANTS.
     COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 PROCEDURE DIVISION.
* -------------------------------------------------------*
   ⋮
*     Open queue and build message.
   ⋮
```

Figure 67 (Part 1 of 2). Using the MQPUT call (COBOL)

```
*
* Set the message descriptor and put-message options to
* the values required to create the message.
* Set the length of the message.
*
  MOVE MQMT-REQUEST        TO MQMD-MSGTYPE.
  MOVE MQCI-NONE           TO MQMD-CORRELID.
  MOVE MQMI-NONE           TO MQMD-MSGID.
  MOVE W02-TEMPORARY-Q     TO MQMD-REPLYTOQ.
  MOVE SPACES              TO MQMD-REPLYTOQMGR.
  MOVE 5                   TO MQMD-PRIORITY.
  MOVE MQPER-NOT-PERSISTENT TO MQMD-PERSISTENCE.
  COMPUTE MQPMO-OPTIONS    = MQPMO-NO-SYNCPOINT +
                            MQPMO-DEFAULT-CONTEXT.
  MOVE LENGTH OF CSQ4BIIM-MSG TO W03-BUFFLEN.
*
     CALL 'MQPUT' USING W03-HCONN
                        W03-HOBJ-INQUIRY
                        MQMD
                        MQPMO
                        W03-BUFFLEN
                        W03-PUT-BUFFER
                        W03-COMPCODE
                        W03-REASON.
     IF W03-COMPCODE NOT = MQCC-OK
   :
     END-IF.
```

Figure 67 (Part 2 of 2). Using the MQPUT call (COBOL)

Putting a message using MQPUT1

Figure 68 demonstrates how to use the MQPUT1 call. This extract is taken from the Credit Check sample application (program CSQ4CVB5) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   :
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*    W03 - MQM API fields
*
 01  W03-HCONN        PIC S9(9) BINARY VALUE ZERO.
 01  W03-OPTIONS      PIC S9(9) BINARY.
 01  W03-COMPCODE     PIC S9(9) BINARY.
 01  W03-REASON       PIC S9(9) BINARY.
 01  W03-BUFFLEN      PIC S9(9) BINARY.
*
 01  W03-PUT-BUFFER.
     05 W03-CSQ4BQRM.
     COPY CSQ4VB4.
```

Figure 68 (Part 1 of 2). Using the MQPUT1 call (COBOL)

```
*
*    API control blocks
*
 01  MQM-OBJECT-DESCRIPTOR.
     COPY CMQODV.
 01  MQM-MESSAGE-DESCRIPTOR.
     COPY CMQMDV.
 01  MQM-PUT-MESSAGE-OPTIONS.
     COPY CMQPMOV.
*
* CMQV contains constants (for filling in the
* control blocks) and return codes (for testing
* the result of a call).
*
 01  MQM-MQV.
 COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 PROCEDURE DIVISION.
* -------------------------------------------------------*
   :
*    Get the request message.
   :
* -------------------------------------------------------*
 PROCESS-QUERY SECTION.
* -------------------------------------------------------*
   :
*    Build the reply message.
   :
*
* Set the object descriptor, message descriptor and
* put-message options to the values required to create
* the message.
* Set the length of the message.
*
  MOVE MQMD-REPLYTOQ    TO MQOD-OBJECTNAME.
  MOVE MQMD-REPLYTOQMGR TO MQOD-OBJECTQMGRNAME.
  MOVE MQMT-REPLY       TO MQMD-MSGTYPE.
  MOVE SPACES           TO MQMD-REPLYTOQ.
  MOVE SPACES           TO MQMD-REPLYTOQMGR.
  MOVE LOW-VALUES       TO MQMD-MSGID.
  COMPUTE MQPMO-OPTIONS = MQPMO-SYNCPOINT +
                        MQPMO-PASS-IDENTITY-CONTEXT.
  MOVE W03-HOBJ-CHECKQ  TO MQPMO-CONTEXT.
  MOVE LENGTH OF CSQ4BQRM-MSG TO W03-BUFFLEN.
*
     CALL 'MQPUT1' USING W03-HCONN
                        MQOD
                        MQMD
                        MQPMO
                        W03-BUFFLEN
                        W03-PUT-BUFFER
                        W03-COMPCODE
                        W03-REASON.
     IF W03-COMPCODE NOT = MQCC-OK
         MOVE 'MQPUT1'         TO M02-OPERATION
         MOVE MQOD-OBJECTNAME  TO M02-OBJECTNAME
         PERFORM RECORD-CALL-ERROR
         PERFORM FORWARD-MSG-TO-DLQ
     END-IF.
*
```

Figure 68 (Part 2 of 2). Using the MQPUT1 call (COBOL)

Getting a message

Figure 69 demonstrates how to use the MQGET
call to remove a message from a queue. This
extract is taken from the Credit Check sample
application (program CSQ4CVB1) supplied with
MQSeries for MVS/ESA. For the names and
locations of the sample applications on other
platforms, see Chapter 30, "Sample programs (all
platforms except MVS/ESA)" on page 313.

```
     ⋮
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*    W03 - MQM API fields
*
 01 W03-HCONN           PIC S9(9) BINARY VALUE ZERO.
 01 W03-HOBJ-RESPONSE   PIC S9(9) BINARY.
 01 W03-OPTIONS         PIC S9(9) BINARY.
 01 W03-BUFFLEN         PIC S9(9) BINARY.
 01 W03-DATALEN         PIC S9(9) BINARY.
 01 W03-COMPCODE        PIC S9(9) BINARY.
 01 W03-REASON          PIC S9(9) BINARY.
*
 01 W03-GET-BUFFER.
    05 W03-CSQ4BAM.
    COPY CSQ4VB2.
*
*    API control blocks
*
 01 MQM-MESSAGE-DESCRIPTOR.
    COPY CMQMDV.
 01 MQM-GET-MESSAGE-OPTIONS.
    COPY CMQGMOV.
*
*    MQV contains constants (for filling in the
*    control blocks) and return codes (for testing
*    the result of a call).
*
 01 MQM-CONSTANTS.
    COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 A-MAIN SECTION.
* -------------------------------------------------------*
  ⋮
*   Open response queue.
  ⋮
* -------------------------------------------------------*
 PROCESS-RESPONSE-SCREEN SECTION.
* -------------------------------------------------------*
*                                                       *
*  This section gets a message from the response queue. *
*                                                       *
*  When a correct response is received, it is           *
*  transferred to the map for display; otherwise        *
*  an error message is built.                           *
*                                                       *
* -------------------------------------------------------*
```

*Figure 69 (Part 1 of 2). Using the MQGET call
(COBOL)*

```
*
*    Set get-message options
*
  COMPUTE MQGMO-OPTIONS = MQGMO-SYNCPOINT +
                      MQGMO-ACCEPT-TRUNCATED-MSG +
                      MQGMO-NO-WAIT.
*
* Set msgid and correlid in MQMD to nulls so that any
* message will qualify.
* Set length to available buffer length.
*
    MOVE MQMI-NONE TO MQMD-MSGID.
    MOVE MQCI-NONE TO MQMD-CORRELID.
    MOVE LENGTH OF W03-GET-BUFFER TO W03-BUFFLEN.
*
    CALL 'MQGET' USING W03-HCONN
                       W03-HOBJ-RESPONSE
                       MQMD
                       MQGMO
                       W03-BUFFLEN
                       W03-GET-BUFFER
                       W03-DATALEN
                       W03-COMPCODE
                       W03-REASON.
    EVALUATE TRUE
        WHEN W03-COMPCODE NOT = MQCC-FAILED
  ⋮
*        Process the message
  ⋮
        WHEN (W03-COMPCODE = MQCC-FAILED AND
           W03-REASON = MQRC-NO-MSG-AVAILABLE)
             MOVE M01-MESSAGE-9 TO M00-MESSAGE
             PERFORM CLEAR-RESPONSE-SCREEN
*
        WHEN OTHER
            MOVE 'MQGET '     TO M01-MSG4-OPERATION
            MOVE W03-COMPCODE TO M01-MSG4-COMPCODE
            MOVE W03-REASON   TO M01-MSG4-REASON
            MOVE M01-MESSAGE-4 TO M00-MESSAGE
            PERFORM CLEAR-RESPONSE-SCREEN
    END-EVALUATE.
```

*Figure 69 (Part 2 of 2). Using the MQGET call
(COBOL)*

Getting a message using the wait option

Figure 70 demonstrates how to use the MQGET call with the wait option and accepting truncated messages. This extract is taken from the Credit Check sample application (program CSQ4CVB5) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
     ⋮
* -----------------------------------------------------*
 WORKING-STORAGE SECTION.
* -----------------------------------------------------*
*
*    W00 - General work fields
*
 01 W00-WAIT-INTERVAL   PIC S9(09) BINARY VALUE 30000.
*
*    W03 - MQM API fields
*
 01 W03-HCONN           PIC S9(9) BINARY VALUE ZERO.
 01 W03-OPTIONS         PIC S9(9) BINARY.
 01 W03-HOBJ-CHECKQ     PIC S9(9) BINARY.
 01 W03-COMPCODE        PIC S9(9) BINARY.
 01 W03-REASON          PIC S9(9) BINARY.
 01 W03-DATALEN         PIC S9(9) BINARY.
 01 W03-BUFFLEN         PIC S9(9) BINARY.
*
 01 W03-MSG-BUFFER.
    05 W03-CSQ4BCAQ.
    COPY CSQ4VB3.
*
*    API control blocks
*
 01 MQM-MESSAGE-DESCRIPTOR.
    COPY CMQMDV.
 01 MQM-GET-MESSAGE-OPTIONS.
    COPY CMQGMOV.
*
*    CMQV contains constants (for filling in the
*    control blocks) and return codes (for testing
*    the result of a call).
*
 01 MQM-MQV.
 COPY CMQV SUPPRESS.
* -----------------------------------------------------*
 PROCEDURE DIVISION.
* -----------------------------------------------------*
   ⋮
*    Open input queue.
   ⋮
```

Figure 70 (Part 1 of 2). Using the MQGET call with the wait option (COBOL)

```
*
*    Get and process messages.
*
  COMPUTE MQGMO-OPTIONS = MQGMO-WAIT +
                          MQGMO-ACCEPT-TRUNCATED-MSG +
                          MQGMO-SYNCPOINT.
  MOVE LENGTH OF W03-MSG-BUFFER TO W03-BUFFLEN.
  MOVE W00-WAIT-INTERVAL TO MQGMO-WAITINTERVAL.
  MOVE MQMI-NONE TO MQMD-MSGID.
  MOVE MQCI-NONE TO MQMD-CORRELID.
*
*    Make the first MQGET call outside the loop.
*
     CALL 'MQGET' USING W03-HCONN
                        W03-HOBJ-CHECKQ
                        MQMD
                        MQGMO
                        W03-BUFFLEN
                        W03-MSG-BUFFER
                        W03-DATALEN
                        W03-COMPCODE
                        W03-REASON.
*
*    Test the output of the MQGET call using the
*    PERFORM loop that follows.
*
*    Perform whilst no failure occurs
*      - process this message
*      - reset the call parameters
*      - get another message
*    End-perform
*
   ⋮
*
*    Test the output of the MQGET call.  If the call
*    fails, send an error message showing the
*    completion code and reason code, unless the
*    completion code is NO-MSG-AVAILABLE.
*
     IF (W03-COMPCODE NOT = MQCC-FAILED) OR
        (W03-REASON NOT = MQRC-NO-MSG-AVAILABLE)
         MOVE 'MQGET '          TO M02-OPERATION
         MOVE MQOD-OBJECTNAME   TO M02-OBJECTNAME
                PERFORM RECORD-CALL-ERROR
     END-IF.
   ⋮
```

Figure 70 (Part 2 of 2). Using the MQGET call with the wait option (COBOL)

Getting a message using signaling

Signaling is available only with MQSeries for MVS/ESA.

Figure 71 demonstrates how to use the MQGET call with signaling. This extract is taken from the Credit Check sample application (program CSQ4CVB2) supplied with MQSeries for MVS/ESA.

```
    ⋮
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*    W00 - General work fields
  ⋮
 01 W00-WAIT-INTERVAL   PIC S9(09) BINARY VALUE 30000.
*
*    W03 - MQM API fields
*
 01 W03-HCONN           PIC S9(9) BINARY VALUE ZERO.
 01 W03-HOBJ-REPLYQ     PIC S9(9) BINARY.
 01 W03-COMPCODE        PIC S9(9) BINARY.
 01 W03-REASON          PIC S9(9) BINARY.
 01 W03-DATALEN         PIC S9(9) BINARY.
 01 W03-BUFFLEN         PIC S9(9) BINARY.
  ⋮
 01 W03-GET-BUFFER.
    05 W03-CSQ4BQRM.
    COPY CSQ4VB4.
*
    05 W03-CSQ4BIIM REDEFINES W03-CSQ4BQRM.
    COPY CSQ4VB1.
*
    05 W03-CSQ4BPGM REDEFINES W03-CSQ4BIIM.
    COPY CSQ4VB5.
  ⋮
*    API control blocks
*
 01 MQM-MESSAGE-DESCRIPTOR.
    COPY CMQMDV.
 01 MQM-GET-MESSAGE-OPTIONS.
    COPY CMQGMOV.
  ⋮
*    MQV contains constants (for filling in the
*    control blocks) and return codes (for testing
*    the result of a call).
*
 01 MQM-MQV.
 COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 LINKAGE SECTION.
* -------------------------------------------------------*
 01 L01-ECB-ADDR-LIST.
    05 L01-ECB-ADDR1      POINTER.
    05 L01-ECB-ADDR2      POINTER.
```

Figure 71 (Part 1 of 3). Using the MQGET call with signaling (COBOL)

```
*
 01 L02-ECBS.
    05 L02-INQUIRY-ECB1    PIC S9(09) BINARY.
    05 L02-REPLY-ECB2      PIC S9(09) BINARY.
 01 REDEFINES L02-ECBS.
    05                     PIC X(02).
    05 L02-INQUIRY-ECB1-CC PIC S9(04) BINARY.
    05                     PIC X(02).
    05 L02-REPLY-ECB2-CC   PIC S9(04) BINARY.
*
* -------------------------------------------------------*
 PROCEDURE DIVISION.
* -------------------------------------------------------*
    ⋮
* Initialize variables, open queues, set signal on
* inquiry queue.
    ⋮
* -------------------------------------------------------*
 PROCESS-SIGNAL-ACCEPTED SECTION.
* -------------------------------------------------------*
*  This section gets a message with signal.  If a       *
*  message is received, process it.  If the signal      *
*  is set or is already set, the program goes into      *
*  an operating system wait.                            *
*  Otherwise an error is reported and call error set.   *
* -------------------------------------------------------*
*
   PERFORM REPLYQ-GETSIGNAL.
*
   EVALUATE TRUE
       WHEN (W03-COMPCODE = MQCC-OK AND
             W03-REASON = MQRC-NONE)
         PERFORM PROCESS-REPLYQ-MESSAGE
*
       WHEN (W03-COMPCODE = MQCC-WARNING AND
             W03-REASON = MQRC-SIGNAL-REQUEST-ACCEPTED)
          OR
            (W03-COMPCODE = MQCC-FAILED AND
             W03-REASON = MQRC-SIGNAL-OUTSTANDING)
         PERFORM EXTERNAL-WAIT
*
       WHEN OTHER
         MOVE 'MQGET SIGNAL'  TO M02-OPERATION
         MOVE MQOD-OBJECTNAME TO M02-OBJECTNAME
         PERFORM RECORD-CALL-ERROR
         MOVE W06-CALL-ERROR  TO W06-CALL-STATUS
   END-EVALUATE.
*
 PROCESS-SIGNAL-ACCEPTED-EXIT.
*    Return to performing section
    EXIT.
    EJECT
*
* -------------------------------------------------------*
 EXTERNAL-WAIT SECTION.
* -------------------------------------------------------*
*  This section performs an external CICS wait on two   *
*  ECBs until at least one is posted.  It then calls    *
*  the sections to handle the posted ECB.               *
* -------------------------------------------------------*
```

Figure 71 (Part 2 of 3). Using the MQGET call with signaling (COBOL)

```
        EXEC CICS WAIT EXTERNAL
            ECBLIST(W04-ECB-ADDR-LIST-PTR)
            NUMEVENTS(2)
        END-EXEC.
*
* At least one ECB must have been posted to get to this
* point.  Test which ECB has been posted and perform
* the appropriate section.
*
        IF L02-INQUIRY-ECB1 NOT = 0
            PERFORM TEST-INQUIRYQ-ECB
        ELSE
            PERFORM TEST-REPLYQ-ECB
        END-IF.
*
 EXTERNAL-WAIT-EXIT.
*
*    Return to performing section.
*
        EXIT.
        EJECT
  :
* -------------------------------------------------------*
 REPLYQ-GETSIGNAL SECTION.
* -------------------------------------------------------*
*                                                       *
* This section performs an MQGET call (in syncpoint with *
* signal) on the reply queue.  The signal field in the  *
* MQGMO is set to the address of the ECB.               *
* Response handling is done by the performing section.  *
*                                                       *
* -------------------------------------------------------*
*
        COMPUTE MQGMO-OPTIONS        = MQGMO-SYNCPOINT +
                                       MQGMO-SET-SIGNAL.
        MOVE W00-WAIT-INTERVAL      TO MQGMO-WAITINTERVAL.
        MOVE LENGTH OF W03-GET-BUFFER TO W03-BUFFLEN.
*
        MOVE ZEROS                  TO L02-REPLY-ECB2.
        SET MQGMO-SIGNAL1 TO ADDRESS OF L02-REPLY-ECB2.
*
*    Set msgid and correlid to nulls so that any message
*    will qualify.
*
        MOVE MQMI-NONE TO MQMD-MSGID.
        MOVE MQCI-NONE TO MQMD-CORRELID.
*
        CALL 'MQGET' USING W03-HCONN
                           W03-HOBJ-REPLYQ
                           MQMD
                           MQGMO
                           W03-BUFFLEN
                           W03-GET-BUFFER
                           W03-DATALEN
                           W03-COMPCODE
                           W03-REASON.
*
 REPLYQ-GETSIGNAL-EXIT.
*
*    Return to performing section.
*
        EXIT.
        EJECT
*
  :
```

Figure 71 (Part 3 of 3). Using the MQGET call with signaling (COBOL)

Inquiring about the attributes of an object

Figure 72 demonstrates how to use the MQINQ call to inquire about the attributes of a queue. This extract is taken from the Queue Attributes sample application (program CSQ4CVC1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
  :
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*    W02 - MQM API fields
*
 01 W02-SELECTORCOUNT    PIC S9(9) BINARY VALUE 2.
 01 W02-INTATTRCOUNT     PIC S9(9) BINARY VALUE 2.
 01 W02-CHARATTRLENGTH   PIC S9(9) BINARY VALUE ZERO.
 01 W02-CHARATTRS        PIC X     VALUE LOW-VALUES.
 01 W02-HCONN            PIC S9(9) BINARY VALUE ZERO.
 01 W02-HOBJ             PIC S9(9) BINARY.
 01 W02-COMPCODE         PIC S9(9) BINARY.
 01 W02-REASON           PIC S9(9) BINARY.
 01 W02-SELECTORS-TABLE.
    05 W02-SELECTORS     PIC S9(9) BINARY OCCURS 2 TIMES
 01 W02-INTATTRS-TABLE.
    05 W02-INTATTRS      PIC S9(9) BINARY OCCURS 2 TIMES
*
*    CMQODV defines the object descriptor (MQOD).
*
 01 MQM-OBJECT-DESCRIPTOR.
    COPY CMQODV.
*
* CMQV contains constants (for setting or testing field
* values) and return codes (for testing the result of a
* call).
*
 01 MQM-CONSTANTS.
 COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 PROCEDURE DIVISION.
* -------------------------------------------------------*
*
*    Get the queue name and open the queue.
*
  :
*
*    Initialize the variables for the inquiry call:
*    - Set W02-SELECTORS-TABLE to the attributes whose
*    status is required
*    - All other variables are already set
*
        MOVE MQIA-INHIBIT-GET TO W02-SELECTORS(1).
        MOVE MQIA-INHIBIT-PUT TO W02-SELECTORS(2).
```

Figure 72 (Part 1 of 2). Using the MQINQ call (COBOL)

```
*
*    Inquire about the attributes.
*
     CALL 'MQINQ' USING W02-HCONN,
                        W02-HOBJ,
                        W02-SELECTORCOUNT,
                        W02-SELECTORS-TABLE,
                        W02-INTATTRCOUNT,
                        W02-INTATTRS-TABLE,
                        W02-CHARATTRLENGTH,
                        W02-CHARATTRS,
                        W02-COMPCODE,
                        W02-REASON.
*
* Test the output from the inquiry:
*
* - If the completion code is not OK, display an error
*   message showing the completion and reason codes
*
* - Otherwise, move the correct attribute status into
*   the relevant screen map fields
*
     IF W02-COMPCODE NOT = MQCC-OK
        MOVE 'MQINQ'     TO M01-MSG4-OPERATION
        MOVE W02-COMPCODE TO M01-MSG4-COMPCODE
        MOVE W02-REASON   TO M01-MSG4-REASON
        MOVE M01-MESSAGE-4 TO M00-MESSAGE
*
     ELSE
*       Process the changes.
  ⋮
          END-IF.
  ⋮
```

