



F E D E R A L
S T U D E N T A I D
We Help Put America Through School

**FSA Modernization Partner
Enterprise Application Integration (EAI)
Core Release 3.0**

Application Enablement Guide (Final)

Version 1.0

**Task Order 80
Deliverable 80.1.4b**

October 30, 2002

TABLE OF CONTENTS

1	EXECUTIVE SUMMARY	5
1.1	PURPOSE.....	5
1.2	APPROACH.....	5
1.3	DESCRIPTION OF SECTIONS.....	5
1.4	MAPPING OF ENABLEMENT GUIDE OBJECTIVES TO SECTIONS.....	7
1.5	SCOPE.....	7
1.6	INTENDED AUDIENCE.....	8
2	MQSERIES ARCHITECTURE CONVENTIONS AND GUIDELINES	9
2.1	PROCESS STEPS FOR BUILDING AN EAI INTERFACE.....	9
2.2	MQSERIES NAMING GUIDELINES.....	10
2.3	MQSERIES APPLICATION MESSAGING INTERFACE (AMI) NAMING GUIDELINES	19
2.4	USING A MQSERIES OBJECT.....	20
2.5	MQSERIES MESSAGING IMPLEMENTATION GUIDELINES.....	22
2.6	MQSERIES CLUSTER DESIGN GUIDELINES.....	22
2.7	MQSERIES CLUSTER IMPLEMENTATION GUIDELINES.....	26
2.8	FSA CLUSTER SPECIFICS.....	27
2.9	MQSERIES SECURITY STANDARDS.....	33
2.10	EAI MQSERIES SECURITY IMPLEMENTATION GUIDELINES.....	34
2.11	MQSERIES WEBSHERE DESIGN GUIDELINES.....	35
3	DATA INTEGRATOR ARCHITECTURE CONVENTIONS AND GUIDELINES	39
3.1	DATA INTEGRATOR STANDARDS.....	39
3.2	DATA INTEGRATOR IMPLEMENTATION.....	39
4	MQSERIES INTEGRATOR ARCHITECTURE CONVENTIONS AND GUIDELINES	41
4.1	MQSI NAMING STANDARDS.....	41
4.2	MQSI IMPLEMENTATION GUIDELINES.....	52
5	FSA APPLICATION ENABLEMENT GUIDELINES	53

5.1	APPLICATION PROGRAMS AND MESSAGING.....	53
5.2	APPLICATION USAGE GUIDELINES FOR MQSERIES	53
5.3	APPLICATION USAGE GUIDELINES FOR MQSERIES APPLICATION MESSAGING INTERFACE (AMI)	59
5.4	APPLICATION INTERFACE PROGRAMMING OPTIONS FOR MESSAGE QUEUE INTERFACE (MQI)	61
5.5	EAI COMMON ERROR HANDLING GUIDELINES.....	63
5.6	TRIGGERED QUEUES AND APPLICATIONS	64
5.7	APPLICATION USAGE GUIDELINES FOR DATA INTEGRATOR	67
5.8	APPLICATION USAGE GUIDELINES FOR DATA INTEGRATOR STATUS UTILITY....	69
5.9	APPLICATION USAGE GUIDELINES FOR MQSERIES INTEGRATOR.....	69
6	APPLICATION CONNECTIVITY (ADAPTERS AND BRIDGES)	78
6.1	MQSERIES APPLICATION ADAPTER	78
6.2	ADAPTER CLASSIFICATIONS.....	78
6.3	MQSERIES-CICS/ESA BRIDGE	79
6.4	RUNNING CICS DPL PROGRAMS	80
7	APPLICATION INTEGRATION EXAMPLES	83
7.1	EXAMPLE OF AN INTERFACE USING DATA INTEGRATOR	83
7.2	EXAMPLE OF AN INTERFACE USING MQSERIES INTEGRATOR.....	84
7.3	CUSTOM EAI ADAPTER.....	86
8	REUSEABLE EAI FUNCTIONS	89
8.1	EAI COMMON LOG FUNCTION	89
9	COMMITTING AND BACKING OUT UNITS OF WORK	96
9.1	COMMITTING AND BACKING OUT.....	96
9.2	SYNCPOINT COORDINATION, SYNCPOINT, UNIT OF WORK.....	96
9.3	SYNCPOINT GUIDELINES.....	96
9.4	MQSERIES SYNCPOINT CALLS FOR OS/390.....	100
9.5	MQSERIES SYNCPOINT CALLS ON WINDOWS NT AND UNIX SYSTEMS.....	100

9.6	SINGLE-PHASE COMMIT	101
9.7	TWO-PHASE COMMIT	101
10	APPENDIX A: REFERENCE MATERIAL	102
11	APPENDIX B: GLOSSARY	104
12	APPENDIX C: CORE SERVICES QUESTIONNAIRE	118
12.1	HIGH LEVEL SYSTEM AND INTERFACE INFORMATION.....	118
12.2	APPLICATION	118
12.3	INTERFACES	118
12.4	DETAIL SYSTEM AND INTERFACE OVERVIEW	122
12.5	DETAIL SYSTEM INFORMATION	122
12.6	DETAIL INTERFACE INFORMATION	125

1 EXECUTIVE SUMMARY

1.1 PURPOSE

The EAI Application Enablement Guide was developed in support of the Department of Education's Federal Student Aid (FSA) Modernization Program, to provide an overview of the MQSeries Messaging functionality being implemented as part of the Enterprise Application Integration (EAI) project. The EAI provides a standard reusable architecture for connecting disparate, heterogeneous systems through common middleware architecture. The EAI architecture is built using the following middleware products: MQSeries Messaging, MQSeries Integrator, and Data Integrator. This deliverable defines the guidelines for enabling FSA application developers to design and implement applications utilizing the features of the EAI Core architecture. This includes work done in Release 1.0, 2.0, and 3.0 of EAI.

The deliverable also serves to provide a high level overview of the features and capabilities of the FSA EAI Messaging infrastructure architecture and product capabilities. This deliverable should be the initial reading for all application developers who will be developing applications to interface to the EAI Bus at FSA.

The document is intended to be a living document and a repository of MQSeries best practices and guidelines, which can be adopted by FSA for the implementation of EAI applications.

1.2 APPROACH

The following approach was used to develop the EAI Application Enablement Guide deliverable:

- Review and modify industry best practices to meet the FSA EAI Core Architecture requirements
- Incorporate additional steps required for applications to integrate and utilize the FSA EAI Core Architecture
- Update Release 2.0 Enablement Guide sections where necessary
- Add sections pertinent to Release 3.0

1.3 DESCRIPTION OF SECTIONS

This deliverable is divided into the following sections:

- Section 1 – Executive Summary

This section provides an introduction and overview of the EAI Application Enablement Guide.

- Section 2 – MQSeries Architecture Conventions and Guidelines

This section will provide guidance on naming conventions for using MQSeries in the FSA EAI architecture. The guidelines provide guidance in defining and implementing MQSeries objects.

- Section 3 – Data Integrator Architecture Conventions and Guidelines

This section will discuss the Data Integrator standards that have been developed and provide suggestions for design and implementation of Data Integrator.

- Section 4 – MQSeries Integrator Architecture Conventions and Guidelines

This section will provide guidance on naming conventions for using MQSeries Integrator in the FSA EAI architecture.

- Section 5 – FSA Application Enablement Guidelines

This section provides an overview of messaging and provides specific steps an application needs to perform in order to connect to a queue manager and to send and receive messages.

- Section 6 – Application Connectivity (Adapters and Bridges)

This section discusses the use of adapters and bridges. Adapters handle data inbound-to and outbound-from the application or environment. A bridge is a software component that moves data between a message on a queue and an application or environment.

- Section 7 – Application Integration Examples

This section will provide guidance on integrating FSA Applications to utilize the EAI Core Architecture through illustrative examples.

- Section 8 – Reusable EAI Functions

This section describes reusable EAI functions that can be utilized by applications integrated with the EAI Core Architecture.

- Section 9 – Committing and Backing Out Units of Work

This section describes how to commit and back out any recoverable get and put operations. It also describes applications and their use of operating under syncpoint control.

- Section 10 – Appendix A: Reference Material

This section provides URL links to on-line documentation referenced within this document.

- Section 11 – Appendix B: Glossary

This section provides a glossary of MQSeries related terms and abbreviations. In addition, it includes terms and abbreviations found in this document.

- Section 12 – Appendix C: Core Services Questionnaire

This section serves as a questionnaire for application teams to provide information regarding all the interfaces for each EAI initiative.

1.4 MAPPING OF ENABLEMENT GUIDE OBJECTIVES TO SECTIONS

This section was added to provide the user of this document with a reference section between the objectives of this document and the document contents.

Objective	Maps To:	Description/Comments
Defines the procedures to enable FSA business applications to connect to the EAI Core Architecture.	Section 2.1	This section was added to provide a high-level view of the process steps involved in developing an interface.
	Section 6.0	This section describes connecting applications to the EAI Bus using adapters.
	Section 8.0	This section describes how to take advantage of application logging functions. These functions assist in troubleshooting to the EAI Bus.
Documents the steps to design and build interfaces between FSA business applications and legacy systems.	Section 5.0 Section 7.0 Section 9.0	These sections contain representative examples of interfaces using each of the EAI middleware products thereby providing guidance on integrating FSA applications with the EAI Core Architecture.
Identifies the procedures to identify the business rules for the interface between FSA business applications and legacy systems.	Appendix C	Core Services Questionnaire

1.5 SCOPE

The scope of this deliverable is to provide guidelines and best practices for designing and implementing interfaces between applications using the EAI Bus (EAI core architecture). The guidelines defined in this deliverable are based on best practices. They provide a structured approach for defining a consistent and maintainable environment.

1.6 INTENDED AUDIENCE

The EAI Application Enablement Guide document is intended for application teams who need to understand the services and capabilities provided by the EAI Core Architecture. The contents of this document should be utilized and built upon in accordance with requirements for applications integrating with the EAI Core Architecture.

2 MQSERIES ARCHITECTURE CONVENTIONS AND GUIDELINES

Prior to Modernization Partner's EAI implementation, FSA had not previously utilized MQSeries as part of its existing middleware infrastructure. During Release 1.0 of this implementation standards were implemented and documented in the Release 1.0 version of the Enablement Guide. With each subsequent EAI Release, the Enablement Guide has been updated. This section provides guidance on naming conventions for using MQSeries in the FSA EAI architecture. These guidelines are meant to provide guidance in defining and implementing MQSeries objects.

These standards have been developed in conjunction with the AIS group from Computer Science Corporation, which will be responsible for monitoring FSA queue managers.

2.1 Process Steps for Building an EAI Interface

The steps for building an EAI interface vary greatly. There are a number of factors that impact the involvement and necessary steps. Key factors are:

- Install Platform/OS
- Type of interface (Data Integrator, MQ Server, MQ Client, Custom)
- EAI team involvement (EAI does not always have access to all systems/processes for interface)
- Type of processing (batch/transactional/pub-sub/real-time)
- Business Logic Required (Transformations, Workflow, Error Handling)

However, there are core steps that can be taken to ensure an interface can be built in a repeatable fashion.

The following procedures are used to enable FSA business applications to connect to the EAI Core Architecture:

1. EAI conducts a core kick-off meeting with the Application Team.
2. Application Team completes the EAI Questionnaire. (Please see Appendix C: EAI Core Services Questionnaire)
3. EAI Team reviews the EAI Questionnaire with the Application Team as input for the Interface Estimate.
4. EAI Team completes the Interface Estimate.
5. EAI Team and Application Team review the Interface Estimate.
6. EAI Team completes the Interface Partner Agreement (IPA).
7. EAI Team and Application Team sign the IPA.
8. EAI Team completes the design documentation.
 - a. EAI Team completes Interface Control Document (ICD)
 - b. EAI Team completes Internal Interface Design (IID)
9. EAI Team installs the MQSeries Infrastructure on the application system in the development environment.
 - a. MQSeries software is installed and configured on the application system

- b. Connectivity test is performed between the EAI Bus and the application system.
 - c. Queue Managers are defined on the application system. (Please see section 2.2.2 and 2.4.2 for Queue Manager details)
 - d. Channels are defined between the application system and the EAI Bus. (Please see section 2.2.11 and 2.4.1 for Channel details)
 - e. Local Queues are defined on the application system. (Please see section 2.2.3 and 2.4.2 for Local Queue details)
 - f. Message delivery test is performed between the EAI Bus and the application system using sample programs provided by the MQSeries software.
10. EAI Team builds application interface based on the design documentation.
 - a. MQSeries Code is baselined in ClearCase repository
 - b. Code review is performed by EAI Team
 11. EAI Team installs the MQSeries Infrastructure on the application system in the test environments (see step 8).
 12. EAI Team tests the application interface.
 13. EAI Team logs testing Issues and reviews/resolves them with the Application Team.
 14. Using ClearQuest, the EAI Team coordinates the deployment procedures and timeline with the application host.
 - a. EAI Team installs the MQSeries Infrastructure on the application system in the production environment (see step 8).
 - b. Code is deployed to production
 - c. Application Team performs Production Readiness Test

2.2 MQSeries Naming Guidelines

This section defines MQSeries Messaging naming guidelines for MQSeries objects within FSA's enterprise technical architecture.

2.2.1 Common Rules

All MQSeries names should follow MQSeries naming conventions, rather than the standard for object names on each supported platform. Key standards and guidelines:

- Use all upper case letters (some platforms default text to upper case and MQSeries names are case sensitive)
 - MQSeries allows both upper and lower case letters in its names. However, MQSeries names are case-sensitive. Using lower and uppercase characters for object names is a common source for naming errors.
- Refrain from using % in names

- This character is valid in all MQSeries names, although it is not commonly used in other names across platforms.
- Limit names to alpha–numeric characters
 - Exceptions are the special characters [_ / .]
- Choose meaningful names within the constraint of the standard.
 - Using meaningful names aids the MQSeries Administrator in maintaining the MQSeries environment.
 - There is no required structure, or hierarchy, in an object name, such as may be found on many systems' file names. MQSeries only compares the name strings.
 - These standards recommend using hierarchical names under certain conditions. One such example is to use a suffix where there are multiple “instances” of an object.
- Document object names and always include a description.
 - All objects have a DESCR attribute for this purpose. MQSeries does not act on the value, but it provides additional information as to the function of the queue.
- Choose meaningful names for new MQSeries interfaces.
 - Each application to be integrated using MQSeries creates one or more MQSeries interfaces. The MQSeries interface defines or exposes some application to the outside world. Implied in an interface is a level of reliability and performance commonly referred to as a contract. Any other component can request and receive a service by awareness and compliance with a defined interface. The application does not need to know how or where the service is performed. The interface becomes a DMZ between an application and the outside world, so changes to the interface may cause repercussions across all users of the interface. XML has become one solution to the static nature of interfaces because it allows for self-defining and extensible interfaces. Still XML does not solve all issues and problems with interface definitions.
- Name an interface for what it does and is, because MQSeries interface names tend to surface in the naming of MQSeries components related to the interface.
- Save the definitions

There are a number of reasons for saving the definitions:

 - In the case of a system failure, objects may need to be recreated. To perform this function, the definitions need to be saved separately from the queue manager.
 - They can be used to reset the attributes to a known state. For example if triggering has been turned off, or GET or PUT disabled, it is helpful to be able to restore the objects to their initial state.
 - The definitions can supplement the MQSeries documentation.

2.2.2 Queue Manager

A queue manager provides the messaging and queuing services to application programs through Message Queue Interface (MQI) program calls. Queue manager names are created at the sole discretion of MQSeries administrators. The following guidelines should be followed when naming queue managers:

- Assign unique names to all queue managers
 - This recommendation can often cause significant problems if queue manager names are not unique. (On MVS, the queue manager name must also be distinct from other subsystem names on the same MVS.)
 - A queue manager can be understood as a “container” for queues and related objects. There is typically one per system, but additional queue managers can be defined.
 - Queue Managers with the same name can be configured to exchange messages - by using Queue Manager aliases. This is strongly discouraged. There are some examples where this can lead to ambiguity, and messages can then be sent to the wrong queue manager.
 - If ReplyToQMgr is left blank in the Message Descriptor, MQSeries inserts the actual local Queue Manager name, not its alias.
 - Dead Letter Queue messages identify the real Queue Manager, not any alias.

- Do not copy documentation examples
 - Copying the documentation examples provided with the installation files is an easy way to produce queue managers with duplicate names. Plan for the names of queue managers ahead of time.

- Keep the queue manager name short and meaningful

A recommendation would be to make queue manager names the same as the network host name. However, keep the following points in mind:

- On MVS, the queue manager name corresponds to the MVS subsystem name. Therefore, the queue manager name is restricted to four characters.
 - Many queue managers use the first eight characters when generating unique message identifiers.
 - Channel names, which by convention are derived from queue manager names, are limited to 20 characters.
 - If there were no obvious name, most users would adopt a convention for constructing a queue manager name. Make sure that the convention provides for further expansion, particularly where the restricted names on MVS are concerned.
- For a Queue Manager alias, use the naming conventions for the specific platform
 - This feature is usually related to defining multiple channels between a pair of queue managers.

2.2.2.1 Naming Convention for MQSeries Queue Manager for Mainframe (CPS and NSLDS on OS/390)

Naming examples for MQSeries queue managers on the OS/390 are illustrated below. OS/390 queue manager names are limited to 4 characters in length.

Examples:

QMP1

QM – Indication that STC(Started Task) is for a queue manager

P – Production (D(development), P(production), or T(test))

1 – First instance

QMP2

QM – Indication that STC(Started Task) is for a queue manager

T – Test (D(development), P(production), or T(test))

2 – Second instance

2.2.2.2 Naming convention of the MQSeries Queue Managers for all other platforms

On Non-MVS platforms the queue manager name should not exceed 8 characters. Queue manager names on distributed platforms will be based on the nature of the work performed, with indicators for environment and distance. For example, EAIBUSP1 is the first instance of a production queue manager on the EAI Bus. PEPSD1 would be the first instance of a queue manager in the PEPS development environment.

Examples:

SAIGT1

SAIG – Student Aid Internet Gateway queue manager

T – Test(D(development), P(production), or T(test))

1 – First instance

Queue Manager Names can also have aliases. This adds another layer of “insulation and abstraction” from the underlying object name. Message routing using alias queue manager names is an example of their use. Consolidation of multiple queue managers to one queue manager is also a way to make use of queue manager alias names to minimize the impact of the consolidation on MQSeries application programs and the MQ Administrator. Although queue manager alias objects are defined via remote queue definitions, they should be named according to queue manager naming guidelines.

2.2.3 Local Queues

As a rule, applications will never reference local queues directly but will always access them via alias queues.

A local queue object defines a local queue belonging to the queue manager to which applications are connected. The following guidelines should be adhered to when naming local queues:

- Local queue names can be up to 48 characters long. They should be short, but long enough to be meaningful.
- Local queue names should not include the name of the queue manager or an indication of the platform used.
- Local queue names should not indicate that the queue is local.

- Local queue names should not include the words local or queue (unless relevant in the context of the application).
- Local queue names should be of the form:
FIRSTNODE.[SECONDNODE].[THIRDNODE].FOURTHNODE
 - The first node is five or six characters indicating the name of the system that owns the object. This will be useful when applications from multiple business units share the same machine/queue manager.
 - The second node is optional and may contain five or six characters. This may be used to define which system the queue is going to or from or some other detail of the interface this queue supports.
 - The third node is optional and may contain five or six characters. This may be used to define which system the queue is going to or from or some other detail of the interface this queue supports.
 - The fourth node is any number of characters, such that entire queue name does not exceed 48 characters in length, that is a unique and descriptive term for the application or business-specific function performed by the queue.

Examples:

SAIG.GETMAIL
SAIG.ONLINE.COD.GETMAIL
SAIG.COD.GETMAIL

2.2.4 Remote Queues

As a rule, applications will never reference remote queues directly but will always access them via alias queues.

A remote queue object identifies a queue belonging to another queue manager. The remote queue is usually given a local definition. The definition specifies the name of the remote queue manager where the queue exists as well as the name of the remote queue itself. The information specified when defining a remote queue object enables the queue manager to find the remote queue manager, so that any messages destined for the remote queue go to the correct queue manager. The following guidelines should be adhered to when naming remote queues:

- Remote queue names can be up to 48 characters long. They should be short, but long enough to be meaningful.
- Remote queue names should be of the form:
TARGETQM.TARGETLOCALQUEUE
 - The first node indicates which queue manager owns the local queue that it references.
 - The second node is the name of the local queue referenced by this remote queue.

This is done to provide operations with a clear view of message flow. Since applications never reference remote queues directly, a change in remote queue name or properties would not have any adverse effect nor require any modifications.

Examples:

SAIGP1.SAIG.GETMAIL
SAIGP1.SAIG.ONLINE.COD.GETMAIL
SAIGP1.SAIG.COD.GETMAIL

2.2.5 Alias Queues

An alias queue object enables applications to access queues by referring to them indirectly in MQI calls. When an alias queue name is used in an MQI call the name is resolved to the name of a message queue at run time. This enables changes to the queues that applications use without changing the application itself in any way. The following guidelines should be adhered to when naming alias queues:

- Alias queue names can be up to 48 characters long. They should be short, but long enough to be meaningful.
- Alias queue names should not include the name of the queue manager or an indication of the platform used.
- Alias queue names should not indicate that the queue is an alias.
- Alias queue names should not include the words alias or queue (unless relevant in the context of the application).
- Alias queue names can be of the form:

TARGETQUEUE.[MODE]

- The first node is the name of the local or remote queue referenced by this alias queue.
- The second node is an indicator or whether this queue is to be enqueued (.PUT) or dequeued (.GET).

Examples:

SAIG.GETMAIL.PUT
SAIG.ONLINE.COD.GETMAIL.GET
SAIG.COD.GETMAIL.PUT

Alias queues which are to be used to enqueue will be GET(DISABLED), while alias queues which are to be used to dequeue will be PUT(DISABLED).

2.2.6 Model and Dynamic Queues

The model queue object defines a set of queue attributes that are used as a template for a dynamic queue. The queue manager creates dynamic queues when an application makes an open queue request specifying a queue that is a model queue. The dynamic queue that is created in this way is a local queue whose name is specified by the application and whose attributes are the same as the model queue.

2.2.6.1 Model Queue Naming Conventions

Generally, model queue names should be of the form:

FIRSTNODE.[SECONDNODE].[THIRDNODE].FOURTHNODE

- The first node is five or six characters indicating the name of the system that owns the object. This will be useful when applications from multiple business units share the same machine/queue manager.
- The second node is optional and may contain five or six characters. This may be used to define which system the queue is going to or from or some other detail of the interface this queue supports.
- The third node is optional and may contain five or six characters. This may be used to define which system the queue is going to or from or some other detail of the interface this queue supports.
- The fourth node is any number of characters, such that entire queue name does not exceed 48 characters in length, that is a unique and descriptive term for the application or business-specific function performed by the queue.

2.2.7 Transmission Queues

A transmission queue temporarily stores messages that are destined for a remote queue manager. Transmission queues must be defined for each remote queue manager that a local queue manager will send messages to. It is possible to associate several transmission queues with different characteristics with a remote queue manager. This allows different classes of transmission service. The following guidelines should be adhered to when naming transmission queues:

- Transmission queue names will include the name of the adjacent (i.e. directly connected) queue manager. The transmission queue name will be the name of the destination queue manager only in the case where the destination queue manager is directly connected with the sending queue manager. Otherwise, the transmission queue name will be the name of some other queue manager that will play the middle party in a multi-hop message transfer to the destination queue manager.
- If there is only one channel to the queue manager, use the exact name of the adjacent queue manager.
- If there will be multiple channels to the queue manager, use the adjacent queue manager name followed by a dot and some class of service.
- If the exact queue manager name is not used, appropriate queue manager alias definitions need to be provided to allow MQSeries to perform queue manager name resolution.
- Transmission queue names should be of the form:

AdjacentQueueManagerName[.ClassOfService]

Examples:

SAIGP1
QMT1
PEPSP2.B

The only class of service defined at this time is batch which is indicated by a ‘.B’ suffixed to the queue name. The class of service will provide a mechanism for separating message traffic by type and service level required. For FSA, any traffic not batch in nature will use the default transmission queue and associated channels.

2.2.8 Dead Letter Queues

A dead-letter queue (also known as an undelivered-message queue) receives messages that cannot be routed to their correct destinations. This occurs when, for example:

- The destination queue is full
- The message cannot be put on the destination queue
- The sender is not authorized to use the destination queue
- The destination queue does not exist

The following guidelines should be adhered to when naming dead-letter queues:

SYSTEM.DEAD.LETTER.QUEUE will always be used.

2.2.9 Initiation Queues

An initiation queue receives trigger messages, which indicate that a trigger event has occurred. A trigger event is caused by a message that satisfies the specified conditions being put onto a queue. Messages are read from the initiation queue by a trigger monitor application that then starts the appropriate application to process the message. If triggers are active, at least one initiation queue must be defined for each queue manager. The following guidelines should be adhered to when naming initiation queues:

- Initiation queue names should be of the form:

FIRSTNODE.SECONDNODE.THIRDNODE.

- The first node should contain the system name.
- Use of the second node is dependent on the system name.
- The third node should be INIT or INITQ, literal standing for the initiation queue.