Figure 72 (Part 2 of 2). Using the MQINQ call (COBOL)

Setting the attributes of a queue

Figure 73 demonstrates how to use the MQSET call to change the attributes of a queue. This extract is taken from the Queue Attributes sample application (program CSQ4CVC1) supplied with MQSeries for MVS/ESA. For the names and locations of the sample applications on other platforms, see Chapter 30, "Sample programs (all platforms except MVS/ESA)" on page 313.

```
   ⋮
* -------------------------------------------------------*
 WORKING-STORAGE SECTION.
* -------------------------------------------------------*
*
*    W02 - MQM API fields
*
 01  W02-SELECTORCOUNT   PIC S9(9) BINARY VALUE 2.
 01  W02-INTATTRCOUNT    PIC S9(9) BINARY VALUE 2.
 01  W02-CHARATTRLENGTH  PIC S9(9) BINARY VALUE ZERO.
 01  W02-CHARATTRS       PIC X     VALUE LOW-VALUES.
 01  W02-HCONN           PIC S9(9) BINARY VALUE ZERO.
 01  W02-HOBJ            PIC S9(9) BINARY.
 01  W02-COMPCODE        PIC S9(9) BINARY.
 01  W02-REASON          PIC S9(9) BINARY.
 01  W02-SELECTORS-TABLE.
     05  W02-SELECTORS   PIC S9(9) BINARY OCCURS 2 TIMES.
 01  W02-INTATTRS-TABLE.
     05  W02-INTATTRS    PIC S9(9) BINARY OCCURS 2 TIMES.
*
*    CMQODV defines the object descriptor (MQOD).
*
 01  MQM-OBJECT-DESCRIPTOR.
     COPY CMQODV.
*
* CMQV contains constants (for setting or testing
* field values) and return codes (for testing the
* result of a call).
*
 01  MQM-CONSTANTS.
 COPY CMQV SUPPRESS.
* -------------------------------------------------------*
 PROCEDURE DIVISION.
* -------------------------------------------------------*
```

Figure 73 (Part 1 of 2). Using the MQSET call (COBOL)

COBOL examples

```
*
*    Get the queue name and open the queue.
*
   :
*
*
* Initialize the variables required for the set call:
* - Set W02-SELECTORS-TABLE to the attributes to be set
* - Set W02-INTATTRS-TABLE  to the required status
* - All other variables are already set
*
     MOVE MQIA-INHIBIT-GET   TO W02-SELECTORS(1).
     MOVE MQIA-INHIBIT-PUT   TO W02-SELECTORS(2).
     MOVE MQQA-GET-INHIBITED TO W02-INTATTRS(1).
     MOVE MQQA-PUT-INHIBITED TO W02-INTATTRS(2).
*
*    Set the attributes.
*
     CALL 'MQSET' USING W02-HCONN,
                        W02-HOBJ,
                        W02-SELECTORCOUNT,
                        W02-SELECTORS-TABLE,
                        W02-INTATTRCOUNT,
                        W02-INTATTRS-TABLE,
                        W02-CHARATTRLENGTH,
                        W02-CHARATTRS,
                        W02-COMPCODE,
                        W02-REASON.
*
* Test the output from the call:
*
*  - If the completion code is not OK, display an error
*    message showing the completion and reason codes
*
*  - Otherwise, move 'INHIBITED' into the relevant
*    screen map fields
*
     IF W02-COMPCODE NOT = MQCC-OK
        MOVE 'MQSET'       TO M01-MSG4-OPERATION
        MOVE W02-COMPCODE  TO M01-MSG4-COMPCODE
        MOVE W02-REASON    TO M01-MSG4-REASON
        MOVE M01-MESSAGE-4 TO M00-MESSAGE
     ELSE
*
*       Process the changes.
   :
     END-IF.
```

Figure 73 (Part 2 of 2). Using the MQSET call (COBOL)

Appendix E. System/390 assembler-language examples

The extracts in this appendix are mostly taken from the MQSeries for MVS/ESA sample applications.

The examples in this appendix demonstrate the following techniques:

489

Connecting to a queue manager

Figure 74 demonstrates how to use the MQCONN call to connect a program to a queue manager in MVS batch. This extract is taken from the Browse sample program (CSQ4BAA1) supplied with MQSeries for MVS/ESA.

```
      .
      .
WORKAREA DSECT
*
PARMLIST CALL ,(0,0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
COMPCODE DS   F            Completion code
REASON   DS   F            Reason code
HCONN    DS   F            Connection handle
         ORG
PARMADDR DS   F            Address of parm field
PARMLEN  DS   H            Length of parm field
*
MQMNAME  DS   CL48         Queue manager name
*
*
************************************************************
*   SECTION NAME : MAINPARM                               *
************************************************************
MAINPARM DS   0H
         MVI  MQMNAME,X'40'
         MVC  MQMNAME+1(L'MQMNAME-1),MQMNAME
*
* Space out first byte and initialize
*
*
* Code to address and verify parameters passed omitted
*
*
PARM1MVE DS   0H
         SR   R1,R3        Length of data
         LA   R4,MQMNAME   Address for target
         BCTR R1,R0        Reduce for execute
         EX   R1,MOVEPARM  Move the data
*
************************************************************
* EXECUTES                                                *
************************************************************
MOVEPARM MVC  0(*-*,R4),0(R3)
*
         EJECT
```

Figure 74 (Part 1 of 2). Using the MQCONN call (Assembler language)

```
************************************************************
*   SECTION NAME : MAINCONN                               *
************************************************************
*
*
MAINCONN DS   0H
         XC   HCONN,HCONN       Null connection handle
*
         CALL MQCONN,                          X
              (MQMNAME,                         X
              HCONN,                            X
              COMPCODE,                         X
              REASON),                          X
              MF=(E,PARMLIST),VL
*
         LA   R0,MQCC_OK      Expected compcode
         C    R0,COMPCODE     As expected?
         BER  R6              Yes .. return to caller
*
         MVC  INF4_TYP,=CL10'CONNECT   '
         BAL  R7,ERRCODE              Translate error
         LA   R0,8                    Set exit code
         ST   R0,EXITCODE             to 8
         B    ENDPROG                 End the program
*
```

Figure 74 (Part 2 of 2). Using the MQCONN call (Assembler language)

Disconnecting from a queue manager

Figure 75 demonstrates how to use the MQDISC call to disconnect a program from a queue manager in MVS batch. This extract is not taken from the sample applications supplied with MQSeries.

```
      .
      .
*
*      ISSUE MQI DISC REQUEST USING REENTRANT FORM
*      OF CALL MACRO
*
*      HCONN WAS SET BY A PREVIOUS MQCONN REQUEST
*      R5 = WORK REGISTER
*
DISC     DS   0H
         CALL MQDISC,                    X
              (HCONN,                    X
              COMPCODE,                  X
              REASON),                   X
              VL,MF=(E,CALLLST)
*
         LA   R5,MQCC_OK
         C    R5,COMPCODE
         BNE  BADCALL
      .
      .
```

Figure 75 (Part 1 of 2). Using the MQDISC call (Assembler language)

```
BADCALL DS    0H
  ⋮
*                CONSTANTS
*
       CMQA
*
*      WORKING STORAGE (RE-ENTRANT)
*
WEG3    DSECT
*
CALLLST CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
HCONN   DS    F
COMPCODE DS   F
REASON  DS    F
*
*
LEG3    EQU  *-WKEG3
        END
```

Figure 75 (Part 2 of 2). Using the MQDISC call (Assembler language)

Creating a dynamic queue

Figure 76 demonstrates how to use the MQOPEN call to create a dynamic queue. This extract is not taken from the sample applications supplied with MQSeries.

```
  ⋮
*
*    R5 = WORK REGISTER.
*
OPEN    DS    0H
*
MVC WOD_AREA,MQOD_AREA INITIALIZE WORKING VERSION OF
*                              MQOD WITH DEFAULTS
MVC WOD_OBJECTNAME,MOD_Q   COPY IN THE MODEL Q NAME
MVC WOD_DYNAMICQNAME,DYN_Q COPY IN THE DYNAMIC Q NAME
L   R5,=AL4(MQOO_OUTPUT)   OPEN FOR OUTPUT AND
A   R5,=AL4(MQOO_INQUIRE)  INQUIRE
ST  R5,OPTIONS
```

Figure 76 (Part 1 of 2). Using the MQOPEN call to create a dynamic queue (Assembler language)

```
*
* ISSUE MQI OPEN REQUEST USING REENTRANT
* FORM OF CALL MACRO
*
        CALL MQOPEN,                        X
             (HCONN,                        X
             WOD,                           X
             OPTIONS,                       X
             HOBJ,                          X
             COMPCODE,                      X
             REASON),VL,MF=(E,CALLLST)
*
  LA R5,MQCC_OK              CHECK THE COMPLETION CODE
  C  R5,COMPCODE             FROM THE REQUEST AND BRANCH
  BNE BADCALL                TO ERROR ROUTINE IF NOT MQCC_OK
*
  MVC  TEMP_Q,WOD_OBJECTNAME  SAVE NAME OF TEMPORARY Q
*                             CREATED BY OPEN OF MODEL Q
*
  ⋮
BADCALL DS    0H
  ⋮
*
*
*   CONSTANTS:
*
MOD_Q DC   CL48'QUERY.REPLY.MODEL'  MODEL QUEUE NAME
DYN_Q DC   CL48'QUERY.TEMPQ.*'      DYNAMIC QUEUE NAME
*
     CMQODA DSECT=NO,LIST=YES  CONSTANT VERSION OF MQOD
     CMQA                      MQI VALUE EQUATES
*
*   WORKING STORAGE
*
       DFHEISTG
HCONN   DS F                  CONNECTION HANDLE
OPTIONS DS F                  OPEN OPTIONS
HOBJ    DS F                  OBJECT HANDLE
COMPCODE DS F                 MQI COMPLETION CODE
REASON  DS F                  MQI REASON CODE
TEMP_Q  DS CL(MQ_Q_NAME_LENGTH)  SAVED QNAME AFTER OPEN
*
WOD     CMQODA DSECT=NO,LIST=YES  WORKING VERSION OF MQOD
*
CALLLST CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L  LIST FORM
*                                                 OF CALL
*                                                 MACRO
  ⋮
        END
```

Figure 76 (Part 2 of 2). Using the MQOPEN call to create a dynamic queue (Assembler language)

Opening an existing queue

Figure 77 demonstrates how to use the MQOPEN call to open a queue that has already been defined. It shows how to specify two options. This extract is not taken from the sample applications supplied with MQSeries.

```
   ⋮
*
*   R5 = WORK REGISTER.
*
OPEN    DS   0H
*
   MVC  WOD_AREA,MQOD_AREA  INITIALIZE WORKING VERSION OF
*                             MQOD WITH DEFAULTS
   MVC  WOD_OBJECTNAME,Q_NAME   SPECIFY Q NAME TO OPEN
   LA   R5,MQOO_INPUT_EXCLUSIVE  OPEN FOR MQGET CALLS
*
        ST   R5,OPTIONS
*
* ISSUE MQI OPEN REQUEST USING REENTRANT FORM
* OF CALL MACRO
*
        CALL MQOPEN,                     X
             (HCONN,                     X
             WOD,                        X
             OPTIONS,                    X
             HOBJ,                       X
             COMPCODE,                   X
             REASON),VL,MF=(E,CALLLST)
*
   LA R5,MQCC_OK        CHECK THE COMPLETION CODE
   C  R5,COMPCODE       FROM THE REQUEST AND BRANCH
   BNE BADCALL          TO ERROR ROUTINE IF NOT MQCC_OK
*
   ⋮
BADCALL  DS   0H
   ⋮
*
*
*   CONSTANTS:
*
Q_NAME   DC   CL48'REQUEST.QUEUE'  NAME OF QUEUE TO OPEN
*
        CMQODA DSECT=NO,LIST=YES  CONSTANT VERSION OF MQOD
        CMQA                      MQI VALUE EQUATES
*
*   WORKING STORAGE
*
        DFHEISTG
HCONN    DS F            CONNECTION HANDLE
OPTIONS  DS F            OPEN OPTIONS
HOBJ     DS F            OBJECT HANDLE
COMPCODE DS F            MQI COMPLETION CODE
REASON   DS F            MQI REASON CODE
*
WOD  CMQODA DSECT=NO,LIST=YES  WORKING VERSION OF MQOD
*
CALLLST  CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L  LIST FORM
                                              OF CALL
*                                              MACRO
   ⋮
        END
```

Figure 77. Using the MQOPEN call to open an existing queue (Assembler language)

Closing a queue

Figure 78 demonstrates how to use the MQCLOSE call to close a queue. This extract is not taken from the sample applications supplied with MQSeries.

```
   ⋮
*
* ISSUE MQI CLOSE REQUEST USING REENTRANT FROM OF
* CALL MACRO
*
*       HCONN WAS SET BY A PREVIOUS MQCONN REQUEST
*       HOBJ  WAS SET BY A PREVIOUS MQOPEN REQUEST
*       R5 = WORK REGISTER
*
CLOSE   DS    0H
        LA    R5,MQCO_NONE      NO SPECIAL CLOSE OPTIONS
        ST    R5,OPTIONS            ARE REQUIRED.
*
        CALL  MQCLOSE,                 X
              (HCONN,                  X
              HOBJ,                    X
              OPTIONS,                 X
              COMPCODE,                X
              REASON),                 X
              VL,MF=(E,CALLLST)
*
        LA    R5,MQCC_OK
        C     R5,COMPCODE
        BNE   BADCALL
*
   ⋮
BADCALL DS    0H
   ⋮
*                CONSTANTS
*
        CMQA
*
*       WORKING STORAGE (REENTRANT)
*
WEG4    DSECT
*
CALLLST CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
HCONN    DS   F
HOBJ     DS   F
OPTIONS  DS   F
COMPCODE DS   F
REASON   DS   F
*
*
LEG4    EQU  *-WKEG4
        END
```

Figure 78. Using the MQCLOSE call (Assembler language)

Putting a message using MQPUT

Figure 79 demonstrates how to use the MQPUT call to put a message on a queue. This extract is not taken from the sample applications supplied with MQSeries.

```
     ⋮
*      CONNECT TO QUEUE MANAGER
*
CONN    DS  0H
     ⋮
*
*      OPEN A QUEUE
*
OPEN    DS  0H
     ⋮
*
*      R4,R5,R6,R7 = WORK REGISTER.
*
PUT DS  0H
    LA   R4,MQMD          SET UP ADDRESSES AND
    LA   R5,MQMD_LENGTH   LENGTH FOR USE BY MVCL
    LA   R6,WMD           INSTRUCTION, AS MQMD IS
    LA   R7,WMD_LENGTH    OVER 256 BYES LONG.
    MVCL R6,R4            INITIALIZE WORKING VERSION
*                         OF MESSAGE DESCRIPTOR
*
    MVC  WPMO_AREA,MQPMO_AREA  INITIALIZE WORKING MQPMO
*

    LA   R5,BUFFER_LEN   RETRIEVE THE BUFFER LENGTH
    ST   R5,BUFFLEN      AND SAVE IT FOR MQM USE
*
    MVC  BUFFER,TEST_MSG    SET THE MESSAGE TO BE PUT
*
*   ISSUE MQI PUT REQUEST USING REENTRANT FORM
*   OF CALL MACRO
*
*      HCONN WAS SET BY PREVIOUS MQCONN REQUEST
*      HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST
*
       CALL MQPUT,                   X
            (HCONN,                  X
            HOBJ,                    X
            WMD,                     X
            WPMO,                    X
            BUFFLEN,                 X
            BUFFER,                  X
            COMPCODE,                X
            REASON),VL,MF=(E,CALLLST)
*
       LA   R5,MQCC_OK
       C    R5,COMPCODE
       BNE BADCALL
*
     ⋮
BADCALL DS  0H
     ⋮
```

Figure 79 (Part 1 of 2). Using the MQPUT call (Assembler language)

```
*
*      CONSTANTS
*
   CMQMDA DSECT=NO,LIST=YES,PERSISTENCE=MQPER_PERSISTENT
   CMQPMOA DSECT=NO,LIST=YES
   CMQA
TEST_MSG DC CL80'THIS IS A TEST MESSAGE'
*
*      WORKING STORAGE DSECT
*
WORKSTG  DSECT
*
COMPCODE DS F
REASON   DS F
BUFFLEN  DS F
OPTIONS  DS F
HCONN    DS F
HOBJ     DS F
*
BUFFER   DS CL80
BUFFER_LEN EQU *-BUFFER
*
WMD      CMQMDA DSECT=NO,LIST=NO
WPMO     CMQPMOA DSECT=NO,LIST=NO
*
CALLLST  CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
     ⋮
         END
```

Figure 79 (Part 2 of 2). Using the MQPUT call (Assembler language)

Putting a message using MQPUT1

Figure 80 demonstrates how to use the MQPUT1 call to open a queue, put a single message on the queue, then close the queue. This extract is not taken from the sample applications supplied with MQSeries.

```
     ⋮
*
*      CONNECT TO QUEUE MANAGER
*
CONN    DS  0H
     ⋮
*
*      R4,R5,R6,R7 = WORK REGISTER.
*
PUT     DS  0H
*
 MVC  WOD_AREA,MQOD_AREA     INITIALIZE WORKING VERSION OF
*                                MQOD WITH DEFAULTS
 MVC  WOD_OBJECTNAME,Q_NAME  SPECIFY Q NAME FOR PUT1
*
 LA   R4,MQMD          SET UP ADDRESSES AND
 LA   R5,MQMD_LENGTH   LENGTH FOR USE BY MVCL
 LA   R6,WMD           INSTRUCTION, AS MQMD IS
 LA   R7,WMD_LENGTH    OVER 256 BYES LONG.
 MVCL R6,R4            INITIALIZE WORKING VERSION
*                      OF MESSAGE DESCRIPTOR
```

Figure 80 (Part 1 of 2). Using the MQPUT1 call (Assembler language)

```
*
 MVC  WPMO_AREA,MQPMO_AREA        INITIALIZE WORKING MQPMO
*

 LA   R5,BUFFER_LEN              RETRIEVE THE BUFFER LENGTH
 ST   R5,BUFFLEN                 AND SAVE IT FOR MQM USE
*
 MVC  BUFFER,TEST_MSG            SET THE MESSAGE TO BE PUT
*
* ISSUE MQI PUT REQUEST USING REENTRANT FORM OF CALL MACRO
*
*      HCONN WAS SET BY PREVIOUS MQCONN REQUEST
*      HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST
*
       CALL MQPUT1,                             X
             (HCONN,                            X
              LMQOD,                            X
              LMQMD,                            X
              LMQPMO,                           X
              BUFFERLENGTH,                     X
              BUFFER,                           X
              COMPCODE,                         X
              REASON),VL,MF=(E,CALLLST)
*
       LA R5,MQCC_OK
       C  R5,COMPCODE
       BNE BADCALL
*
   ⋮
BADCALL DS  0H
  ⋮
*
*      CONSTANTS
*
 CMQMDA DSECT=NO,LIST=YES,PERSISTENCE=MQPER_PERSISTENT
 CMQPMOA DSECT=NO,LIST=YES
 CMQODA DSECT=NO,LIST=YES
 CMQA
*
TEST_MSG DC CL80'THIS IS ANOTHER TEST MESSAGE'
Q_NAME   DC CL48'TEST.QUEUE.NAME'
*
*      WORKING STORAGE DSECT
*
WORKSTG  DSECT
*
COMPCODE DS F
REASON   DS F
BUFFLEN  DS F
OPTIONS  DS F
HCONN    DS F
HOBJ     DS F
*
BUFFER   DS CL80
BUFFER_LEN EQU *-BUFFER
*
WOD   CMQODA DSECT=NO,LIST=YES     WORKING VERSION OF MQOD
WMD      CMQMDA DSECT=NO,LIST=NO
WPMO     CMQPMOA DSECT=NO,LIST=NO
*
CALLLST  CALL ,(0,0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
  ⋮
       END
```

*Figure 80 (Part 2 of 2). Using the MQPUT1 call
(Assembler language)*

Getting a message

Figure 81 demonstrates how to use the MQGET
call to remove a message from a queue. This
extract is not taken from the sample applications
supplied with MQSeries.

```
  ⋮
*
*      CONNECT TO QUEUE MANAGER
*
CONN    DS  0H
  ⋮
*
*      OPEN A QUEUE FOR GET
*
OPEN    DS  0H
  ⋮
*
*      R4,R5,R6,R7 = WORK REGISTER.
*
GET DS 0H
    LA   R4,MQMD              SET UP ADDRESSES AND
    LA   R5,MQMD_LENGTH       LENGTH FOR USE BY MVCL
    LA   R6,WMD               INSTRUCTION, AS MQMD IS
    LA   R7,WMD_LENGTH        OVER 256 BYES LONG.
    MVCL R6,R4                INITIALIZE WORKING VERSION
*                                 OF MESSAGE DESCRIPTOR
*
    MVC  WGMO_AREA,MQGMO_AREA  INITIALIZE WORKING MQGMO
*

    LA   R5,BUFFER_LEN        RETRIEVE THE BUFFER LENGTH
    ST   R5,BUFFLEN           AND SAVE IT FOR MQM USE
*
*
* ISSUE MQI GET REQUEST USING REENTRANT FORM OF CALL MACRO
*
*      HCONN WAS SET BY PREVIOUS MQCONN REQUEST
*      HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST
*
       CALL  MQGET,                       X
             (HCONN,                      X
              HOBJ,                       X
              WMD,                        X
              WGMO,                       X
              BUFFLEN,                    X
              BUFFER,                     X
              DATALEN,                    X
              COMPCODE,                   X
              REASON),                    X
              VL,MF=(E,CALLLST)
*
       LA R5,MQCC_OK
       C  R5,COMPCODE
       BNE BADCALL
*
   ⋮
BADCALL DS  0H
  ⋮
```

*Figure 81 (Part 1 of 2). Using the MQGET call
(Assembler language)*

```
*
*      CONSTANTS
*
       CMQMDA DSECT=NO,LIST=YES
       CMQGMOA DSECT=NO,LIST=YES
       CMQA
*
*      WORKING STORAGE DSECT
*
WORKSTG  DSECT
*
COMPCODE DS F
REASON   DS F
BUFFLEN  DS F
DATALEN  DS F
OPTIONS  DS F
HCONN    DS F
HOBJ     DS F
*
BUFFER   DS CL80
BUFFER_LEN EQU *-BUFFER
*
WMD      CMQMDA DSECT=NO,LIST=NO
WGMO     CMQGMOA DSECT=NO,LIST=NO
*
CALLLST  CALL ,(0,0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
   ⋮
         END
```

*Figure 81 (Part 2 of 2). Using the MQGET call
(Assembler language)*

Getting a message using the wait option

Figure 82 demonstrates how to use the wait
option of the MQGET call. This code accepts
truncated messages. This extract is not taken
from the sample applications supplied with
MQSeries.

```
   ⋮
*    CONNECT TO QUEUE MANAGER
CONN   DS  0H
   ⋮
*    OPEN A QUEUE FOR GET
OPEN   DS  0H
   ⋮
*    R4,R5,R6,R7 = WORK REGISTER.
GET DS  0H
    LA  R4,MQMD          SET UP ADDRESSES AND
    LA  R5,MQMD_LENGTH    LENGTH FOR USE BY MVCL
    LA  R6,WMD            INSTRUCTION, AS MQMD IS
    LA  R7,WMD_LENGTH     OVER 256 BYES LONG.
    MVCL R6,R4            INITIALIZE WORKING VERSION
*                            OF MESSAGE DESCRIPTOR
```

*Figure 82 (Part 1 of 3). Using the MQGET call with
the wait option (Assembler language)*

```
*
    MVC  WGMO_AREA,MQGMO_AREA   INITIALIZE WORKING MQGMO
    L    R5,=AL4(MQGMO_WAIT)
    A    R5,=AL4(MQGMO_ACCEPT_TRUNCATED_MSG)
    ST   R5,WGMO_OPTIONS
    MVC  WGMO_WAITINTERVAL,TWO_MINUTES  WAIT UP TO TWO
*                                       MINUTES BEFORE
*                                       FAILING THE
*                                       CALL
*
    LA  R5,BUFFER_LEN   RETRIEVE THE BUFFER LENGTH
    ST  R5,BUFFLEN      AND SAVE IT FOR MQM USE
*
*  ISSUE MQI GET REQUEST USING REENTRANT FORM OF CALL MACRO
*
*      HCONN WAS SET BY PREVIOUS MQCONN REQUEST
*      HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST
*
       CALL  MQGET,                       X
             (HCONN,                       X
              HOBJ,                        X
              WMD,                         X
              WGMO,                        X
              BUFFLEN,                     X
              BUFFER,                      X
              DATALEN,                     X
              COMPCODE,                    X
              REASON),                     X
             VL,MF=(E,CALLLST)
*
  LA  R5,MQCC_OK               DID THE MQGET REQUEST
  C   R5,COMPCODE              WORK OK?
  BE  GETOK                    YES, SO GO AND PROCESS.
  LA  R5,MQCC_WARNING          NO, SO CHECK FOR A WARNING.
  C   R5,COMPCODE              IS THIS A WARNING?
  BE  CHECK_W                  YES, SO CHECK THE REASON.
*
  LA  R5,MQRC_NO_MSG_AVAILABLE  IT MUST BE AN ERROR.
*                               IS IT DUE TO AN EMPTY
  C   R5,REASON                QUEUE?
  BE  NOMSG                    YES, SO HANDLE THE ERROR
  B   BADCALL                  NO, SO GO TO ERROR ROUTINE
*
CHECK_W DS 0H
       LA  R5,MQRC_TRUNCATED_MSG_ACCEPTED   IS THIS A
*                                            TRUNCATED
       C   R5,REASON                         MESSAGE?
       BE  GETOK                 YES, SO GO AND PROCESS.
       B   BADCALL               NO, SOME OTHER WARNING
*
NOMSG   DS 0H
  ⋮
GETOK   DS 0H
  ⋮
BADCALL DS 0H
  ⋮
*
*      CONSTANTS
*
       CMQMDA DSECT=NO,LIST=YES
       CMQGMOA DSECT=NO,LIST=YES
       CMQA
*
TWO_MINUTES DC F'120000'     GET WAIT INTERVAL
*
*      WORKING STORAGE DSECT
```

*Figure 82 (Part 2 of 3). Using the MQGET call with
the wait option (Assembler language)*

```
*
WORKSTG  DSECT
*
COMPCODE DS F
REASON   DS F
BUFFLEN  DS F
DATALEN  DS F
OPTIONS  DS F
HCONN    DS F
HOBJ     DS F
*
BUFFER   DS CL80
BUFFER_LEN EQU *-BUFFER
*
WMD      CMQMDA DSECT=NO,LIST=NO
WGMO     CMQGMOA DSECT=NO,LIST=NO
*
CALLLST  CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
   ⋮
         END
```

Figure 82 (Part 3 of 3). Using the MQGET call with the wait option (Assembler language)

Getting a message using signaling

Figure 83 demonstrates how to use the MQGET call to set a signal so that you are notified when a suitable message arrives on a queue. This extract is not taken from the sample applications supplied with MQSeries.

```
   ⋮
*
*    CONNECT TO QUEUE MANAGER
*
CONN    DS  0H
   ⋮
*
*    OPEN A QUEUE FOR GET
*
OPEN    DS  0H
   ⋮
*
*    R4,R5,R6,R7 = WORK REGISTER.
*
GET DS  0H
    LA   R4,MQMD            SET UP ADDRESSES AND
    LA   R5,MQMD_LENGTH     LENGTH FOR USE BY MVCL
    LA   R6,WMD             INSTRUCTION, AS MQMD IS
    LA   R7,WMD_LENGTH      OVER 256 BYES LONG.
    MVCL R6,R4              INITIALIZE WORKING VERSION
*                              OF MESSAGE DESCRIPTOR
```

Figure 83 (Part 1 of 3). Using the MQGET call with signaling (Assembler language)

```
*
    MVC  WGMO_AREA,MQGMO_AREA   INITIALIZE WORKING MQGMO
    LA   R5,MQGMO_SET_SIGNAL
    ST   R5,WGMO_OPTIONS
    MVC  WGMO_WAITINTERVAL,FIVE_MINUTES   WAIT UP TO FIVE
                                          MINUTES BEFORE
*                                         FAILING THE CALL
*
    XC   SIG_ECB,SIG_ECB   CLEAR THE ECB
    LA   R5,SIG_ECB        GET THE ADDRESS OF THE ECB
    ST   R5,WGMO_SIGNAL1   AND PUT IT IN THE WORKING
*                              MQGMO
*

    LA   R5,BUFFER_LEN     RETRIEVE THE BUFFER LENGTH
    ST   R5,BUFFLEN        AND SAVE IT FOR MQM USE
*
*
*  ISSUE MQI GET REQUEST USING REENTRANT FORM OF CALL MACRO
*
*    HCONN WAS SET BY PREVIOUS MQCONN REQUEST
*    HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST
*
    CALL  MQGET,                          X
          (HCONN,                         X
          HOBJ,                           X
          WMD,                            X
          WGMO,                           X
          BUFFLEN,                        X
          BUFFER,                         X
          DATALEN,                        X
          COMPCODE,                       X
          REASON),                        X
          VL,MF=(E,CALLLST)
*
    LA R5,MQCC_OK        DID THE MQGET REQUEST
    C  R5,COMPCODE       WORK OK?
    BE GETOK             YES, SO GO AND PROCESS.
    LA R5,MQCC_WARNING   NO, SO CHECK FOR A WARNING.
    C  R5,COMPCODE       IS THIS A WARNING?
    BE CHECK_W           YES, SO CHECK THE REASON.
    B  BADCALL           NO, SO GO TO ERROR ROUTINE
*
```

Figure 83 (Part 2 of 3). Using the MQGET call with signaling (Assembler language)

```
CHECK_W  DS  0H
         LA  R5,MQRC_SIGNAL_REQUEST_ACCEPTED
         C   R5,REASON   SIGNAL REQUEST SIGNAL SET?
         BNE BADCALL     NO, SOME ERROR OCCURRED
         B   DOWORK      YES, SO DO SOMETHING
*                             ELSE
*
CHECKSIG DS  0H
         CLC SIG_ECB+1(3),=AL3(MQEC_MSG_ARRIVED)
                              IS A MESSAGE AVAILABLE?
         BE  GET              YES, SO GO AND GET IT
*
         CLC SIG_ECB+1(3),=AL3(MQEC_WAIT_INTERVAL_EXPIRED)
                              HAVE WE WAITED LONG ENOUGH?
         BE  NOMSG       YES, SO SAY NO MSG AVAILABLE
         B   BADCALL     IF IT'S ANYTHING ELSE
*                        GO TO ERROR ROUTINE.
*
DOWORK   DS  0H
    ⋮
         TM  SIG_ECB,X'40'  HAS THE SIGNAL ECB BEEN POSTED?
         BO  CHECKSIG       YES, SO GO AND CHECK WHY
         B   DOWORK         NO, SO GO AND DO MORE WORK
*
NOMSG    DS  0H
    ⋮
GETOK    DS  0H
    ⋮
BADCALL  DS  0H
    ⋮
*
*        CONSTANTS
*
         CMQMDA DSECT=NO,LIST=YES
         CMQGMOA DSECT=NO,LIST=YES
         CMQA
*
FIVE_MINUTES DC F'300000'      GET SIGNAL INTERVAL
*
*        WORKING STORAGE DSECT
*
WORKSTG  DSECT
*
COMPCODE DS F
REASON   DS F
BUFFLEN  DS F
DATALEN  DS F
OPTIONS  DS F
HCONN    DS F
HOBJ     DS F
SIG_ECB  DS F
*
BUFFER   DS CL80
BUFFER_LEN EQU *-BUFFER
*
WMD      CMQMDA DSECT=NO,LIST=NO
WGMO     CMQGMOA DSECT=NO,LIST=NO
*
CALLLLST CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
*
    ⋮
         END
```

Figure 83 (Part 3 of 3). Using the MQGET call with signaling (Assembler language)

Inquiring about and setting the attributes of a queue

Figure 84 demonstrates how to use the MQINQ call to inquire about the attributes of a queue and to use the MQSET call to change the attributes of a queue. This extract is taken from the Queue Attributes sample application (program CSQ4CAC1) supplied with MQSeries for MVS/ESA.

```
    ⋮
DFHEISTG DSECT
    ⋮
OBJDESC  CMQODA LIST=YES   Working object descriptor
*
SELECTORCOUNT    DS F      Number of selectors
INTATTRCOUNT     DS F      Number of integer attributes
CHARATTRLENGTH   DS F      char attributes length
CHARATTRS        DS C      Area for char attributes
*
OPTIONS  DS  F             Command options
HCONN    DS  F             Handle of connection
HOBJ     DS  F             Handle of object
COMPCODE DS  F             Completion code
REASON   DS  F             Reason code
SELECTOR DS  2F            Array of selectors
INTATTRS DS  2F            Array of integer attributes
    ⋮
OBJECT   DS  CL(MQ_Q_NAME_LENGTH)   Name of queue
    ⋮
CALLLIST CALL ,(0,0,0,0,0,0,0,0,0,0,0),VL,MF=L
**********************************************************
*              PROGRAM EXECUTION STARTS HERE        *
    ⋮
CSQ4CAC1 DFHEIENT CODEREG=(R3),DATAREG=(R13)
    ⋮
*        Initialize the variables for the set call
*
         SR  R0,R0             Clear register zero
         ST  R0,CHARATTRLENGTH Set char length to zero
         LA  R0,2              Load to set
         ST  R0,SELECTORCOUNT  selectors add
         ST  R0,INTATTRCOUNT   integer attributes
*
         LA  R0,MQIA_INHIBIT_GET Load q attribute selector
         ST  R0,SELECTOR+0     Place in field
         LA  R0,MQIA_INHIBIT_PUT Load q attribute selector
         ST  R0,SELECTOR+4     Place in field
*
UPDTEST  DS  0H
         CLC ACTION,CINHIB     Are we inhibiting?
         BE  UPDINHBT          Yes branch to section
*
         CLC ACTION,CALLOW     Are we allowing?
         BE  UPDALLOW          Yes branch to section
*
         MVC M00_MSG,M01_MSG1  Invalid request
         BR  R6                Return to caller
*
```

Figure 84 (Part 1 of 2). Using the MQINQ and MQSET calls (Assembler language)

Assembler-language examples

```
UPDINHBT DS   0H
         MVC  UPDTYPE,CINHIBIT     Indicate action type
         LA   R0,MQQA_GET_INHIBITED Load attribute value
         ST   R0,INTATTRS+0        Place in field
         LA   R0,MQQA_PUT_INHIBITED Load attribute value
         ST   R0,INTATTRS+4        Place in field
         B    UPDCALL              Go and do call
*
UPDALLOW DS   0H
         MVC  UPDTYPE,CALLOWED     Indicate action type
         LA   R0,MQQA_GET_ALLOWED  Load attribute value
         ST   R0,INTATTRS+0        Place in field
         LA   R0,MQQA_PUT_ALLOWED  Load attribute value
         ST   R0,INTATTRS+4        Place in field
         B    UPDCALL              Go and do call
*
UPDCALL  DS   0H
         CALL MQSET,                              C
              (HCONN,                             C
              HOBJ,                               C
              SELECTORCOUNT,                      C
              SELECTOR,                           C
              INTATTRCOUNT,                       C
              INTATTRS,                           C
              CHARATTRLENGTH,                     C
              CHARATTRS,                          C
              COMPCODE,                           C
              REASON),                            C
              VL,MF=(E,CALLLIST)
*
         LA   R0,MQCC_OK    Load expected compcode
         C    R0,COMPCODE   Was set successful?
  ⋮
*  SECTION NAME : INQUIRE                                *
*  FUNCTION     : Inquires on the objects attributes     *
*  CALLED BY    : PROCESS                                 *
*  CALLS        : OPEN, CLOSE, CODES                      *
*  RETURN       : To Register 6                           *
INQUIRE  DS   0H
  ⋮
*        Initialize the variables for the inquire call
*
         SR   R0,R0              Clear register zero
         ST   R0,CHARATTRLENGTH  Set char length to zero
         LA   R0,2               Load to set
         ST   R0,SELECTORCOUNT   selectors add
         ST   R0,INTATTRCOUNT    integer attributes
*
         LA   R0,MQIA_INHIBIT_GET Load attribute value
         ST   R0,SELECTOR+0       Place in field
         LA   R0,MQIA_INHIBIT_PUT Load attribute value
         ST   R0,SELECTOR+4       Place in field
         CALL MQINQ,                              C
              (HCONN,                             C
              HOBJ,                               C
              SELECTORCOUNT,                      C
              SELECTOR,                           C
              INTATTRCOUNT,                       C
              INTATTRS,                           C
              CHARATTRLENGTH,                     C
              CHARATTRS,                          C
              COMPCODE,                           C
              REASON),                            C
              VL,MF=(E,CALLLIST)
         LA   R0,MQCC_OK         Load expected compcode
         C    R0,COMPCODE        Was inquire successful?
  ⋮
```

*Figure 84 (Part 2 of 2). Using the MQINQ and
MQSET calls (Assembler language)*

Appendix F. PL/I examples

The use of PL/I is supported by MQSeries for MVS/ESA, AIX, OS/2 Warp, and Windows NT only.

The examples demonstrate the following techniques:

Connecting to a queue manager

Figure 85 demonstrates how to use the MQCONN call to connect a program to a queue manager in MVS batch. This extract is not taken from the sample applications supplied with MQSeries.

```
%INCLUDE SYSLIB(CMQP);
%INCLUDE SYSLIB(CMQEPP);
  :
/****************************************************/
/* STRUCTURE BASED ON PARAMETER INPUT AREA (PARAM) */
/****************************************************/
DCL 1 INPUT_PARAM      BASED(ADDR(PARAM)),
      2 PARAM_LENGTH   FIXED BIN(15),
      2 PARAM_MQMNAME  CHAR(48);
  :
/****************************************************/
/* WORKING STORAGE DECLARATIONS                 */
/****************************************************/
DCL MQMNAME                 CHAR(48);
DCL COMPCODE                BINARY FIXED (31);
DCL REASON                  BINARY FIXED (31);
DCL HCONN                   BINARY FIXED (31);
  :
/****************************************************/
/* COPY QUEUE MANAGER NAME PARAMETER            */
/* TO LOCAL STORAGE                             */
/****************************************************/
MQMNAME = ' ';
MQMNAME = SUBSTR(PARAM_MQMNAME,1,PARAM_LENGTH);
  :
/****************************************************/
/* CONNECT FROM THE QUEUE MANAGER               */
/****************************************************/
CALL MQCONN (MQMNAME,      /* MQM SYSTEM NAME     */
             HCONN,        /* CONNECTION HANDLE   */
             COMPCODE,     /* COMPLETION CODE     */
             REASON);      /* REASON CODE         */

/****************************************************/
/* TEST THE COMPLETION CODE OF THE CONNECT CALL.   */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE   */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE. */
/****************************************************/
IF COMPCODE ¬= MQCC_OK
   THEN DO;
  :
      CALL ERROR_ROUTINE;
   END;
```

Figure 85. Using the MQCONN call (PL/I)

Disconnecting from a queue manager

Figure 86 demonstrates how to use the MQDISC call to disconnect a program from a queue manager in MVS batch. This extract is not taken from the sample applications supplied with MQSeries.

```
%INCLUDE SYSLIB(CMQP);
%INCLUDE SYSLIB(CMQEPP);
  :
/****************************************************/
/* WORKING STORAGE DECLARATIONS                 */
/****************************************************/
DCL COMPCODE                BINARY FIXED (31);
DCL REASON                  BINARY FIXED (31);
DCL HCONN                   BINARY FIXED (31);

  :
/****************************************************/
/* DISCONNECT FROM THE QUEUE MANAGER            */
/****************************************************/
CALL MQDISC (HCONN,        /* CONNECTION HANDLE   */
             COMPCODE,     /* COMPLETION CODE     */
             REASON);      /* REASON CODE         */
/*************************************************************/
/* TEST THE COMPLETION CODE OF THE DISCONNECT CALL.       */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE          */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE.       */
/*************************************************************/
IF COMPCODE ¬= MQCC_OK
   THEN DO;
  :
      CALL ERROR_ROUTINE;
   END;
```

Figure 86. Using the MQDISC call (PL/I)