Example:

CPS.BATCH.INIT

CPT1.CICSDEV2.INITQ

2.2.10 Processes

A process definition object defines an application to an MQSeries queue manager. Typically in MQSeries, an application puts or gets messages from one or more queues and processes them. A process definition object is used for defining applications to be started by a trigger monitor. The definition includes the application ID, the application type, and application specific data. A process may only be used to service a single local queue.

The following guidelines should be adhered to when naming processes:

- Process names should not include the name of the queue manager or an indication of the platform used.
- All process names should be of the form:

LOCALQUEUE.PRC

- The first node is the local queue served by this process
- The second node is the ‘PRC’ literal indicating this MQSeries object is a process definition.

Examples:

SAIG.GETMAIL.PRC
SAIG.ONLINE.COD.GETMAIL.PRC
SAIG.COD.GETMAIL.PRC

2.2.11 Channels

A channel provides a communication path. There are two types of channels, message channels and MQI channels. A message channel provides a communication path between two queue managers on the same, or different, platforms. The message channel is used for the transmission of messages from one queue manager to another, and shields the application programs from the complexities of the underlying networking protocols. A message channel can transmit messages in only one direction. If two-way communication is required between two queue managers, two message channels are required.

An MQI channel connects an MQSeries client to a queue manager on a server machine. It is for the transfer of MQI calls and responses only and is bi-directional. A channel definition exists for each end of the link. The following guidelines should be adhered to when naming channels:

- Channel names can be up to 20 characters long.
- Channel names should be of the form:

SendingQM.ReceivingQM[.ClassOfService]

- *SendingQM* is the name of the sending queue manager (without the *_QM*).
- *ReceivingQM* is the name of the receiving queue manager (without the *_QM*).
- *ClassOfService* is optional and is used to distinguish between different classes of service between the same two queue managers. The only class of service defined at this time is batch which is indicated by a ‘B’ suffixed to the channel name. The class of service will provide a mechanism for separating message traffic by type and service level required.

Based on the above channel-naming convention, channel names can always be interpreted as *FromQueueManager.ToQueueManager* without ambiguity.

Examples:

SAIGP1.QMP1
EAIBUSP1.CODP1.B

2.3 MQSeries Application Messaging Interface (AMI) Naming Guidelines

FSA has standardized on the use of Application Messaging Interface (AMI) as a programming API. The AMI is a higher-level programming interface and abstracts many of the messaging specific details into external repositories, removing them from the programmer's responsibility. AMI is organized into three major categories: Services, Policies, and Messages. That is: "Where", "How", and "What".

The OAG OAMAS messaging standard has been implemented by IBM, resulting in the Application Messaging Interface (AMI). AMI has three major components requiring naming standards to be applied. AMI objects exposed to the applications are highly abstracted. Consequently AMI object naming will be highly logical, exposing no implementation specific details. AMI objects are maintained in external repositories. In the interest of maintaining the sanity of MQSeries administrators, a single AMI repository will be used requiring objects to be qualified by the system that uses them. This will ensure the capability to provide different options to different applications requesting the same service.

2.3.1 Service Points

Services are AMI objects that describe the "what" of the request. A service definition contains queue name, queue manager and other details related to what queues are to be used for the request and reply.

Service point names should be of the format:

Calling System.Application Details.Extension

- Calling system is the name of the system invoking AMI for this request
- Application details describe the function performed by the service
- Extension describes the action within the dialog and can be one of the following:
 - REQSDR

Request Sender: This indicates that this service point is used to send requests for a given service.

- REQRCVR

Request Receiver: This indicates that this service point is used to receive requests for a given service

- REPRCVR

Reply Receiver: This indicates that this service point is used to receive replies to request for a given service.

Examples of service point names are:

COD.GETMAIL.REQSDR

This is the service that would be used by COD to request mail from a SAIG mailbox.

SAIG.GETMAIL.REQRCVR

This is the service that would be used by SAIG to receive requests for mailbox data.

SAIG.COD.GETMAIL.REQRCVR

This is the service that would be used by SAIG to receive requests for mailbox data from COD.

2.3.2 Policies

Policies are objects that contain “how” the request to AMI is to be executed. Policy objects contain clauses for connection requests, send and receive requests, as well as publish, subscribe, and policy handler details. It should be possible to create only one policy per application named per that application. If further granularity is required, this will be revisited and this section revised.

Examples of policy names are:

COD

This is the policy used by COD for all calls to AMI.

SAIG

This is the policy used by SAIG for all calls to AMI.

2.4 Using a MQSeries Object

This section supplements the build steps outlined in 2.1. It provides detailed information about two critical components of MQSeries infrastructure, the channels and queues.

2.4.1 Channels

In order for two machines to communicate via MQSeries, a channel must exist. If two systems must exchange messages, then two channels are required. Channels are created by system administrators or dynamically by the MQSeries queue manager. Although used by the MQSeries queue manager to move messages from one system to another, channels are of little interest to the application developer.

2.4.2 Queues

MQSeries system queues are simple FIFO disk-resident buffers that hold messages. Queues can be divided into local queues and remote queues. Local queues reside on the local system and remote queues reside on a remote system. If messages are destined for a remote system, then a remote queue should be used. Messages destined for applications on the local system are sometimes referred to as destination queues, application queues, or as local queues. Local queues are usually looked upon as queues from which applications GET messages. Queues should be created based on application needs and used when messages need to move between systems or between applications on the same system. Local queues were used on each FSA legacy system.

Another type of queue is a transmission queue. Messages destined for remote queue managers are placed in special queues called transmission queues. Messages reside in the transmission queue until they can be delivered to the remote system via the sender channel. From the perspective of the local system,

transmission queues hold outbound messages. Again, transmission queues are created by the system administrator and could be considered background objects. Transmission queues are used when messages are PUT to a remote queue; the application developer does not write them to directly. At least one transmission queue must be defined for each remote queue manager to whom the local queue manager is to send messages directly. Transmission queues were used on each FSA legacy system.

Remote queues and alias queues are alternative logical names, which can be used to address an MQSeries system queue instead of using the actual queue name. In the case of the remote queue definition, a single name is provided for use by an application that relieves the application of needing to know the location (queue manager name) of the destination queue. Remote queues are used when sending messages to a destination queue defined on a remote queue manager. Both remote queues and alias queues are used by the application developer to get and put messages. Remote queues were used on each FSA legacy system.

Alias queues provide a simple one-to-one name substitution capability. An alias associates an alternative (alias) name with an already defined queue. By defining an alias, the MQSeries system administrator has the ability to redirect message traffic. By using alias queue definitions, the programmer is insulated from changing their application code to fit the changing needs of the network. An alias queue is not a queue, but an object that one can use to access another queue.

Initiation queues are queues that are used in triggering. A queue manager puts a trigger message on an initiation queue when a trigger event occurs. A trigger event is a logical combination of conditions that is detected by a queue manager. Initiation queues are defined by the system administrator for the use of triggering. Initiation queues are not used for the get and put of messages by the application developer. They are used by the queue manager. Initiation queues were defined and used on each FSA legacy system.

A dead letter queue is a queue that stores messages that cannot be routed to their correct destinations. There should only be one dead letter queue defined on each queue manager. The dead letter queue is defined by the system administrator at the time the queue manager is created. Applications can also use the queue for messages they cannot deliver. Dead letter queues were created on each FSA legacy system.

A model queue defines a set of queue attributes that are used as a template for creating a dynamic queue. Dynamic queues are created by the queue manager when an application issues a MQOPEN request specifying a queue name that is the name of a model queue. The dynamic queue that is created in this way is a local queue whose attributes are taken from the model queue definition. Dynamic queues do not survive product restarts; use dynamic queues with caution. Model and dynamic queues are used based on application needs. These were not used for FSA

Processes allow an application to be started without the need for operator intervention. 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. Processes are usually associated with a trigger event: when the trigger event conditions are met, the application associated with the process is initiated. For FSA, processes were used to start the adapters.

2.5 MQSeries Messaging Implementation Guidelines

The following is a list of suggestions for MQSeries design and administration:

- The MQSeries Administrator is responsible for defining and maintaining MQSeries objects such as queues, queue managers, channels, and processes.
- The configuration values of MQSeries objects should be selected carefully to satisfy the requirements of each application. The default value is usually the recommended value. It should not be changed without careful evaluation.
- Include a Dead Letter Queue for every implementation.
- Avoid trigger types “DEPTH” and “EVERY”. These triggering methods have the potential to overload the system.
- Long running units of work are detrimental to the performance of the network. Break the work into small pieces; this tends to have the additional benefit of improved restart capability.
- Use verified network port addresses. Every queue manager needs a listener port in order to negotiate communications and manage the various queues. The default port address is 1414. Check with the network engineers to avoid any port address conflicts during implementation.
- Always evaluate using clusters of queues for redundancy and load balancing. Clusters provide a means to distribute the work in a queue among multiple processes. These processes may be on the same or different physical machines, and the machines may be located in the same or different locations. The only restriction on the locations of the members is that the members must be able to communicate via TCP/IP. Communications between the queue managers participating in each cluster enable the sending queue manager to route the message to the appropriate queue manager based on the default load balancing method or user defined cluster workload exit routine.

2.6 MQSeries Cluster Design Guidelines

2.6.1 Selecting Queue Managers to Hold Repositories

In each cluster, select at least one, preferably two, or possibly more of the queue managers to hold repositories. A cluster could work quite adequately with only one repository but using two improves availability. The repository queue managers are interconnected by defining cluster-sender channels between them. A repository is a collection of information about queue managers that are members of a cluster. This information includes queue manager names, their locations, their channels, what queues they host, and so on. Typically, two queue managers in a cluster hold a full repository. The other queue managers in a cluster inquire on the information in the full repositories and build up their own subsets of this information in partial repositories.

The cluster is configured to include the Websphere Application Server and the two Sun Solaris Servers. The Sun Servers were selected to be the repositories for the cluster.

- The most important consideration is that the queue managers chosen to hold repositories need to be reliable and well managed.

- Consider the location of the queue managers and choose ones that are in a central position geographically or perhaps ones that are located on the same system as a number of other queue managers in the cluster.
- Another consideration might be whether a queue manager already holds the repositories for other clusters. If a queue manager were a repository for one cluster, it would be wise to use the same queue manager as a repository for other clusters of which it is a member.

When a queue manager sends out some information about itself, or requests some information about another queue manager, the information or request is sent to two or more repositories. A repository handles the request whenever possible but if the chosen repository is not available another repository is used. When the first repository becomes available again, it collects the latest new and changed information from the others so that the queue managers are kept in synch. The repository queue managers send messages to each other to be sure that they are both kept up to date with new information about the cluster. The automatic updating of repositories by queue managers is part of the behavior that is inherent to clusters and is done behind the scenes without any intervention by the user.

The following cluster-sender and cluster-receiver definitions were taken directly from the IBM MQSeries Queue Manager Clusters Manual:

“A cluster-sender (CLUSSDR) channel definition defines the sending end of a channel on which a cluster queue manager can send cluster information to one of the full repositories. The cluster-sender channel is used to notify the repository of any changes to the queue manager’s status, for example the addition or removal of a queue. It is also used to transmit messages. The repository queue managers themselves have cluster-sender channels that point to each other. They use them to communicate cluster status changes to each other.”

“A cluster-receiver channel (CLUSRCVR) channel definition defines the receiving end of a channel on which a cluster queue manager can receive messages from other queue managers in a cluster. A cluster-receiver channel can also carry information about the cluster-information destined for the repository. The definition of a cluster-receiver channel has the effect of advertising that a queue manager is available to receive messages. You need at least one cluster-receiver channel for each cluster queue manager.”

If all the repository queue managers go out of service at the same time, queue managers continue to work using the information contained in their partial repositories. New information and requests for updates cannot be processed. When the repository queue managers reconnect to the network, messages are exchanged to bring all repositories (both full and partial) back up to date.

2.6.2 Organizing a cluster

Having selected some queue managers to hold repositories, decide which queue managers should link to which repository. The CLUSSDR channel definition links a queue manager to a repository from which it finds out about the other repositories in the cluster. From then on, the queue manager sends messages to any two or more repositories, but it always tries to use the one to which it has a CLUSSDR channel definition first. It is not significant which repository is chosen.

It is not advisable to use a repository queue manager on an OS/390 system as the repository queue manager because MQSeries for OS/390 does not have a command server. To ensure that a particular repository queue manager is not used by the MQSeries Explorer, include the string ‘%NOREPOS%’ in

the description field of its cluster-receiver channel definition. When the explorer is choosing which repository to link to, it ignores those channel description containing ‘%NOREPOS%’, and treats them as though the queue manager did not hold a repository for the cluster. If there are a large number of repositories or they are spread over a large area, it would be advisable to make a second CLUSSDR channel definition.

Choosing names

When setting up a new cluster, consider a naming convention for the queue managers. Every queue manager must have a different name, but it may help to remember which queue managers are grouped where if given a set of similar names. The queue naming convention of a cluster queue manager follows the same naming convention of any other queue manager. Please refer to section 2.1.2 for queue manager naming conventions. It is recommended that the cluster name be descriptive of the function the cluster is performing. The cluster name is limited in length to 48 characters. For example, the name given to the MQSeries cluster for FSA was “EAI”.

- Every cluster-receiver channel must have a unique name. One possibility is to use the queue-manager name preceded by the preposition ‘TO’. The name would be of the form:

FIRSTNODE.SECONDNODE.

Where:

- FIRSTNODE is replace with the literal TO.
- SECONDDNODE is replaced with the queue manager name.

Example:

TO.SU35E16
TO.SU35E17

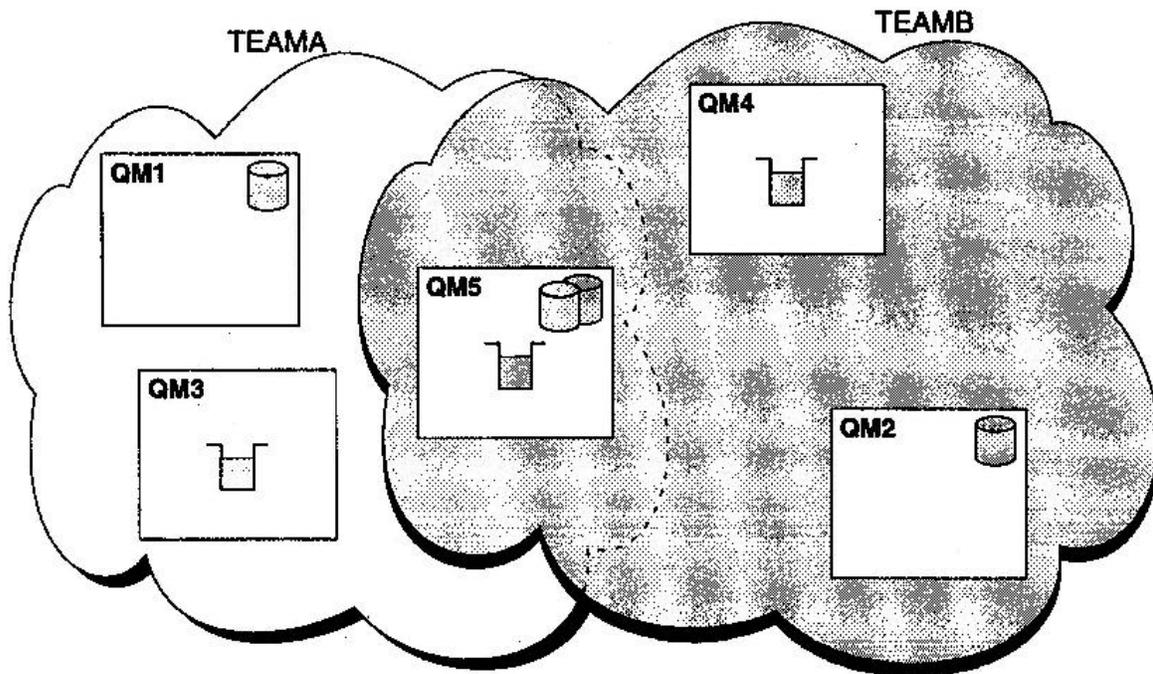
Remember that all cluster-sender channels have the same name as their corresponding cluster-receiver channel.

- Do not use generic connection names on cluster-receiver definitions. If a CLUSRCVR is defined with a generic CONNAME there is no guarantee that the CLUSSDR channels will point to the queue managers intended. The initial CLUSSDR may end up pointing to any queue manager in the queue-sharing group, not necessarily one that hosts a repository. Furthermore, if a channel goes to retry status, it may reconnect to a different queue manager with the same generic name and the flow of messages will be disrupted. Basically, the CONNAME should be the network address of the machine the queue manager resides on.

2.6.3 Overlapping clusters

Create clusters that overlap. There are a number of reasons to do this, for example:

- To allow different organizations to have their own administration.
- To allow independent applications to be administered separately.
- To create classes of service.
- To create test and production environments.



In the figure above, the queue manager QM5 is a member of both the clusters illustrated.

If there is more than one cluster in the network, it is essential to give them different names. If two clusters with the same name are ever merged, it will not be possible to separate them again. When defining a cluster, the following objects are included in the set of default objects defined when creating a queue manager on V5.X of Sun Solaris and Windows NT, and in the customization samples for MQSeries for OS/390.

- SYSTEM.CLUSTER.REPOSITORY QUEUE
- SYSTEM.CLUSTER.COMMAND.QUEUE
- SYSTEM.CLUSTER.TRANSMIT.QUEUE
- SYSTEM.DEF.CLUSSDR
- SYSTEM.DEF.CLUSRCVR

Do **not** alter the default queue definitions. This could alter the default channel definitions in the same way as any other channel definition, using MQSC or PCF commands.

2.6.4 In the Unlikely Event of a Repository Failure

Cluster information is carried to repositories (whether full or partial) on a local queue called SYSTEM.CLUSTER.COMMAND.QUEUE. If this queue should fill up, perhaps because the queue manager has stopped working, the cluster-information messages are routed to the dead-letter queue. If this is observed from the messages on the queue-manager log or OS/390 system console, an application will need to be executed to retrieve the messages from the dead-letter queue and reroute them to the correct destination.

If errors occur on a repository queue manager, messages will appear defining what error has occurred and how long the queue manager will wait before trying to restart. On MQSeries for OS/390 the

SYSTEM.CLUSTER.COMMAND.QUEUE is get-disabled. After identifying and resolving the error, get-enable the SYSTEM.CLUSTER.COMMAND.QUEUE so that the queue manager will be able to restart successfully.

In the unlikely event of a queue manager's repository running out of storage, storage allocation errors will appear on the queue-manager log or OS/390 system console. If this happens, stop and then restart the queue manager. When the queue manager is restarted, more storage is automatically allocated to hold all the repository information.

2.6.5 Cluster channels

Although using clusters relieves the need to define channels (because MQSeries defines them by default), the same channel technology used in distributed queuing is used for communication between queue managers in a cluster. To understand about cluster channels, become familiar with matters such as:

- How channels operate
- How to find their status
- How to use channel exits

These topics are all discussed in the *MQSeries Intercommunication* book.

When defining cluster-sender channels and cluster-receiver channels, do not set the “disconnect interval” too low (less than about 10 seconds). If it is set too low, the channel may close down between sending a request to a repository queue manager and receiving the response.

If the cluster-sender end of a channel fails and subsequently tries to restart, the restart is rejected if the cluster-receiver end of the channel has remained active. To avoid this problem, arrange for the cluster-receiver channel to be terminated and restarted, when a cluster-sender channel attempts to restart.

On V5.X of MQSeries for Sun Solaris and Windows NT

Control this using the AdoptNewMCA, AdoptNewMCATimeout, and AdoptNewMCACheck attributes in the qm.ini file or the Windows NT Registry. See the *MQSeries System Administration* book for more information.

On MQSeries for OS/390

Control this using the ADOPTMCA and ADOPTCHK parameters of CSQ6CHIP. See the *MQSeries for OS/390 System Setup Guide* for more information.

All documentation referenced above can be found in appendix A

2.7 MQSeries Cluster Implementation Guidelines

- On OS/390 clustering cannot be used if the system is using CICS for distributed queuing. In order to get the most benefit out of using clusters, the queue managers in the network need to be on a platform that supports clusters. Until all the systems are migrated to a platform that supports clusters, the system may have queue managers outside a cluster that are not able to access the cluster queues without extra manual definitions. The clustering facility is available to queue managers on the following platforms:

MQSeries for AIX V5.1

MQSeries for AS/400 V5.1
MQSeries for HP-UX V5.1
MQSeries for OS/2 Warp V5.1
MQSeries for OS/390 V2.1
MQSeries for Sun Solaris V5.1
MQSeries for Windows NT V5.1

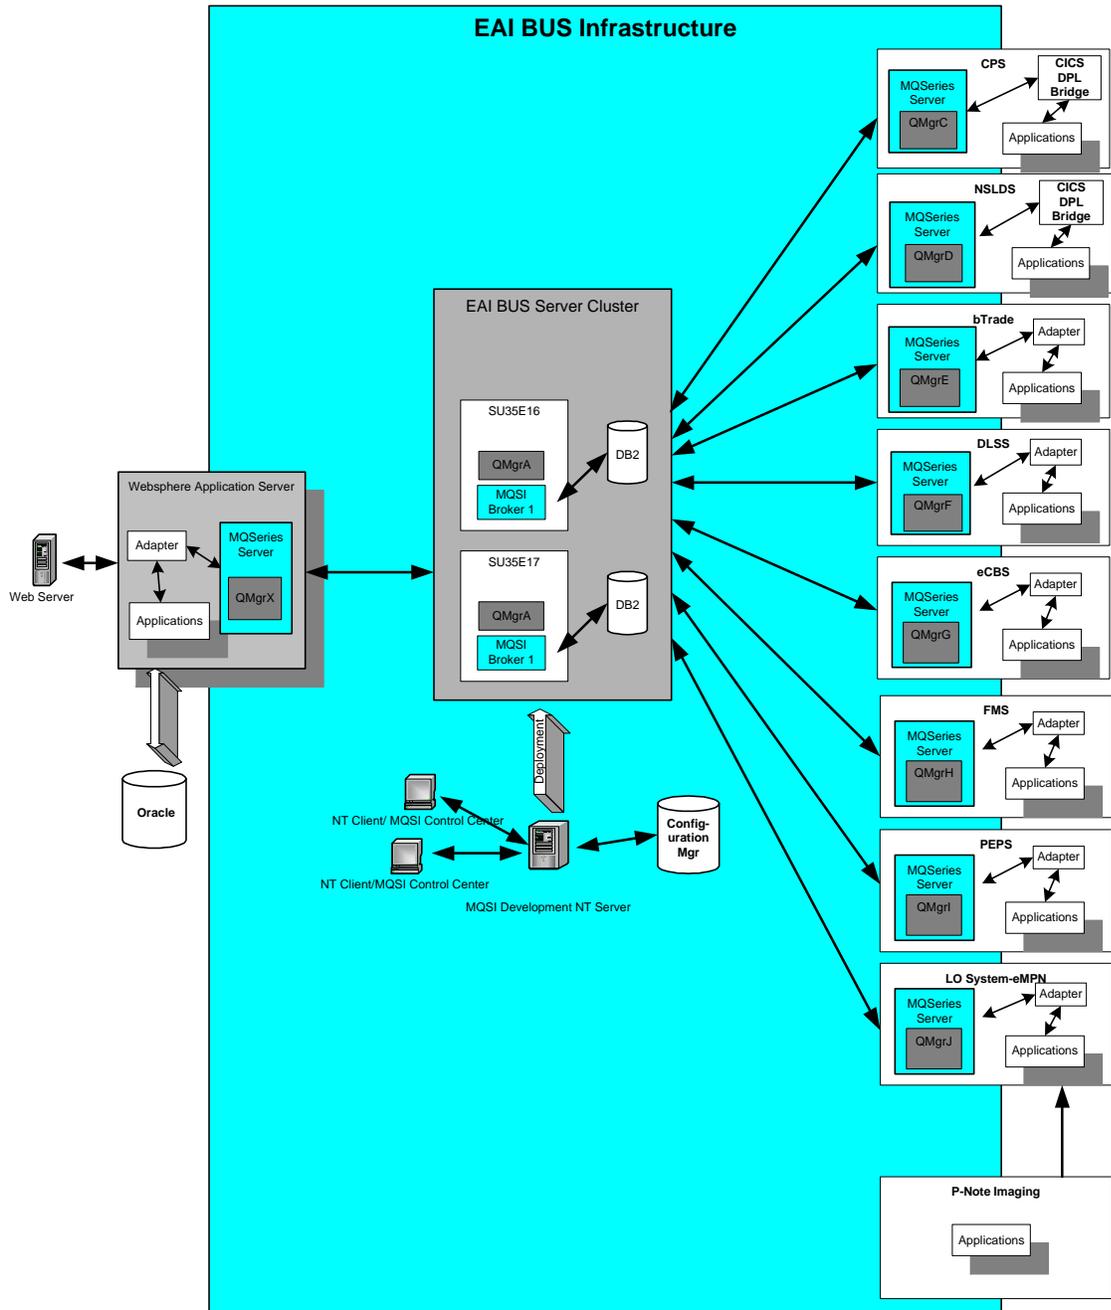
- If two clusters with the same name were merged, it would not be possible to separate them again. Therefore, it is advisable to give all clusters a unique name.
- If a message arrives at a queue manager but there is no queue there to receive it, the message is put to the dead-letter queue as usual. (If there is no dead-letter queue, the channel fails and retries, as described in “Dead-letter queue Guidelines” in the MQSeries Intercommunication book.)
- Using clusters reduces system administration. Clusters make it easy to connect larger networks with many more queue managers than would be possible to contemplate using distributed queuing. However, as with distributed queuing, there is a risk that the system may consume excessive network resources if attempting to enable communication between every queue manager in a cluster.
- The purpose of distribution lists, which are supported on V5.1 of MQSeries for Sun Solaris and Windows NT, is to use a single MQPUT command to send the same message to multiple destinations. Distribution lists can be used in conjunction with queue manager clusters. However, in a clustering environment all the messages are expanded at MQPUT time and so the advantage, in terms of network traffic, is not so great as in a non-clustering environment. The advantage of distribution lists, from the administrator’s point of view, is that the numerous channels and transmission queues do not need to be defined manually.
- If using clusters to achieve workload balancing, first examine the applications to see whether the applications require messages to be processed by a particular queue manager or in a particular sequence. Such applications are said to have message affinities. Applications may need to be modified before being used in complex clusters.
- It is not advisable to use clustering in an environment where IP addresses change on an unpredictable basis such as on machines where Dynamic Host Configuration Protocol (DHCP) is being used.

2.8 FSA Cluster Specifics

2.8.1 Physical layout of the cluster

The hardware architecture implemented at FSA is shown in the diagram below.

EAI BUS Architecture Overview (Development/Test)



This diagram is an operational model of the EAI Bus and trading partner systems. Referenced from left to right this diagram shows 4 logical areas consisting of the WAS Server, EAI Bus Cluster, MQSI Configuration area, and the trading partner systems. The WAS server can be used as a front end for testing of interfaces with trading partner systems. In addition, two trading partner systems may interface with each other via the Bus.

2.8.1.1 Cluster configuration – Development/Test

The FSA EAI cluster consists of 2 Sun Solaris Servers named SU35E16 and SU35E17. The Sun Servers are the repository queue managers for the cluster.

The steps used in creating the cluster are:

1. Install MQSeries on the system.
2. Create the queue managers and the default objects with the `crtmqm` command.
3. Start the channel initiator and the channel listener. The channel initiator monitors the system-defined initiation queue `SYSTEM.CHANNEL.INITQ` which is the initiation queue for all transmission queues. The channel listener must be run on each system. A channel listener program ‘listens’ for incoming network requests and starts the appropriate receiver channel when it is needed.
4. Decide upon the cluster name, in the case of FSA the name of EAI was chosen for the cluster.
5. Determine which queue managers should hold full repositories. For FSA, both nodes SU35E16 and SU35E17 were chosen to hold full repositories.
6. Alter the queue manager definitions to add repository definitions. The command `ALTER QMGR REPOS(EAI)` was executed on both SU35E16(Development) and SU35E17(Test).
7. Define the `CLUSRCVR` channels. For each queue manager in a cluster you need to define a cluster receiver channel on which the queue manager can receive messages. The command was executed on SU35E5, SU35E16(Development), and SU35E17(Test) with the command:

For example:

```
On SU35E5: DEFINE CHANNEL(TO.SU35E5) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)  
CONNAME(ip address of SU35E5) CLUSTER(EAI)
```

```
On SU35E16 Development: DEFINE CHANNEL(TO.SU35E16(Development))  
CHLTYPE(CLUSRCVR) TRPTYPE(TCP) CONNAME(ip address of SU35E16(Development))  
CLUSTER(EAI)
```

```
On SU35E17 Test: DEFINE CHANNEL(TO.SU35E17(Test)) CHLTYPE(CLUSRCVR)  
TRPTYPE(TCP) CONNAME(ip address of SU35E17(Test)) CLUSTER(EAI)
```

8. Define the CLUSSDR channels. On every queue manager in a cluster, you need to define one cluster-sender channel on which the queue manager can send messages to one of the repository queue managers.

On SU35E5: DEFINE CHANNEL(TO.SU35E16(Development)) CHLTYPE(CLUSSDR)
TRPTYPE(TCP) CONNAME(ip address of SU35E16) CLUSTER (EAI)

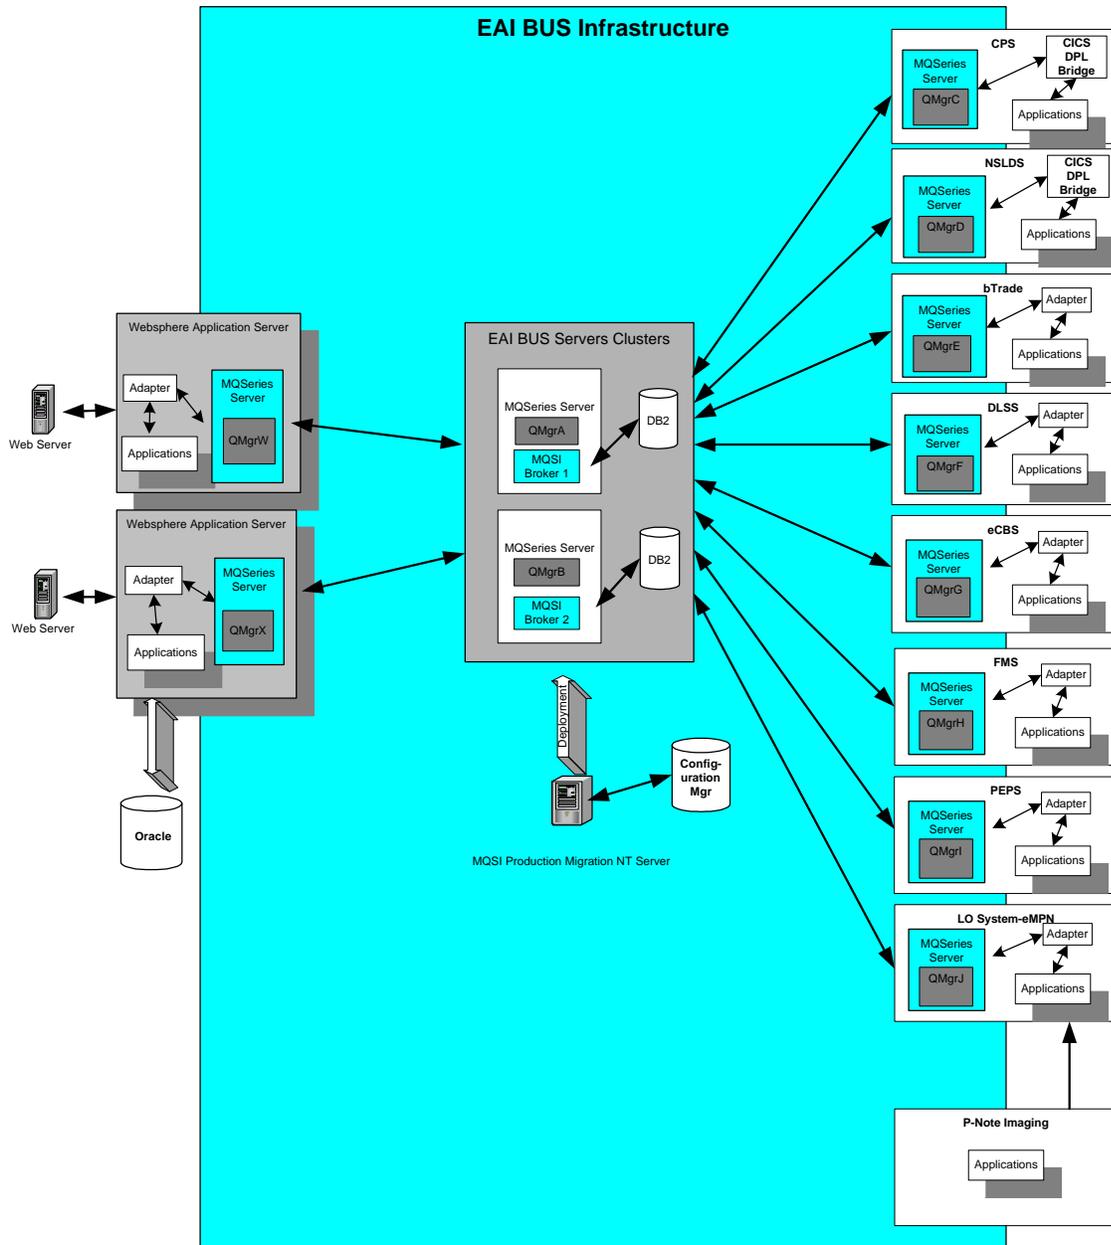
On SU35E16 Development: DEFINE CHANNEL(TO.SU35E17(Test)) CHLTYPE(CLUSSDR)
TRPTYPE(TCP) CONNAME(ip address of SU35E17(Test)) CLUSTER (EAI)

On SU35E17 Test: DEFINE CHANNEL(TO.SU35E16(Development)) CHLTYPE(CLUSSDR)
TRPTYPE(TCP)
CONNAME(ip address of SU35E16(Development)) CLUSTER(EAI)

Once the queue manager has definitions for both a cluster-receiver channel and a cluster-sender channel in the same cluster, the cluster-sender channel is started.

9. Define any cluster queues. For example:
On SU35E16 Development: DEFINE QLOCAL(EAI.FROM.WAS.LOAN) CLUSTER(EAI)

EAI BUS Architecture Overview (Production)



2.8.1.2 Cluster configuration – Production

The FSA EAI cluster “EAIPROD” consists of 2 Sun Solaris machines named SU35E3 and SU35E14. These two machines are the repository queue managers for the cluster.

The steps used in creating the cluster are:

1. Install MQSeries on the system.
2. Create the queue managers and the default objects with the crtmqm command.
3. Start the channel initiator and the channel listener. The channel initiator monitors the system-defined initiation queue SYSTEM.CHANNEL.INITQ which is the initiation queue for all transmission queues. The channel listener must be run on each system. A channel listener program 'listens' for incoming network requests and starts the appropriate receiver channel when it is needed.
4. Decide upon the cluster name, in the case of FSA the name of EAIPROD was chosen for the cluster.
5. Determine which queue managers should hold full repositories. For FSA, both nodes SU35E3 and SU35E14 were chosen to hold full repositories.
6. Alter the queue manager definitions to add repository definitions. The command ALTER QMGR REPOS(EAIPROD) was executed on both SU35E3 and SU35E14.
7. Define the CLUSRCVR channels. For each queue manager in a cluster you need to define a cluster receiver channel on which the queue manager can receive messages. The command was executed on SU35E3, SU35E9, SU35E13, and SU35E14 with the command:

For example:

On SU35E3: DEFINE CHANNEL(TO.SU35E3) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(ip address of SU35E3) CLUSTER(EAIPROD)

On SU35E9: DEFINE CHANNEL(TO.SU35E9) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(ip address of SU35E9) CLUSTER(EAIPROD)

On SU35E13: DEFINE CHANNEL(TO.SU35E13) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(ip address of SU35E13) CLUSTER(EAIPROD)

On SU35E14: DEFINE CHANNEL(TO.SU35E14) CHLTYPE(CLUSRCVR) TRPTYPE(TCP)
CONNAME(ip address of SU35E14) CLUSTER(EAIPROD)

8. Define the CLUSSDR channels. On every queue manager in a cluster, you need to define one cluster-sender channel on which the queue manager can send messages to one of the repository queue managers.

On SU35E3: DEFINE CHANNEL(TO.SU35E14) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(ip address of SU35E14) CLUSTER (EAIPROD)

On SU35E9: DEFINE CHANNEL(TO.SU35E14) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(ip address of SU35E14) CLUSTER (EAIPROD)

On SU35E13: DEFINE CHANNEL(TO.SU35E14) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(ip address of SU35E14) CLUSTER (EAIPROD)

On SU35E14: DEFINE CHANNEL(TO.SU35E3) CHLTYPE(CLUSSDR) TRPTYPE(TCP)
CONNAME(ip address of SU35E3) CLUSTER (EAIPROD)

Once the queue manager has definitions for both a cluster-receiver channel and a cluster-sender channel in the same cluster, the cluster-sender channel is started.

9. Define any cluster queues. For example:
On SU35E3: DEFINE QLOCAL(EAI.FROM.WAS.LOAN) CLUSTER(EAIPROD)

2.9 MQSeries Security Standards

In general, security can be addressed at many different levels in a messaging environment. These security levels are identified as the application, middleware, operating system, network, and link levels. Another way to look at security is to think about providing access control, confidentiality, authentication, non-repudiation and integrity functionality.

EAI is committed to protecting the EAI framework from illegal or damaging actions by individuals, either knowingly or unknowingly. The EAI framework is to be used for business purposes in serving the interests of Federal Student Aid, its clients and customers in the course of normal operations. Effective security is a team effort involving the participation and support of every Federal Student Aid employee and affiliate who deals with information and/or information systems.

EAI complies with FSA/CSC policies and procedures. This consists of file-level security and use of a restricted number of non-published tcp/ip ports. Only users who belong to the mqm group can execute EAI Services. There is only one communication port used by each system for EAI communication purposes.

Therefore:

- 1) All EAI binaries will be run behind the DMZ. It means that no messages will be transmitted in clear text across a public (CSC, TSYS, ACS). This reduces the security exposure of the EAI to the outside world is not bigger than the likelihood of hacker accessing other FSA resources.
Note: See Data Encryption policy for exception(s).
- 2) Use of EAI is restricted to only the users who belong to the mqm group. It serves 2 purposes. First of all it requires explicit action on behalf of every EAI user to be added to the mqm group. Second, malicious users will not be able to get access to EAI. EAI users who failed to go through the step of being added to mqm group as well as malicious users trying to communicate with the "EAI BUS" will be returned a MQSeries 2035 error reason code, which means "Not authorized".
- 3) On each production system there is only one tcp/ip port open to allow incoming communication.

2.9.1.1 Data Encryption Policy (internal to FSA)

The data encryption policy defines requirements for encryption algorithms used within the organization.

EAI data that resides on hardware managed by FSA (CSC) will not be encrypted. This policy is subject to change depending upon application requirements

2.9.1.2 Data Encryption Policy (external to FSA)

The data encryption policy defines requirements for encryption algorithms used with external trading partners. Hardware level encryption is used when exchanging data between FSA managed servers and external trading partner servers; i.e. servers not hosted by the CSC.

Application level encryption will also be considered depending on future requirements.

2.10 EAI MQSeries Security Implementation Guidelines

Queue Managers need to ensure that they exchange messages with the correct partner Queue Managers. Note that it is just as important for the sending Queue Manager to be sure of the receiver's identity as it is for the receiving Queue Manager to be sure of the sender's identify. Such an environment is called mutual authentication. EAI will develop, test, and implement MQSeries channel exits for mid-tier servers to authenticate connectivity from authorized servers. Only servers in the authorized list will be permitted to connect/utilize EAI resources.

2.10.1 EAI Application Consideration

The following table is a guideline of security services that can be provided depending upon application requirements, the category of requirement (NR – not required, O – optional, M –mandatory, V-Varies by application) and which levels might be used to satisfy each service (AL-Application Level, OS-Operating System Level, MW-Middleware Level, LL- Link Level, NW- Network Level, PP-Policy and Procedures).

Service	Category	Applicable Approaches
❖ Identification & Authentication (I & A)		
➤ End User	M	OS, AL
➤ Application Processes	M	OS
➤ Channel Agent	V	OS, MW
❖ Authorization (Access Control)		
➤ Application Processes		
▪ Queues	O	OS, MW
▪ Message Headers	V	OS, MW
➤ System Processes		
▪ Queues	O	OS, MW
▪ Message Headers	V	MW
➤ Restricted Commands	M	OS, MW, AL
➤ Resource Definitions	M	OS, MW, AL
❖ Message Integrity		
➤ Modification Detection	O	LL, NW, MW, AL
❖ Message Non-Repudiation		
➤ Sender	O	NW, MW, AL
➤ Receiver	O	NW, MW, AL
❖ Message Privacy		

Service	Category	Applicable Approaches
➤ Entire Message	O	LL, NW, MW, AL
➤ Selected Fields	O	MW, AL
❖ Logging		
➤ Changes to security information	M	OS, MW, AL
➤ Security-related events		
▪ Failed access attempts	M	NW, OS, MW, AL
▪ Failed message content validation	O	NW, MW, OS, AL

2.11 MQSeries Websphere Design Guidelines

The Web Application Server communicates with the EAI bus to retrieve and put information to different legacy data sources. WebSphere Application Server is the standard Java Application Server in the Integrated Technical Architecture (ITA) at FSA. The WebSphere Server will host Web Based applications that act as middleware between the client browser and FSA's Legacy Systems via the EAI bus. Using Java Server Pages, Servlets and Enterprise Java Beans, WebSphere implements FSA's business application logic through Java based Applications. Several methods exist to enable communication between a WebSphere hosted application and the EAI bus.

2.11.1 WebSphere Connectors

WebSphere Common Connector Framework Classes - IBM java classes provided with WebSphere Application Server which provide a highly abstract view of multiple middleware products.

Application Messaging Interface Java Classes - IBM provided java classes which implement the Open Applications Group Open Applications Messaging Standard (OAG OAMAS). These classes provide a "services" view of middleware, shielding developers from the underlying messaging semantics.

IBM Java class for MQSeries - IBM provided java classes which provide a thin java native interface (JNI) wrapper around the MQSeries native libraries. These classes are provided with the base MQSeries product.

IBM Java Messaging Service classes - IBM provided java classes which provide an implementation of the Sun Java Messaging Service specification. These classes allow applications to be developed using a vendor neutral interface which would allow for the messaging layer to be replaced or the programs moved to other platforms where other JMS services could be used interchangeably.

EAI Messaging components - EAI developed java components to provide an RPC-like interface to FSA middleware. Two implementations of these components exist, one using AMI and one using the base MQSeries java classes.

2.11.1.1 Common Connector Framework

The Common Connector Framework is a standard for developing applications using E-Business Patterns. When a Web Application Server needs to access a Backend Enterprise Information System, whether it is a middleware messaging system, Enterprise Database system, or a System 390 Transaction Management System, several common communication procedures must take place. These procedures may include starting a transaction, processing data, passing status, and closing the transaction. Whether the backend system is CICS, IMS, DB2 or an Oracle RDBMS, the actual commands and parameters may be different but the high level procedures are common. Since these procedures and backend systems have already been identified, prebuilt java classes can be written to communicate with these systems. This requires a change to the parameters and data that is passed to the backend systems.

The Common Connector Framework (CCF) is actually implemented within IBM's java development tool, Visual Age for Java (VAJ). The needed classes that implement the binding between the Web Application Server and MQSeries are included within VAJ's Enterprise Access Builder, which is part of VAJ Enterprise Edition. Programs written using the MQSeries CCF connector classes can communicate with MQSeries Applications using the standard MQSeries Programming Interface or the MQSeries Client classes for java interface. A programmer can use the SmartGuide Wizard within VAJ to build a program shell that will communicate with MQSeries and all that is required is to add the application business logic that will make decisions.

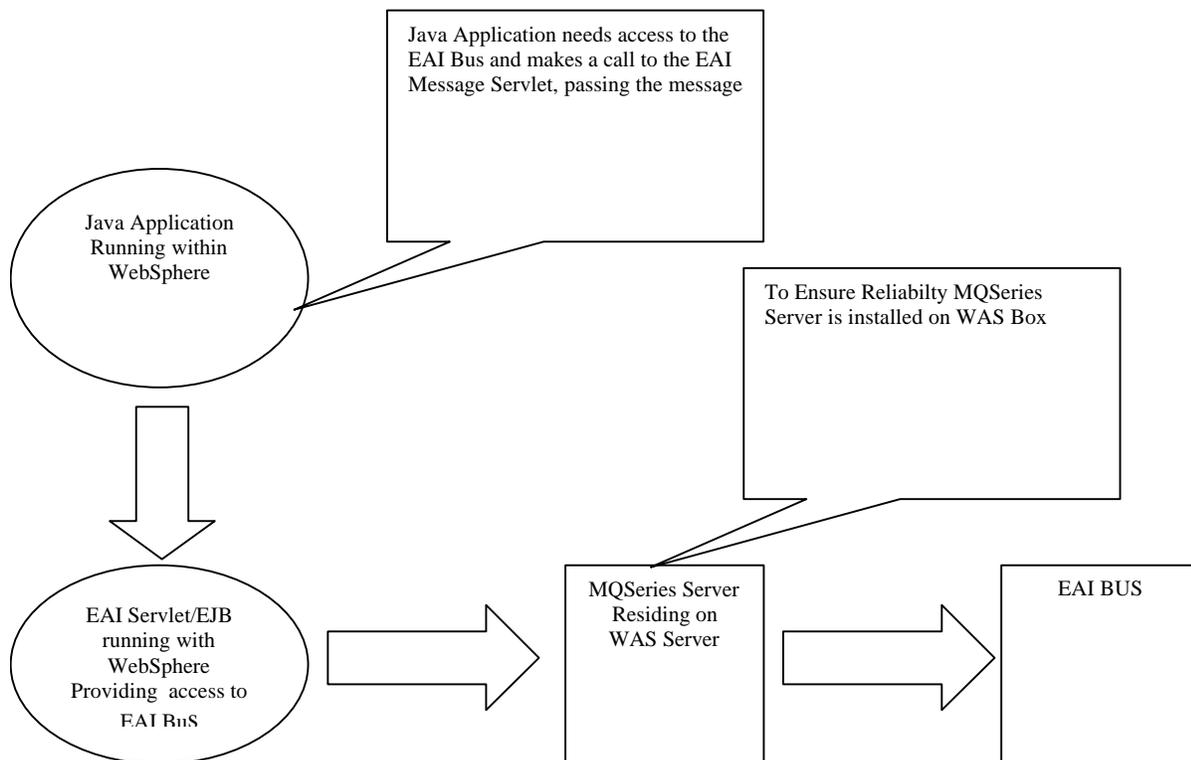
2.11.1.2 Build Your Own Connector

Using MQSeries client classes for Java, a programmer can develop their own interface to MQSeries. This option should only be used by very experienced programmers that have previously implemented Java interfaces to messaging systems. This option is not recommended because the CCF framework is so readily available.

2.11.2 Architecture look and feel

Whether using the CCF framework or building a custom connector, it is important to have a standard, reusable application component within the Java Application Server that enforces communication and data transfer standards between the Application Server and EAI Bus. This reusable component can exist as a Servlet or an Enterprise Java Bean on the WebSphere Application Server. When other applications require access to the EAI Bus, the applications would make a call to the servlet/EJB, which then forwards the message to the EAI bus. This reusable servlet/EJB enforces naming standards, queue names and cluster names before sending message data to the EAI bus. This servlet/EJB would also control the number of connections to MQSeries and allow a central place to tune and manage the web application interface to MQSeries.

To provide reliability and availability of the EAI Bus, the MQSeries Server component should be installed on all WebSphere Application Server (WAS). If an active MQSeries Server with defined Queue Managers are installed on the WAS Server, this ensures assured delivery of all messages to the target destination. If the Queue Manager on the target destination server goes down the sending MQSeries Server will retain the message data and send once connectivity to the target Queue Manager is restored.



2.11.3 FSA EAI WebSphere Reusable Component

The EAI Core Architecture team has developed a reusable WebSphere Java component as an aid to FSA application developers in connecting Internet applications to the EAI Bus. This reusable component is written in Java. It provides a class file for putting and getting messages into and out of a MQSeries message queue. The application incorporates this Java class within the application code to provide a transparent mechanism for putting messages into a queue to be sent to the EAI Bus for processing, and to retrieve message data from a queue upon return. The application specific logic must be built into the application to pass the required message data and to process the message data upon receipt.

2.11.3.1 WebSphere MQ Adapter Overview

The WebSphere MQ Adapter is a Java component that provides a Class file to put message data into a MQSeries message queue and to get message data from a MQSeries message queue. The adapter utilizes MQSeries MQI calls to perform this functionality. In addition, the MQ Adapter provides XML translation capability. This transforms the input message from the WebSphere server application into the application specific XML format. The input data can be of any format, as long as the XML mapping is defined in the application specific MQ Adapter. The message can be passed to the EAI Bus for transformation by MQSI. This functionality was provided for the PEPS and bTrade validation test.

The developed EAI MQ Adapter resides in the ClearCase repository. Any FSA application development team can utilize this functionality for putting messages into a MQSeries message queue, transforming into XML format, and getting the returned message data from the legacy system for processing by the application.

3 DATA INTEGRATOR ARCHITECTURE CONVENTIONS AND GUIDELINES

Data Integrator (DI) is a COTS product that uses MQ messaging to transfer data between different systems. When sending a file, DI automatically splits it up into MQ messages, sends it across a channel(s), and reassembles it at the destination machine. FSA has not previously utilized Data Integrator as part of its existing middleware infrastructure therefore no standards currently exist. This section will provide guidance on Data Integrator standards that have been developed in the FSA EAI architecture.

3.1 Data Integrator Standards

3.1.1 Bi-Directional Sending

Bi-directional sending means that messages can be sent in both directions between queue managers. Every trading partner that uses Data Integrator is set up to send and receive messages in both directions, even in cases where the ability to send in one direction is all that is required.