Creating a dynamic queue

Figure 87 demonstrates how to use the MQOPEN
call to create a dynamic queue. This extract is not
taken from the sample applications supplied with
MQSeries.

```
    %INCLUDE SYSLIB(CMQP);
    %INCLUDE SYSLIB(CMQEPP);
    :
/***********************************************************/
/* WORKING STORAGE DECLARATIONS                            */
/***********************************************************/
DCL COMPCODE              BINARY FIXED (31);
DCL REASON                BINARY FIXED (31);
DCL HCONN                 BINARY FIXED (31);
DCL HOBJ                  BINARY FIXED (31);
DCL OPTIONS               BINARY FIXED (31);
   :
DCL MODEL_QUEUE_NAME    CHAR(48) INIT('PL1.REPLY.MODEL');
DCL DYNAMIC_NAME_PREFIX CHAR(48) INIT('PL1.TEMPQ.*');
DCL DYNAMIC_QUEUE_NAME  CHAR(48) INIT(' ');
   :
/***********************************************************/
/* LOCAL COPY OF OBJECT DESCRIPTOR                         */
/***********************************************************/
DCL 1 LMQOD  LIKE MQOD;
   :
/***********************************************************/
/* SET UP OBJECT DESCRIPTOR FOR OPEN OF REPLY QUEUE        */
/***********************************************************/
LMQOD.OBJECTTYPE =MQOT_Q;
LMQOD.OBJECTNAME = MODEL_QUEUE_NAME;
LMQOD.DYNAMICQNAME = DYNAMIC_NAME_PREFIX;
OPTIONS = MQOO_INPUT_EXCLUSIVE;

    CALL MQOPEN (HCONN,
                 LMQOD,
                 OPTIONS,
                 HOBJ,
                 COMPCODE,
                 REASON);

/***********************************************************/
/* TEST THE COMPLETION CODE OF THE OPEN CALL.              */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE           */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE.        */
/* IF THE CALL HAS SUCCEEDED THEN EXTRACT THE NAME OF      */
/* THE NEWLY CREATED DYNAMIC QUEUE FROM THE OBJECT         */
/* DESCRIPTOR.                                             */
/***********************************************************/
    IF COMPCODE ¬= MQCC_OK
       THEN DO;

  :

         CALL ERROR_ROUTINE;
       END;
       ELSE
         DYNAMIC_QUEUE_NAME = LMQOD_OBJECTNAME;
```

Figure 87. Using the MQOPEN call to create a
dynamic queue (PL/I)

Opening an existing queue

Figure 88 demonstrates how to use the MQOPEN
call to open an existing queue. This extract is not
taken from the sample applications supplied with
MQSeries.

```
%INCLUDE SYSLIB(CMQP);
%INCLUDE SYSLIB(CMQEPP);
:
/***********************************************************/
/* WORKING STORAGE DECLARATIONS                            */
/***********************************************************/
DCL COMPCODE              BINARY FIXED (31);
DCL REASON                BINARY FIXED (31);
DCL HCONN                 BINARY FIXED (31);
DCL HOBJ                  BINARY FIXED (31);
DCL OPTIONS               BINARY FIXED (31);
   :
DCL QUEUE_NAME         CHAR(48) INIT('PL1.LOCAL.QUEUE');
   :
/***********************************************************/
/* LOCAL COPY OF OBJECT DESCRIPTOR                         */
/***********************************************************/
DCL 1 LMQOD  LIKE MQOD;
   :
/***********************************************************/
/* SET UP OBJECT DESCRIPTOR FOR OPEN OF REPLY QUEUE        */
/***********************************************************/
LMQOD.OBJECTTYPE = MQOT_Q;
LMQOD.OBJECTNAME = QUEUE_NAME;
OPTIONS = MQOO_INPUT_EXCLUSIVE;

CALL MQOPEN (HCONN,
             LMQOD,
             OPTIONS,
             HOBJ,
             COMPCODE,
             REASON);


/***********************************************************/
/* TEST THE COMPLETION CODE OF THE OPEN CALL.              */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE           */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE.        */
/***********************************************************/
    IF COMPCODE ¬= MQCC_OK
       THEN DO;

  :

        CALL ERROR_ROUTINE;
       END;
```

Figure 88. Using the MQOPEN call to open an existing
queue (PL/I)

Closing a queue

Figure 89 demonstrates how to use the MQCLOSE call. This extract is not taken from the sample applications supplied with MQSeries.

```
%INCLUDE SYSLIB(CMQP);
%INCLUDE SYSLIB(CMQEPP);
 :
/*********************************************************/
/* WORKING STORAGE DECLARATIONS                          */
/*********************************************************/
DCL COMPCODE              BINARY FIXED (31);
DCL REASON                BINARY FIXED (31);
DCL HCONN                 BINARY FIXED (31);
DCL HOBJ                  BINARY FIXED (31);
DCL OPTIONS               BINARY FIXED (31);
  :
/*********************************************************/
/* SET CLOSE OPTIONS                                     */
/*********************************************************/
OPTIONS=MQCO_NONE;

/*********************************************************/
/* CLOSE QUEUE                                           */
/*********************************************************/
   CALL MQCLOSE (HCONN,      /* CONNECTION HANDLE      */
                 HOBJ,       /* OBJECT HANDLE          */
                 OPTIONS,    /* CLOSE OPTIONS          */
                 COMPCODE,   /* COMPLETION CODE        */
                 REASON);    /* REASON CODE            */

/*********************************************************/
/* TEST THE COMPLETION CODE OF THE CLOSE CALL.           */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE         */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE.      */
/*********************************************************/
    IF COMPCODE ¬= MQCC_OK
       THEN DO;

  :

         CALL ERROR_ROUTINE;
        END;
```

Figure 89. Using the MQCLOSE call (PL/I)

Putting a message using MQPUT

Figure 90 demonstrates how to use the MQPUT call using context. This extract is not taken from the sample applications supplied with MQSeries.

```
%INCLUDE SYSLIB(CMQP);
%INCLUDE SYSLIB(CMQEPP);
 :
/*********************************************************/
/* WORKING STORAGE DECLARATIONS                          */
/*********************************************************/
DCL COMPCODE              BINARY FIXED (31);
DCL REASON                BINARY FIXED (31);
DCL HCONN                 BINARY FIXED (31);
DCL HOBJ                  BINARY FIXED (31);
DCL OPTIONS               BINARY FIXED (31);
DCL BUFFLEN               BINARY FIXED (31);
DCL BUFFER                CHAR(80);
  :
DCL PL1_TEST_MESSAGE      CHAR(80)
INIT('*****    THIS IS A TEST MESSAGE    *****');
  :
*********************************************************/
/* LOCAL COPY OF MESSAGE DESCRIPTOR                    */
/* AND PUT MESSAGE OPTIONS                             */
/*********************************************************/
DCL 1 LMQMD  LIKE MQMD;
DCL 1 LMQPMO LIKE MQPMO;
  :
/*********************************************************/
/* SET UP MESSAGE DESCRIPTOR                            */
/*********************************************************/
LMQMD.MSGTYPE = MQMT_DATAGRAM;
LMQMD.PRIORITY = 1;
LMQMD.PERSISTENCE = MQPER_PERSISTENT;
LMQMD.REPLYTOQ = ' ';
LMQMD.REPLYTOQMGR = ' ';
LMQMD.MSGID = MQMI_NONE;
LMQMD.CORRELID = MQCI_NONE;

/*********************************************************/
/* SET UP PUT MESSAGE OPTIONS                           */
/*********************************************************/
LMQPMO.OPTIONS = MQPMO_NO_SYNCPOINT;

/*********************************************************/
/* SET UP LENGTH OF MESSAGE BUFFER AND THE MESSAGE      */
/*********************************************************/
BUFFLEN = LENGTH(BUFFER);
BUFFER = PL1_TEST_MESSAGE;
/*********************************************************/
/*                                                       */
/* HCONN WAS SET BY PREVIOUS MQCONN REQUEST.             */
/* HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST.              */
/*                                                       */
/*********************************************************/
CALL MQPUT (HCONN,
            HOBJ,
            LMQMD,
            LMQPMO,
            BUFFLEN,
            BUFFER,
            COMPCODE,
            REASON);
```

Figure 90 (Part 1 of 2). Using the MQPUT call (PL/I)

```
/********************************************************/
/* TEST THE COMPLETION CODE OF THE PUT CALL.          */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE      */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE.   */
/********************************************************/
    IF COMPCODE ¬= MQCC_OK
       THEN DO;

  ⋮

          CALL ERROR_ROUTINE;
       END;
```

Figure 90 (Part 2 of 2). Using the MQPUT call (PL/I)

Putting a message using MQPUT1

Figure 91 demonstrates how to use the MQPUT1 call. This extract is not taken from the sample applications supplied with MQSeries.

```
%INCLUDE SYSLIB(CMQEPP);
%INCLUDE SYSLIB(CMQP);
:
/********************************************************/
/* WORKING STORAGE DECLARATIONS                       */
/********************************************************/
DCL COMPCODE           BINARY FIXED (31);
DCL REASON             BINARY FIXED (31);
DCL HCONN              BINARY FIXED (31);
DCL OPTIONS            BINARY FIXED (31);
DCL BUFFLEN            BINARY FIXED (31);
DCL BUFFER             CHAR(80);
  ⋮
DCL REPLY_TO_QUEUE     CHAR(48) INIT('PL1.REPLY.QUEUE');
DCL QUEUE_NAME         CHAR(48) INIT('PL1.LOCAL.QUEUE');
DCL PL1_TEST_MESSAGE   CHAR(80)
  INIT('***** THIS IS ANOTHER TEST MESSAGE *****');
  ⋮
/********************************************************/
/* LOCAL COPY OF OBJECT DESCRIPTOR, MESSAGE DESCRIPTOR */
/* AND PUT MESSAGE OPTIONS                            */
/********************************************************/
DCL 1 LMQOD  LIKE MQOD;
DCL 1 LMQMD  LIKE MQMD;
DCL 1 LMQPMO LIKE MQPMO;
  ⋮
/********************************************************/
/* SET UP OBJECT DESCRIPTOR AS REQUIRED.              */
/********************************************************/
LMQOD.OBJECTTYPE = MQOT_Q;
LMQOD.OBJECTNAME = QUEUE_NAME;


/********************************************************/
/* SET UP MESSAGE DESCRIPTOR AS REQUIRED.             */
/********************************************************/
LMQMD.MSGTYPE = MQMT_REQUEST;
LMQMD.PRIORITY = 5;
LMQMD.PERSISTENCE = MQPER_PERSISTENT;
LMQMD.REPLYTOQ = REPLY_TO_QUEUE;
LMQMD.REPLYTOQMGR = ' ';
LMQMD.MSGID = MQMI_NONE;
LMQMD.CORRELID = MQCI_NONE;
```

Figure 91 (Part 1 of 2). Using the MQPUT1 call (PL/I)

```
/********************************************************/
/* SET UP PUT MESSAGE OPTIONS AS REQUIRED              */
/********************************************************/
     LMQPMO.OPTIONS = MQPMO_NO_SYNCPOINT;

/********************************************************/
/* SET UP LENGTH OF MESSAGE BUFFER AND THE MESSAGE     */
/********************************************************/
     BUFFLEN = LENGTH(BUFFER);
     BUFFER = PL1_TEST_MESSAGE;

     CALL MQPUT1 (HCONN,
                  LMQOD,
                  LMQMD,
                  LMQPMO,
                  BUFFLEN,
                  BUFFER,
                  COMPCODE,
                  REASON);

/********************************************************/
/* TEST THE COMPLETION CODE OF THE PUT1 CALL.          */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE SHOWING */
/* THE COMPLETION CODE AND THE REASON CODE.            */
/********************************************************/
     IF COMPCODE ¬= MQCC_OK
        THEN DO;

  ⋮

           CALL ERROR_ROUTINE;
        END;
```

Figure 91 (Part 2 of 2). Using the MQPUT1 call (PL/I)

Getting a message

Figure 92 demonstrates how to use the MQGET call to remove a message from a queue. This extract is not taken from the sample applications supplied with MQSeries.

```
     %INCLUDE SYSLIB(CMQP);
     %INCLUDE SYSLIB(CMQEPP);
     :
/********************************************************/
/* WORKING STORAGE DECLARATIONS                        */
/********************************************************/
     DCL COMPCODE              BINARY FIXED (31);
     DCL REASON                BINARY FIXED (31);
     DCL HCONN                 BINARY FIXED (31);
     DCL HOBJ                  BINARY FIXED (31);
     DCL BUFFLEN               BINARY FIXED (31);
     DCL DATALEN               BINARY FIXED (31);
     DCL BUFFER                CHAR(80);

  ⋮
```

Figure 92 (Part 1 of 2). Using the MQGET call (PL/I)

```
/********************************************************/
/* LOCAL COPY OF MESSAGE DESCRIPTOR AND                */
/* GET MESSAGE OPTIONS                                 */
/********************************************************/
     DCL 1 LMQMD  LIKE MQMD;
     DCL 1 LMQGMO LIKE MQGMO;

  ⋮

/********************************************************/
/* SET UP MESSAGE DESCRIPTOR AS REQUIRED.              */
/* MSGID AND CORRELID IN MQMD SET TO NULLS SO FIRST    */
/* AVAILABLE MESSAGE WILL BE RETRIEVED.                */
/********************************************************/
     LMQMD.MSGID = MQMI_NONE;
     LMQMD.CORRELID = MQCI_NONE;

/********************************************************/
/* SET UP GET MESSAGE OPTIONS AS REQUIRED.             */
/********************************************************/
     LMQGMO.OPTIONS = MQGMO_NO_SYNCPOINT;

/********************************************************/
/* SET UP LENGTH OF MESSAGE BUFFER.                    */
/********************************************************/
     BUFFLEN = LENGTH(BUFFER);
/********************************************************/
/*                                                     */
/* HCONN WAS SET BY PREVIOUS MQCONN REQUEST.           */
/* HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST.            */
/*                                                     */
/********************************************************/

     CALL MQGET (HCONN,
                 HOBJ,
                 LMQMD,
                 LMQGMO,
                 BUFFERLEN,
                 BUFFER,
                 DATALEN,
                 COMPCODE,
                 REASON);

/********************************************************/
/* TEST THE COMPLETION CODE OF THE GET CALL.           */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE       */
/* SHOWING THE COMPLETION CODE AND THE REASON CODE.    */
/********************************************************/
     IF COMPCODE ¬= MQCC_OK
        THEN DO;
          :
          :
          :
          CALL ERROR_ROUTINE;
        END;
```

Figure 92 (Part 2 of 2). Using the MQGET call (PL/I)

Getting a message using the wait option

Figure 93 demonstrates how to use the MQGET call with the wait option and accepting truncated messages. This extract is not taken from the sample applications supplied with MQSeries.

```
    %INCLUDE SYSLIB(CMQP);
    %INCLUDE SYSLIB(CMQEPP);
       :
/********************************************************/
/* WORKING STORAGE DECLARATIONS                         */
/********************************************************/
    DCL COMPCODE                  BINARY FIXED (31);
    DCL REASON                    BINARY FIXED (31);
    DCL HCONN                     BINARY FIXED (31);
    DCL HOBJ                      BINARY FIXED (31);
    DCL BUFFLEN                   BINARY FIXED (31);
    DCL DATALEN                   BINARY FIXED (31);
    DCL BUFFER                    CHAR(80);

      :
/********************************************************/
/* LOCAL COPY OF MESSAGE DESCRIPTOR AND GET MESSAGE     */
/* OPTIONS                                              */
/********************************************************/
    DCL 1 LMQMD  LIKE MQMD;
    DCL 1 LMQGMO LIKE MQGMO;

      :
/********************************************************/
/* SET UP MESSAGE DESCRIPTOR AS REQUIRED.               */
/* MSGID AND CORRELID IN MQMD SET TO NULLS SO FIRST     */
/* AVAILABLE MESSAGE WILL BE RETRIEVED.                 */
/********************************************************/
    LMQMD.MSGID = MQMI_NONE;
    LMQMD.CORRELID = MQCI_NONE;

/********************************************************/
/* SET UP GET MESSAGE OPTIONS AS REQUIRED.              */
/* WAIT INTERVAL SET TO ONE MINUTE.                     */
/********************************************************/
    LMQGMO.OPTIONS = MQGMO_WAIT +
                     MQGMO_ACCEPT_TRUNCATED_MSG +
                     MQGMO_NO_SYNCPOINT;
    LMQGMO.WAITINTERVAL=60000;

/********************************************************/
/* SET UP LENGTH OF MESSAGE BUFFER.                     */
/********************************************************/
    BUFFLEN = LENGTH(BUFFER);
```

Figure 93 (Part 1 of 2). Using the MQGET call with the wait option (PL/I)

```
/********************************************************/
/*                                                      */
/* HCONN WAS SET BY PREVIOUS MQCONN REQUEST.            */
/* HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST.             */
/*                                                      */
/********************************************************/
    CALL MQGET (HCONN,
                HOBJ,
                LMQMD,
                LMQGMO,
                BUFFERLEN,
                BUFFER,
                DATALEN,
                COMPCODE,
                REASON);

/********************************************************/
/* TEST THE COMPLETION CODE OF THE GET CALL.            */
/* TAKE APPROPRIATE ACTION BASED ON COMPLETION CODE AND */
/* REASON CODE.                                         */
/********************************************************/

    SELECT(COMPCODE);
      WHEN (MQCC_OK) DO;   /* GET WAS SUCCESSFUL */

      :

      END;
      WHEN (MQCC_WARNING) DO;
        IF REASON = MQRC_TRUNCATED_MSG_ACCEPTED
          THEN DO;          /* GET WAS SUCCESSFUL */

      :

          END;
          ELSE DO;

      :

            CALL ERROR_ROUTINE;
          END;
      END;
      WHEN (MQCC_FAILED) DO;

      :

          CALL ERROR_ROUTINE;
        END;
      END;
      OTHERWISE;
    END;
```

Figure 93 (Part 2 of 2). Using the MQGET call with the wait option (PL/I)

Getting a message using signaling

Signaling is available only with MQSeries for MVS/ESA.

Figure 94 demonstrates how to use the MQGET call with signaling. This extract is not taken from the sample applications supplied with MQSeries.

```
    %INCLUDE SYSLIB(CMQP);
    %INCLUDE SYSLIB(CMQEPP);
      :
/********************************************************/
/* WORKING STORAGE DECLARATIONS                       */
/********************************************************/
    DCL COMPCODE              BINARY FIXED (31);
    DCL REASON                BINARY FIXED (31);
    DCL HCONN                 BINARY FIXED (31);
    DCL HOBJ                  BINARY FIXED (31);
    DCL DATALEN               BINARY FIXED (31);
    DCL BUFFLEN               BINARY FIXED (31);
    DCL BUFFER                CHAR(80);

   :
    DCL ECB_FIXED        FIXED BIN(31);
    DCL 1 ECB_OVERLAY BASED(ADDR(ECB_FIXED)),
        3 ECB_WAIT    BIT,
        3 ECB_POSTED  BIT,
        3 ECB_FLAG3_8 BIT(6),
        3 ECB_CODE  PIC'999';

   :
/********************************************************/
/* LOCAL COPY OF MESSAGE DESCRIPTOR AND GET MESSAGE   */
/* OPTIONS                                            */
/********************************************************/
    DCL 1 LMQMD  LIKE MQMD;
    DCL 1 LMQGMO LIKE MQGMO;

   :
/********************************************************/
/* CLEAR ECB FIELD.                                   */
/********************************************************/
    ECB_FIXED = 0;

   :
/********************************************************/
/* SET UP MESSAGE DESCRIPTOR AS REQUIRED.             */
/* MSGID AND CORRELLID IN MQMD SET TO NULLS SO FIRST  */
/* AVAILABLE MESSAGE WILL BE RETRIEVED.               */
/********************************************************/
    LMQMD.MSGID = MQMI_NONE;
    LMQMD.CORRELID = MQCI_NONE;
/********************************************************/
/* SET UP GET MESSAGE OPTIONS AS REQUIRED.            */
/* WAIT INTERVAL SET TO ONE MINUTE.                   */
/********************************************************/
    LMQGMO.OPTIONS = MQGMO_SET_SIGNAL +
                     MQGMO_NO_SYNCPOINT;
    LMQGMO.WAITINTERVAL=60000;
    LMQGMO.SIGNAL1 = ADDR(ECB_FIXED);
```

Figure 94 (Part 1 of 3). Using the MQGET call with signaling (PL/I)

```
/********************************************************/
/* SET UP LENGTH OF MESSAGE BUFFER.                   */
/* CALL MESSGE RETRIEVAL ROUTINE.                     */
/********************************************************/
    BUFFLEN = LENGTH(BUFFER);
    CALL GET_MSG;

/********************************************************/
/* TEST THE COMPLETION CODE OF THE GET CALL.          */
/* TAKE APPROPRIATE ACTION BASED ON COMPLETION CODE AND */
/* REASON CODE.                                       */
/********************************************************/
    SELECT;
      WHEN ((COMPCODE = MQCC_OK) &
            (REASON = MQCC_NONE)) DO

   :
        CALL MSG_ROUTINE;

   :
      END;
      WHEN ((COMPCODE = MQCC_WARNING) &
            (REASON = MQRC_SIGNAL_REQUEST_ACCEPTED)) DO;

   :
        CALL DO_WORK;

   :
      END;
      WHEN  ((COMPCODE = MQCC_FAILED) &
            (REASON = MQRC_SIGNAL_OUTSTANDING)) DO;

   :
        CALL DO_WORK;

   :
      END;
      OTHERWISE DO;         /* FAILURE CASE */
/********************************************************/
/* ISSUE AN ERROR MESSAGE SHOWING THE COMPLETION CODE  */
/* AND THE REASON CODE.                               */
/********************************************************/

   :
          CALL ERROR_ROUTINE;

   :
      END;
    END;

   :
```