3.1.2 Pool Architecture

Pooled queues are used to support large file transfers. Local queues have size limitations of 1 GB, so pooling 50 local queues allows each trading partner to receive files up to 50 GB in size. All interfaces on the Bus use 50 pooled queues, and the default data pool used for sending via the EAI Bus is EAIPOOL.

3.1.3 Scripts

Unix scripts that send data between interfaces using Data Integrator are generally named in the following format: <Destination Name>interface.sh (ex. A script on FMS called PEPSinterface.sh would send data from FMS to PEPS).

All scripts for Data Integrator reside in the \$EAIDIR/ftf directory

3.1.4 Configuration file

The configuration file contains default settings for Data Integrator, including defaults for its components (discussed later), logging, and pooled queues. An example is the DefaultPool value of 'EAIPOOL', as mentioned above. There is one .ini configuration file (ftfconfig.ini) that is used across all trading partners. The .ini files used for all trading partners have the same options specified.

3.2 Data Integrator Implementation

The following is a list of suggestions for the design and administration of Data Integrator.

- Unix environment variables (i.e. \$LQM) should be used whenever possible.
- Avoid the use of 'dirmon' (directory monitoring) whenever possible. The script that is supposed to be "kicked off" should be called directly instead.
- When scripts are called via cron:
 - Ensure that the owner of the cronjob has permission to execute the DI scripts.

- Ensure that `ftfenvlist.sh` has been sourced in from the EAI build.
- Trading partner data for the COD interface is written out to the `$EAIDATA` directory. If another interface would prefer data written out to a different directory, ensure that the appropriate users have read/write access to that directory.

4 MQSERIES INTEGRATOR ARCHITECTURE CONVENTIONS AND GUIDELINES

MQSeries Integrator provides data transformation and message routing capability. This section will provide guidance on naming conventions for using MQSeries Integrator in the FSA EAI architecture. These guidelines are meant to provide guidance in defining and implementing MQSeries Integrator objects.

4.1 MQSI Naming Standards

4.1.1 Common Rules

For FSA, there are rules that must be adhered to when providing names or identifiers for message flows, message sets, messages, message processing nodes, brokers, and execution groups in the MQSI broker domain:

- Use all upper case letters
- Version all message flows, message sets, and messages names.
- Underscore should be used instead of spaces when naming all message flows, message sets, and messages.
- MRM object identifiers must match the object name. MRM objects include categories, element qualifiers, elements, element lengths, messages, types and element valid values. The reason for this is that the objects are referenced by their identifiers and not their names.
- A corresponding description should be provided for any objects created.

4.1.2 Brokers

The broker is a MQSI resource that hosts and controls business processes defined in message flows. The following guidelines should be followed when naming brokers:

- The broker name must be unique within the MQSI domain.
- Associate each broker with a separate MQSeries Queue Manager. There is a one-to-one correlation between a broker and a queue manager.
- The broker name should have the same name as the Queue Managers they are associated with.
- For example a Queue Manager called “MQSI” would have a broker “MQSI” associated with it.

4.1.3 Execution Groups

An execution group provides an isolated execution environment within the broker, and is started as a separate operating system process. The following guidelines should be followed when naming execution groups:

- The execution group must be unique within a broker.

- Execution group names should be of the form:

FIRSTNODE_SECONDNODE

- □ The first node should indicate the name of the system that messages are coming from or sent to. This means that message flows are grouped easily within the broker based on the systems it communicates with.
- □ The second node should indicate the name of the system that messages are coming from or sent to. This means that message flows are grouped easily within the broker based on the systems it communicates with.

Examples:

COD_FMS

4.1.4 Message Flows

A message flow is a sequence of operations on a message, performed by a series of message processing nodes. The actions are defined in terms of the message format, its content, and the results of individual actions along the message flow. The following guidelines should be followed when naming message flows:

- The name for a message flow must be unique within a broker domain.
- Message flow names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE_FORTHNODE

- □ The first node is the name of the system that messages are coming from.
- □ The second node is the name of the system that messages are sent to.
- □ The third node is a short description of as to the function of the message flow.
- □ The fourth node is the current version of the message flow being used.

Examples:

FMS_COD_RESPONSE_1

EAI_COD_ERRORS_1

4.1.5 Message Sets

A message set is a logical grouping of related messages. The following guidelines should be followed when naming message sets:

- The name for a message set must be unique within a broker domain.
- Message set names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node is the name of the system that messages are coming from or sent to.
- □ The second node is the name of the system that messages are coming from or sent to.

- □ The third node is the current version of the message flow being used.

Examples:

COD_FMS_1

4.1.6 Messages

A message is the logical representation of data that is exchanged between client applications and MQSeries Integrator. The following guidelines should be followed when naming messages:

- The name for a message must be unique within a message set.
- Message names should be in the following format:

FIRSTNODE_SECONDNODE_THRIDNODE_FORTHNODE

- □ The first node is the name of the system that owns that messages structure.
- □ The second node is a short description of as to the function of the message.
- □ The third node represents the how the message will be used. Valid values are “INPUT” and “OUTPUT”. If the message will be used as both input and output, then this should be left blank.
- □ The fourth node is the current version of the message being used.

Examples:

COD_FINANCIAL_1

COD_VENDOR_INPUT_1

FMS_RESPONSE_INPUT_1

4.1.7 Message Flow Nodes

A message processing node is a point in the message that represents a well defined processing stage. It can be one of several primitive types or can represent a sub flow. There is no hard and fast way that message flow nodes should be named. However there are some guidelines that can be presented here in order to make the message flow clearer for people trying to understand its business purpose.

4.1.7.1 Check

The Check node compares the format of a message with a message-type specification that you supply when you configure the Check node. The message-type specification comprises any combination of the message domain, message set, and message type. The following guidelines should be followed when naming Check nodes:

- Check node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE_FORTHNODE

- □ The first node must be CHECK
- □ The second node must be “DOMAIN”, “SET”, and/or “TYPE”.
- □ The third node is the short description for the function of the node.

- □ The fourth node is optional and is the instance the Check node within the message flow.

Examples:

CHECK_DOMAIN_XML

4.1.7.2 Compute

The Compute node constructs a new message or modifies elements of an existing message, or its associated destination or exception list, or both. These components of the message can be based on elements of both the input message and data from an external database. The following guidelines should be followed when naming Compute nodes:

- Compute node names should be in the following format:

FIRSTNODE.SECONDNODE.THIRDNODE.FOURTHNODE

- □ The first node is the name of the input message format that is sent to the node.
- □ The second node is the name of the output message format that is sent from the node.
- □ The third node is the short description of as to the function of the node.
- □ The fourth node is optional and is the instance the Compute node within the message flow.

Examples:

COD_VENDOR_INPUT_1.COD_FINANCIAL_1.TRANSFORMATION

COD_FINANCIAL_1.COD_FINANCIAL_1.DBLOOKUP.2

4.1.7.3 Database

The Database node allows a database operation in the form of an ESQl statement to be applied to the specified ODBC data source. The following guidelines should be followed when naming Database nodes:

- Database node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE_FOURTHNODE

- □ The first node represents the how the database function that will be performed. Valid values are “INSERT_INTO”, “UPDATE_IN”, and “DELETE_FROM”.
- □ The second node is the database name that the node connects to.
- □ The third node is the table name that the node references.
- □ The fourth node is optional and is the instance the Database node within the message flow.

Examples:

INSERT_INTO_DEV_COD_1

UPDATE_IN_DEV_COD_2

DELETE_FROM_COD_1

4.1.7.4 DataDelete

The DataDelete node is a specialized form of the Database node that allows deletion of one or more rows from a table in the specified ODBC data source. The following guidelines should be followed when naming DataDelete nodes:

- DataDelete node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE_FOURTHNODE

- □ The first node must be “DELETE_FROM”.
- □ The second node is the database name that the node connects to.
- □ The third node is the table name that the node references.
- □ The fourth node is optional and is the instance the DataDelete node within the message flow.

Examples:

DELETE_FROM_COD_1

4.1.7.5 DataInsert

The DataInsert node is a specialized form of the Database node that allows insertion of one or more rows into a table in the specified ODBC data source. The following guidelines should be followed when naming DataInsert nodes:

- DataInsert node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE_FOURTHNODE

- □ The first node must be “INSERT_INTO”.
- □ The second node is the database name that the node connects to.
- □ The third node is the table name that the node references.
- □ The fourth node is optional and is the instance the DataInsert node within the message flow.

Examples:

INSERT_INTO_DEV_COD_1

4.1.7.6 DataUpdate

The DataUpdate node is a specialized form of the Database node that allows the modification of one or more rows in a table in specified ODBC data source. The following guidelines should be followed when naming DataUpdate nodes:

- DataUpdate node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE_FOURTHNODE

- □ The first node must be “UPDATE_IN”.
- □ The second node is the database name that the node connects to.
- □ The third node is the table name that the node references.
- □ The fourth node is optional and is the instance the DataUpdate node within the message flow.

Examples:

UPDATE_IN_DEV_COD_1

4.1.7.7 Extract

The Extract node derives an output message from those elements of the input message that you specify for inclusion. The following guidelines should be followed when naming Extract nodes:

- Extract node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “EXTRACT”.
- □ The second node is the short description of as to the function of the node.
- □ The third node is optional and is the instance the Extract node within the message flow

Examples:

EXTRACT_HEADER

4.1.7.8 Filter

The Filter node routes a message according to message content using a filter expression specified in SQL. The following guidelines should be followed when naming Filter nodes:

- Filter node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node is a question title that represents the functionality of the node.
- □ The second node is optional and is the instance the Filter node within the message flow.

Examples:

IS_GL_TRANSACTION

IS_AP_TRANSACTION_1

4.1.7.9 FlowOrder

The FlowOrder node enables you to control the order in which a message is processed by a message flow.

- FlowOrder node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “ORDER”.
- □ The second node is the short description of as to the function of the node.
- □ The third node is optional and is the instance the FlowOrder node within the message flow.

Examples:

ORDER_DB_INSERT

4.1.7.10 Input Terminal

The Input Terminal provides an in terminal for an embedded message flow. The following guidelines should be followed when naming Input Terminal nodes:

- Input Terminal node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node is the message flow name.
- □ The second node must be “IN”.

Examples:

COD_ERROR_IN

4.1.7.11 Label

The Label node is a named destination for a message processed by a RouteToLabel node. The Label node is identified by an entry in a destination list of the message when it is processed by a RouteToLabel node. The following guidelines should be followed when naming Label nodes:

- Label node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node is the function of the adjacent nodes that are associated with the Label node.
- □ The second node must be “LABEL”.

Examples:

DB_INSERT_LABEL

FMS_LABEL

DEFAULT_LABEL

4.1.7.12 MQInput

The MQInput node reads a message from an MQSeries message queue defined on the broker's queue manager, and establishes the processing environment for the message. The following guidelines should be followed when naming MQInput nodes:

- The name must be the same name as the underlying MQSeries queue that it references.

Examples:

EAI.COD.FINANCIAL.GET

EAI.FMS.FINANCIAL.RESPONSE.GET

EAI.COD.VENDOR.GET

4.1.7.13 MQOutput

The MQOutput node writes messages to an MQSeries message queue defined on any queue manager accessible by the broker's queue manager, or to the destinations identified in the destination list associated with the message. The following guidelines should be followed when naming MQOutput nodes:

- The name must be the same name as the underlying MQSeries queue that it references.

Examples:

FMS.COD.FINANCIAL.RESPONSE.PUT
COD.FMS.FINANCIAL.RESPONSE.PUT
COD.ERRORS.PUT

4.1.7.14 MQReply

The MQReply node is a specialized form of the MQOutput node that sends a response to the originator of the message by putting a message to the MQSeries queue identified by the ReplyToQueue field of the message header. The following guidelines should be followed when naming MQReply nodes:

- MQReply node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be "REPLY_TO".
- □ The second node is the name of the system that messages are sent to.
- □ The third node is optional and is the instance the MQReply node within the message flow.

Examples:

REPLY_TO_COD
REPLY_TO_FMS_1

4.1.7.15 NeonFormatter

The NEONFormatter node is used to transform a message from a known input format to a specified output format. The message definition and transformations are defined using the NEON Formatter graphical utility. The following guidelines should be followed when naming NeonFormatter nodes:

- NeonFormatter node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be "TRANSFORM_TO_NEON".
- □ The second node is the output format that is defined in the NEON Formatter.
- □ The third node is optional and is the instance the NEONFormatter node within the message flow

Examples:

TRANSFORM_TO_NEON_FMS

TRANSFORM_TO_NEON_COD_1

4.1.7.16 NeonRules

The NEONRules node provides an encapsulation of the NEON Rules engine within the MQSeries Integrator Version 2 environment. The following guidelines should be followed when naming NeonRules nodes:

- NeonRules node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node must be “NEONRULES”.
- □ The second node is optional and is the instance the NeonRules node within the message flow.

Examples:

NEONRULES_1

4.1.7.17 Output Terminal

The Output Terminal provides an out terminal for an embedded message flow. The following guidelines should be followed when naming Output Terminal nodes:

- Output Terminal node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node is the message flow name.
- □ The second node must be “OUT”.

Examples:

COD_ERROR_OUT

4.1.7.18 Publication

The Publication node filters and transmits the output from a message flow to subscribers who have registered an interest in a particular set of topics. The following guidelines should be followed when naming Publication nodes:

- Publication node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “PUBLISH_TO”.
- □ The second node is optional and represents the subscription point.
- □ The third node is optional and is the instance the Publication node within the message flow.

Examples:

PUBLISH_TO_XML

PUBLISH_TO_FIXED_1

4.1.7.19 ResetContentDescriptor

The ResetContentDescriptor node takes the bit stream of the input message and reparses it using a different message template from the same or a different message dictionary. The node can reset any combination of message domain, set, type, and format. The following guidelines should be followed when naming ResetContentDescriptor nodes:

- ResetContentDescriptor node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “RESET_TO”.
- □ The second node is the new message template that the message will follow.
- □ The third node is optional and is the instance the ResetContentDescriptor node within the message flow.

Examples:

RESET_TO_XML

RESET_TO_NEON_1

4.1.7.20 RouteToLabel

The RouteToLabel node provides a routing facility based on the contents of the destination list associated with the message. The destination list contains the identity of one or more target Label nodes. The following guidelines should be followed when naming RouteToLabel nodes:

- RouteToLabel node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node must be “ROUTETOLABEL”.
- □ The second node is optional and is the instance the RouteToLabel node within the message flow.

Examples:

ROUTETOLABEL_1

4.1.7.21 Throw

The Throw node provides a mechanism for throwing an exception within a message flow. The following guidelines should be followed when naming Throw nodes:

- Throw node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “THROW”.
- □ The second node is a short description of the exception that is thrown.
- □ The third node is optional and is the instance the Throw node within the message flow.

Examples:

THROW_INVALID_TRANSACTION_TYPE
THROW_COD_ERROR_1

4.1.7.22 Trace

The Trace node generates trace records that can incorporate text, message content, and date and time information, to help you to monitor the behavior of the message flow. The following guidelines should be followed when naming Trace nodes:

- Trace node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “TRACE”.
- □ The second node is a short description of the function of the trace node.
- □ The third node is optional and is the instance the Trace node within the message flow.

Examples:

TRACE_MESSAGE
TRACE_MESSAGE_HEADERS_1

4.1.7.23 TryCatch

The TryCatch node provides a special handler for exception processing. The following guidelines should be followed when naming TryCatch nodes:

- TryCatch node names should be in the following format:

FIRSTNODE_SECONDNODE

- □ The first node must be “TRYCATCH”.
- □ The second node is optional and is the instance the TryCatch node within the message flow.

Examples:

TRYCATCH

4.1.7.24 Warehouse

The Warehouse node is a specialized form of the Database node that stores the entire message, or parts of the message, or both, to the specified ODBC data source. The following guidelines should be followed when naming Warehouse nodes:

- Warehouse node names should be in the following format:

FIRSTNODE_SECONDNODE_THIRDNODE

- □ The first node must be “HOUSE_IN”.
- □ The second node is the database name that the node connects to.

- □ The third node is the table name that the node references.

Examples:

HOUSE_IN_DEV_COD

4.2 MQSI Implementation Guidelines

The following is a list of suggestions for designing message flows and message sets in MQSeries Integrator.

- Develop message flows to be as concise as possible. There is a cost associated with passing through each of the primitive nodes, so it is best to develop a flow in as few nodes as possible.
- Use RouteToLabel and Label nodes for message flows that routes messages in several (more than 3) directions, since they are cheaper than several Filter nodes.
- Combine Filter and Compute nodes into a single Compute node with an IF THEN ELSE structure when possible.
- Use non-persistent messages instead of persistent messages. Persistent messages cause a great deal more logging and therefore the speed of the flow is bound by the speed of the disk hardware.
- Set Transaction Mode to automatic as a property of the MQInput node, to allow persistent messages to be treated as transactions while non-persistent messages are not.
- Refrain from using nesting loops inside each other since they offer no significant gain because the extra overhead of the additional ESQL counters the small gains from navigating more quickly through the message structure.
- Minimize converting between message XML, MRM and BLOB formats.
- Reusable logic should be placed in sub-flow, so that other flows have access to it. Examples would include a common error handling route.
- Message elements, element lengths and types should be automatically created by using the MQSI importer. The MQSI importer allows for the importing of C structures and COBOL copybooks.

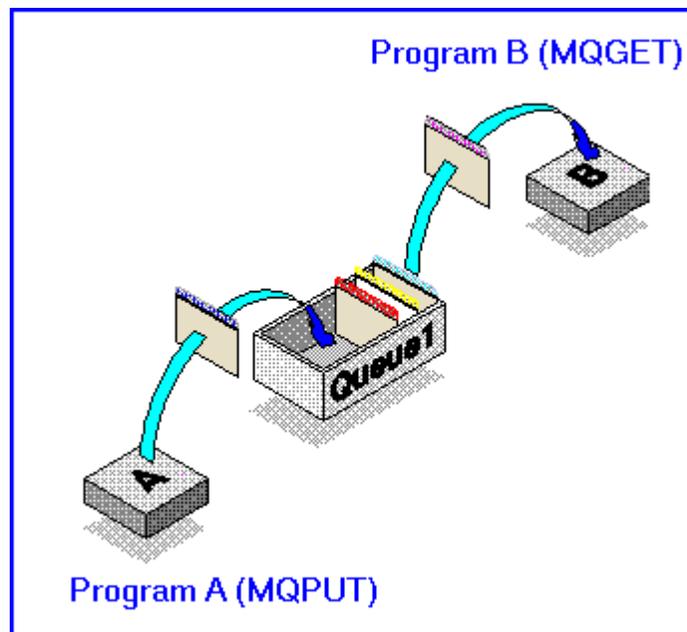
5 FSA APPLICATION ENABLEMENT GUIDELINES

5.1 Application Programs and Messaging

The IBM MQSeries range of products provides application-programming services that enable application programs to communicate with each other using messages and queues. This form of communication is referred to as commercial messaging. It provides assured, once-only delivery of messages. Using MQSeries means that you can separate application programs, so that the program sending a message can continue processing without having to wait for a reply from the receiver. If the receiver, or the communication channel to it, is temporarily unavailable, the message can be forwarded later. MQSeries also provides mechanisms for providing acknowledgements of messages received.

The programs that comprise a MQSeries application can be running on different computers, on different operating systems, and at different locations. The applications are written using a common programming interface known as the Message Queue Interface (MQI), so that applications developed on one platform can be transferred to another.

This figure shows that when two applications communicate using messages and queues, one application puts a message on a queue, and the other application gets that message from the queue.



5.2 Application Usage Guidelines For MQSeries

A queue is a MQSeries object owned by a queue manager, upon which applications can put or retrieve messages. Applications access a queue by using the Message Queue Interface (MQI). Before a message can be put on a queue, the queue must already exist. Each queue must have a name that is unique to the owning queue manager. Before an application can use a queue, it must open the queue, specifying what it wants to do with it. For example, the application can open a queue to:

- Browse messages only (do not delete them)
- Retrieve messages

- Put messages on the queue
- Inquire about the attributes of the queue
- Set the attributes of the queue

For a complete list of the options related to opening a queue, see the description of the MQOPEN call in the *MQSeries Application Programming Reference* manual.

There are different types of queues. These types include:

- *Local*: a local queue is managed by the queue manager to which the application is connected
- *Remote*: a remote queue is managed by a queue manager other than the one to which the application is connected
- *Alias*: an alias queue points to another queue
- *Model*: a model queue is a template for queue definition
- *Dynamic*: a dynamic queue is a temporary queue defined based on a model queue

In FSA's technical environment, the use of alias queues is discouraged, unless a business need dictates its use (e.g. limiting security access to certain queues). Applications putting messages to remote queues will use the remote queue definition. This allows the application to only specify the remote queue name and not be required to know the remote queue manager name. Model and dynamic queues should be used only when a business need dictates their use.

5.2.1 Identifying an Application for a Queue Manager

Any MQSeries application must make a successful connection to a queue manager before it can make any other MQI calls. When the application successfully makes the connection, the queue manager returns a connection handle. This is an identifier that the application must specify each time it issues a MQI call. An application can connect to only one queue manager at a time* (known as its local queue manager), so only one connection handle is valid (for that particular application) at a time. When the application has connected to a queue manager, that queue manager processes all the MQI calls that the application issues until the application issues another MQI call to disconnect from that queue manager. Each adapter written for FSA performs the task of connecting to the queue manager.

* When an application connects to a queue manager, it issues a MQCONN call. The scope of a MQCONN call is limited to the thread that issued it within all of the following:

- MQSeries for AS/400
- MQSeries for Compaq (Digital) OpenVMS
- MQSeries for OS/2 Warp
- MQSeries on UNIX systems
- MQSeries for Windows
- MQSeries for Windows NT

That is, the connection handle returned from a MQCONN call is valid only within the thread that 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 the application has multiple threads and each wishes to use MQSeries calls, each one must individually issue MQCONN. Each thread can connect to a different queue manager on OS/2 and Windows NT, but not on OS/400 or UNIX. If the application is running as a

client, it may connect to more than one queue manager within a thread. This does not apply if the application is not running as a client.

5.2.2 Opening and Closing Queues

Before opening a queue using the MQOPEN call, the application must connect to a queue manager. The application can then use the MQOPEN call to open a queue. The application can also then use the MQCLOSE call to close a queue. When an application opens a queue, the application receives an object handle for that queue. This handle is used in subsequent calls to get or put messages. The same queue can be opened more than once; each open call creates a new object handle. However, most applications will only need to open a given queue once.

Once an application has opened a queue, the application has access to that queue until it closes the queue. The MQOPEN call is costly in terms of time, so once an application has opened a queue and plans to use it in the future, keep the queue open, except when an application only needs to ‘put’ one message. The MQPUT1 call was designed for this case: this call opens a queue, puts the message, and closes the queue, eliminating the need to use the MQOPEN and MQCLOSE calls.

Queues are automatically closed when an application closes its connection to the queue manager. However, it is a good practice to close all queues before disconnecting from the queue manager.

Each adapter written for FSA performed MQOPEN and MQCLOSE calls.

It is recommended to use the FAIL_IF QUIESCING open option for the MQOPEN call. This will allow the MQSeries administrators more control of the system.

5.2.2.1 MQOPEN Call

As input to the MQOPEN call, the application must supply:

- A connection handle, using the connection handle returned by the MQCONN call.
- A description of the object 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 access to the queue. Use this as input to any subsequent MQI calls for this queue.
- A modified object-descriptor structure, if the application is creating a dynamic queue.
- A completion code.
- A reason code.

Always verify the completion code. If the call is unsuccessful, inspect the reason code for an indication as to why the call failed.

5.2.2.2 MQCLOSE Call

As input to the MQCLOSE call, the application must supply:

- A connection handle, using the same connection handle used to open the queue.
- The handle of the queue to close. This comes from the output of the MQOPEN call.

The output from MQCLOSE is:

- A completion code.

- A reason code.

Always verify the completion code. If the call is unsuccessful, inspect the reason code for an indication as to why the call failed.

5.2.3 Putting Messages On A Queue

To put messages on a queue, an application must use the MQOO_OUTPUT option when issuing the MQOPEN call. After the queue has been opened using this option, the application can issue a MQPUT call to put a message on the open queue. If the application is only putting one message and will not use the queue again, use the MQPUT1 call.

It is recommended to use the FAIL_IF QUIESCING put-message option for the MQPUT and MQPUT1 calls. This will allow the MQSeries administrators more control of the system.MQPUT call.

As input to the MQPUT call, the application must supply:

- A connection handle, using the connection handle that was returned when the application issued the MQCONN call.
- A queue handle, using the queue handle that was returned when the application issued the MQOPEN call.
- A description of the message the application is putting on the queue. This is in the form of a message descriptor structure.
- Control information, in the form of a put-message options structure. This options structure needs to be redefined for every MQPUT call.
- The length of the application data contained within the message.
- The application data itself.

The output from the MQPUT call is:

- A reason code.
- A completion code.
- If the call completes successfully, it also returns the put-message options structure and the message descriptor structure. One or both structures may have modified attributes within them. For more detail, look at the *MQSeries Application Programming Guide*.

Always verify the completion code. If the call is unsuccessful, inspect the reason code for an indication as to why the call failed.