Figure 94 (Part 2 of 3). Using the MQGET call with signaling (PL/I)

```
DO_WORK: PROC;

  ⋮
    IF ECB_POSTED
      THEN DO;
        SELECT(ECB_CODE);
          WHEN(MQEC_MSG_ARRIVED) DO;

  ⋮

            CALL GET_MSG;

  ⋮

          END;
          WHEN(MQEC_WAIT_INTERVAL_EXPIRED) DO;

  ⋮

            CALL NO_MSG;

  ⋮

          END;
          OTHERWISE DO;          /* FAILURE CASE */
/**********************************************************/
/* ISSUE AN ERROR MESSAGE SHOWING THE COMPLETION CODE  */
/* AND THE REASON CODE.                                */
/**********************************************************/

  ⋮

            CALL ERROR_ROUTINE;

  ⋮

          END;

        END;

      END;

  ⋮
END DO_WORK;

GET_MSG: PROC;
/**********************************************************/
/*                                                      */
/* HCONN WAS SET BY PREVIOUS MQCONN REQUEST.            */
/* HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST.             */
/* MD AND GMO SET UP AS REQUIRED.                       */
/*                                                      */
/**********************************************************/

    CALL MQGET (HCONN,
                HOBJ,
                LMQMD,
                LMQGMO,
                BUFFLEN,
                BUFFER,
                DATALEN,
                COMPCODE,
                REASON);

END GET_MSG;

NO_MSG: PROC;

  ⋮

END NO_MSG;
```

Figure 94 (Part 3 of 3). Using the MQGET call with signaling (PL/I)

Inquiring about the attributes of an object

Figure 95 demonstrates how to use the MQINQ call to inquire about the attributes of a queue. This extract is not taken from the sample applications supplied with MQSeries.

```
    %INCLUDE SYSLIB(CMQP);
    %INCLUDE SYSLIB(CMQEPP);
    :
/**********************************************************/
/* WORKING STORAGE DECLARATIONS                         */
/**********************************************************/
    DCL COMPCODE              BINARY FIXED (31);
    DCL REASON                BINARY FIXED (31);
    DCL HCONN                 BINARY FIXED (31);
    DCL HOBJ                  BINARY FIXED (31);
    DCL OPTIONS               BINARY FIXED (31);
    DCL SELECTORCOUNT         BINARY FIXED (31);
    DCL INTATTRCOUNT          BINARY FIXED (31);
    DCL 1 SELECTOR_TABLE,
        3 SELECTORS(5)        BINARY FIXED (31);
    DCL 1 INTATTR_TABLE,
        3 INTATTRS(5)         BINARY FIXED (31);
    DCL CHARATTRLENGTH        BINARY FIXED (31);
    DCL CHARATTRS             CHAR(100);

  ⋮

/**********************************************************/
/* SET VARIABLES FOR INQUIRE CALL                       */
/* INQUIRE ON THE CURRENT QUEUE DEPTH                   */
/**********************************************************/

    SELECTORS(01) = MQIA_CURRENT_Q_DEPTH;

    SELECTORCOUNT  = 1;
    INTATTRCOUNT   = 1;

    CHARATTRLENGTH = 0;
/**********************************************************/
/*                                                      */
/* HCONN WAS SET BY PREVIOUS MQCONN REQUEST.            */
/* HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST.             */
/*                                                      */
/**********************************************************/
    CALL MQINQ (HCONN,
                HOBJ,
                SELECTORCOUNT,
                SELECTORS,
                INTATTRCOUNT,
                INTATTRS,
                CHARATTRLENGTH,
                CHARATTRS,
                COMPCODE,
                REASON);
```

Figure 95 (Part 1 of 2). Using the MQINQ call (PL/I)

```
/******************************************************/
/* TEST THE COMPLETION CODE OF THE INQUIRE CALL.      */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE SHOWING */
/* THE COMPLETION CODE AND THE REASON CODE.           */
/******************************************************/
     IF COMPCODE ¬= MQCC_OK
        THEN DO;

   :

          CALL ERROR_ROUTINE;
        END;
```
Figure 95 (Part 2 of 2). Using the MQINQ call (PL/I)

Setting the attributes of a queue

Figure 96 demonstrates how to use the MQSET
call to change the attributes of a queue. This
extract is not taken from the sample applications
supplied with MQSeries.

```
     %INCLUDE SYSLIB(CMQP);
     %INCLUDE SYSLIB(CMQEPP);
     :
/******************************************************/
/* WORKING STORAGE DECLARATIONS                       */
/******************************************************/
     DCL COMPCODE              BINARY FIXED (31);
     DCL REASON                BINARY FIXED (31);
     DCL HCONN                 BINARY FIXED (31);
     DCL HOBJ                  BINARY FIXED (31);
     DCL OPTIONS               BINARY FIXED (31);
     DCL SELECTORCOUNT         BINARY FIXED (31);
     DCL INTATTRCOUNT          BINARY FIXED (31);
     DCL 1 SELECTOR_TABLE,
         3 SELECTORS(5)        BINARY FIXED (31);
     DCL 1 INTATTR_TABLE,
         3 INTATTRS(5)         BINARY FIXED (31);
     DCL CHARATTRLENGTH        BINARY FIXED (31);
     DCL CHARATTRS             CHAR(100);

   :


/******************************************************/
/* SET VARIABLES FOR SET CALL                         */
/* SET GET AND PUT INHIBITED                          */
/******************************************************/

     SELECTORS(01) = MQIA_INHIBIT_GET;
     SELECTORS(02) = MQIA_INHIBIT_PUT;

     INTATTRS(01) = MQQA_GET_INHIBITED;
     INTATTRS(02) = MQQA_PUT_INHIBITED;

     SELECTORCOUNT  = 2;
     INTATTRCOUNT   = 2;

     CHARATTRLENGTH = 0;
```
Figure 96 (Part 1 of 2). Using the MQSET call (PL/I)

```
/******************************************************/
/*                                                    */
/* HCONN WAS SET BY PREVIOUS MQCONN REQUEST.          */
/* HOBJ WAS SET BY PREVIOUS MQOPEN REQUEST.           */
/*                                                    */
/******************************************************/
     CALL MQSET (HCONN,
                 HOBJ,
                 SELECTORCOUNT,
                 SELECTORS,
                 INTATTRCOUNT,
                 INTATTRS,
                 CHARATTRLENGTH,
                 CHARATTRS,
                 COMPCODE,
                 REASON);

/******************************************************/
/* TEST THE COMPLETION CODE OF THE SET CALL.          */
/* IF THE CALL HAS FAILED ISSUE AN ERROR MESSAGE SHOWING */
/* THE COMPLETION CODE AND THE REASON CODE.           */
/******************************************************/
     IF COMPCODE ¬= MQCC_OK
        THEN DO;

   :

          CALL ERROR_ROUTINE;
        END;
```
Figure 96 (Part 2 of 2). Using the MQSET call (PL/I)

Appendix G. MQSeries data definition files

MQSeries provides data definition files to assist you with the writing of your applications. Data definition files are also known as:

Language	Data definitions
C	Include files or header files
Visual Basic	Include files or header files
COBOL	Copy files
Assembler	Macros
PL/I	Include files
TAL	Header files

See Appendix A, "Language compilers and assemblers" on page 433 for the compilers that are supported and suitable for use with these data definition files.

The data definition files to assist with the writing of channel exits are described in the MQSeries Intercommunication book.

The data definition files to assist with the writing of installable services exits are described in the *MQSeries Programmable System Management* book.

For data definition files supported on C++, see the *MQSeries Using C++* book.

For data definition files supported on RPG, see the *MQSeries for AS/400 Application Programming Reference (RPG)* book.

The names of the data definition files have the prefix CMQ, and a suffix that is determined by the programming language:

Suffix	Language
a	Assembler language
b	Visual Basic
c	C
l	COBOL (without initialized values)
p	PL/I
v	COBOL (with default values set)

Notes:

1. MQSeries for MVS/ESA supports all the above programming languages.
2. MQSeries for AT&T GIS UNIX supports C only.
3. MQSeries for Windows supports C and Visual Basic.
4. Other platforms support C and COBOL.
5. MQSeries for AIX, MQSeries for MVS/ESA, MQSeries for OS/2 Warp, and MQSeries for Windows NT support PL/I.
6. This does not apply to MQSeries for Digital OpenVMS.

509

> ┌─ **Installation directory and library** ─────────────────┐
>
> Throughout this appendix, the name **mqmtop** is used to represent the name of
> the installation directory on UNIX systems. On Digital OpenVMS, the actual
> directory is /mqm; on SunOS, the actual directory is /usr/mqm; on other UNIX
> platforms, the actual directory is /opt/mqm. The name **thlqual** is the high-level
> qualifier of the installation library on MVS/ESA.

C language include files

The MQSeries C include files are listed in Table 45. They are installed in the
following directories:

Platform	Installation directory or library
AIX	/usr/lpp/mqm/inc/
Other UNIX platforms	/**mqmtop**/inc/
OS/2 and Windows NT	\mqm\tools\c\include
Windows V2.0	\MQW\INCLUDE
Windows V2.1	\Program Files\MQSeries for Windows\Lib
MVS/ESA	**thlqual**.SCSQC370
Tandem NSK	$volume.zmqslib

Notes:

1. For UNIX platforms (not including Digital OpenVMS), the include files are
   symbolically linked into /usr/include.

2. For AS/400, include files are supplied as members of H in library QMQM.

For more information on the structure of directories, see the *MQSeries System
Administration* for MQSeries for AIX, HP-UX, OS/2, Sun Solaris, and Windows NT;
for other platforms, see the appropriate *System Management Guide*.

Table 45. C include files for MQSeries	
File name	**Contents**
<cmqc.h>	Call prototypes, data types, structures, return codes, and constants
<cmqcfc.h>(1)	Definitions for programmable commands
<cmqxc.h>	Definitions for channel exits and data-conversion exits
<cmqzc.h> (1, 2)	Definitions for installable services exits

Notes: The files are protected against multiple declaration, so you can include
them many times.

1. MQSeries for MVS/ESA and MQSeries for Windows do not provide this
   include file.
2. Not provided on MQSeries for AS/400.
3. On Tandem NSK filenames cannot contain a period (.) so the header
   filenames are <cmqch> and so on.

Visual Basic header files

MQSeries for Windows Version 2.0 provides two Visual Basic header files. They are listed in Table 46 and installed in \MQW\INCLUDE.

Table 46. Visual Basic header files for MQSeries for Windows V2.0	
File name	**Contents**
CMQB3.BAS	Call declarations, data types, and named constants for the 16-bit MQI.(1)
CMQB4.BAS	Call declarations, data types, and named constants for both the 16-bit and 32-bit MQI.(2)
Notes:	

1. Use this with Microsoft Visual Basic Version 3.
2. Use this with Microsoft Visual Basic Version 4.

MQSeries for Windows Version 2.1 provides two Visual Basic header files. They are listed in Table 47 and installed in \Program Files\MQSeries for Windows\Lib.

Table 47. Visual Basic header files for MQSeries for Windows V2.1	
File name	**Contents**
CMQB.BAS	Call declarations, data types, and named constants for the main MQI.
CMQB4.BAS	Call declarations, data types, and named constants for the channel exits.

Note: In a default installation, the form files (.BAS) are supplied in the \Program Files\MQSeries for Windows\Include subdirectory.

COBOL copy files

For COBOL, MQSeries provides separate copy files containing the named constants, and two copy files for each of the structures. There are two copy files for each structure because each is provided both with and without initial values:

- In the WORKING-STORAGE SECTION of a COBOL program, use the files that initialize the structure fields to default values. These structures are defined in the copy files that have names suffixed with the letter "V" (values).

- In the LINKAGE SECTION of a COBOL program, use the structures without initial values. These structures are defined in copy files that have names suffixed with the letter "L" (linkage).

The MQSeries COBOL copy files are listed in Table 48 on page 513. They are installed in the following directories:

Platform	Installation directory or library
AIX	/usr/lpp/mqm/inc/
Other UNIX platforms	/**mqmtop**/inc/
OS/2 and Windows NT	\mqm\tools\cobol\copybook (for Micro Focus COBOL) \mqm\tools\cobol\copybook\VAcobol (for IBM VisualAge for COBOL)
MVS/ESA	**thlqual**.SCSQCOBC

Notes: For AS/400, they are supplied in the library QMQM:

For OPM, they are supplied as members of the file QLBLSRC.

For ILE, they are supplied as members of the file QCBLLESRC.

For Tandem NSK, all the sections are contained in one ENSCRIBE file CMCPCOBOL.

Table 48. COBOL copy files		
File name (with initial values)	**File name (without initial values)**	**Contents**
CMQBOV (not AS/400)	CMQBOL (not AS/400)	Begin options structure (MQBO)
CMQCNOV	CMQCNOL	Connect options structure (MQCNO)
CMQDHV	CMQDHL	Distribution header structure (MQDH)
CMQDLHV	CMQDLHL	Dead-letter (undelivered-message) header structure (MQDLH)
CMQDXPV	CMQDXPL	Data-conversion exit parameter structure (MQDXP)
CMQGMOV	CMQGMOL	Get-message options structure (MQGMO)
CMQIIHV	CMQIIHL	IMS header structure (MQIIH)
CMQMDV	CMQMDL	Message descriptor structure (MQMD)
CMQMDEV	CMQMDEL	Message descriptor extension structure (MQMDE)
CMQODV	CMQODL	Object descriptor structure (MQOD)
CMQORV	CMQORL	Object record structure (MQOR)
CMQPMOV	CMQPMOL	Put-message options structure (MQPMO)
CMQRRV	CMQRRL	Response record structure (MQRR)
CMQTMCV	CMQTMCL	Trigger-message structure (character format)
CMQTMC2V	CMQTMC2L	Trigger-message structure (character format) (MQTMC)
CMQTMV	CMQTML	Trigger-message structure (MQTM)
CMQV	not applicable	Named constants for the MQI
CMQXV	not applicable	Named constants for exits
CMQXQHV	CMQXQHL	Transmission-queue header structure (MQXQH)

Include in your program only those files you need. Do this with one or more COPY statements after a level-01 declaration. This means you can include multiple versions of the structures in a program if necessary. However, note that CMQV is a large file.

COBOL copy files

Here is an example of COBOL code for including the CMQMDV copy file:

```
01 MQM-MESSAGE-DESCRIPTOR.
   COPY CMQMDV.
```

Each structure declaration begins with a level-10 item; this means you can declare several instances of the structure by coding the level-01 declaration followed by a COPY statement to copy in the remainder of the structure declaration. To refer to the appropriate instance, use the IN keyword.

Here is an example of COBOL code for including two instances of CMQMDV:

```
* Declare two instances of MQMD
 01  MY-CMQMD.
     COPY CMQMDV.
 01  MY-OTHER-CMQMD.
     COPY CMQMDV.
*
* Set MSGTYPE field in MY-OTHER-CMQMD
     MOVE MQMT-REQUEST TO MQMD-MSGTYPE IN MY-OTHER-CMQMD.
```

The structures should be aligned on 4-byte boundaries. If you use the COPY statement to include a structure following an item that is not the level-01 item, try to ensure that the structure is a multiple of 4-bytes from the start of the level-01 item. If you do not do this, you may get a reduction in the performance of your application.

The structures are described in the *Application Programming Reference* manual. The descriptions of the field in the structures show the names of fields without a prefix. In COBOL programs you must prefix the field names with the name of the structure followed by a hyphen, as shown in the COBOL declarations. The fields in the structure copy files are prefixed this way.

The field names in the declarations in the structure copy files are in uppercase. You can use mixed case or lowercase instead. For example, the field *StrucId* of the MQGMO structure is shown as MQGMO-STRUCID in the COBOL declaration and in the copy file.

The V-suffix structures are declared with initial values for all of the fields, so you need to set only those fields where the value required is different from the initial value.

System/390 assembler-language macros

MQSeries for MVS/ESA provides two assembler-language macros containing the named constants, and one macro to generate each structure. They are listed in Table 49 and installed in **thlqual**.SCSQMACS.

Table 49. System/390 assembler-language macros	
Macro	**Contents**
CMQA	Values of the return codes for the API calls Constants for filling in the option fields Constants for each object attribute, used by the MQINQ and MQSET calls
CMQDLHA	Definition of the MQDLH structure
CMQDXPA	Definition of the MQDXP structure
CMQGMOA	Definition of the MQGMO structure
CMQIIHA	Definition of the MQIIH structure
CMQMDA	Definition of the MQMD structure
CMQODA	Definition of the MQOD structure
CMQPMOA	Definition of the MQPMO structure
CMQTMA	Definition of the MQTM structure
CMQTMC2A	Definition of the MQTMC2 structure
CMQXA	Constants for exits
CMQXPA	Definition of the MQXP structure
CMQXQHA	Definition of the MQXQH structure

These macros are called using code like this:

```
MY_MQMD   CMQMDA EXPIRY=0,MSGTYPE=MQMT_DATAGRAM
```

PL/I include files

MQSeries for MVS/ESA, AIX, OS/2 Warp, and Windows NT provide include files that contain all the definitions you need when you write MQSeries applications in PL/I. They are listed in Table 50 on page 516. They are installed in the following directories:

Platform   **Installation directory or library**
AIX   /usr/lpp/mqm/inc/
OS/2 and Windows NT   \mqm\tools\pli\include
MVS/ESA   **thlqual**.SCSQPLIC

PL/I include files

Table 50. PL/I include files

Include file	Contents
CMQCFP	Definitions for programmable commands
CMQEPP	Entry point definitions for the API calls.
CMQP	Definitions of all the constants and return codes, data types and structures, and constants to initialize the structures.
CMQXP	Definitions for channel exits and data-conversion exits on MVS/ESA. Named constants related to PCF on AIX, OS/2 Warp, and Windows NT.

Include these files in your program if you are going to link the MQ stub to your program (see "Preparing your program to run" on page 267). Include only CMQP if you intend to link the MQ calls dynamically (see "Dynamically calling the MQSeries stub" on page 270). Dynamic linking can be performed for batch and IMS programs only.

Appendix H. List of MQSeries products

At the time of publication of this document, MQSeries products were available for the hardware and software platforms listed below.

Some of the products in the list were renamed in November 1994. For those products, the name by which the product used to be known is given, as well as the current product name.

These products fall into two broad categories, Level 1 and Level 2, in respect of the function they support. These are referred to in Appendix I, "MQSeries platforms - functional comparisons" on page 519.

Level 1 products

Table 51. MQSeries products, Level 1	
Product name	**Previous name**
MQSeries for SCO UNIX V1.4	ezBRIDGE Transact on SCO UNIX for MQSeries Release 3.0
MQSeries for UnixWare V1.4.1	ezBRIDGE Transact on UnixWare for MQSeries Release 3.0
MQSeries for VSE/ESA V1.4	ezBRIDGE Transact on VSE/ESA for MQSeries Release 3.0

Level 2 products

Table 52 (Page 1 of 2). MQSeries products, Level 2	
Product name	**Previous name**
MQSeries for AIX V5.0 (clients and servers)	None (new version)
MQSeries for AT&T GIS UNIX V2.2 (clients and servers)	None (new release)
MQSeries for Digital OpenVMS V2.2 (clients and servers)	None (new release)
MQSeries for HP-UX V5.0 (clients and servers)	None (new version)
MQSeries for MVS/ESA V1.2	None (new level)
MQSeries for OS/2 Warp V5.0 (clients and servers)	None (new version)
MQSeries for AS/400 V4R2	None (new release)
MQSeries for SINIX and DC/OSx V2.2 (clients and servers)	None (new product)
MQSeries for SunOS V2.2 (clients and servers)	None (new product)
MQSeries for Sun Solaris V5.0 (clients and servers)	None (new version)
MQSeries for Tandem NonStop Kernel V2.2	None (new release)
MQSeries Three Tier for AIX	None (new product)
MQSeries Three Tier for OS/2	None (new product)
MQSeries for Windows V2.0	None (new product)

MQSeries products

Table 52 (Page 2 of 2). MQSeries products, Level 2	
Product name	**Previous name**
MQSeries for Windows V2.1	None (new release)
MQSeries for Windows NT V5.0 (clients and servers)	None (new version)

Appendix I. MQSeries platforms - functional comparisons

This describes the functional differences between MQSeries queue manager products, concentrating on the functionality relating to the Message Queue Interface (MQI).

You may need this information for the following:

- Program migration when the queue-manager is upgraded

  Programs written to work with one MQSeries queue manager will migrate unchanged, with only a relink, when the queue manager is upgraded to one which offers more function, operating on the same platform. You may, alternatively, want to plan for additional functionality.

- Program migration from one platform to another

  You may need to write programs in such a way that they are readily ported to a different platform on which there is another MQSeries queue manager.

- Personnel migration

  You may need to become familiar with the functionality of more than one MQSeries queue manager.

- Message movement

  Most of the functional differences are not relevant to the progress of messages from one platform (or queue manager) to another. However a few are relevant, for example differences in the maximum message size that is supported. See "Message properties and protocols" on page 530 for differences which are relevant in this area.