5.2.4 Getting Messages From A Queue

To open a queue so that the messages on that particular queue can be browsed (does not remove the message from the queue), use the MQOPEN call with the MQOO_BROWSE option. To get (and remove) messages from a queue, an application must use the MQOO_INPUT_AS_Q_DEF, MQOO_INPUT_SHARED, or MQOO_INPUT_EXCLUSIVE option when issuing the MQOPEN call. Selection of one of these three options is used to specify if the application opens the queue in exclusive, or shared, mode. See the *MQSeries Application Programming Guide* for more information. After the queue has been opened using one of these options, the application can issue a MQGET call to get a message from the open queue.

By specifying the MsgId and/or CorrelId fields in the message descriptor structure, the application can search the queue for a particular message. If the application uses MQGET call more than once (for example, to step through the messages in the queue), it must set the MsgId and CorrelId fields of this

structure to null after each call. This prevents the call from filling these fields with the identifiers of the message that were retrieved, and therefore getting messages with the same identifiers as the previous message.

If the fields in the message descriptor structure are not specified to search for a particular message, the MQGET call will retrieve the first message in the queue.

It is recommended to use the FAIL_IF_QUIESCING get-message option for the MQGET call. This will allow the MQSeries administrators more control of the system.

5.2.4.1 MQGET Call

As input to the MQGET call, the application must supply:

- A connection handle, using the connection handle that was returned when the application issued the MQCONN call.
- A queue handle, using the queue handle that was returned when the application issued the MQOPEN call.
- A description of the message the application wants to get from the queue. This is in the form of a message descriptor structure.
- Control information in the form of a get-message options structure. This control information describes if the application is browsing or removing messages. The control information also describes if the MQI call waits (and how long it waits) for a message or if the call returns immediately.
- The size of the buffer you have assigned to hold the message.
- The address of the storage location in which the message must be put.

The output from the MQGET call is:

- A reason code
- A completion code
- The message in the buffer area specified, if the call completed successfully
- The options structure, modified to show the name of the queue from which the message was retrieved.
- The message descriptor structure, with the contents of the fields modified to describe the message that was retrieved
- The length of the message
- Always verify the completion code. If the call is unsuccessful, inspect the reason code for an indication as to why the call failed

5.2.5 Queue Manager Connectivity Guidelines

A queue manager supplies applications with MQSeries services. An application must have a connection to a queue manager before it can use the services of that queue manager. An application can make this connection explicitly (using the MQCONN call), or the connection can be made implicitly. For example, CICS for MVS/ESA and CICS/MVS programs do not need to explicitly connect to a queue manager, because the CICS system itself is connected to a queue manager. However, for portability it is recommended that CICS for MVS/ESA and CICS/MVS programs use the MQCONN and MQDISC calls.

5.2.6 Connecting To and Disconnecting From a Queue Manager

To connect to a queue manager, an application must use the MQCONN call. To disconnect from a queue manager, an application must use the MQDISC call.

MQCONN Call

As input to the MQCONN call, the application must supply a queue manager name. To connect to the default queue manager, specify a queue manager name consisting entirely of blanks or starting with a null character.

The output from MQCONN is:

- A connection handle, using this handle in subsequent MQI calls associated with this queue manager.
- A completion code.
- A reason code.

Always verify the completion code. If the call is unsuccessful, inspect the reason code for an indication as to why the call failed. 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 MQCONN call fails if the queue manager is in a queuing state when issuing the call, or if the queue manager is shutting down.

MQDISC Call

As input to the MQDISC call, the application must supply the connection handle that was returned by MQCONN when the application connected to the queue manager.

The output from MQDISC is:

- A completion code.
- A reason code.

Always verify the completion code. If the call is unsuccessful, inspect the reason code for an indication as to why the call failed.

All adapters written for FSA had to connect to the queue manager, open a queue, perform a MQGET or MQPUT, close a queue, and disconnect from the queue manager. Each future adapter written for FSA will also need to perform each of the above in order to get and put messages on a queue.

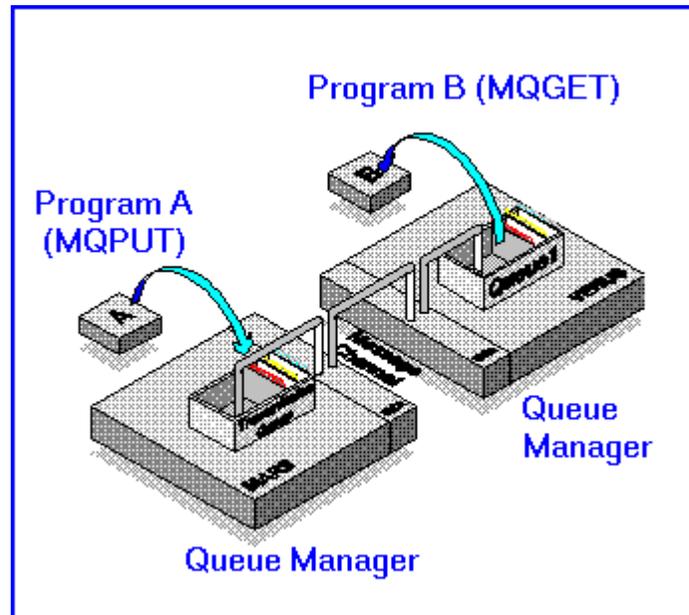
5.2.7 Pass the Connection Name as a Program Parameter

This allows a program to run unchanged on any Queue Manager. This provides the capability for multiple concurrent instances; or a queue driven application could be moved to a different queue manager without impacting the application code.

5.2.8 Messaging Using More Than One Queue Manager

This arrangement is not typical for a real messaging application because both programs are running on the same computer, and connected to the same queue manager. In a commercial application, the putting and getting programs would probably be on different computers, and so connected to different queue managers.

This figure shows how messaging works when the program putting the message and the program getting the message are on the different computers, and are connected to different queue managers.



In this situation, it is necessary to create message channels to carry MQSeries messages between the queue managers.

5.3 Application Usage Guidelines for MQSeries Application Messaging Interface (AMI)

AMI is a highly abstracted interface to MQSeries that externalizes much of the complexity associated with MQSeries usage into an external repository. Understanding its organization is key to its use. AMI is organized into three major levels: Service Points, Policies, and Messages. Service Points contain information related to queues. Policies contain information related to connections, queue interaction, publish and subscribe and AMI user exits. Messages are not abstracted into the AMI repository and are the containers that hold the application data to be placed to or received from queues.

The AMI is object oriented. All errors are reported in the form of thrown exceptions that are caught and evaluated by the application.

The AMI repository is created, updated, and managed by MQSeries administrators who in each case will ensure that objects match application requirements and options are appropriate. The use of an external repository dramatically reduces the amount of middleware knowledge application programmers are required to possess. Comparing the MQI and AMI guidelines demonstrate this conclusively.

5.3.1 AMI Connectivity Guidelines

Any AMI application must establish a Session before it can make any other AMI calls. An AMI Session is a container object that holds the queue manager connection information.

5.3.2 Establishing and Terminating AMI Sessions

In order to establish an AMI session it is necessary to create an AMI Session object. This object will be used to establish connections to the underlying queue manager as well as provide the context for other AMI objects. These Session objects are created with a logical name that must be unique within the application.

Connecting to a Queue Manager

Connecting to a queue manager is a result of running the “open” method of the previously created Session object. An AMI Policy is provided as input. The “Initialization” section of the referenced Policy is used to determine queue manager name, whether to use client or server binding and whether to run as a trusted/fastpath application. When successfully opened, the Session contains an active connection to a queue manager.

Disconnecting from a Queue Manager

Disconnecting from a queue manager is a result of running the “close” method of the previously created Session object. An AMI policy is provided as input. Information from the policy is used in the case where users exits are required. All related objects become invalid after having closed the AMI Session through which they were created.

5.3.3 AMI Sender and AMI Receiver Objects

In order to put messages to and get messages from queues it is necessary to create AMI Sender and AMI Receiver objects. These objects contain queue information and are used to direct interaction with those queues.

5.3.3.1 Using AMI Sender objects

When creating a Sender object a Service Point name is provided as input. This is a reference to a Service Point in the AMI repository. The Service Point contains the queue name that is to be used to put messages.

Once created, the “open” method is used to establish a handle to the target MQSeries queue. A Policy is provided as input. The “Send” section of the policy is used to determine the options related to the placement of messages including priority, persistence, expiry interval, report options and more.

To then send data using this Sender, the “send” method is used providing a Policy and message data as input.

When complete, using the “close” method of the Sender invalidates its handle to the underlying MQSeries queue and closes it.

5.3.3.2 Using AMI Receiver objects

When creating a Receiver object a Service Point name is provided as input. This is a reference to a Service Point in the AMI repository. The Service Point contains the queue name that is to be used to get messages.

Once created, the “open” method is used to establish a handle to the target MQSeries queue. A Policy is provided as input. The “Receive” section of the policy is used to determine the options related to the receipt of messages including wait interval, message conversion and more.

To then receive data using this Receiver, the “receive” method is used providing a Policy and message buffer as input.

When complete, using the “close” method of the Receiver invalidates its handle to the underlying MQSeries queue and closes it.

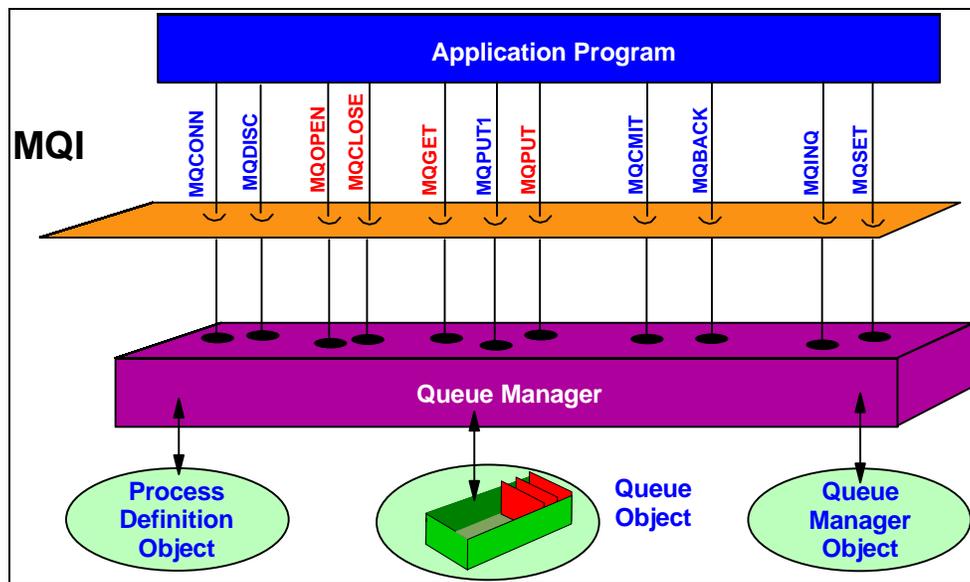
5.4 Application Interface Programming Options for Message Queue Interface (MQI)

There is a wide range of options for communicating with MQSeries programs including new interfaces for message content as well as message delivery. Programs written using any of these message delivery styles can communicate with each other and with programs written in any of the other MQSeries delivery styles.

5.4.1 Message Delivery

5.4.1.1 Message Queue Interface (MQI)

The Message Queue Interface (MQI) is the common API across all platforms. The calls made by the applications running on each platform are common. This allows application programmers to focus on the business logic of the application, rather than the interface differences of each platform. This makes it much easier to write and maintain applications, as well as facilitate migration of applications from one platform to another as required by changing business needs. Each adapter written for FSA made use of a majority of the MQI function calls as shown below. The following figure represents the MQI.



Message Queue Interface

5.4.1.2 Java Message Service (JMS)

Java Message Service (JMS) is supported by a MQSeries implementation of this Java standard API for Enterprise Messaging Services. Using JMS, applications can communicate with other MQSeries JMS applications, with applications written to the MQI, or to the Application Message Interface (AMI).

5.4.1.3 Application Messaging Interface (AMI)

The Application Messaging Interface (AMI) provides a simpler and higher-level programming interface than the MQI. Although it has some limitations compared with the MQI, its function should be sufficient for the majority of users. The AMI supports both point-to-point and publish/subscribe messaging models. The AMI eliminates the need for application programmers to understand all of the options and functions available in the MQI. This was not used at FSA, but is mentioned for future use if the need arises.

The MQSeries AMI can be used to build client applications, and the AMI will automatically build any required headers as specified using the AMI, including the new RFH2 headers. The AMI is designed to simplify the task of the application programmer, while enabling the more advanced functions and message broker facilities to be used.

AMI is a high level API that moves many functions normally performed by messaging applications into the middleware layer, where a set of policies defined by the enterprise is applied on the application's behalf. Policies hold details of how messages are to be handled, for example, priority, confirmation of delivery, timed expiry.

5.4.2 Message Content

5.4.2.1 Extensible Markup Language (XML)

Extensible Markup Language (XML) is an industry-wide standard for self-defining messages. It enables diverse systems and databases to understand each other's data (for example, to identify fields) by indicating both the content and the role of the data.

XML is supported in MQSeries Integrator Version 2 and MQSeries Workflow Version 3.2; XML will be supported within MQSeries Messaging via the Common Messaging Interface.

For FSA, all messages passed into MQSeries Integrator were in XML. IBM is not advocating the use of XML and the adoption of XML as a standard is outside the scope of this document.

Sample XML Message:

```
<?xml version="1.0"?>
<!DOCTYPE Message SYSTEM "C:\TestEnvironment\XMLFiles\LifeQuote.dtd">
<!--Generated by XML Authority.-->
<Message issuedTime = "string" Authorisation = "string" sessionID = "string" creationTime = "string"
issueProgram =
"string" issueUser = "string" ID = "id1" issueSystem = "string" txnScope = "string" eventID = "string"
zoneOffset = "string"
language = "string"><!-- (Command.valueQuoteRequest* , Command.valueQuoteResponse* )-->
<Command.valueQuoteRequest responseDTD = "string" echoBack = "string" cmdMode = "always" ID =
"id2"><!--
(%CustomizeAgreement , SystemInfo )-->
<LifeAgreement ID = "id3" REFID = "string" status = "string" UUID = "UUID"><!-- (%Agreement ,
```

```
Product
)-->
<policyNumber>only text</policyNumber>
<effectiveFromDate>only text</effectiveFromDate>
<companyCode>only text</companyCode>
<ratingCompany>only text</ratingCompany>
<policyType>only text</policyType>
<renewalDate>only text</renewalDate>
<paymentPlan>only text</paymentPlan>
<agreementState>only text</agreementState>
<lineOfBusinessCode>only text</lineOfBusinessCode>
<effectiveFromDate>only text</effectiveFromDate>
<agentOfRecord>only text</agentOfRecord>
<agentCommission>only text</agentCommission>
<PolicyMessage/>
<MoneyObligation ID="id4" REFID="string" status="string" UUID="UUID"><!--(type, amount,
frequency)-->
<type>only text</type>
<amount>only text</amount>
<frequency>only text</frequency>
</MoneyObligation>
<Discount-Surcharge/>
<Underwriting/>
<Applicant ID = "id5" REFID = "string" status = "string" UUID = "UUID"><!-- (Person )-->
<Person ID = "id6" REFID = "string" status = "string" UUID = "UUID"><!-- (%Party , Body
, PartyActivity* , Residency , PartyContactPointUsage? )-->
<id>only text</id>
<uuid>only text</uuid>
<FamilyName/>
<!-- <UnstructuredName>only text</UnstructuredName> -->
<Body ID = "id7" REFID = "string" status = "string" UUID = "UUID"><!-- (gender ,
height , weight , birthdate , MedicalCondition+ )-->
<gender>Female</gender>
<height>6.2</height>
<weight>250</weight>
<birthdate>01/01/1980</birthdate>
<MedicalCondition ID = "id8" REFID = "string" status = "string" UUID =
"UUID"><!-- (description , response )-->
<description>High Blood Pressure</description>
<response>Yes</response>
</MedicalCondition>
<MedicalCondition ID = "id9" REFID = "string" status = "string" UUID =
"UUID"><!-- (description , response )-->
<description>Heart Disease</description>
<response>No</response>
</MedicalCondition>
</Body>
```

5.5 EAI Common Error Handling Guidelines

Whenever possible, the queue manager returns any errors as soon as a MQI call is made. The three most common errors that the queue manager can report immediately are described in this section.

5.5.1 Failure of a MQI Call

An example of a MQI call failure is being unable to put a message to a queue because the queue is full. The completion code and return code of the MQI call specify the nature of the failure. Applications should inspect these codes for every MQI call and be able to handle all possible return codes.

5.5.2 System Interruption

The queue manager is an example of a system component needed by the application and when the queue manager is interrupted, the application encounters an error. Applications must ensure no data is lost due to this sort of interruption. To ensure no data loss, applications will get and put messages under syncpoint. This syncpoint activity can be controlled by the queue manager or by some external resource coordinator (e.g. CICS, Encina, etc.).

5.5.3 Unable to Process Messages

Messages containing data that cannot be processed successfully are known as poisoned messages. When applications operate under syncpoint, if the application cannot successfully process a message, 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 for MQGET calls which DO NOT use any of the Browse type get message options. Messages whose backout counts increase over time are being repeatedly rejected by the application – the application should be designed to handle such situations. There are many different tactics to handling poisoned messages. One method would be to write the messages to a file and a common “poison message application” attempt to process them at a later point in time. Another method is to have the application itself deal with the message. Messages could also be written to the dead letter queue and then be processed by a dead letter handler. Based on your application requirements a method should be adopted.

5.5.4 Responding to Errors

Applications should respond in a similar manner to errors returned by MQI calls. One possible way to implement this common error handling methodology is to provide error-handling routines for the application developer. Use of these common error-handling routines ensures that all application programmers handle MQSeries errors in the same way and do not have to write their own error handling routines.

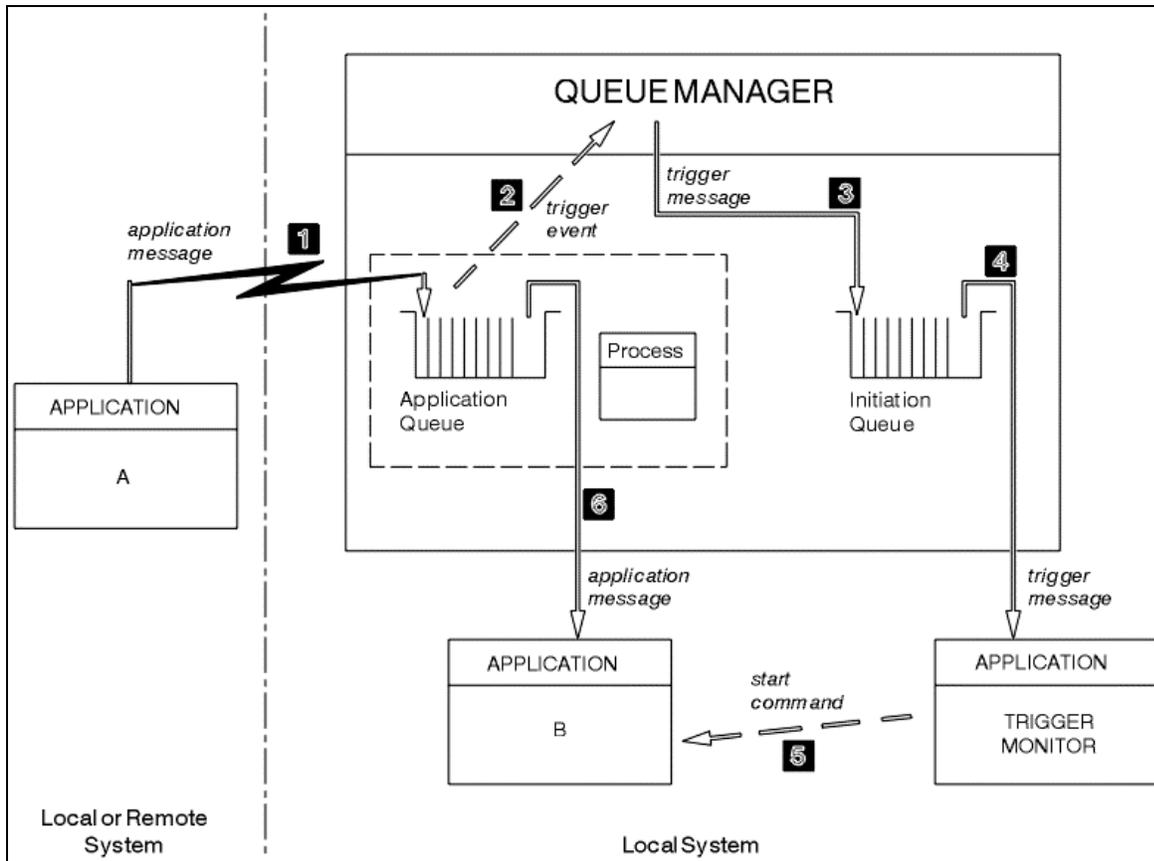
Note: Refer to Section 8.1 - Reusable EAI Functions: EAI Common Log Component for additional information regarding common error handling. The EAI Common Log Component interface enables applications to record events to local and centralized logs.

5.6 Triggered queues and applications

5.6.1 Designing MQSeries Applications

Some MQSeries applications that serve queues run continuously, and 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*.



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 relevant only to trigger type FIRST and DEPTH.
2. The queue manager checks to see if the conditions are met under which it has to generate a trigger event. If so, a trigger event is generated. Information that is held within the associated process definition object is used when creating the trigger message.
3. The queue manager creates a trigger message and puts it on the initiation queue associated with this application queue, but only if an application (trigger monitor) has the initiation queue open for input.
4. The trigger monitor retrieves the trigger message from the initiation queue.
5. The trigger monitor issues a command to start application B (the server application).
6. Application B opens the application queue and retrieves the message.

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 since the queue is already being served.
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. When using triggering for channels, you are recommended to use trigger type FIRST or DEPTH.

Each adapter created for FSA utilized triggering. The MQSeries object definitions can be seen by viewing each system script file contained in the Clearcase repository. Specifically, you want to look for the objects with the “trigger” attribute.

5.6.2 Starting MQSeries Applications

Trigger messages created because of trigger events that are not part of a unit of work are:

- put on the initiation queue,
- put outside any unit of work, with no dependence on any other messages
- available for retrieval by the trigger monitor immediately

Trigger messages created because of trigger events that are a 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.

Notes:

1. The queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist.

With triggering of type FIRST or DEPTH, trigger messages are made available even if the unit of work is backed out so that a trigger message is always available when the required conditions are met. An example is 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. Rather, 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. 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 to use the wait option with the MQGET call, setting the *WaitInterval* to a suitable value.

2. For local shared queues (that is, shared queues in a queue-sharing group) the queue manager counts committed messages only.

For FSA, the adapters were triggered on the trigger type of “FIRST”, the queues were then read until empty.

5.7 Application Usage Guidelines for Data Integrator

5.7.1 Data Integrator Components

Three components must be running on each machine that will be using Data Integrator: the Manager, Sender, and Receiver. Each component performs a different task related to the transfer of data using MQ Series. Each component performs common duties such as monitoring its input queue, logging events and sending status messages. When two components are located on different queue managers, the destination will be resolved to the appropriate transmission queue and sent to its destination.

5.7.1.1 e-Adapter Manager

The originating queue manager in a transaction acts as the Manager, and is responsible for starting and ending the transfer unit of work and managing all transfers. The Manager determines the appropriate Sender, and the message is sent to that Sender's input queue.

At the end of a transaction, the Manager returns one of the following completion codes: Request completed successfully, request failed, request expired, or request canceled. In the event of a failure, additional information, including error codes, is returned as well.

5.7.1.2 e-Adapter Sender

The source queue manager in a transaction acts as the Sender, and is responsible for transforming the data to be sent into MQSeries messages. Depending on the nature of the transaction, the Sender may send the message(s) directly to the specified Receiver or move them to a staging area. Once its work is complete, the Sender will report back to the Manager of the transaction.

5.7.1.3 e-Adapter Receiver

The destination queue manager in a transaction acts as the Receiver, and is responsible for processing the transfer request from the Sender and transforming the MQSeries messages into the target data. Once finished it will send a reply back to the Manager of the transaction.

5.7.2 Common Script Arguments

A data transfer using Data Integrator can be performed directly from the command line or typed into a command script. There are a number of arguments that can be used with the 'FTF' command to accomplish this. The most commonly used ones are listed below, along with a brief description.

5.7.2.1 Queue Manager Arguments

- **lqm** – Local Queue Manager – the queue manager from which the command is issued.
- **oqm** – Originating Queue Manager – the queue manager where the Manager will operate, defaults to the lqm.
- **sqm** – Source Queue Manager – the queue manager where the Sender will operate, defaults to the lqm.
- **dqm** - Destination Queue Manager – the queue manager where the Receiver will operate.

5.7.2.2 Source/Target File Arguments

- **spath** – Source Path – the full path and filename of the source file to be transmitted.
- **dpath** – Destination Path - the full path and filename of the destination file.

5.7.2.3 Process Arguments

- **immed** – this argument will force the transfer request to be processed synchronously between the Sender and Receiver. The Receiver will begin processing immediately upon receipt of the first message, instead of waiting until all messages have arrived (the default method). This technique requires less queue storage on the destination machine.