What is included here

Enough information about the functions is given here to allow you to understand the differences between the products. For detailed information, see the *MQSeries Application Programming Reference* manual.

The categories used here are:

- Syncpoint (see page 521)
- Triggering (see page 525)
- Input (see page 527)
- Message properties (see page 530)
- Threading (see page 534)
- Security (see page 535)
- Object types (see page 539)
- Administration (see page 541)
- Programming languages (see page 543)
- Connection (see page 544)

Within each category, a summary table of differences is given, and a brief explanation of the features involved. Differences between support by the various products are discussed, and recommendations made for ways in which migration can be made easier.

Functional comparison

The various products are represented as follows:

A   MQSeries for Digital OpenVMS V2.2

C   MQSeries for MVS/ESA V1.2

D   MQSeries for OS/2 Warp V5.0

E   MQSeries for AS/400 V4R2

F   MQSeries for SCO UNIX V1.4

G   MQSeries for Tandem NonStop Kernel V2.2

H   MQSeries for UnixWare V1.4.1

I   MQSeries for VSE/ESA V1.4

J   MQSeries for Windows V2.0

K   MQSeries for Windows V2.1

L   MQSeries for Windows NT V5.0

M   MQSeries V2 for UNIX systems. This includes MQSeries for AT&T GIS UNIX V2.2, MQSeries for SINIX and DC/OSx V2.2, and MQSeries for SunOS V2.2.

N   MQSeries V5 for UNIX systems. This includes MQSeries for AIX V5.0, MQSeries for HP-UX V5.0, and MQSeries for Sun Solaris V5.0.

Syncpoint

Table 53. Syncpoint support - function summary

Product	Syncpoint available	Coordination with other RMs	Coordination of other RMs	No syncpoint options (1)	Syncpoint options (2)	Default	Browse with lock	Backout count	Mark skip backout
A	Yes	No	No	Accepted (checked)	Accepted (checked)	No syncpoint	Yes	Yes	No
C	Yes	CICS or IMS resources	No	Accepted (checked)	Accepted (checked)	Syncpoint	No	Yes	Yes
D	Yes	CICS or other resources	Yes	Accepted (checked)	Accepted (checked)	No syncpoint	Yes	Yes	No
E	Yes	CICS and OS/400 resources	No	Accepted (checked)	Accepted (checked)	No syncpoint	Yes	Yes	No
F	No	No	No	Accepted	Rejected	No syncpoint	Yes	No	No
G	Yes	TMF resources	No	Rejected	Accepted	Syncpoint	No	No	No
H	No	No	No	Accepted	Rejected	No syncpoint	Yes	No	No
I	Yes	CICS resources	No	Rejected	Accepted	Syncpoint	Yes	No	No
J	Yes	No	No	Accepted	Accepted	No syncpoint	No	Yes	No
K	Yes	No	No	Accepted	Accepted	No syncpoint	No	Yes	No
L	Yes	CICS or other resources	Yes	Accepted (checked)	Accepted (checked)	No syncpoint	Yes	Yes	No
M	Yes	CICS or other resources	No	Accepted (checked)	Accepted (checked)	No syncpoint	Yes	Yes	No
N	Yes	CICS or other resources	Yes	Accepted (checked)	Accepted (checked)	No syncpoint	Yes	Yes	No

Note:

1. MQGMO_NO_SYNCPOINT for MQGET, and MQPMO_NO_SYNCPOINT for MQPUT and MQPUT1
2. MQGMO_SYNCPOINT for MQGET, and MQPMO_SYNCPOINT for MQPUT and MQPUT1
3. Not required since shared input is not supported

Syncpoint coordination for MQSeries allows coordination of the effect of a number of get or put, or both, actions issued from the same application within a single unit of work, so that either all of these actions take place automatically, or none of them does. Usually other actions, such as file or database updates, can also be coordinated within the same unit of work.

A resource manager that participates in syncpoint is a component which owns a particular resource, such as a file system. The resource manager controls updates to that resource so that updates made from a particular unit of work are not visible from outside that unit of work until a syncpoint is reached. At that time, depending

upon whether the syncpoint action is to commit or backout, all of the updates made from within the unit of work are either made permanent, and visible to other units of work, or backed out, so that other units of work never see them. If the updates are made, they are made automatically. This will happen even if the application or the system fails at any point. In order to carry out this action, the resource manager must log changes, and lock uncommitted changes to prevent other units of work from seeing them.

If the updates from more than one resource manager are to be coordinated, then the services of a syncpoint coordinator are required. If the updates from all of the resource managers are to take place automatically, then two-phase commit is needed; otherwise, or if only one resource manager is involved, single-phase commit is all that is needed. On MQSeries Version 5 products products, MQSeries can act as a syncpoint coordinator coordinating other resource managers.

Note the variation between platforms in the calls to declare the syncpoint, and indicate the application's decision to commit or backout. Also, in some cases, there is a variation between platforms in calls to declare the start of a new unit of work. For MQSeries for MVS/ESA batch, MQSeries for OS/2 Warp V5.0, MQSeries Version 2 for UNIX systems, and MQSeries for Windows NT applications, these calls are provided through the MQI, otherwise they are provided by the environment. MQSeries for OS/2 Warp V5.0, MQSeries Version 2 for UNIX systems and MQSeries for Windows NT support both internally-coordinated units of work (using the MQI) and externally-coordinated units of work.

In some cases these calls need not be issued at all by the application. For example, if an application terminates abnormally before committing its unit of work, the unit of work will be backed out automatically. Similarly, if a CICS transaction terminates normally, it is as if it had committed the unit of work.

MQSeries for Tandem Guardian V1 supports a verb MQSYNC with a parameter to commit or roll back. This does not actually perform the function but rather coordinates resident statistics with commits and roll backs performed by the application. The actual syncpoint support is provided by the system environment. For example, an application will start a unit of work, perform all operations, complete the unit with a commit, issue MQSYNC/commit to update counters to reflect the committed work.

Level 1 product range support

The products in this range never act as a resource manager. An underlying file or database system is always used to hold the message queues. The resource manager is the owner of the appropriate file or database system. This means that the underlying system controls whether or not the MQI gets and puts participate in a syncpoint operation. In some cases the underlying system is never coordinated, so in these cases MQI gets and puts are never coordinated.

In other cases, whether the underlying system is coordinated depends upon how the user has set it up. The setup can specify that actions on some queues are always coordinated, while those on others are never coordinated.

A product in this subset will not be informed of this, and therefore does not have any record of whether the MQI request will be coordinated. If it is coordinated, however, it will be fully coordinated with any other updates to the underlying system, since only a single resource manager is involved. Additionally, if this

resource manager acts as a participant with other resource managers, under the control of a syncpoint coordinator, MQI requests will also be coordinated with updates to the resources owned by these other resource managers. Where an MQI call might or might not be coordinated depending upon the setup, the queue manager accepts both syncpoint and no-syncpoint options, without being able to validate that the option will be honored.

Level 1–browse-with-lock support

Where there is no syncpoint support, an application can be written in order to avoid the possibility of losing incoming messages if the system or application crashes:

```
Do forever
   MQGET with MQGMO_BROWSE_FIRST
   Process message, and make updates as necessary
   MQGET with MQGMO_MSG_UNDER_CURSOR
End
```

If there is a failure before the updates are complete, the message is still available on the input queue after restart. The second MQGET deletes the message from the queue.

However, if more than one application can process the queue, this technique is not adequate, since several applications might browse the message at the top of the queue and begin making updates on its behalf.

To overcome this, a browse-with-lock feature is supported on platforms that don't have any syncpoint capability but do support shared input. This enables an application to browse the first message on the queue, locking it at the same time so that it is not available to other applications. The application is written in a similar way, except that the MQGMO_LOCK option is added for the initial browse:

```
Do forever
   MQGET with MQGMO_BROWSE_FIRST and MQGMO_LOCK
   Process message, and make updates as necessary
   MQGET with MQGMO_MSG_UNDER_CURSOR
End
```

Level 2 product range support

In Level 2 products, management of the message queues is not delegated to an underlying file system. These products always act as a resource manager, and include full coordination with the other resource managers, except for:

- MVS/ESA batch applications, and
- MQSeries Version 2 for UNIX systems, MQSeries for OS/2 V2.0.1, and MQSeries for Windows NT V2.0 if there is no external coordinator available,

where syncpoint facilities only apply to MQI requests.

The full MQI support enables applications to specify, on a particular get or put request, whether or not this request is to be coordinated within a unit of work. If they do not specify this, a default is taken.

The full MQI also provides a *BackoutCount*, which informs the getter of a message how many times this message has been retrieved and subsequently backed out. The count may be inaccurate, unless the *HardenGetBackout* local queue attribute for the queue specifies that it should be hardened. This does carry a resulting

performance penalty. This count can be used as an indication that a message might be causing the application to terminate abnormally and that the application should not attempt to process it again. Of course, the message might have been backed out for reasons unconnected with the message itself - for example, a system crash.

On the MQSeries for MVS/ESA queue manager only, the option MQGMO_MARK_SKIP_BACKOUT can be issued with an MQGET call. This is an advanced option, which can be used if the application might make some updates as a result of an input message, but might then discover that these updates should be backed out without reinstating the input message on the queue (since it would only give the same result when it was processed the next time).

Migration recommendations

1. If your application is not sensitive to whether the gets and puts are within syncpoint scope, do not specify the syncpoint or no_syncpoint options with MQGET, MQPUT, and MQPUT1.

2. If your application is designed to run without syncpoint, always specify the no_syncpoint options (MQGMO_NO_SYNCPOINT and MQPMO_NO_SYNCPOINT) with MQGET, MQPUT, and MQPUT1. With queue managers that accept these options but do not validate them, ensure that the underlying file or database system is set up so that operations on the relevant objects do not participate in syncpoint.

   If you want to port this application to a platform that always has syncpoint, the application will have to be modified.

3. If your application is designed to run with syncpoint, always specify the syncpoint options (MQGMO_SYNCPOINT and MQPMO_SYNCPOINT) with MQGET, MQPUT, and and MQPUT1. With queue managers that accept these options but do not validate them, ensure that the underlying file or database system is set up so that operations on the relevant objects do participate in syncpoint.

   If you want to port this application to a platform that does not support syncpoint, the application will have to be modified. Use the browse technique outlined above, with lock support if multiple applications can service the queue.

   If you have already written the application in this way, and now wish to port it to a platform that does support syncpoint, it will have to be changed to use the syncpoint support rather than the browse-with-lock technique. However, MQSeries Version 2 for UNIX systems and MQSeries for OS/2 Warp V5.0 do support the lock option to allow you to move your application without making the changes. If you are writing a new application on either of these platforms you are recommended to use syncpoint support.

4. Try to avoid MQGMO_MARK_SKIP_BACKOUT if portability is important. If the database is designed so that the application can establish, before making any updates, that the transaction can be processed successfully, this option will not be necessary.

5. Avoid using the *BackoutCount* if portability beyond the Level 2 queue managers is important.

Triggering

Table 54. Triggering support - function summary

Product	Triggering	Depth, priority triggering	Type of applications started	Passed to started application (1)
A	Yes	Yes	VMS system applications	MQTMC2 structure
C	Yes	Yes	CICS or IMS (2)	MQTM structure for CICS; MQTMC2 structure for IMS
D	Yes	Yes	OS/2 applications (including CICS)	MQTMC2 structure
E	Yes	Yes	OS/400 applications (including CICS)	MQTMC2 structure
F	No	No	-	-
G	Yes	Yes	Tandem NSK applications	MQTMC2 structure
H	No	No	-	-
I	Yes	No	CICS only	MQTM structure
J	No	No	-	-
K	No	No	-	-
L	Yes	Yes	Windows NT applications (including CICS)	MQTMC2 structure
M	Yes	Yes	UNIX systems applications (including CICS)	MQTMC2 structure
N	Yes	Yes	UNIX systems applications (including CICS)	MQTMC2 structure

Note:

1. MQTMC is a version of the trigger message (MQTM) structure, in which the two binary fields in the structure are replaced with character fields of the same length. MQTMC2 is a version of the trigger message structure, in which the two binary fields in the structure are replaced with character fields of the same length and to which the queue manager name has been added.

2. There is currently no automatic support supplied for starting MVS/ESA batch applications. However, this could be added by a user.

If triggering is supported, a trigger message is written to the initiation queue when the trigger conditions for the queue are satisfied. Various conditions can be specified as described in "Conditions for a trigger event" on page 204. Briefly, the conditions include the following:

- When any message arrives on the queue (*every* triggering)

- When the queue goes from being empty to having messages on it (*first* triggering)

- When a specified number of messages have arrived on the queue (*depth* triggering)

With any of these conditions, a minimum priority can also be specified. Only messages with this or a higher priority are included when evaluating whether the trigger conditions have been met (*priority* triggering).

There are also a number of unusual conditions which cover cases in which a restart or other failure has occurred.

When the application is started, it is passed a structure which contains the name of the queue which caused the trigger event, and from which the application is now expected to retrieve messages. Other data in this structure includes information which the administrator can add to the process definition, which can be used to control the operation of the application.

The MQTM, MQTMC, and MQTMC2 structures are very similar; the difference between them is that some binary fields in the MQTM structure are replaced by character equivalents (of the same length) in the MQTMC and MQTMC2 structure. (MQTMC2 also contains the name of the queue manager.) If these fields are not referenced, portability is increased. However, one of the fields is the structure version, and it is safer to always check that, at the cost of a possible small change to the application in an area which would probably need change anyway.

Level 1 product range support

Only MQSeries for VSE/ESA supports triggering. In this case there is no actual initiation queue; the application is started directly by the queue manager. User information in the trigger message is not supported for MQSeries for VSE/ESA.

On MQSeries for VSE/ESA two forms of triggering are supported:

- In one form, the started application must issue a CICS RETRIEVE call in order to get the trigger message. This form is compatible with triggering for MVS/ESA.

- In the other form, the trigger message is passed in COMMAREA.

Depth and priority triggering are not supported by this subset of the product range.

Level 2 product range support

These products provide full support. There is a real initiation queue, and the trigger monitor which reads trigger messages from this queue and starts applications, could be replaced by a customer-written version.

The trigger monitor on MQSeries Version 2 for UNIX systems, MQSeries for OS/2 Warp V5.0, and MQSeries for Windows NT V5.0 can run either on a server or a client, with the exception of clients in DOS and Windows 3.1 environments.

Migration recommendations

1. At present, triggering is not available on several platforms. Alternative techniques must be used, for example on some platforms a long-running application which does a get with signal may be used. This may migrate better between the platforms of interest. See "Input" on page 527.

2. If migration between CICS/VSE and CICS/MVS or CICS/ESA is important, use the form of triggering on VSE/ESA which is compatible with the MVS/ESA form.

3. Note that Level 1 products do not include any user information in the trigger message.

Input

Table 55. Input features - function summary

Product	Shared input	Exclusive input	Input default to queue definition	Browse with lock	Full browse (1)	Signal	Selective get by MsgId / CorrelId (2)	Data conversion
A	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
C	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
D	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
E	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
F	Yes	Yes	No	Yes	No (3)	No	No	No
G	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes
H	Yes	Yes	No	Yes	No (3)	No	No	No
I	Yes	Yes	No	Yes	Yes	No	No	No
J	Yes	Yes	Yes	No	Yes	No	Yes	No
K	Yes	Yes	Yes	No	Yes	Yes	Yes	No
L	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
M	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
N	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes

Note:

1. Other than Browse-with-lock

2. All products support generating a *MsgId* if zero is specified at put time, conveying both *MsgId* and *CorrelId* with the message, and delivering them to the getter. Not all products support selective get, however, which is the ability of the getter to specify the *MsgId* or *CorrelId* or both that the retrieved message must have.

3. Only the first message on a queue can be browsed.

Shared and exclusive input

These options (MQOO_INPUT_EXCLUSIVE and MQOO_INPUT_SHARED) control whether more than one application is allowed to have the queue open for input at the same time.

If the application is not sensitive to this, MQOO_INPUT_AS_Q_DEF can be specified, and exclusivity is determined by the queue definition. This feature is available only with queue managers within Level 2 products.

On MQSeries for Digital VMS VAX V1.5 only exclusive input is possible. The effect of this is that each server must have its own queue. If on this platform it is necessary to have messages on a single queue that are to be processed by different applications (that are active at the same time), then a server must be written which reads the messages on the single queue and passes each one to the relevant application's queue.

Another feature available only with queue managers within the Level 2 products, is a queue attribute for *Shareability*. If this is set to MQQA_NOT_SHAREABLE, only one application is allowed to have the queue open for input at any time (the MQOO_INPUT_SHARED open option for this queue is effectively treated as if it were MQOO_INPUT_EXCLUSIVE).

MQOO_INPUT_EXCLUSIVE is supported on all platforms and is the safest to choose from a portability point of view. However, if an application is not sensitive to exclusivity, and portability to Digital Vax platform is not an issue, MQOO_INPUT_SHARED, is a better choice.

Note that on MQSeries for VSE/ESA, a single MQOPEN cannot request both input and output functions. If input and output for the same queue are required, two MQOPEN calls must be issued, one for each. Other queue managers do not have this restriction, so bear it in mind if your application might need to migrate to MQSeries for VSE/ESA in the future.

Browse-with-lock, and full browse

For more information about browse-with-lock support see "Syncpoint" on page 521. With the products in Level 1, only the first message on the queue can be browsed. Browse-with-lock is supported only by queue managers for which no syncpoint support is available, and where shared input is possible, with the exception of MQSeries for AS/400 V4R2, MQSeries Version 2 for UNIX systems, and MQSeries for OS/2 Warp V5.0 where it is supported to enable you to migrate from Version 1 to Version 2.

Full browse support allows an application to inspect each of the messages on the queue in turn, without retrieving them destructively. A browse cursor marks the position on the queue that the browse scan has reached. Full browse is supported by the queue managers of Level 2 products, and by some queue managers in the Level 1 product range. Others in the Level 1 product range offer a limited form of browse, in which only the first message on the queue can be browsed.

Browse can be used to select a message to get destructively. It can also be used in order to monitor or display the messages on the queue.

Signal

Get using an operating-system signal (option MQGMO_SET_SIGNAL) is used to allow the application to continue with other work, without blocking, while waiting for a suitable message to arrive. The signal is an operating-system dependent object (for example, in MVS it is an ECB). If the operating system supports waiting on multiple signals, as MVS does, then the application can effectively wait for a message arriving on one of several queues.

This feature is only available on platforms where there is a suitable facility provided in the operating environment. If applications that use this feature are ported between platforms, some change will be needed even if signal support is available on both platforms, because the signal itself is specific to the operating environment. If signal is not available, the application logic will have to change; MQGMO_WAIT can be used to provide a blocking wait (MQGMO_WAIT is available with all queue managers).

Get by MsgId and CorrelId

This allows an application to specify a non-zero *MsgId* or *CorrelId* or both when issuing an MQGET (either a browse or a destructive get). If this is done, only messages which have a matching *MsgId* or *CorrelId*, or both, are eligible to be satisfied by the get. Setting either value to zero means that no matching is required for that identifier.

This can be used for retrieving a group of messages, all of which have the same identifier, usually the *CorrelId*, even though they are not all together at the top of the queue. Getting for a specific *MsgId*, and possibly also for a specific *CorrelId* as well, can be used to retrieve a specific message. Note however that there is no guarantee that *MsgId*s are unique within a queue. If a message that is to be retrieved has been found by issuing browse requests, it is much safer to retrieve the message using the MQGMO_MSG_UNDER_CURSOR option than by quoting its *MsgId*.

This facility is provided only by the products in Level 2 queue managers. Level 1 queue managers check that *MsgId* and *CorrelId* have been set to zero for an MQGET, so applications migrating from Level 1 to Level 2 queue managers will need no change in this area. Migrating the other way will require the application to search for the required message itself by browsing, where this is supported. If browse is not supported, more radical changes may be needed.

However, as previously mentioned, all MQSeries products support the putting of a message with a non-zero *MsgId* and *CorrelId*. Therefore, in a client-server type of environment, the client might MQPUT or MQPUT1 *MsgId* and/or *CorrelId* to a server system, such as MQSeries for MVS/ESA, or MQSeries for AS/400, where MQGET by *MsgId* / *CorrelId* is supported. The recomposition of several related messages could then be done at the server or host with the results message then sent back to the client.

Alternatively, each group of messages could request replies to unique reply-to queue names; therefore facilitating *recomposition* on the same or another system, even if the specific MQGET by *CorrelId* function is not available on that system.

Data conversion

See Chapter 11, "Data-conversion exits" on page 157 for a list of conditions under which a data-conversion exit is invoked.

User-written exits can be provided to perform data conversion. A utility is supplied with these products which accepts a C structure, and produces source code from which a data-conversion exit suitable for use by the queue manager can be generated.

Some queue managers do not support data conversion (see Table 55 on page 527). However, when sending messages from a queue manager that does support data conversion to one that doesn't, a channel attribute can be set at the sending queue manager to request that text is converted *before* each message is transmitted.

Message properties and protocols

Table 56. Message properties - function summary

Product	Default maximum message length	Maximum priority	Nonpersistent messages	Exception reports	MsgId, CorrelId response options	Extra message types	COA, COD, Expiry	PAN, NAN	Protocols
A	4MB - 4000	9	Yes	Yes	Yes	Yes	Yes	No	LU 6.2 TCP/IP DECnet
C	4 MB - 4000 (1)	9	Yes	Yes	Yes	Yes	Yes	Yes	LU 6.2 TCP/IP
D	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	Yes	LU 6.2 TCP/IP NetBIOS SPX
E	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	Yes	LU 6.2 TCP/IP
F	64000	0	No	No	No	No	No	No	LU 6.2 TCP/IP
G	4MB - 4000	0	No	No	No	No	No	No	LU 6.2 TCP/IP
H	64000	0	No	No	No	No	No	No	LU 6.2 TCP/IP
I	30000	0	No	No	No	No	No	No	LU 6.2
J	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	No	TCP/IP
K	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	No	TCP/IP
L	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	Yes	LU 6.2 TCP/IP NetBIOS SPX
M	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	No	LU 6.2 TCP/IP
N	4 MB - 4000	9	Yes	Yes	Yes	Yes	Yes	Yes	LU 6.2 TCP/IP

Note:

1. The CICS message mover on the MVS Queue Manager has a maximum segment size of 32K. Because Level 1 queue managers do not support segmentation, the maximum transmitted message length is 32K to a CICS mover.

Message context is another message property. See "Security" on page 535.

Maximum message length

The lowest maximum size of message (excluding headers) supported is 10 000 bytes. You should limit your message text size to this if you want to enable your message to go to any platform.

For Level 2 products, the maximum size is given by the queue manager attribute *MaxMsgLength*, minus MQ_MSG_HEADER_LENGTH. This currently gives 4 MB - 4000 bytes by default. This formula can be used with these queue managers for messages that are not going to travel to, or through, a product in Level 1. If there is any possibility of this, messages should be restricted appropriately.

If longer messages have to be split up for this (or any other) reason, a sequence number or relative starting address should be carried in the message data, so that the receiving application can reconstruct the original text correctly. In some cases it cannot be guaranteed that messages will be delivered in the same order as that in which they were put. An example of this is where multiple hops may be necessary to reach the final destination.

Maximum priority

The queue managers of Level 2 products recognize 10 message priority levels, in the range 0 through 9. A message's priority controls the position at which it is enqueued when it arrives at a queue, and hence the order in which it will be delivered to an application issuing MQGET.

The queue managers of Level 2 products also support a local queue attribute *MsgDeliverySequence*, that allows this to be overridden. If this attribute is set to FIFO only, all messages are enqueued at the same priority. However, the priority value carried with the message still records its intended priority.

Level 1 queue managers treat all messages as having equal priority, so effectively the maximum message priority for them is zero. However, to allow messages of higher priority to pass through, they do not reject an MQPUT of a message where the priority is greater than zero. In this case they issue a warning. So, although the priority does not affect the speed at which the message passes through the system, it can still be honored when the message reaches its destination at a Level 2 queue manager.

However, application program writers must ensure that the priority value is in the range 0 through 9 when a put is issued. Values greater than 9 are rejected by the MQSeries for MVS/ESA queue manager.

Nonpersistent messages

The queue managers of Level 2 products support both persistent and nonpersistent messages. Persistent messages are ones which are guaranteed not to be lost if the system crashes. Messages on them are written to DASD, either as part of the queue backup storage, or as part of the log, so that they can be recovered if necessary. Nonpersistent messages, however, are guaranteed *not* to be recovered after a restart of the queue manager which owns the queue; even if they do happen to be found on DASD, they are discarded.

Persistent messages are vital if the message represents a business action that has to be performed. In order for messages not to be lost, as guaranteed by MQSeries products, they must be marked as persistent. All MQSeries queue managers support persistent messages.

Nonpersistent messages are supported only by Level 2 queue managers. They offer important performance advantages because the message does not have to be written to DASD at all if the capacity of the main-storage buffers permits, and it is processed sufficiently quickly. Nonpersistent messages can be used for queries, where the query will simply be repeated if necessary, and for non-critical notifications (for example, an update to an airport flight information screen).

The queue managers of Level 2 products also provide a third option, which allows the persistence property of the message to be defaulted to the queue, rather than specified by the putting application. In general a given queue can contain a mixture of persistent and nonpersistent messages.

Although the Level 1 queue managers reject an attempt by an application to put a nonpersistent message, they do not prevent a nonpersistent message from passing through (or arriving at) a queue manager. On MQSeries for VSE/ESA the Message Channel Agent changes any nonpersistent message that it receives to be persistent, and marks it as such. So even if the ultimate target is a Level 2 queue manager, the message will have become persistent if it has passed through the queue manager of a Level 1 product. On MQSeries for Digital VMS VAX V1.5, MQSeries for AIX V1, and MQSeries for OS/2 V1.3 the message remains marked as nonpersistent, although it is treated as persistent while at a queue manager of a Level 1 product.

If a program that runs with a queue manager of an Level 2 product specifies nonpersistent when putting a message, the program will need to be changed to run with a queue manager of a Level 1 product.

Exception reports

With Level 2 queue managers, exception reports can be requested when a message is put. In this case, if there is a problem with the message, for example, it cannot be delivered because the destination queue is full, an exception report is generated and sent to the reply-to queue. The report carries a code indicating the nature of the problem.

In addition to requesting exception reports, applications can ask that the first 100 bytes of the original message or all of the original message be included with the report. This is convenient if state information is held at the front of the message data; it helps the originator to identify the message which had a problem.

An exception report is not generated if the putting application can be notified of the problem synchronously, by a reason code at the time of the put.

Level 1 queue managers and transmission programs do not generate exception reports, nor do they allow reports to be requested. They do, however, allow an application to generate a report message. Events that would cause an exception report to be generated are notified instead, for example with a console message, at the queue manager where the problem occurs. There is, however, no direct notification to the application.

MsgId and CorrelId options for responses

With Level 2 queue managers, it is possible to request how the *MsgId* and *CorrelId* of the report should be constructed. It is beneficial if applications honor these requests, to facilitate requests from applications which use different schemes to manage business unit-of-work correlation.

However, these options are not supported for applications connected to Level 1 queue managers. Applications have to choose, based on the correlation scheme in use, how to generate the *MsgId* and *CorrelId* of their replies.

Extra message types

With Level 2 queue managers the *MsgType* field in the message descriptor allows a wide range of values to be accepted. These values are not accepted by queue managers of Level 1 products, and should only be used with messages that do not pass through, or to, applications connected to Level 1 queue managers.

COA, COD, and Expiry message

MQSeries for AS/400 V4R2, MQSeries Version 2 for UNIX systems, MQSeries for OS/2 Warp V5.0, MQSeries for Windows NT V5.0, and MQSeries for MVS/ESA V1.2 allow a program putting a message to request:

- That a confirm-of-arrival report (COA) is generated when the message has been put on its final destination queue.

- That a confirm-of-delivery report (COD) is generated when the message has been retrieved from its final destination queue.

- That the message is discarded if it has not been retrieved within a specified time, and optionally a report is generated if this happens.

These functions are not supported on other MQSeries queue managers.

Level 1 queue managers (and earlier versions of MQSeries for MVS/ESA) do not allow a message requesting any of these options to pass through them.

PAN and NAN messages

MQSeries for AS/400 V4R2, MQSeries for MVS/ESA V1.2, MQSeries for AIX V5.0, MQSeries for Sun Solaris V5.0, MQSeries for HP-UX V5.0, MQSeries for OS/2 Warp V5.0, and MQSeries for Windows NT V5.0, allow a program putting a message to request:

- That a positive action notification (PAN) is generated when the request has been successfully serviced

- That a negative action notification (NAN) is generated when the request has not been successfully serviced

These functions are not supported on other MQSeries queue managers.

Threading

Table 57. Threading - function summary	
Product	**Multi-threaded applications**
J	No
K	Yes
L	Yes
M	(1)
N	Yes
Note:	
1. MQSeries for SINIX supports multi-threaded applications whereas MQSeries for DC/OSx does not.	

Threaded applications on UNIX platforms

On UNIX, an application can connect to a single queue manager only; it **cannot** connect different threads to different queue managers.

A multi-threaded MQSeries application must ensure that MQDISC is executed in each thread. This can be done using a thread exit handler. Do not use detached threads.

Those platforms that support multithreaded applications have some limitations:

AIX, Sun Solaris, and HP-UX

In addition to the considerations for "UNIX signal handling" on page 85, a multi-threaded MQSeries application must not establish signal handlers for SIGSEGV and SIGBUS while any thread has an open connection to MQSeries. These signal handlers should be set up process wide before MQSeries is called. MQSeries' handlers will then attempt to invoke the application's handlers when possible and appropriate.

SINIX

In addition to the considerations for "UNIX signal handling" on page 85, a multi-threaded MQSeries application must not use SIGALRM, SIGSEGV, or SIGBUS.

Security

Table 58. Security - function summary					
Product	**Authorization checks**	**Message context**	**Alternate user ID checks**	**Transmission program exits**	**API crossing exit**
A	Yes	Yes	Yes	Yes	No
C	Yes	Yes	Yes	Yes	Yes (2)
D	Yes	Yes	Yes	Yes	No
E	Yes	Yes	Yes	Yes	No
F	Yes	No (1)	No	No	No
G	Yes	No (1)	No	No	No
H	Yes	No (1)	No	No	No
I	Yes	No (1)	No	No	No
J	No	No	No	No	No
K	No	No	No	No	No
L	Yes	Yes	Yes	Yes	No
M	Yes	Yes	Yes	Yes	No
N	Yes	Yes	Yes	Yes	No
Note:					
1. Context information is accepted and passed on, but not validated. 2. CICS applications only					

Authorization checks

Authorization checks are carried out for each resource, for example a queue, that is accessed. The user identifier under which the application is running is used to validate access to the resource.

Level 1 queue managers provide no direct support for this. However, queues are always implemented as files or database objects, and so when a read or write to the underlying object takes place as a result of a request from a Level 1 product, authority to perform this operation can be validated by the underlying system.

Level 2 queue managers provide more direct support, in that they request security validation for each function and object for which the application attempts an open. However, Level 2 queue managers do not themselves maintain the security tables which specify the operations each user is authorized for, although they may keep temporary caches of this information. In MVS/ESA these tables are maintained by an external security manager such as, RACF, and in OS/400 by the native OS/400 support. MQSeries Version 2 on UNIX systems and MQSeries for Windows NT provide an Object Authority Manager (OAM) that controls access to queue manager resources by user and group IDs. MQSeries for OS/2 Warp V5.0 provides an interface that allows users to provide controlled access to queue manager resources by user and group IDs.

At this level, there is no impact on the MQI statements in a program, and so there is no direct effect on program migration. However, users may need to review the security capabilities on each platform to determine how they will map.

Message context

Level 2 queue managers support various fields that are transmitted with each message, collectively referred to as the context information for the message. This information includes the user identifier and accounting token of the user who originated the message, and the name of the program that generated it. There are also fields which can be set by suitably-authorized applications or exits (see "Channel exits" on page 537) to contain information such as whether the other context information has been validated at this queue manager.

Progressively higher levels of authority are required to alter these fields. At the lowest level of authority, an application is unable to set any information in them; the queue manager sets them all to default values.

Within the Level 2 products the *UserIdentifier* context field holds a user identifier that can be used by an application for validating authority to open resources.

Level 1 queue managers do not support context fields directly. That is, they do not set information in them on behalf of an application, nor do they check that an application is suitably authorized to set them. Whatever information an application sets in these fields when putting a message is carried with the message, and made available to the recipient of the message. A message passing through one of these queue managers will therefore retain any context information that it contains.

Applications connected to Level 1 queue managers can set user identifier and other context information, and can carry out their own checks when receiving messages, bearing in mind that with these queue managers no system security functions have been used to authorize the setting of this information. Note, however, that if applications do set the context information themselves, they will not migrate unchanged to Level 2 queue managers unless they are given an unusually high level of authority.

If messages are received at a Level 2 product queue manager from a Level 1 product queue manager, the context fields can be set, changed or verified in the transmission program's Message Exit (see "Channel exits" on page 537).

Alternate user ID checks

When a resource of a Level 2 product, such as a queue, is opened, the authorization check that is carried out by default uses the user identifier under which the application is running. Level 2 queue managers support an option which allows the user identifier which is used to be specified explicitly by the application. An application, for example a server, can use this feature to specify the user identifier taken from the context field of a message that it has retrieved, and for which it is currently doing work.

This feature is not supported by Level 1 queue managers.

Channel exits

Channels (also called Message Channel Agents), of Level 2 products support exits that can be used for:

- Security exit
- Message exit
- Send exit
- Receive exit

    On MQSeries for AIX V2.2.1, MQSeries for AS/400 V4R2, MQSeries for Digital OpenVMS, MQSeries for OS/2 Warp V5.0, MQSeries for Tandem NSK, and MQSeries for Windows NT there is also support for:

- Message retry exit

Each end of the channel, if it is connected to a Level 2 queue manager can have each of the four types of exit. For some purposes, for example encryption, it is necessary that there is a corresponding exit at the other end of the channel.

Security exits are given control when the channel is initiated. They can initiate an exchange of security messages, of a format that is defined by the exits themselves, Only when each exit, if present, is satisfied can transmissions proceed along the channel.

Message exits receive control:

- At the sending end, after a message has been retrieved from the transmission queue, and before it is segmented if necessary, and transmitted, and

- At the receiving end, just before a message is put to its destination queue.

In particular, these exits can manipulate the context information. For example they can:

- Validate user identifiers

- Translate user identifiers, if the messages are entering a new authority domain

- Set the user context fields, to show other programs the degree to which the message is trusted.

Send and receive exits can be used for journaling and message encryption.

Message retry exits can be used to reduce the number of messages on a dead-letter queue. After a message has failed to be put on a queue, the exit will make additional attempts to put the message on the target queue. The message is only added to the dead-letter queue if this fails.

With the exception of MQSeries for Digital VMS VAX V1.5, transmission program exits are not available with Level 1 products.

API-crossing exit

The MQSeries for MVS/ESA queue manager supports an exit, for CICS for MVS/ESA applications only, which is executed both before and after every API call. The exit can cause the call to be rejected, and can set, or change, any of the parameters being returned to the application. It can therefore be used to implement additional security checks.

Application programs do not have to be changed to cause this exit to be invoked, so in general there are no issues of program migration involved with it, although if it is used for security purposes migration of this function may be difficult.

This exit could also be used to perform processing essential to the working of the application, for example, data conversion (see "Data conversion" on page 529). If this is done, difficulty will again occur if the application has to be migrated.

Object types

Table 59. Object types - function summary

Product	Model and dynamic queues	Namelists	Queue manager object	Distribution lists
A	Yes	No	Yes	No
B	No	No	No	No
C	Yes	Yes	Yes	No
D	Yes	No	Yes	Yes
E	Yes	No	Yes	(1)
F	No	No	No	No
G	Yes	Yes	Yes	No
H	No	No	No	No
I	No	No	No	No
J	Yes	No	Yes	No
K	Yes	No	Yes	No
L	Yes	No	Yes	Yes
M	Yes	No	Yes	No
N	Yes	No	Yes	Yes

Note:

1. Supported on MQSeries for AS/400 V4R2 only

Model and dynamic queues

The MQSeries for AS/400 V4R2, MQSeries for MVS/ESA, MQSeries Version 2 for UNIX systems, MQSeries for OS/2 Warp V5.0, MQSeries for Tandem NSK, and MQSeries for Windows NT queue managers support a type of queue called a model queue. Opening a model queue causes a dynamic queue to be created.

One of the main uses of this is for a requester to create a reply queue dynamically, avoiding the need for it to be predefined. This is convenient if there needs to be several instances of such requesters. Using dynamic queues means that there need be no fixed limit on the total number of them.

However, this feature is currently available only with some queue managers, and so applications making use of it will need to be changed if they are to be migrated to other queue managers.

Namelists

MQSeries for MVS/ESA and MQSeries for Tandem NSK support a type of object called a namelist. Although this is intended to hold a list of queue names, it can in fact hold a list of any names (subject to a maximum of 48 characters per name), since there is no check that the objects named do in fact exist.

Namelists can be used by monitoring applications to inquire a list of the queues to be monitored. The benefit is that the list of queues in the namelist is maintained by the systems administrator, independently of the application.

However, namelists are supported by only these two platforms, so applications making use of them will not migrate to other platforms.

Queue-manager object

With Level 2 queue managers, a queue-manager object can be opened. This allows the application to inquire the values of various queue-manager attributes, including the maximum message length (*MaxMsgLength*) supported by the queue manager. This feature is not available with Level 1 queue managers, and applications making use of it will not migrate to them unchanged.

Administration

There are substantial differences in the way administration is carried out for the various queue managers, in particular between the Level 1 queue managers, as a whole, and each of the Level 2 queue managers. All Level 2 products provide single point of control for their object administration. In addition, the administration of MQSeries for OS/400 V3R1 or later, MQSeries Version 2 for UNIX systems, MQSeries for Windows NT, and MQSeries for OS/2 Warp V5.0 can be done remotely, from any one of MQSeries for OS/400 V2R3 or later, MQSeries Version 2 for UNIX systems, or MQSeries for OS/2 Warp V5.0.

With Level 1 products, each queue manager must be administered individually. The only factors that are discussed here are those which affect program migration. For other details see the appropriate publication for the queue manager.

Table 60. Administration - function summary

Product	Command queue	Handles broken	MQSET supported
A	Yes	Yes	Yes
C	Yes	Yes	Yes
D	Yes	Yes	Yes
E	Yes	Yes	Yes
F	No	No	No
G	Yes	Yes	Yes
H	No	No	No
I	No	No	No
J	No	Yes	Yes
K	Yes	Yes	Yes
L	Yes	Yes	Yes
M	Yes	Yes	Yes
N	Yes	Yes	Yes

Command queues

Level 2 queue managers support a command input queue to which command messages can be sent. If the user identifier in the message's context is authorized for the command, the command is executed and a reply message sent giving the results of the command.

However, the format of these commands is different between the MQSeries for MVS/ESA queue manager, and that for the other Level 2 queue managers. For MVS/ESA, the command messages contain character information, in the same format as is supported at the command line. The other Level 2 queue managers support messages containing Programmable Command Formats, which include binary information. Replies to these messages are in a similar format, which is easier for programs, as opposed to humans, to parse.

The queue managers of the Level 1 products do not provide command queues. Administration is only carried out using the supplied utility.

Handles broken

When changes are made to the configuration information, for example a new queue is defined or the attributes of a queue are changed, the immediacy with which this takes effect varies.

With Level 2 queue managers, changes generally take immediate effect. Any change to an attribute which affects the name resolution of a queue (for example, a change to the name of the queue to which an alias queue resolves), causes any handle for that queue to be marked as *broken* by the queue manager. Any subsequent call (other than MQCLOSE) using such a handle fails, with an MQRC_OBJECT_CHANGED reason code. The application has to close the handle, and reopen the queue, whereupon the change will take effect. Changes to other attributes generally take immediate effect for any new messages arriving on the queue.

With Level 1 queue managers, handles are not broken when changes to the configuration occur. The application does not normally see the change until it closes and reopens the queue; the MQRC_OBJECT_CHANGED reason code is never issued. Changes to the Inhibit-get and Inhibit-put attributes do take immediate effect, because these attributes are checked at get or put time, not at open time.

Some Level 1 queue managers, however, do not permit changes to the configuration which would affect currently-open handles.

It is likely that only long-running applications need to be sensitive to changes in the configuration tables. If portability is a strong consideration for long-running applications, it may be better to close and reopen queues periodically, rather than check for MQRC_OBJECT_CHANGED.

MQSET support

Level 2 queue managers support the MQSET call to set a limited number of queue attributes. This call is not supported by Level 1 queue managers.

Programming languages

Support here means that the queue manager has been tested with the language, and that header files or copybooks as appropriate are supplied, together with relevant documentation. It is possible that other languages can be used, providing the calling interface is (or can be made to be) the same as that generated by one of the supported languages.

Table 61. Programming languages supported - function summary

Product	C	C++	Visual Basic	COBOL	RPG	Assembler	PL/I
A	Yes	No	No	Yes	No	No	No
C	Yes	No	No	Yes	No	Yes	Yes
D	Yes	Yes	No	Yes	No	No	Yes
E	Yes	Yes	No	Yes	Yes	No	No
F	Yes	No	No	No	No	No	No
G	Yes	No	No	No	No	Yes (TAL)	No
H	Yes	No	No	No	No	No	No
I	No	No	No	Yes	No	No	No
J	Yes	No	Yes	No	No	No	No
K	Yes	No	Yes	No	No	No	No
L	Yes	Yes	No	Yes	No	No	Yes
M	Yes	No	No	(1)	No	No	No
N	Yes	(2)	No	Yes	No	No	(2)

Note:

1. Not supported on MQSeries for AT&T GIS UNIX
2. Supported on MQSeries for AIX only

No language is supported on every platform, but C is the one most widely supported. It is supported on all platforms except VSE/ESA.