5.7.2.4 User Exit Arguments

These arguments are used when a process needs to be started immediately following the completion of a data transfer. A Unix script, specified using the user exit arguments, will be run once the transaction has finished. These arguments must be placed together and in order.

- **exit** – Exit Number – the exit number to be invoked.
- **exitdll** – DLLName – the DLL used to invoke the exit module.
- **exitentry** – Entry Point – the name of the function in the DLL that contains the exit module.
- **exitdata** – Data Value – the command-line argument that will be executed.

5.7.2.5 Data Specification Arguments

- **compress** – will cause the data to be compressed before it is sent.
- **pool** – Pool Name – the name of the data pool that will be used for transferring between the Sender and Receiver. This pool must be defined in the configuration file, and will default to the default pool specified in this file.

5.7.2.6 OS/390 Arguments

These arguments are only necessary when dealing with a target machine running OS/390.

- **blksize** – Block Size – the block size for the target file, it is usually specified .
- **lrecl** – Logical Record Length – the logical record length for the target file, it is usually specified.

5.7.2.7 Additional Script Arguments

Additional script arguments may be found in the e-Adapter Technical Reference.

5.8 Application Usage Guidelines for Data Integrator Status Utility

The Data Integrator product can be configured to generate status and diagnostic messages for each Data Integrator file transfer. These messages are in XML format and can be used to verify that file transfers occurred and assist in problem determination when a file could not be transferred.

5.8.1 Storing the status messages

When a file is transferred using Data Integrator, the related status messages are directed to an MQSeries queue. An MQSeries Integrator message flow reads messages from the queue and places them into an Oracle database. These messages are retained in the database awaiting later retrieval and analysis.

5.8.2 Retrieving the status messages

Messages are retrieved from the Oracle database via an MQSeries Integrator message flow. Requests for information are in XML format and they may contain one of several search criteria, such as a date/time range or a unique Data Integrator file transfer identifier. Using MQSeries Integrator as the retrieval mechanism allows messages to be requested from any platform that has MQSeries access to the EAI bus.

5.9 Application Usage Guidelines for MQSeries Integrator

5.9.1 Defining Messages

MQSeries Integrator provides a message brokering function that can transform messages from one format to another. The brokers that manage these transformations need to interpret the structure and content of the messages they receive to perform the full range of transformation functions available with MQSeries Integrator.

5.9.1.1 Message domains

The messages supported by MQSeries Integrator are of three broad types that are identified by a property of the message called the message domain:

- A message can be unstructured: its message domain must be set to BLOB.
- A message can be self-defining: its message domain must be set to XML.

Two additional domains are included in this category to support JMS messages: the domain JMSMap can be used for jms_map messages and the domain JMSStream can be used for jms_stream messages.

- A message can be predefined. Its message domain must be set to one of:
 - □ MRM
 - □ NEON

A predefined message has a logical structure and a physical structure:

- □ The logical structure of a predefined message is a tree structure that demonstrates the hierarchical relationships between the components of a message.

- □ The physical structure of a message, which is also referred to as its wire format, is just a string of bits and bytes. Without the logical structure, the physical structure (the bit-stream) has no intrinsic meaning.

5.9.1.2 Unstructured messages in the BLOB domain

An unstructured message must have a message domain of *BLOB*. It has no known (or defined structure). These messages can be processed and routed by MQSeries Integrator, but the manipulation that you can perform is very limited. You can perform some simple manipulation at the message level, and take other actions on the whole message.

5.9.1.3 Self-defining messages in the XML domain

A self-defining message must have a message domain of *XML*. It carries the information about its content and structure within the message. Its definition is not held anywhere else. When a self-defining message is received by the broker, it is handled by the XML parser, and a tree is created according to the XML definitions contained within that message. A self-defining message is also known as a *generic XML message*. It does not have a recorded format.

A self-defining message can be handled by every IBM-supplied message processing node. The whole message can be stored in a database, and headers can be added to or removed from the message as it passes through the message flow. The message can also be manipulated, constructed, and reformatted by nodes in the message flow, using a specialized form of standard database Structured Query Language (SQL). This specialized form is known as Extended SQL, or ESQL, and supports MQSeries Integrator processing of the message structure. This means that although you do not have to define the message structure to the Control Center, you do have to understand the definition to be able to construct valid ESQL for message manipulation.

5.9.1.4 Predefined messages in the MRM domain

A predefined message in the MRM message domain must have its message domain set to *MRM*. It must be defined to the *Message Repository Manager*, a component of the Configuration Manager. You can define messages to the MRM domain using the Control Center (Message Sets view). The MRM maintains these messages in the message repository. You can also predefine a message to the MRM in the XML message domain. If you define a message to the XML domain, you can use all the facilities available to MRM domain messages to manipulate and reference the message in the nodes within your message flows in the Control Center.

However, you are not expected to assign these message sets to a broker, nor to deploy them. Because the domain is set to *XML*, the XML parser is invoked by the broker and does not reference any external message definition. An MRM message can be handled by every IBM-supplied message processing node. The whole message, or parts of the message, can be stored in a database, and headers can be added to or removed from the message as it passes through the message flow. The message can be manipulated using ESQL defined within all message processing nodes that support manipulation (for example, compute and filter).

You can also transform any message in the MRM domain into any other format defined to the MRM using ESQL (in most cases, just one line of ESQL). This includes code page and encoding conversion. It provides the significant benefit that data conversion exists in MQSeries and therefore applications are not required to provide this function.

Messages with a message domain of MRM have three other characteristics for further classification:

- Message format

Three message formats are supported by the MRM:

- □ A message can have a message format of CWF (Custom Wire Format).

These messages are MRM representations of legacy data structures created in the C or COBOL programming language, and imported into the MRM using the Control Center facilities. You can also create new messages using this format.

- □ A message can have a message format of PDF.

This is a specialized format used predominantly in the finance industry. It does not have any connection with the Portable Document Format defined by Adobe (also known as PDF). If you already use messages of this format, you can continue to use them and process them by specifying this format in the definitions.

- □ A message can have a message format of XML.

These messages are represented as XML documents. They conform to an XML DTD (Document Type Definition) that can be generated by the Control Center for documentation purposes.

- Message set

This identifies the message set to which each message belongs. This is specified as the message set *identifier*, not the message set name. When you define a message in the MRM message domain, you must define a message set that contains it. A message set can contain one or more related messages.

- Message type

The message type identifies the message definition within the set. It is the unique identifier for each message of this particular content and format.

5.9.1.4.1 Predefined messages in the NEON domain

A predefined message in the NEON message domain must have its message domain set to *NEON*. It must be defined using the MQSeries Integrator Version 1 graphical utilities that are supplied with MQSeries Integrator Version 2. You can create new messages and use existing messages defined to the NEON domain. A NEON message can be handled by every IBM-supplied message processing node. The whole message can be stored in a database, and headers can be added to or removed from the message as it passes through the message flow. The NEONFormatter node can be used to transform a NEON message. No other node can manipulate the message contents.

5.9.2 Designing Message Flows

A message flow is a sequence of operations on a message, performed by a series of message processing nodes. The actions are defined in terms of the message format, its content, and the results of individual actions along the message flow. MQSeries Integrator includes a range of message processing nodes, called primitives, that provide most of the function that you will need in most situations. A message flow and the message processing nodes it contains describes the transformation and routing applied to an incoming message to transform it into outgoing messages. These actions form the rules by which the message is processed. A message flow can also be made up of a sequence of other message flows, that

are joined together. This function allows you to define a message flow containing a specific sequence of message processing nodes, and reuse that message flow in other message flows wherever that action is needed. When you complete the creation of your message flow, you can assign it for execution to one or more brokers. When you do this, the message flow must be operationally complete. That is, it must contain at least one MQInput node (one of the primitives). Most message flows will also contain at least one MQOutput or one Publication node, although this is not required (both of these nodes are also primitives).

5.9.2.1 Message flows and units of work

A message flow is transactional. You can define your message flows to perform all processing within a single unit of work. Therefore the receipt of every message by the input node, and the database operations performed as a result of that message being received and processed by the message flow, are coordinated. If an error occurs within a transactional message flow, the transaction is rolled back and the message will be handled according to normal error handling rules. You can also define a message flow to work outside of a unit of work if you do not want this support.

5.9.2.2 Parallel processing of message flow instances

When you define, assign, and deploy a message flow, the broker automatically starts an instance of the message flow for each input node (one or more). This is the default behavior. Each instance retrieves a message from the input node, and runs in parallel with other instances that retrieve a message from other input nodes. If you want to further increase the throughput of this message flow, you can set a property of the assigned message flow that defines how many additional instances are to be started by the broker for that message flow. You can set properties of the input node to exercise control over the order in which messages are processed.

You can also increase message flow throughput by assigning more than one copy of the message flow to the same broker. However, this is only appropriate if the message order is not important, because the multiple copies of the message flow are handled independently by the broker, with no correlation between them. Therefore, if more than one copy of the same message flow is active within the broker, each copy can be processing a message at the same time, from the same queue. It is possible for the processing time of a message flow to vary, and multiple message flows accessing the same queue could therefore read the messages from the queue in a random order. Also, the order of messages produced by the message flows might not correspond to the order of the original messages. You can influence the order in which the input node removes messages from the queue (using the Order Mode property). You are therefore recommended to increase the instances of a single copy of the message flow if you want to increase throughput and parallel processing but wish to have control over the message order.

5.9.2.3 Transformation

Most enterprises have applications that have been developed over many years, on different systems, using different programming languages, and different methods of communication. Standard message queuing technology can bridge differences like these, but applications still need to be aware of, and negotiate, the format in which the messages flow. With MQSeries Integrator the knowledge of each application is stored just once in the broker and each message is translated into the receiving application's format. Because the broker knows the requirements of each application, it can transform the message to the correct format without the sending or receiving application needing any modification.

A message flow can completely rebuild a message, convert it from one format to another (whether format means order of fields, byte order, language, and so on), remove content from the message, or introduce specific data into it.

5.9.2.4 Intelligent routing

Intelligent routing encapsulates business knowledge of how information should be distributed between sending and receiving applications throughout the enterprise. This knowledge is stored in the broker as a set of rules that are applied to each message as it passes through the broker. Routing is independent of the requirement for message transformation, although you will usually define sets of rules (as message flows) that combine the two in some way. Messages are distributed according to criteria applied to the values of fields within the message.

You can also establish a more dynamic routing option by building additional routing information into the message when it is processed. Optional sets of rules are set up to receive messages according to values (destinations) set into the message. You can establish these rules such that a message is processed by one or more of the optional sets of rules, in an order determined by the added message content. You can create, modify, and use these rules to develop a very flexible approach to the distribution of information. New ideas and requirements can be stated clearly, and turned into new or changed rules in the broker, and your business goals are met. You don't have to rework your applications. Your business processes range from the simple to the very complex. You can create rules to cover every case, building new rules, and reusing and combining existing ones to develop even the most complex solution.

5.9.2.5 Enriching message content

When a message is processed by a message flow, it is possible to update and add to the message content. This allows you to add value between sender and receiver in any way you choose. A typical way in which you can enhance the message content is by adding data from a database. This can be done by appending fields to the message, or merging information from the two sources. For example, a new field value can be calculated using the database information.

5.9.3 Using Message Processing Nodes

Message flow nodes are the key components of a message flow. A message processing node is a stand-alone procedure defined within a message flow that receives a message, performs a specific action against it, and outputs zero or more messages as a result of the action it has taken. This section describes the types of nodes, using the primitives included in MQSeries Integrator to illustrate the function they provide. You can create additional message processing nodes to provide enhanced or replacement function if you choose, except where noted.

5.9.3.1 MQSI Primitives

The MQSeries Integrator 2.0.1 Control Center provides a number of message flow nodes. The table below identifies the message flow nodes supplied with MQSeries Integrator, which are known as the IBM Primitives.

IBM Primitive	Function
Check node	Compares the format of an incoming message with a predefined message specification.

Compute node	Derives an output message from an input message and, optionally, from data taken from a external database. A computation can be applied to each element of the input message before the output message is constructed.
Database node	Combines database operations with message processing.
DataDelete node	Deletes one or more rows from a database table.
DataInsert node	Inserts one or more rows in a database table.
DataUpdate node	Updates the contents of one or more rows in a database table.
Extract node	Derives an output message from the fields in an input message.
Filter node	Evaluates an input message against an ESQL expression.
FlowOrder node	Determines a specific order for the processing of a message.
Input Terminal	Provides an in terminal for an embedded message flow.
Label node	Receives a message from a RouteToLabel node.
MQInput node	Reads MQSeries messages from a specified message queue.
MQOutput node	Writes MQSeries messages to a specified message queue.
MQReply node	Sends a response message to the originator of the message that caused this message flow to be invoked.
NEONFormatter node	Transforms an input message using the NEON Formatter engine.
NEONRules node	Passes an input message to the NEON Rules engine for evaluation.
Output Terminal	Provides an out terminal for an embedded message flow.
Publication node	Publishes a message to subscribers.
ResetContentDescriptor node	Reparses the bit stream of an input message.
RouteToLabel node	Routes a message to one or more specific destinations that are identified in the message.
Throw node	Throws an exception within a message flow.
Trace node	Generates a trace record.
TryCatch node	Catches any exceptions that are thrown by nodes further on in the message flow.
Warehouse node	Stores message data in a data repository.

5.9.3.2 Common node characteristics

Every message processing node has a fixed number of input points and output points. These points are known as terminals. Each node normally has one input terminal (on which it receives messages), and multiple output terminals to handle a variety of situations. Output terminals are defined according to the

characteristics of the individual node. For example, a filter node has true, false, failure, and unknown output terminals.

A Connector joins an output terminal of one node to an input terminal of the next node in the message flow. You can leave an output terminal unconnected, or you can connect a single output terminal to more than one target node. After a node has finished processing a message, the connectors defined from the node's output terminals determine which node(s), process the message next. If a node has more than one output terminal connected to a target node, the node determines the order in which the different execution paths are executed. If a single output terminal has more than one connector to a target node, the broker determines the order in which the different execution paths are executed. You cannot change the order of processing determined by the node or broker.

A node does not always produce an output message for every output terminal. Often it produces one output for a specific terminal depending on the message received. For example, a filter node will typically send a message on either the true terminal, or the false terminal, but not both. When the processing determined by one connector has been completed, the node issues the message again to the next connector, until all possible paths have been completed. Updates to a message are never propagated to previously executed nodes, only to nodes following the node in which the update has been made. The message flow can only accept a new message for processing when all paths through the message flow (that is, all connected nodes from all output terminals, as appropriate) have been completed.

5.9.3.3 Input and output nodes

Some message nodes have special characteristics. They define points in the message flow to which clients send messages (input nodes or MQInput), or from which clients receive messages (output nodes or MQOutput). These special nodes represent MQSeries queues. Client applications interact with these nodes by putting messages to, or getting messages from, these queues. A message flow has a set of (one or more) input nodes to which senders can post their messages, and a set of output nodes from which receivers can pick up messages.

If a message is being processed under transactional control, the output node only puts the message to the destination queue when all processing by the message flow has been successfully completed, unless the output node is set up to put the message outside the global (message flow) transaction. Before you can use a message flow, the input nodes must be associated with queues that represent the sources of messages. An output node must also be associated with a queue in most cases. However, you can set an output node property that causes the node to put the message to every queue in a *destination list*, which is contained within the message itself. You must use the primitive MQInput node for every message flow input node. You cannot replace it with one of your own. You can replace the output node if you choose.

Publication nodes are a special type of output node that use the queues identified by current subscribers whose subscriptions match the characteristics of the current message. Subscribers provide the identity of the queue on which they want to receive all matching publications.

5.9.3.4 Processing messages

All nodes other than the input and output nodes receive an input message from the previous node in the message flow and transform it into zero or more output messages to be made available to the next node (or nodes) in the message flow. Messages passing between nodes are not put to an intermediate queue: each message is held in local memory. These nodes can perform any kind of processing on a message. For example, they can:

- Reformat the message (NEONFormatter).

- Transform the message (Compute).
- Subset the data within the message (Extract).
- Route the message to one or more targets (NEONRules).
- Archive the message in a message warehouse (Warehouse).
- Update database information from the message content (Database).

5.9.3.5 Error handling

All primitive message processing nodes have a failure output terminal, to which a message is transferred if an error is detected within the node. If the failure terminal is not connected to a target node, an exception is generated and propagated back towards the MQInput node:

- If a TryCatch node is encountered before the exception reaches the MQInput node, the flow of control proceeds down the catch terminal. The message that is propagated through the catch terminal is the message originally received by the TryCatch node: any changes made to the message by later nodes in the message flow are not preserved. However, any external processing (for example, updates to a database through a Database node) are preserved. It is not possible to rollback these database updates from within the message flow.

Before the TryCatch node passes on the message to the node connected to the catch terminal, it adds the exception information to the ExceptionList item in the message tree. Existing information in the ExceptionList field in the message is written to the local error log, and then overwritten with the new exception information.

- If the message reaches the input node:
 - If the input node's catch terminal is connected to another node, the message is propagated to that node. In this case, an error is not recorded in the local error log.
 - If the input node's catch terminal is not connected, and the message is being processed under transactional control, the message is returned to the input queue. An error is recorded in the local error log. The MQInput node will then read the message again for retry. It first checks to see if the backout count for this message has now exceeded the backout threshold:

- If the backout count has not exceeded the threshold, the message processing is retried.
- If the backout count has exceeded the threshold, and the failure terminal is connected to another node, the message is propagated to that node.

If the failure terminal is not connected, the message is put on the backout queue, if one is defined for this input queue, or the queue manager's dead-letter queue (DLQ), if a backout queue does not exist.

If the queue manager does not have a DLQ defined, the message is left on the input queue. (If the broker's queue manager has been created by the create broker command mqsicreatebroker, a DLQ has been defined and enabled for this queue manager.)

- If the catch terminal is not connected and the message is not being processed under transactional control, the message is discarded.

You can provide a minimum level of error handling within every message flow you define if you choose. This minimum level might include:

- Define a dead-letter queue (DLQ) on the broker's queue manager (or use the default supplied DLQ).
- Change the queue manager's attributes to use this DLQ.

5.9.3.6 Adding or enhancing message processing nodes

MQSeries Integrator provides an external interface that allows you to add new capabilities to the broker by implementing new node types. The interface comprises a set of calls implemented in the C language. These calls are of two kinds:

- Calls that the broker makes to the node, for example to initialize the node.
- Calls that the node makes to the broker, for example, to inquire about the content of the message being processed.

5.9.4 Assigning and Deploying Resources to Brokers

The complete process of assigning resources and deploying them to a broker is too lengthy to be included in this guide. The steps are completely documented in Chapter 7 and 8 in the MQSeries Integrator "Using the Control Center" Manual. The manual can be found at the following url:

<http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/mqsiv202.html>

6 APPLICATION CONNECTIVITY (ADAPTERS AND BRIDGES)

The EAI application will be interfacing with several systems. The interfaces between EAI and other systems may require special mechanisms called adapters and bridges.

An adapter or a bridge is a piece of software that moves data between a message on a queue and an application or environment. Adapters handle data inbound-to and outbound-from the application or environment.

6.1 MQSeries Application Adapter

MQSeries provides a mechanism for assured delivery of messages, which can be sent even when the target is disconnected. It can be used to distribute work around a large number of disparate systems in an environment where trying to propagate transactional two-phase commit is not practical.

6.2 Adapter Classifications

6.2.1 Type of Message

Adapters may be classified by the type of message that will be processed:

- **Request/Reply**
An incoming XML request message from the front-end is posted to the back-end. In response, the adapter always synchronously routes the back-end results in the form of a valid XML document.
- **Fire & Forget**
An incoming XML request from the front-end is posted to the back-end and no response is required.
- **Notification**
The adapter routes an incoming message from the back-end to the front-end in the form of a valid XML message. This may be the reply to a message received.

All adapters written for FSA were of the Request/Reply type.

6.2.2 Interface Type

Adapters may be classified by interface type:

- **Java Object** - Creates Java objects that corresponds to the XML message elements.
- **Host structure** -
 1. Converts data from valid XML values to valid host values. Uses tables for simple cases and code for complex transformations.
 2. Creates host objects that correspond to the host data structures and maps the values from the XML objects to the host objects
- **XML Message** – The input data and the output data are both in XML format. The adapter may add the standard header and perform other functions, but does not need to transform the message

6.3 MQSeries-CICS/ESA Bridge

The MQSeries-CICS/ESA Bridge enables an application, not running in a CICS environment, to run a program or transaction on CICS/ESA and get a response back. This non-CICS application can be run from any environment that has access to a MQSeries network that encompasses MQSeries for MVS/ESA.

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

A transaction is a CICS transaction designed to run on a 3270 terminal. This transaction can use BMS or TC commands. It can be conversational or part of a pseudo conversation. It is permitted to issue syncpoints.

6.3.1 Using the CICS Bridge

Only FSA applications that use a CICS commarea to communicate can utilize the CICS Bridge; any applications that use terminal I/O CICS commands can use the CICS DPL Bridge.

The CICS Bridge allows an application to run a single CICS program or a 'set' of CICS programs (often referred to as a unit of work). The adapter written for the CPS system utilizes the CICS Bridge. For more information on the CPS adapter please reference the Technical Specification document. The CICS Bridge works with the application that waits for a response to come back before it runs the next CICS program (synchronous processing). It also works with the application that requests one or more CICS programs to run, but doesn't wait for a response (asynchronous processing).

The CICS Bridge also allows an application to run a 3270-based CICS transaction, without knowledge of the 3270 data stream. The CICS Bridge uses standard CICS and MQSeries security features. It can be configured to authenticate, trust, or ignore the requestor's user ID.

With this flexibility, there are many instances where the CICS Bridge can be used. For example,

- To write a new MQSeries application that needs access to logic or data (or both) that reside on your CICS server.
- Enabling a Lotus Notes application to run CICS programs.
- To be able to access CICS applications from a MQSeries Java client application or a web browser using the MQSeries Internet gateway.

6.3.2 CICS Bridge at Work

This section explains how the CICS Bridge works and the options available when deciding what level of security to use.

With respect to system setup, note the following:

- Ensure that the MQSeries-CICS adapter is enabled.
- The CICS Bridge requires that both MQSeries and CICS are running in the same MVS image.
- The MQSeries request queue must be local to the CICS Bridge, however the response queue can be local or remote.
- The CICS bridge tasks must run in the same CICS as the bridge monitor. The user programs can be in the same or a different CICS system.

6.4 Running CICS DPL programs

Data necessary to run the program is provided in the MQSeries message. The bridge builds a COMMAREA from this data, and runs the program using EXEC CICS LINK.

The following shows the components and data flow to run a CICS DPL program.

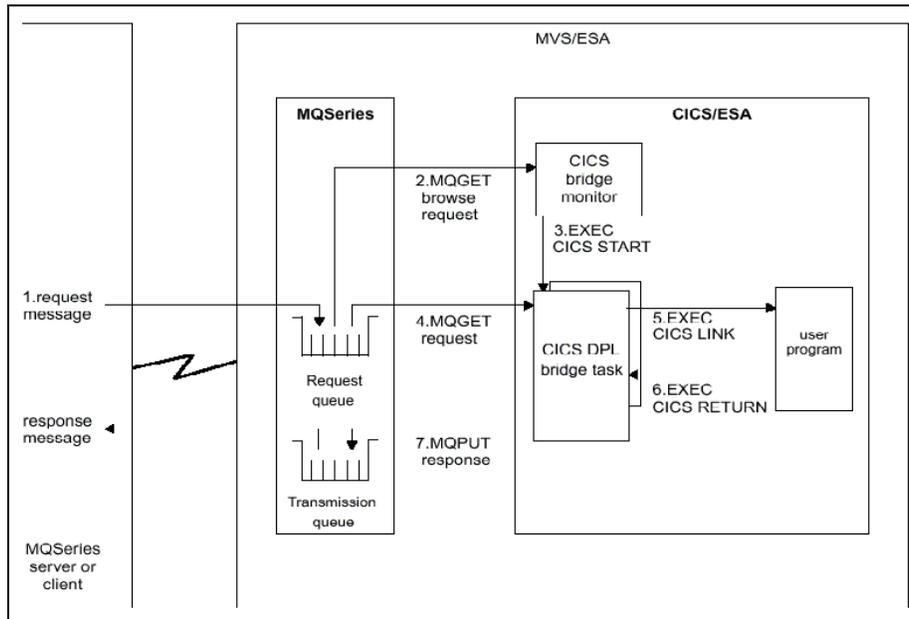


Figure 1 – CICS DPL Transaction

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

1. A message, with a request to run a CICS program, is put on the request queue.
2. The CICS Bridge monitor task, which is constantly browsing the queue, recognizes that a 'start unit of work' message is waiting (*CorrelId=MQCI_NEW_SESSION*).
3. Relevant authentication checks are made, and a CICS DPL Bridge task is started with the appropriate authority.
4. The CICS DPL Bridge task removes the message from the request queue.
5. The CICS DPL Bridge task builds a COMMAREA from the data in the message and issues an EXEC CICS LINK for the program requested in the message.
6. The program returns the response in the COMMAREA used by the request.
7. The CICS DPL Bridge task reads the COMMAREA, creates a message, and puts it on the reply-to queue specified in the request message. All response messages (normal and error, requests and replies) are put to the reply-to queue with default context.
8. The CICS DPL bridge task ends.