Header files

The header files and copybooks supplied with queue managers include, for example, constants for options, that are supported by that queue manager. However, some constants (reason codes for example), that are usually part of a test rather than being set within a program, are also included if they cannot be generated on a particular platform. This makes it easier to port applications which do need to test for reason codes.

Constant values are not reused for different purposes across queue managers. For example, MQRC_OBJECT_CHANGED has the value 2041. There is no risk that the value 2041 will be used for a different reason code on a queue manager that never generates MQRC_OBJECT_CHANGED.

Connection

Table 62. Default connection to a queue manager - function summary	
Product	**Default connection to a queue manager**
A	No
B	No
C	Yes (1)
D	No
E	Yes
F	No
G	No
H	No
I	Yes (1)
J	No
K	No
L	No
M	No
N	No
Note:	
1. CICS applications only	

Default connection

| With MQSeries for AS/400 and with MQSeries for VSE/ESA there is a default
| connection to a queue manager. It is not essential for applications to issue the
MQCONN and MQDISC verbs. However, it is recommended that even on these
platforms applications do always issue these verbs, in order to enhance their
portability.

Multiple queue managers on AS/400

MQSeries for OS/400 V1 allows multiple queue manager instances to be defined
on one AS/400 system (although an application cannot connect to more than one at
a time). MQSeries for OS/400 V2R3 and subsequent releases do not support this
capability. To avoid migration problems, it is therefore recommended that even for
MQSeries for OS/400 V1 only a single queue manager instance is defined on any
one AS/400.

Appendix J. 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. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service 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.

This publication contains sample programs. Permission is hereby granted to copy and store the sample programs into a data processing machine and to use the stored copies for internal study and instruction only. No permission is granted to use the sample programs for any other purpose.

Programming interface information

This book documents General-use Programming Interface and Associated Guidance Information and Product-sensitive Programming Interface and Associated Guidance Information provided by MQSeries for AIX V5.0, MQSeries for AS/400 V4R2, 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 SINIX and DC/OSx V2.2, MQSeries for SunOS V2.2, MQSeries for Sun Solaris V5.0, MQSeries for Tandem NonStop Kernel V2.2, MQSeries for Windows V2.0, MQSeries for Windows V2.1, and MQSeries for Windows NT V5.0.

General-use programming interfaces allow the customer to write programs that obtain the services of these products.

General-use Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a chapter or section or by the following marking:

General-use programming interface

General-use Programming Interface and Associated Guidance Information...

End of General-use programming interface

Product-sensitive programming interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of these products. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive programming interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.

Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a chapter or section or by the following marking:

Product-sensitive programming interface

Product-sensitive Programming Interface and Associated Guidance Information...

End of Product-sensitive programming interface

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

AD/Cycle	AIX	AIX/6000
AS/400	BookManager	CICS
CICS/ESA	CICS/MVS	CICS/VSE
COBOL/370	COBOL/400	Common User Access
C Set++	C/370	C/400
DB2	DFS	FFST
First Failure Support Technology	IBM	IBMLink
ILE	IMS	IMS/ESA
MQ	MQSeries	MQSeries Three Tier
MVS	MVS/ESA	OpenEdition
OS/2	OS/400	RACF
RPG/400	SAA	Systems Application Architecture
System/390 400	VisualAge	VSE/ESA

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, Windows NT, and the Windows logo are 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.

Notices

Part 6. 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

abend reason code. A 4-byte hexadecimal code that uniquely identifies a problem with MQSeries for MVS/ESA. A complete list of MQSeries for MVS/ESA abend reason codes and their explanations is contained in the *MQSeries for MVS/ESA Messages and Codes* manual.

active log. See *recovery log*.

adapter. An interface between MQSeries for MVS/ESA and TSO, IMS, CICS, or batch address spaces. An adapter is an attachment facility that enables applications to access MQSeries services.

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.

address space. The area of virtual storage available for a particular job.

address space identifier (ASID). A unique, system-assigned identifier for an address space.

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.

alert monitor. In MQSeries for MVS/ESA, a component of the CICS adapter that handles unscheduled events occurring as a result of connection requests to MQSeries for MVS/ESA.

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.

allied address space. See *ally*.

ally. An MVS address space that is connected to MQSeries for MVS/ESA.

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-defined format. In message queuing, application data in a message, which has a meaning defined by the user application. Contrast with *built-in format*.

application environment. The software facilities that are accessible by an application program. On the MVS platform, CICS and IMS are examples of application environments.

application log. In Windows NT, a log that records significant application events.

application queue. A queue used by an application.

archive log. See *recovery log*.

ASID. Address space identifier.

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*.

basic mapping support (BMS). An interface between CICS and application programs that formats input and output display data and routes multiple-page output messages without regard for control characters used by various terminals.

BMS. Basic mapping support.

bootstrap data set (BSDS). A VSAM data set that contains:

- An inventory of all active and archived log data sets known to MQSeries for MVS/ESA
- A wrap-around inventory of all recent MQSeries for MVS/ESA activity

The BSDS is required if the MQSeries for MVS/ESA subsystem has to be restarted.

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.

BSDS. Bootstrap data set.

buffer pool. An area of main storage used for MQSeries for MVS/ESA queues, messages, and object definitions. See also *page set*.

built-in format. In message queuing, application data in a message, which has a meaning defined by the queue manager. Synonymous with *in-built format*. Contrast with *application-defined format*.

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). 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.

CI. Control interval.

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*.

CL. Control Language.

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 prefix (CPF). In MQSeries for MVS/ESA, a character string that identifies the queue manager to which MQSeries for MVS/ESA commands are directed, and from which MQSeries for MVS/ESA operator messages are received.

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.

control interval (CI). A fixed-length area of direct access storage in which VSAM stores records and creates distributed free spaces. The control interval is the unit of information that VSAM transmits to or from direct access storage.

Control Language (CL). In MQSeries for AS/400, a language that can be used to issue commands, either at the command line or by writing a CL program.

controlled shutdown. See *quiesced shutdown*.

CPF. Command prefix.

Cross Systems Coupling Facility (XCF). Provides the MVS coupling services that allow authorized programs in a multisystem environment to communicate with programs on the same or different MVS systems.

D

DAE. Dump analysis and elimination.

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.

data-conversion service. A service that converts application data to the character set and encoding that are required by applications on other platforms.

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.

deferred connection. A pending event that is activated when a CICS subsystem tries to connect to MQSeries for MVS/ESA before MQSeries for MVS/ESA has been started.

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.

| **distribution list**. A list of queues to which a message
| can be put using a single MQPUT or MQPUT1
| statement.

DLQ. Dead-letter queue.

DQM. Distributed queue management.

dual logging. A method of recording MQSeries for MVS/ESA activity, where each change is recorded on two data sets, so that if a restart is necessary and one data set is unreadable, the other can be used. Contrast with *single logging*.

dual mode. See *dual logging*.

dump analysis and elimination (DAE). An MVS service that enables an installation to suppress SVC dumps and ABEND SYSUDUMP dumps that are not needed because they duplicate previously written dumps.

dynamic queue. A local queue created when a program opens a model queue object. See also *permanent dynamic queue* and *temporary dynamic queue*.

E

environment. See *application environment*.

ESM. External security manager.

ESTAE. Extended specify task abnormal exit.

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.

extended specify task abnormal exit (ESTAE). An MVS macro that provides recovery capability and gives control to the specified exit routine for processing, diagnosing an abend, or specifying a retry address.

external security manager (ESM). A security product that is invoked by the MVS System Authorization Facility. RACF is an example of an ESM.

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 AS/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)

forced shutdown. A type of shutdown of the CICS adapter where the adapter immediately disconnects from MQSeries for MVS/ESA, regardless of the state of any currently active tasks. Contrast with *quiesced shutdown*.

format. In message queuing, a term used to identify the nature of application data in a message. See also *built-in format* and *application-defined format*.

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)

FRR. Functional recovery routine.

functional recovery routine (FRR). An MVS recovery/termination manager facility that enables a recovery routine to gain control in the event of a program interrupt.

G

GCPC. Generalized command preprocessor.

generalized command preprocessor (GCPC). An MQSeries for MVS/ESA component that processes MQSeries commands and runs them.

Generalized Trace Facility (GTF). An MVS service program that records significant system events, such as supervisor calls and start I/O operations, for the purpose of problem determination.

get. In message queuing, to use the MQGET call to remove a message from a queue. See also *browse*.

global trace. An MQSeries for MVS/ESA trace option where the trace data comes from the entire MQSeries for MVS/ESA subsystem.

GTF. Generalized Trace Facility.

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*.

in-built format. See *built-in format*.

in-doubt unit of recovery. In MQSeries for MVS/ESA, the status of a unit of recovery for which a syncpoint has been requested but not yet performed.

.ini file. See *configuration file*.

initialization input data sets. Data sets used by MQSeries for MVS/ESA when it starts up.

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 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.

Interactive Problem Control System (IPCS). A component of MVS that permits online problem management, interactive problem diagnosis, online debugging for disk-resident abend dumps, problem tracking, and problem reporting.

Interactive System Productivity Facility (ISPF). An IBM licensed program that serves as a full-screen editor and dialog manager. It is used for writing application programs, and provides a means of generating standard screen panels and interactive dialogues between the application programmer and terminal user.

IPCS. Interactive Problem Control System.

ISPF. Interactive System Productivity Facility.

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

machine check interrupt. An interruption that occurs as a result of an equipment malfunction or error. A machine check interrupt can be either hardware recoverable, software recoverable, or nonrecoverable.

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 format service (MFS). In IMS, and editing facility that allows application programs to deal with simple logical messages, instead of device-dependent data, thus simplifying the application development process. See *message input descriptor* and *message output descriptor*.

| **message group**. A group of logical messages. Logical grouping of messages allows applications to group messages that are similar and to ensure the sequence of the messages.

message input descriptor (MID). In IMS, the MFS control block that describes the format of the data presented to the application program. Contrast with *message output descriptor*.

message output descriptor (MID). In IMS, the MFS control block that describes the format of the output

data produced by the application program. Contrast with *message input descriptor*.

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 segment**. One of a number of segments of | a message that is too large either for the application or | for the queue manager to handle.

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*.

MFS. Message format service.

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

namelist. An MQSeries for MVS/ESA object that contains a list of queue names.

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.

off-loading. In MQSeries for MVS/ESA, an automatic process whereby a queue manager's active log is transferred to its archive log.

Open Transaction Manager Access (OTMA). A transaction-based, connectionless client/server protocol. It functions as an interface for host-based communications servers accessing IMS TM applications through the MVS Cross Systems Coupling Facility (XCF). OTMA is implemented in an MVS sysplex environment. Therefore, the domain of OTMA is restricted to the domain of XCF.

OTMA. Open Transaction Manager Access.

output log-buffer. In MQSeries for MVS/ESA, a buffer that holds recovery log records before they are written to the archive log.

output parameter. A parameter of an MQI call in which the queue manager returns information when the call completes or fails.

P

page set. A VSAM data set used when MQSeries for MVS/ESA moves data (for example, queues and messages) from buffers in main storage to permanent backing storage (DASD).

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.

point of recovery. In MQSeries for MVS/ESA, the term used to describe a set of backup copies of MQSeries for MVS/ESA page sets and the corresponding log data sets required to recover these page sets. These backup copies provide a potential restart point in the event of page set loss (for example, page set I/O error).

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.

recovery log. In MQSeries for MVS/ESA, data sets containing information needed to recover messages, queues, and the MQSeries subsystem. MQSeries for MVS/ESA writes each record to a data set called the *active log*. When the active log is full, its contents are off-loaded to a DASD or tape data set called the *archive log*. Synonymous with *log*.

recovery termination manager (RTM). A program that handles all normal and abnormal termination of tasks by passing control to a recovery routine associated with the terminating function.

reference message. A message that refers to a piece of data that is to be transmitted. The reference message is handled by message exit programs, which attach and detach the data from the message so allowing the data to be transmitted without having to be stored on any queues.

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.

relative byte address (RBA). The displacement in bytes of a stored record or control interval from the beginning of the storage space allocated to the data set to which it belongs.

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.

RESLEVEL. In MQSeries for MVS/ESA, an option that controls the number of CICS user IDs checked for API-resource security in MQSeries for MVS/ESA.

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. Any facility of the computing system or operating system required by a job or task. In MQSeries for MVS/ESA, examples of resources are buffer pools, page sets, log data sets, queues, and messages.

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*.

RTM. Recovery termination manager.

rules table. A control file containing one or more rules that the dead-letter queue handler applies to messages on the DLQ.

S

SAF. System Authorization Facility.

SDWA. System diagnostic work area.

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. (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.

session ID. In MQSeries for MVS/ESA, the CICS-unique identifier that defines the communication link to be used by a message channel agent when moving messages from a transmission queue to a link.

shutdown. See *immediate shutdown*, *preemptive shutdown*, and *quiesced shutdown*.

signaling. In MQSeries for MVS/ESA and MQSeries for Windows 2.1, a feature that allows the operating system to notify a program when an expected message arrives on a queue.

single logging. A method of recording MQSeries for MVS/ESA activity where each change is recorded on one data set only. Contrast with *dual logging*.

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*.

SIT.   System initialization table.

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.

storage class.   In MQSeries for MVS/ESA, a storage class defines the page set that is to hold the messages for a particular queue. The storage class is specified when the queue is defined.

store and forward.   The temporary storing of packets, messages, or frames in a data network before they are retransmitted toward their destination.

subsystem.   In MVS, a group of modules that provides function that is dependent on MVS. For example, MQSeries for MVS/ESA is an MVS subsystem.

supervisor call (SVC).   An MVS instruction that interrupts a running program and passes control to the supervisor so that it can perform the specific service indicated by the instruction.

SVC.   Supervisor call.

switch profile.   In MQSeries for MVS/ESA, a RACF profile used when MQSeries starts up or when a refresh security command is issued. Each switch profile that MQSeries detects turns off checking for the specified resource.

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 Authorization Facility (SAF).   An MVS facility through which MQSeries for MVS/ESA communicates with an external security manager such as RACF.

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.

system diagnostic work area (SDWA).   Data recorded in a SYS1.LOGREC entry, which describes a program or hardware error.

system initialization table (SIT).   A table containing parameters used by CICS on start up.

SYS1.LOGREC.   A service aid containing information about program and hardware errors.

T

target library high-level qualifier (thlqual).   High-level qualifier for MVS/ESA target data set names.

task control block (TCB).   An MVS control block used to communicate information about tasks within an address space that are connected to an MVS subsystem such as MQSeries for MVS/ESA or CICS.

task switching.   The overlapping of I/O operations and processing between several tasks. In MQSeries for MVS/ESA, the task switcher optimizes performance by allowing some MQI calls to be executed under subtasks rather than under the main CICS TCB.

TCB.   Task control block.

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*.

termination notification.   A pending event that is activated when a CICS subsystem successfully connects to MQSeries for MVS/ESA.

thlqual.   Target library high-level qualifier.

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.

trace. In MQSeries, a facility for recording MQSeries activity. The destinations for trace entries can include GTF and the system management facility (SMF). 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.

transaction manager. A software unit that coordinates the activities of resource managers by managing global transactions and coordinating the decision to commit them or roll them back. V5.0 of MQSeries for AIX, HP-UX, OS/2 Warp, Sun Solaris, and Windows NT is a transaction manager.

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.

XCF. Cross Systems Coupling Facility.

Index

Special Characters

Index

Index

coded character sets   31
CodedCharSetId (CCSID)   29
   message data   29
commit
   single-phase   185
   two-phase   185
commitment control considerations   189
committing changes   185
communication
   connectionless   6
   time-independent   7
comparisons, functional   519
comparisons, MQSeries platforms   519
compilers supported   433
compiling
   for MQSeries for AIX   253
   for MQSeries for AS/400   257
   for MQSeries for AT&T GIS UNIX   259
   for MQSeries for Digital OpenVMS   261
   for MQSeries for HP-UX   263
   for MQSeries for MVS/ESA   267
   for MQSeries for OS/2 Warp   279
   for MQSeries for SINIX and DC/OSx   283
   for MQSeries for Sun Solaris   293
   for MQSeries for SunOS   289
   for MQSeries for Windows   301
   for MQSeries for Windows NT   303
completion code   74
confirmation of arrival (COA) report   25
confirmation of delivery (COD) report   25
connecting to a queue manager   88, 89
connection   544
connection handle
   returned from MQCONN   88
   returned from MQCONNX   89
   using with MQGET   120
   what it is   74
connectionless communication   6
constants in COBOL   79
context
   default   108
   identity   36
   message   35
   MQOPEN options   99
   MQPUT options   108
   origin   36
context (Credit Check sample application)   421
Context field   105
control information   21
conversion of report messages   164
conversion processing conventions   159
convert characters call   166
convert message data
   MQGET   121, 148
copy files
   how to use them   512

copying messages   149
correlation identifier   34
CorrelId and MsgId with Lotus Notes   247
CorrelId field   34, 133
creating conversion-exit code   166
credit application manager (CAM)   414
credit check sample (MQSeries for MVS/ESA)   409
CRTMQCVX   166
CSQ4BAA1 sample   394
CSQ4BCA1 sample   394
CSQ4BVA1 sample   394
CSQ4CAC1 sample   400
CSQ4CCB5 sample   418
CSQ4CCC1 sample   400
CSQ4CCG1 sample   396
CSQ4CVB1 sample   414
CSQ4CVB2 sample   414
CSQ4CVB3 sample   417
CSQ4CVB4 sample   418
CSQ4CVB5 sample   418
CSQ4CVC1 sample   400
CSQ4CVD1 sample   406
CSQ4CVD2 sample   407
CSQ4CVD3 sample   407
CSQ4CVD4 sample   408
CSQ4CVD5 sample   408
CSQ4ICB3 sample   423
CSQ4RCD1 sample   402
CSQ4RVD1 sample   402
CSQ4TCD1 sample   406
CSQ4TCD2 sample   407
CSQ4TCD4 sample   408
CSQ4TCD5 sample   408
CSQ4TVD1 sample   406
CSQ4TVD2 sample   407
CSQ4TVD4 sample   408
CSQ4TVD5 sample   408
CSQCAPX sample   223
CSQQTRMN transaction   212, 217
cursor, browse   98, 150
CVTMQMDTA   166

D

data
   application   21
   message   21, 107
data conversion
   amqsvfc2 sample program   351
   AMQSVFC4 sample program   351
   amqsvfcn sample program   351
   amqsvfcx sample program   351
   application   30
   comparison by platform   529
   conversion processing   159
   convert characters call   166

Index

Index

Index

TUXEDO ubbstxcn.cfg example for MQSeries for
  Windows NT 368
TUXEDO ubbstxcx.cfg example for MQSeries on UNIX
  systems 366
two-phase commit 185
type styles in this book xvii
TZ environment variable
  MQSeries for OS/2 Warp only 37

U
ubbstxcn.cfg example for MQSeries for Windows
  NT 368
ubbstxcx.cfg example for MQSeries on UNIX
  systems 366
undelivered messages, IMS bridge 230
undelivered-message queue, using 57
unit of work
  message persistence 33
  syncpoint 185
  triggering 210
UNIX signal handling 85
UNIX-based platforms 517
use of message types 23
user exits 223
user-defined formats 30
UserIdentifier field 36

V
valid syntax
  creating conversion-exit code 168
  input data set 168
Version field 105
  MQGMO structure 121
  MQPMO structure 105
Visual Basic header files
  CMQB.BAS 511
  CMQB3.BAS 511
  CMQB4.BAS 511
  header files 511
  include files 511
Visual Basic language
  include files 511

W
waiting for messages 16, 142
WaitInterval field 122, 142
Windows 3.1 clients
  compilers supported 433
Windows 95 clients
  compilers supported 433
Windows Help xxiii
writing applications 232

writing exit programs
  data conversion
    Digital OpenVMS 174
    MQSeries for AS/400 169
    MQSeries for MVS/ESA 172
    MQSeries for OS/2 Warp 170
    Tandem NSK systems 173
    UNIX systems 174
    Windows NT 178

X
X/Open XA interface support 192
XA resource manager
  name 192
  structure 193
  XA switch location 193
XA transaction manager samples 354

Sending your comments to IBM

MQSeries

Application Programming Guide

SC33-0807-08

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

Application Programming Guide

SC33-0807-08

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

Company or Organization

Telephone

Address

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.).

---

2   Fold along this line

By air mail
Par avion

IBRS/CCRI NUMBER: PHQ - D/1348/SO

NE PAS AFFRANCHIR

NO STAMP REQUIRED

IBM

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 _____

---

4   Fasten here with adhesive tape

IBM®

SC33-0807-08