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

6.4.1 Running CICS 3270 transactions

Data necessary to run the transaction is provided in the MQSeries message. The CICS transaction runs as if it has a real 3270 terminal, but instead uses one or more MQSeries messages to communicate between the CICS transaction and the MQSeries application. Unlike traditional 3270 emulators, the bridge does not work by replacing the VTAM flows with MQSeries messages.

Instead, the message consists of a number of parts called vectors, each of which corresponds to an EXEC CICS request. Therefore, the application is talking directly to the CICS transaction, rather than via an emulator, using the actual data used by the transaction (known as application data structures or ADSs).

The following shows the components and data flows to run a CICS 3270 transaction.

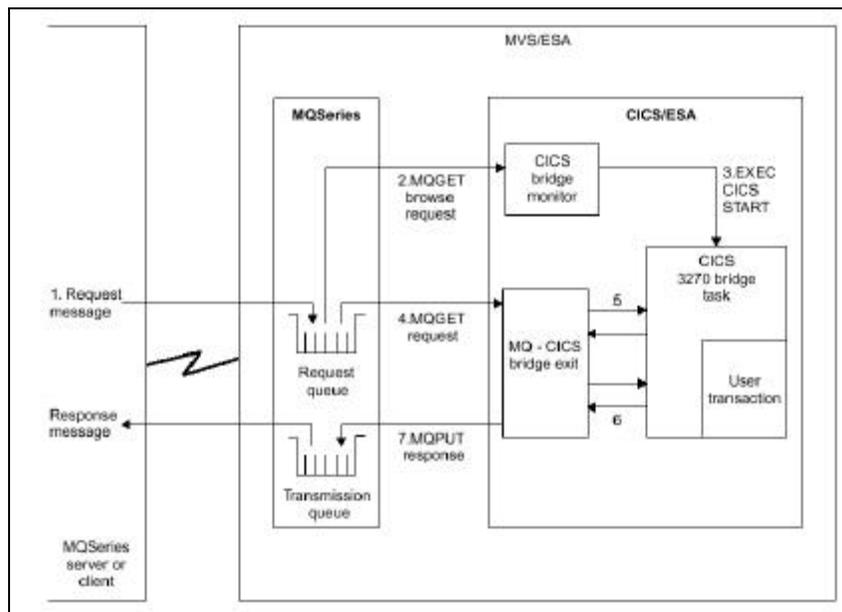


Figure 2 – CICS 3270 Transaction

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

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

8. The CICS 3270 bridge task ends.

A traditional CICS application usually consists of one or more transactions linked together as a pseudo conversation. In general, the 3270 terminal user entering data onto the screen and pressing an AID key starts each transaction. This model of application can be emulated by a MQSeries application. A message is built for the first transaction, containing information about the transaction, and input vectors. This is put on the queue.

The reply message will consist of the output vectors, the name of the next transaction to be run, and a token that is used to represent the pseudo conversation. The MQSeries application builds a new input message, with the transaction name set to the next transaction and the facility token set to the value returned on the previous message. Vectors for this second transaction are added to the message, and the message put on the queue. This process is continued until the application ends.

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

7 APPLICATION INTEGRATION EXAMPLES

This section contains representative examples of interfaces using each of the EAI middleware products thereby providing guidance on integrating FSA applications with the EAI Core Architecture.

7.1 Example of an interface using Data Integrator

The Post-Secondary Education Participants System (PEPS) to eCB (eCampus Based) and COD (Common Origination and Disbursement) interface ensures accuracy of individual, Campus Based accounting transactions. This interface uses Data Integrator product in conjunction with the EAI Bus to enable file transfer between these systems.

There are a series of core steps that every application team must follow to build an interface to the EAI Bus. These steps are supplemented with Data Integrator steps that enable bulk file transfers between different applications via the EAI Bus. These may include additional configuration or custom development, specifying source and target file system location, etc. The steps specific to building a Data Integrator interface are highlighted and italicized.

#	Steps
1	EAI conducts a core kick-off meeting with the Application Team.
2	Application Team completes the EAI Questionnaire. (Please see Appendix C: EAI Core Services Questionnaire)
3	EAI Team reviews the EAI Questionnaire with the Application Team as input for the Interface Estimate.
4	EAI Team completes the Interface Estimate.
5	EAI Team and Application Team review the Interface Estimate.
6	EAI Team completes the Interface Partner Agreement (IPA).
7	EAI Team and Application Team sign the IPA.
8	EAI Team completes the design documentation. <ul style="list-style-type: none"> a. EAI Team completes Interface Control Document (ICD) b. EAI Team completes Internal Interface Design (IID)
9	EAI Team installs the MQSeries Infrastructure on the application system in the development environment. <ul style="list-style-type: none"> d. MQSeries software is installed and configured on the application system e. Connectivity test is performed between the EAI Bus and the application system. f. Queue Managers are defined on the application system. g. Channels are defined between the application system and the EAI Bus. h. Local Queues are defined on the application system. i. Message delivery test is performed between the EAI Bus and the application system using sample programs provided by the MQSeries software.
10	<i>EAI Team installs the Data Integrator Infrastructure on the application system in the development environment.</i> <ul style="list-style-type: none"> <i>a. Data Integrator software is installed and configured on the application system</i>

	<ul style="list-style-type: none"> <i>b. Connectivity test is performed between the EAI Bus and the application system.</i> <i>c. Queue Managers are defined for Data Integrator on the application system.</i> <i>j. Channels are defined for Data Integrator between the application system and the EAI Bus.</i> <i>d. Local Queues are defined for Data Integrator on the application system.</i> <i>e. File delivery test is performed between the EAI Bus and the application system using sample Data Integrator scripts developed by the EAI team.</i>
11	<i>EAI team configures a directory monitor process on PEPS. The directory monitor process (delivered with the Data Integrator software) polls a directory looking for the existence of a new file matching certain naming convention criteria. Data Integrator scripts are developed by the EAI team to send the PEPS file eCB and COD via the EAI Bus.</i>
12	<p><i>EAI Team develops Data Integrator script to send the PEPS file to eCB and COD via the EAI Bus:</i></p> <ul style="list-style-type: none"> <i>a. Specify source and target Queue manager</i> <i>b. Specify source and target file system path</i> <i>c. Specify pre or post processing after delivery of the file</i>
13	<i>EAI team develops a Java transformation program that extracts the PEPS delta file. Once the file arrives on the EAI Bus, the transformation program is triggered</i>
14	<i>Data Integrator is configured allow the same file to be delivered to multiple locations. For instance, after the PEPS file is received on the EAI Bus, but before the transformation program is called, the file is sent to eCB, another Trading Partner.</i>
15	The Data Integrator scripts and transformation code are reviewed and baselined in ClearCase repository by the EAI team.
16	EAI Team installs the MQSeries and <i>Data Integrator</i> Infrastructure on the application system in the test environments (see steps 9 & 10).
17	EAI Team tests the application interface.
18	EAI Team logs testing Issues and reviews/resolves them with the Application Team.
19	<p>Using ClearQuest, the EAI Team coordinates the deployment procedures and timeline with the application host.</p> <ul style="list-style-type: none"> a. EAI Team installs the MQSeries and <i>Data Integrator</i> Infrastructure on the application system in the production environment (see steps 9 &10). b. Code/scripts is deployed to production c. Application Team performs Production Readiness Test

7.2 Example of an interface using MQSeries Integrator

The eCampus Based (eCB) to Financial Management System (FMS) interface ensures consistent and consolidated financial reporting of Campus- Based awards. This interface uses MQSeries Integrator, Data Integrator, and custom adapter in conjunction with the EAI Bus to enable files to be converted into messages and eventually inserted into the FMS database.

There are a series of core steps that every application team must follow to build an interface to the EAI Bus. These steps are supplemented with Data Integrator steps that enable bulk file transfers, MQSeries Integrator to perform file to message transformation, and custom adapter to feed the messages into the database via the EAI Bus. This example will only highlight the MQSeries Integrator steps. These may include additional configuration or custom development, creating MQSeries Integrator messages, flows, etc, etc. The steps specific to building MQSeries Integrator interfaces are highlighted and italicized.

#	Steps
1	EAI conducts a core kick-off meeting with the Application Team.
2	Application Team completes the EAI Questionnaire. (Please see Appendix C: EAI Core Services Questionnaire)
3	EAI Team reviews the EAI Questionnaire with the Application Team as input for the Interface Estimate.
4	EAI Team completes the Interface Estimate.
5	EAI Team and Application Team review the Interface Estimate.
6	EAI Team completes the Interface Partner Agreement (IPA).
7	EAI Team and Application Team sign the IPA.
8	EAI Team completes the design documentation. <ul style="list-style-type: none"> a. EAI Team completes Interface Control Document (ICD) b. EAI Team completes Internal Interface Design (IID)
9	<i>EAI Team creates high-level MQSeries Integrator process flow diagrams based on the business requirements (e.g., transformation, routing, and etc)</i>
10	EAI Team installs the MQSeries Infrastructure on the application system in the development environment. <ul style="list-style-type: none"> a. MQSeries software is installed and configured on the application system b. Connectivity test is performed between the EAI Bus and the application system. c. Queue Managers are defined on the application system. d. Channels are defined between the application system and the EAI Bus. e. Local Queues are defined on the application system. f. Message delivery test is performed between the EAI Bus and the application system using sample programs provided by the MQSeries software.
11	EAI Team installs the Data Integrator Infrastructure on the application system in the development environment. <ul style="list-style-type: none"> a. Data Integrator software is installed and configured on the application system b. Connectivity test is performed between the EAI Bus and the application system. c. Queue Managers are defined for Data Integrator on the application system. d. Channels are defined for Data Integrator between the application system and the EAI Bus. e. Local Queues are defined for Data Integrator on the application system. f. File delivery test is performed between the EAI Bus and the application system using sample Data Integrator scripts developed by the EAI team.

12	EAI Team develops Data Integrator script to send the eCB UTCL file to FMS via the EAI Bus: a. Specify source and target Queue manager b. Specify source and target file system path c. Specify pre or post processing after delivery of the file
13	<i>EAI team develops SQL insert statements based on the input UTCL file that is used to load the UTCL data into FMS. The insert statements are associated with MQSeries message through MQSeries Integrator</i>
14	<i>EAI Team defines the message set(s) and elements in MQSeries Integrator.</i>
	<i>EAI Team defines the message flow(s) in MQSeries Integrator that performs the following transformation on the UTCL file:</i> a. <i>Performs any necessary FMS validations and transformation</i> b. <i>Creates an SQL statement to place the message in the appropriate table</i> c. <i>Places the SQL statement on MQSeries queues bound for the FMS MQSeries queue manager</i>
15	<i>EAI team develops a custom FMS-MQSeries Adapter that will execute the SQL statement contained in the MQSeries Message against the FMS database. All insert SQL statements are performed in a single transaction</i>
16	The Data Integrator scripts and <i>MQSeries Integrator flows</i> are reviewed and baselined in ClearCase repository by the EAI team.
17	EAI Team installs the MQSeries, Data Integrator, and <i>MQSeries Integrator Infrastructure</i> on the application system in the test environments (see steps 9 & 10).
18	EAI Team tests the application interface.
19	EAI Team logs testing Issues and reviews/resolves them with the Application Team.
20	Using ClearQuest, the EAI Team coordinates the deployment procedures and timeline with the application host. a. EAI Team installs the MQSeries, Data Integrator, and MQSeries Integrator Infrastructure on the application system in the production environment (see steps 9 & 10). b. Adapter code scripts are deployed to production c. <i>The Message Sets and Flows are deployed to production</i> d. Application Team performs Production Readiness Test

7.3 Custom EAI Adapter

The eCampus Based (eCB) to Financial Management System (FMS) interface ensures consistent and consolidated financial reporting of Campus- Based awards. This interface uses MQSeries Integrator, Data Integrator, and custom adapter in conjunction with the EAI Bus to enable files to be converted into messages and eventually inserted into the FMS database.

There are a series of core steps that every application team must follow to build an interface to the EAI Bus. These steps are supplemented with Data Integrator steps that enable bulk file transfers, MQSeries

Integrator to perform file to message transformation, and custom adapter to feed the messages into the database via the EAI Bus. This example will only highlight the Custom EAI Adapter steps. These may include installation of development tools, shared libraries, psuedo code, etc. The steps specific to building EAI Custom Adapter interfaces are highlighted and italicized.

#	Steps
1	EAI conducts a core kick-off meeting with the Application Team.
2	Application Team completes the EAI Questionnaire. (Please see Appendix C: EAI Core Services Questionnaire)
3	EAI Team reviews the EAI Questionnaire with the Application Team as input for the Interface Estimate.
4	EAI Team completes the Interface Estimate.
5	EAI Team and Application Team review the Interface Estimate.
6	EAI Team completes the Interface Partner Agreement (IPA).
7	EAI Team and Application Team sign the IPA.
8	EAI Team completes the design documentation. <ul style="list-style-type: none"> a. EAI Team completes Interface Control Document (ICD) b. EAI Team completes Internal Interface Design (IID)
9	<i>Psuedo code are created as a supplement to the design document.</i>
10	<i>EAI Team designs the custom adapter by applying the existing EAI error handling framework or customize it to meet the business requirements</i>
11	EAI Team installs the MQSeries Infrastructure on the application system in the development environment. <ul style="list-style-type: none"> a. MQSeries software is installed and configured on the application system b. Connectivity test is performed between the EAI Bus and the application system. c. Queue Managers are defined on the application system. d. Channels are defined between the application system and the EAI Bus. e. Local Queues are defined on the application system. f. Message delivery test is performed between the EAI Bus and the application system using sample programs provided by the MQSeries software.
12	EAI Team installs the Data Integrator Infrastructure on the application system in the development environment. <ul style="list-style-type: none"> a. Data Integrator software is installed and configured on the application system b. Connectivity test is performed between the EAI Bus and the application system. c. Queue Managers are defined for Data Integrator on the application system. d. Channels are defined for Data Integrator between the application system and the EAI Bus. e. Local Queues are defined for Data Integrator on the application system. f. File delivery test is performed between the EAI Bus and the application system using sample Data Integrator scripts developed by the EAI team.

13	<p><i>EAI Team installs the following library or support pacs on the application system in the development environment.</i></p> <ul style="list-style-type: none"> <i>a. Application Message Interface (AMI) library – provides a layer of MQSeries interface options</i> <i>b. MQSeries client (MA88 support pac)</i> <i>c. Development compiler or engine (e.g., C/C++ compiler, JDK or JVM)</i>
14	<p><i>EAI Team develops the Custom EAI FMS adapter using AMI and the MQSeries client to perform the following actions based on design documents:</i></p> <ul style="list-style-type: none"> <i>a. Connect to a specified Queue Manager</i> <i>b. Put messages on a specified queue</i> <i>c. Get messages from a specified queue</i> <i>d. Listen for messages on a specified queue</i> <i>e. Disconnect from a specified Queue Manager</i> <i>f. Error handling</i> <i>g. Additional database access capabilities are developed to insert records into the FMS database</i>
15	<p>The Custom EAI FMS adapter code is reviewed and baselined in ClearCase repository by the EAI team.</p>
16	<p>EAI Team installs the MQSeries, Data Integrator, and MQSeries Integrator Infrastructure on the application system in the test environments (see steps 11-13).</p>
17	<p>EAI Team tests the Custom EAI FMS Adapter.</p>
18	<p>EAI Team logs testing Issues and reviews/resolves them with the Application Team.</p>
19	<p>Using ClearQuest, the EAI Team coordinates the deployment procedures and timeline with the application host.</p> <ul style="list-style-type: none"> <i>d. EAI Team installs the Data Integrator, MQSeries Integrator Infrastructure on the application system in the production environment (see steps 11-13).</i> <i>e. Code/scripts is deployed to production</i> <i>f. Application Team performs Production Readiness Test</i>

8 REUSEABLE EAI FUNCTIONS

Reusable EAI functions described in the following section are application services that can be utilized by applications integrated with the EAI Core Architecture. Additional reusable functions will be included as they are developed and deployed in future EAI Core Architecture efforts.

8.1 EAI Common Log Function

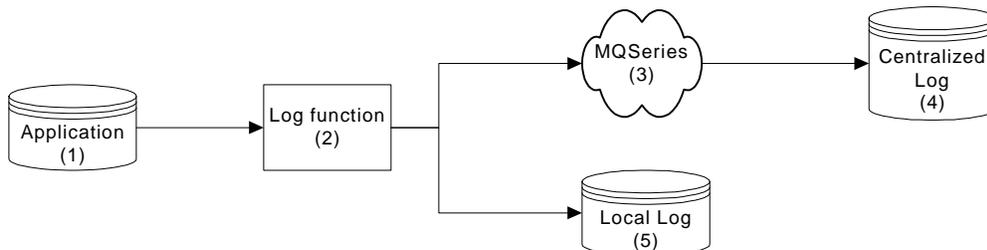
The following outlines design and implementation information required to utilize the EAI Common Log Function.

8.1.1 Interface Design Specification

<i>Interface Name:</i>	EAI Common Log Function
<i>Interface Type:</i>	Uni-Directional
<i>Interface Short Description:</i>	This interface enables applications to record events to the local and centralized logs.
<i>Source Application:</i>	Any
<i>Destination Application:</i>	Local and centralized logs.
<i>Functional Requirement References:</i>	Message Logging
<i>Related Interface Control Document:</i>	N/A
<i>Related Unit Test Document:</i>	TBD
<i>Other Related Interfaces:</i>	N/A

8.1.2 Interface Overview

Flow Diagram:



#	Name	Description
1	Application	The source application
2	Log Function	A library function that sends the log entry to the centralized and/or local logs.
3	MQ Series	The MQ Series transport mechanism.
4	Centralized Log	<i>The centralized log repository.</i>
5	Local Log	The local log file.

8.1.2.1 Detailed Technical Overview

An application (1) generates an event that it needs to record. The application will call the log function (2) according to the specified function signature. The log function creates a message. It then sends the message via MQ Series (3) to the Centralized Log (4). The log function also records the event on the local log (5).

8.1.2.2 Background EAI Logging Objectives

The “logging” framework will help standardize and simplify exception handling for FSA’s application teams. The standardized exception handling will also help reduce the possibility of uncaught exception scenarios.

An exception is a code or language construct that indicates when an unusual or unexpected error condition occurs in an application. Examples of exceptions are hardware, network, I/O, or memory problems. If an exception is “handled” in code, it can be dealt with gracefully and will not necessarily have to cause program termination. Exception handling provides a mechanism for writing robust, resilient code that is capable of dealing with the unexpected.

In addition to exception logging, the following categories were reviewed for consideration:

1. Performance Logging
2. Capture Service Level Agreement Metrics
3. Provide information for system tuning
4. Exception Logging
5. Provide clarity as to where the problem has occurred
6. Debugging/Tracing
7. Aid developers in development and testing
8. Score Card Logging
9. Provide overall transaction status; i.e file X was transferred from server A to server B
10. Alert Logging
11. Provide a mechanism to alert operations of a problem

Empirically it can be observed that information required when satisfying the varied logging requirements overlap. For example information required to “Alert” operations of a problem will also aide in “Problem Identification”.

8.1.2.3 Logging Thresholds Provided via EAI Logging facility

Each message logged within the framework has a severity. A masking of this value determines whether the Logger allows the message to continue to the destination.

The severities allowed within a message are:

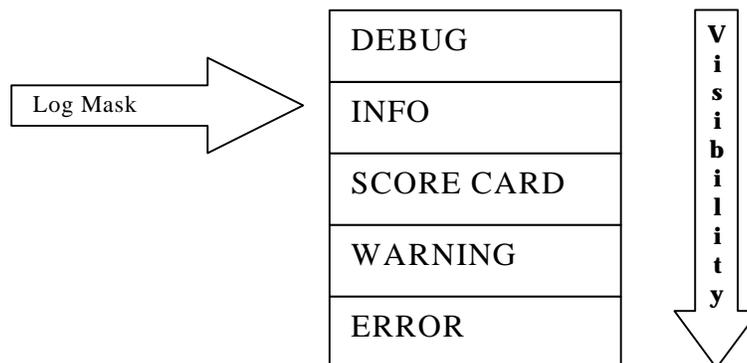
- **Debug Logging**
 - These are debugging messages usually placed by the programmers for the tracing and debugging purposes.

- **Informational Logging**
 - These are useful informational messages about what is occurring.

- **Score Card Logging**
 - Provides an overall status of the interface request; for example a datagram message that originates from NSLDS and terminates at COD would produce logging records for use as an audit mechanism (This feature is not currently implemented).

- **Warning Logging**
 - These messages warn that something abnormal has happened, but that the system will attempt to recover from it. These messages are usually used by programmers to show that something is starting to go wrong.

- **Error Logging**
 - These messages state that something abnormal has occurred, but that it is not severe enough to cause the system to fail in general. A specific task may fail and some users may get an error, but the system will keep going. Exceptions are generally logged at this level.
 - For example, if a Loggers mask is set to INFO, then any message that comes in with a severity that is below INFO will be sent on to the destination. A message that has severity DEBUG will be ignored. With this Log Mask, all info, warning, error, and fatal messages will show up at the destination.



8.1.2.4 Configuration Parameters

The following configuration parameters are required for message logging.

Description	Variable	Example
Environment Variable identifying location of configuration file and name of configuration file	EAILOG_PATH	C:\somedirectory\eailog\eailog.ini
Logging Threshold	LOGGING_THRESHOLD	0 (Debug)
Log file path name	LOGGING_PATH_NAME	C:\somedirectory\eailog\ eailog.yyyymmdd.txt
Remote queue	LOGGING_REMOTE_QUEUE	EAI.LOG

8.1.2.5 Component Model

The following function calls form the public interface of the Error logging subcomponent. These public interfaces will be published on the following platforms:

- Solaris
- HP-UX
- OS/390
- OpenVMS (no AMI)

1. For AMI enabled platforms, logging will be invoked via AMI's "Policy Handler Interface". "Policy Handler" eliminates the need for EAI BUS developers to invoke the logging facility for interactions that utilize MQSeries resources. "Policy Handler Post Transport Request Invocations" will be utilized to execute the logging mechanism.

Post-transport requests:

Post-MQBACK

Post-MQBEGIN

Post-MQCLOSE

Post-MQCMIT

Post-MQCONN

Post-MQCONNX

Post-MQDISC

Post-MQGET

Post-MQINQ

Post-MQOPEN
 Post-MQPUT
 Post-MQPUT1
 Post-MQSET.

2. C/C++ function interface:

```
long EAILog(
  long    lSeverity,
  char    *msgCode,
  char    *msgText,
  char    *interfaceid);
```

3. A Java interface (JNI).

```
public class EAIMSGLOG {
  public native int eaiLog( long severity,
                          String msgCode,
                          String msgText,
                          String interfaceid);
}

package gov.ed.eailog;
```

8.1.3 Design Assumptions

#	ASSUMPTIONS
1	The application is expected to call the log function whenever an event needs to be logged. At a minimum, informational logging will occur post-transport request.
2	The application is expected to call the log function according to the specified function signature.
3	Each application using this API is expected to install and configure MQ Series v5.2 (OpenVMS excluded).
4	EAI BUS File Transfer product includes a logging mechanism.
5	MQSeries 5.2 is not supported on OpenVMS, therefore all logging must be coded by the developer
6	Applications must use this mechanism within EAI adapters; at a minimum this will be called at the start and end of a adapter
7	All servers must have a C/C++ compiler
8	A COTS mqseries monitoring tool will be utilized

8.1.4 Design Dependencies

#	DEPENDENCIES
1	MQSeries 5.2
2	AMI Support Pack

8.1.5 Detailed Technical Design

Component Name: EAI Common Log Function

Related Interface Control Document: N/A

Technical Design Description: Applications will call the EAI Common Log Function according to the previously specified function signature:

Field Descriptions:

Message logging output file description:

	Description	Informational Logging	Exception Logging	Score Card Logging
Version	the version of the EAILogStruct definition being used; currently	1	1	1
Severity	the severity of the message being logged; valid values are: 00 – Debug 04 – Score Card 08 - Informational 12 – warning 16 – error			
msgCode	a freeform field for error codes; typically a MQSeries error code	blank	MQRC=9999	MQRC=9999
MsgText	a freeform field for the error description;	function/method name for informational messages	MQRC_XXX_XXX_XX X	MQRC_XXX_XXX_XX X
interfaceId	interface control document	Interface Control Id	Interface Control Id	Interface Control Id

instance	occurrence of a transaction	Hash value	Hash value	Hash Value
System	hostname of the system generating the error	Hostname	Hostname	Hostname
programId	Program id	Program_id	Program_id	Program_id
ReturnCode	specifies the status of the function upon completion; valid values are: 0 – success 1 – unable to log message			

Message logging functions generate a file delimited as follows:

<Version> <hostname> <program_name> <Instance> <date_time> <severity> <interface_id>
 <message text> <return code>

Error Handling:

#	Type	Reporting/Communication Method	Message
1	Error – continue processing	Return code	Unable to locate EAILOG environment variable.
2	Error – continue processing	Return code	Unable to send message to centralized log.
3	Error – continue processing	Return code	Unable to read control record information.
4	Error – continue processing	Return code	Unable to write to local log.

9 COMMITTING AND BACKING OUT UNITS OF WORK

This section describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work. The following terms, described below, are used in this section:

- Commit
- Back out
- Syncpoint coordination
- Syncpoint
- Unit of work
- Single-phase commit
- Two-phase commit

9.1 Committing and Backing Out

When a program puts a message on a queue within a unit of work, that message is made visible to other programs only when the program *commits* the unit of work. To commit a unit of work, all updates must be successful to preserve data integrity. If the program detects an error and decides that the put operation should not be made permanent, it can *back out* the unit of work. When a program performs a back out, MQSeries restores the queue by removing the messages that were put on the queue by that unit of work. The way in which the program performs the commit and back out operations depends on the environment in which the program is running.

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 by making the messages available to be retrieved by other programs. 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.

9.2 Syncpoint Coordination, Syncpoint, Unit of Work

Syncpoint coordination is the process by which units of work are either committed or backed out with data integrity. The decision to commit or back out the changes is taken, in the simplest case, at the end of a transaction. However, it can be more useful for an application to synchronize data changes at other logical points within a transaction. These logical points are called *syncpoints* (or *synchronization points*) and the period of processing a set of updates between two syncpoints is called a *unit of work*. Several MQGET calls and MQPUT calls can be part of a single unit of work. The maximum number of messages within a unit of work can be controlled by the DEFINE MAXSMSGS command on OS/390, or by the MAXUMSGS attribute of the ALTER QMGR command on other platforms. See the *MQSeries Command Reference* book for details of these commands.

9.3 Syncpoint Guidelines

A MQSeries application can specify on every put and get call whether the call is 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 calling MQPUT. For a get operation, use the MQGMO_SYNCPOINT value in the *Options* field of the MQGMO structure. If not explicitly choosing

an option, the default action depends on the platform. The syncpoint control default on OS/390 and Tandem NSK is 'yes'; for all other platforms, it is 'no'.

If a program issues the MQDISC call while uncommitted requests exist, an implicit syncpoint occurs, except on OS/390 batch with RRS. If the program ends abnormally, an implicit backout occurs. On OS/390, an implicit syncpoint occurs if the program ends normally without first calling MQDISC.

For MQSeries for OS/390 programs, 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).

9.3.1 Syncpoints in MQSeries for Windows NT, MQSeries on UNIX systems

Syncpoint support operates on two types of units of work: local and global. A *local* unit of work is one in which the only resources updated are those of the 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 databases, are also updated. MQSeries can coordinate such units of work itself or the units of work can also be coordinated by an external commitment controller such as another transaction manager.

For full integrity, a two-phase commit procedure must be used. Two-phase commit can be provided by XA-compliant transaction managers and databases such as IBM's TXSeries and UDB. MQSeries Version 5 products (except MQSeries for OS/390) can coordinate global units of work using a two-phase commit process.

9.3.2 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.

9.3.3 Global units of work

Use global units of work when needing to include updates to resources belonging to other resource managers. Here the coordination may be internal or external to the queue manager:

9.3.4 Internal syncpoint coordination

Queue manager coordination of global units of work is supported only on MQSeries Version 5 products except for MQSeries for OS/390. It is not supported in a MQSeries client environment. Here, the coordination is performed by MQSeries. To start a global unit of work, the application issues the MQBEGIN call.

As input to the MQBEGIN call, supply the connection handle (*Hconn*), which is returned by the MQCONN or MQCONNX call. This handle 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 resource 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, Oracle, and Sybase) are firstly all asked to prepare 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. 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, the call starts a unit of work and completes with a warning.

In these cases, the unit of work should 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), a 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. A 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* manual.

9.3.5 External syncpoint coordination

External syncpoint coordination occurs when a syncpoint coordinator other than MQSeries (e.g. CICS, Encina, and Tuxedo) has been selected. MQSeries on a UNIX system or MQSeries for Windows NT will register its interest in the outcome of the unit of work, with the syncpoint coordinator. This happens in order to commit or roll back any uncommitted get or put operations as required. The external syncpoint coordinator determines whether one- or two-phase commitment protocols are provided. When an external coordinator is used 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 programming interface provided by the syncpoint coordinator. An explicit call may, or may not, be required. If an explicit call is required, and the MQPUT call specifying the MQPMO_SYNCPOINT option is specified when a unit of work is not started, the completion code MQRC_SYNCPOINT_NOT_AVAILABLE is returned.

The syncpoint coordinator determines the scope of the unit of work. 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. This is known as a pending disconnect.

9.3.6 Interfaces to external syncpoint managers

MQSeries on UNIX systems 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 supply 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 *CAE Specification Distributed Transaction Processing: The XA Specification*, published by The Open Group.

In an XA configuration, MQSeries on UNIX systems and MQSeries for Windows NT fulfill the role of an XA Resource Manager. An XA syncpoint coordinator can manage a set of XA Resource Managers, and synchronize the commit or backout of transactions in both Resource Managers.

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 calls to update the resources managed by the resource managers associated with the current transaction.
4. The application requests that the syncpoint coordinator either commits or rolls back the transaction.
5. The syncpoint coordinator issues calls to each resource manager using two-phase commit protocols to complete the transaction as requested. The XA specification requires each Resource Manager to provide a structure called an *XA Switch*. This structure declares the capabilities of the Resource Manager, and the functions that are to be called by the syncpoint coordinator.

There are two versions of this structure:

MQRMIXASwitch

Static XA resource management

MQRMIXASwitchDynamic

Dynamic XA resource management

The structure is found in the following libraries:

mqmx.lib

Windows NT XA library for Static resource management

mqmenc.lib

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 (except Sun Solaris) 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. Also, consult the documentation provided by that coordinator to determine how to enable MQSeries to cooperate with the 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 or MQCONNX, and may be blank if the default queue manager is to be used.

9.4 MQSeries Syncpoint Calls for OS/390

MQSeries for OS/390 provides the MQCMIT and MQBACK calls. Use these calls in OS/390 batch programs to tell the queue manager that all the MQGET and MQPUT 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.

RRS use MQCMIT and MQBACK or SRRCMIT and SRRBACK as appropriate.

Note: SRRCMIT and SRRBACK are 'native' RRS commands, and are not MQI calls.

For backward 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.

9.5 MQSeries Syncpoint Calls on Windows NT and UNIX systems

The following products provide the MQCMIT and MQBACK calls:

- MQSeries for Windows NT
- MQSeries on UNIX systems

Use syncpoint calls in programs to tell the queue manager that all the MQGET and MQPUT 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.

9.6 Single-phase Commit

A *single-phase commit* process is one in which a program can commit updates to a queue without coordinating its changes with other resource managers.

9.7 Two-phase Commit

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. 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.

Transaction managers (such as CICS, IMS, Encina, and Tuxedo) can participate in two-phase commit, coordinated with other recoverable resources. This means that the queuing functions provided by MQSeries can be brought within the scope of a unit of work, managed by the transaction manager.

10 APPENDIX A: REFERENCE MATERIAL

For more information on the software and hardware prerequisites for the OS/390, please refer to the “MQSeries for OS/390 v5.2 Program Directory” and the “MQSeries for OS/390 v5.2 Concepts and Planning Guide” books on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on WebSphere Application Server prerequisites, please refer to the “MQSeries for Windows NT and 2000 Quick Beginnings” book on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on EAI BUS prerequisites, please refer to the “MQSeries for Windows NT and 2000 Quick Beginnings” book on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on DLSS prerequisites, please refer to the “MQSeries for Compaq (DIGITAL) OpenVMS System Management” book on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on PEPS prerequisites, please refer to the “MQSeries for HP-UX v5.2 Quick Beginnings” book on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on BTrade prerequisites, please refer to the “MQSeries for HP-UX v5.2 Quick Beginnings” book on the IBM website:

[http://www-4.ibm.com/software/ts/mqseries/library/manualsa.](http://www-4.ibm.com/software/ts/mqseries/library/manualsa/)

For more information on how to customize MQSeries objects for application specific requirements, please refer to the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on MQSeries application error handling, event monitoring and MQSI error handling, please refer to the following books:

“MQSeries Application Programming Reference”

“MQSeries Event Monitoring”

“MQSeries Integrator Introduction and Planning”

on the IBM website: <http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on managing clusters and developing a custom cluster workload exit, please refer to the “MQSeries Queue Manager Clusters” book on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manuals>

For more information on the MQSeries Integrator Control Center and the MQSeries commands and control commands, please refer to the following books:

“MQSeries Integrator Using the Control Center”

“MQSeries MQSC Command Reference”

“MQSeries Systems Administration”

“MQSeries for Compaq (DIGITAL) OpenVMS System Management”

“MQSeries for OS/390 System Administration Guide”

on the IBM website: <http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

For more information on the MQSI configuration manger, please refer to the “MQSeries Integrator Using the Control Center” book on the IBM website:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/>

MQSeries Application Programming Guide can be found at:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/> - Latest family books

MQSeries Application Programming Reference can be found at:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/> - Latest family books

MQSeries Application Messaging Interface manual can be found at:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/> - Latest family books

MQSeries Using C++ manual can be found at:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/> - Latest family books

MQSeries Using Java manual can be found at:

<http://www-4.ibm.com/software/ts/mqseries/library/manualsa/> - Latest family books

11 APPENDIX B: GLOSSARY

A

ACS

Affiliated Computer Systems. The company that manages the DLSS system located in Rockville, MD.

active log

See recovery log.

AIS

Application Information Services

adapter

An adapter is an attachment facility (program) that enables applications to access MQSeries services. More specifically an adapter is used to isolate an application implementing an interface which manages format conversions and application specific behavior.

alias queue object

A MQSeries object, the name of which is an alias for a base queue defined to the local queue manager. When an application or a queue manager uses an alias queue, the alias name is resolved and the requested operation is performed on the associated base queue.

alternate user security

A security feature in which the authority of one user ID can be used by another user ID; for example, to open a MQSeries object.

AMI

Application Message Interface. An MQSeries term.

API

Application Programming Interface.

archive log

See *recovery log*.

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*.

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.

B

BLOB

An MQSI message domain where all unstructured messages are contained.

bootstrap data set (BSDS)

A VSAM data set that contains:

- An inventory of all active and archived log data sets known to MQSeries for OS/390
- A wrap-around inventory of all recent MQSeries for OS/390 activity

The BSDS is required if the MQSeries for OS/390 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.

bTrade

A 3rd party vendor to FSA that provides a product to send files across the internet and allows the data to be compressed and encrypted.

C

CCF

[Custom Connector Framework](#)

channel

See *message channel*.

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

A time when significant information is written on the log. Contrast with *syncpoint*. 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.

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*.

CISC

[Customer Information Control System](#). A subsystem of the OS/390 computing platform.

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.

cluster

A network of queue managers that are logically associated in some way.

CPS

Central Processing System.

COD

Common Origination and Distribution System.

command

In MQSeries, an administration instruction that can be carried out by the queue manager.

command server

The MQSeries component that reads commands from the system-command input queue, verifies them, and passes valid commands to the command processor.

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.

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.

COTS

Custom Off The Shelf. Usually used in reference to software.

CSC

Computer Sciences Corporation. Company which hosts and manages the systems located in Meriden, CT.

CWF

Custom Wire Format

D

data bag

In the MQAI, a bag that allows you to handle properties (or parameters) of objects.

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.

DB2

A relational database marketed by IBM. Also known as UDB or Universal Database.

DCI

Data conversion interface.

dead-letter queue (DLQ)

A queue to which a queue manager or application sends messages that it cannot deliver to their correct destination.

dead-letter queue handler

A MQSeries-supplied utility that monitors a dead-letter queue (DLQ) and processes messages on the queue in accordance with a user-written rules table.

DHCP

Dynamic Host Configuration Protocol

DI

Data Integrator. Product used for file transfer utilizing MQSeries as the message transport.

distributed queue management (DQM)

In message queuing, the setup and control of message channels to queue managers on other systems.

DLSS

Direct Loan Servicing System.

DMZ

Demilitarized Zone

DPL

Distributed Program Load

DTD

Document Type Definition – an MQSI component.

DLQ

Dead-letter queue.

dual logging

A method of recording MQSeries for OS/390 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*.

dynamic queue

A local queue created when a program opens a model queue object. See also *permanent dynamic queue* and *temporary dynamic queue*.

E

EAI

Enterprise Application Integration

ECB

ECampus Based System

EID - Enterprise Integration Domain

One of five domains within IAFeB developed to provide an enterprise-wide scalable framework that allows multiple front-end applications (such as web and call centers) to inter-operate with back-end applications (such as policy administration and claims systems) in an effective and efficient manner.

EJB

Enterprise Java Bean

ESQL

Extended Structured Query Language

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 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.

F

FIFO

First In First Out

FMS

Financial Management System

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)

FSA

Federal Student Aid

G

get

In message queuing, to use the MQGET call to remove a message from a queue. See also *browse*.

H

HACMP

High Availability Cluster Multi-Processing - IBM's high availability offering for AIX platforms to provide dynamic fail-over within a cluster of separate AIX systems.

I

IAA

Insurance Application Architecture. Insurance business object model.

IAFeB

Insurance architecture for e-business. Framework of common insurance specific functionality built on top MQSeries and MQSeries Integrator. Used by insurance companies to build eBusiness/integration systems.

IBM

International Business Machines

in-doubt unit of recovery

In MQSeries, the status of a unit of recovery for which a syncpoint has been requested but not yet confirmed.

initiation queue

A local queue on which the queue manager puts trigger messages.

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.

ITA

Integrated Technical Architecture

J

JMS

Java Messaging Service

L

LDAP

Lightweight directory access protocol.

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

A MQSeries object belonging to a local queue manager.

local definition of a remote queue

A 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.

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 the queue managers are running on the same system as the program.

log

In MQSeries, a file recording the work done by queue managers while the queue managers receive, transmit, and deliver messages. The log file is used to recover in the event of failure.

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.

M

message

In message queuing applications, a communication sent between programs. See also *persistent message* and *nonpersistent message*. 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 at one end and a receiver at the other end) 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. See also *message queue interface*.

message channel interface (MCI)

The MQSeries interface to which customer- or vendor-written programs that transmit messages between a 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 a MQSeries message. The format of the message descriptor is defined by the MQMD structure.

message priority

In MQSeries, an attribute of a message that can affect the order in which messages on a queue are retrieved, and whether a trigger event is generated.

message queue

Synonym for *queue*.

message queue interface (MQI)

The programming interface provided by the MQSeries queue managers. This programming interface allows application programs to access message queuing services.

message queuing

A programming technique in which each program within an application communicates with the other programs by putting messages on queues.

messaging

See *synchronous messaging* and *asynchronous messaging*.

model queue object

A set of queue attributes that act as a template when a program creates a dynamic queue.

MQOD

MQSeries Object Descriptor. The MQOD structure is used to specify an object by name. The structure is an input/output parameter on the MQOPEN and MQPUT1 calls.

The following types of object are valid:

- Queue or distribution list
- Namelist
- Process definition
- Queue manager

MQSeries – Message Queue Series

A family of IBM licensed programs that provides message queuing services across a broad array of operating system platforms and network protocols.

MQSeries Administration Interface (MQAI)

A programming interface to MQSeries.

MQSeries client

Part of a 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)*.

MQSI - MQSeries Integrator

Second generation message broker product that provides basic message routing and data translation capabilities.

MQWF

MQSeries Workflow. A workflow product built to executed long running transactions and other workflow functions over a MQSeries foundation.

MRM

Message Respository Manager. A component of the Configuration Manager that manages MQSI messages.

MVS

Multiple Virtual System

N

namelist

A MQSeries object that contains a list of names, for example, 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.

NEON

New Era of Networks. MQSI interface available through the Control Center.

nonpersistent message

A message that does not survive a restart of the queue manager. Contrast with *persistent message*.

NSLDS

National Student Loan Data System

O

OAMAS

Open Applications Group Middleware API Specifications

OAG

Open Applications Group. The Open Applications Group is a non-profit consortium focusing on best practices and processes based on XML content for eBusiness and Application Integration.

object

In MQSeries, an object is a queue manager, a queue, a process definition, a channel, a namelist, or a storage class (OS/390 only).

object authority manager (OAM)

In MQSeries on UNIX systems, MQSeries for AS/400, 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.

output log-buffer

In MQSeries for OS/390, a buffer that holds recovery log records.

P

page set

A VSAM data set used when MQSeries for OS/390 moves data (for example, queues and messages) from buffers in main storage to permanent backing storage (DASD).

PCF

Programmable Command Format. The purpose of MQSeries programmable command format (PCF) commands is to allow administration tasks to be programmed into an administration program. In this way you can create queues and process definitions, and change queue managers, from a program.

PDF

Specialized message format used in MQSI and predominately found in the finance industry.

PEPS

Post-Secondary Education Participants System

performance event

A category of event indicating that a limit condition has occurred.

persistent message

A message that survives a restart of the queue manager. Contrast with *nonpersistent message*.

platform

In MQSeries, the operating system under which a queue manager is running.

point of recovery

In MQSeries for OS/390, the term used to describe a set of backup copies of MQSeries for OS/390 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).

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

A MQSeries object that contains the definition of a 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*.

Q

queue

A 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.

queue manager

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*. A MQSeries object that defines the attributes of a particular queue manager.

queuing

See message queuing.

R

recovery log

In MQSeries for OS/390, data sets containing information needed to recover messages, queues, and the MQSeries subsystem. MQSeries for OS/390 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*.

remote queue

A queue belonging to a remote queue manager. Programs can put messages on remote queues, but 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.

request message

A type of message used to request a reply from another program.

RESLEVEL

In MQSeries for OS/390, an option that controls the number of CICS user IDs checked for API-resource security in MQSeries for OS/390.

return codes

The collective name for completion codes and reason codes.

RRS

Resource Recovery Service

S

SAIG

Student Aid Internet Gateway

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.

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*.

signaling

In MQSeries for OS/390 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 OS/390 activity where each change is recorded on one data set only. Contrast with *dual logging*.

SQL

Structured Query Language. A database language used to query databases.

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*.

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.

T

TCP/IP

Transmission Control Protocol / Internet Protocol.

thread

In MQSeries, the lowest level of parallel execution available on an operating system platform.

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*.

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.

TSYS

Total System. A vendor which manages the COD environment.

U

undelivered-message queue

See dead-letter queue.

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*.

URL

Uniform Resource Locator.

UTCL

Unpaid Teacher Cancellation Policies.

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.

V

VAJ

Visual Age for Java. A Java programming development environment offered by IBM.

VTAM

Virtual Terminal Access Manager.

W

WAS

WebSphere Application Server.

X

XA

XA interface is a specification that describes the protocol for transaction coordination, commitment, and recovery between a transaction manager and one or more resource managers.

XML

Extensible markup language

12 APPENDIX C: CORE SERVICES QUESTIONNAIRE

This document serves as a questionnaire for application teams to provide information regarding all the interfaces for each EAI initiative. Examples of the initial tasks required for every EAI initiatives are listed below. The [REQUIRED SECTION (S)] must be completed by the application teams before the EAI team can accomplish the specific task.

- EAI Estimates for all the interfaces [REQUIRED SECTION (S)]:
 - *12.1 High Level System and Interface Information*
- Interface Partner Agreement (similar to a statement of work) [REQUIRED SECTION (S)]:
 - *12.1 High Level System and Interface Information*
 - *12.4 Detail System and Interface Overview*
- Detail design for all the interfaces [REQUIRED SECTION (S)]:
 - *12.1 High Level System and Interface Information*
 - *12.4 Detail System and Interface Overview*

12.1 High Level System and Interface Information

This section is intended to outline the minimum set of information that the EAI team needs to produce a high level estimate for work.

12.2 Application

This section lists all the interfacing applications for this EAI initiative.

12.2.1 Application Identification

Application name	
Application acronym	
Provide a detailed explanation of the application's business usage	
Date Questionnaire Completed/Revised	
Data sources	<input type="checkbox"/> <i>Database</i> <input type="checkbox"/> <i>File System</i> <input type="checkbox"/> <i>Other (please specify)</i>

12.3 Interfaces

List all of the system's interfaces to external/internal systems by completing the following sections:

Interface Description – This section captures the high-level information regarding the interfaces. For each interface, the following two sections should be completed.

- *Transaction details – Data* – This section captures different types of data that will be processed by the interface.

- *Transaction details – Transformation/Formatting* – This section captures the data transformation/formatting effort for all the data that will be processed by the interface.

The following scenarios below provide general guidelines for appropriate grouping of the identified interfaces:

For example, twelve external subsystems need to interface with a backend FSA system to process student financial data. The interface(s) requires the external systems to extract information from the data source, transfer the file(s) to FSA system, perform any transformation/formatting, and load the data into the data source residing on the FSA system.

Scenario #1:

If the *twelve external subsystems* all perform similar business functions (e.g., all private collection agencies), and each subsystem have ten unique types of financial data to be loaded into the FSA system. In this scenario, only one interface needs to be documented. This results in the following sections documented:

2.2.1 Interface Description [one instance of this section]
2.2.1.1 Transaction Detail – Data [# of events/messages]
Example: (12 subsystems) x (10 unique events) = 120 total events for this interface.
2.2.1.2 Transaction Detail – Transformation & Formatting
Example: [100] No transformation required
[20] Simple
...

Scenario #2:

If the *twelve external subsystems* all perform unrelated business functions (e.g., government agencies, private collection agencies, Department of Justice, etc), and each subsystem have ten unique types of financial data to be loaded into the FSA system. In this scenario, twelve interfaces and each transaction type needs to be documented. This results in the following sections documented:

2.2.1 Interface Description [#1]
2.2.1.1 Transaction Detail – Data [# of events/messages]
Example: 10 events for this interface.
2.2.1.2 Transaction Detail – Transformation & Formatting
Example: [10] No transformation required
2.2.2 Interface Description [#2]
2.2.2.1 Transaction Detail – Data [# of events/messages]
Example: 10 events for this interface
2.2.2.2 Transaction Detail – Transformation & Formatting
Example: [10] Simple
2.2.3 Interface Description [#3]
2.2.3.1 Transaction Detail – Data [# of events/messages]
Example: [10] No transformation required
2.2.3.2 Transaction Detail – Transformation & Formatting
Example: [10] Medium
.....
.....
2.2.12 Interface Description [#12]
2.2.12.1 Transaction Detail – Data [# of events/messages]
Example: [10] No transformation required
2.2.12.2 Transaction Detail – Transformation & Formatting
Example: [10] Medium

The detail information about each of the documented interfaces should be captured in section **12.6 Detail Interface Information** of this document.

12.3.1 Interface Description [one per interface type]:

The following section should be completed for each interface. This section captures the high-level information regarding the interfaces.

1. Interface name(s)	
2. Interface Short Description	
3. Interface Business Functionality Description	
4. Interfacing System(s)	
5. Interface Mode (online or batch)	<input type="checkbox"/> Online – real-time transactions (e.g., web front end requesting information) <input type="checkbox"/> Batch – scheduled transactions (e.g., feeding daily new students file from schools) <input type="checkbox"/> Other (please specify)
6. Interface Direction	<input type="checkbox"/> Source -> Target - Data flow from Source system to Target system (i.e., request or datagram event/message) <input type="checkbox"/> Target -> Source - Data flow from Target system to Source system (i.e., reply to a request event/message) <input type="checkbox"/> Both
7. Interaction Method	<input type="checkbox"/> Datagram – Send the data to a known target resource without waiting for a reply. <input type="checkbox"/> Request & Reply - Requests a service from a known resource, where the transaction cannot continue without the reply back from the request event (on-line application) <input type="checkbox"/> Publish & Subscribe - The data owner publishes documents without knowing if there are any subscribers to that information. This is similar to publishing magazines, if you like it, subscribe to it. If don't want it you cancel your subscription. It doesn't wait for subscriber acknowledgements <input type="checkbox"/> Publish & Reply - requests a service from unknown resources. Once the source publishes the event, it can continue to process other information. This interaction method eliminates the source application's need to know the target resource's identity <input type="checkbox"/> Other (please specify)
8. Interface Protocol (data transmission protocols used to support interface.)	<input type="checkbox"/> TCP/IP <input type="checkbox"/> FTP <input type="checkbox"/> HTTP <input type="checkbox"/> SNA <input type="checkbox"/> Other (please specify)
9. Interface Encryption (Does the interface support encryption? If so, what algorithm/product is used?)	<input type="checkbox"/> Yes <input type="checkbox"/> No
10. What are the interface's response time requirements?	

11. Interface usage frequency	
12. Interface usage volume? Number of transactions?	
13. Internal/External?	<input type="checkbox"/> Internal (inside the firewall or within FSA) <input type="checkbox"/> External (outside the firewall e.g., Department of Justice)
14. Interface Data Processing (Describe all processing [e.g., copying data, writing to log files, backing up data, etc.] the interface performs on data.)	
15. Are protocols or formats used to ensure message integrity (i.e. message digest, digital signature)	<input type="checkbox"/> Yes <input type="checkbox"/> No
16. Can the interface recover from a crash without permanent data loss?	<input type="checkbox"/> Yes <input type="checkbox"/> No
17. List any vendor(s) and software licensed to support this interface.	

12.3.1.1 Transaction Details - Data

1. For this interface, list the number of events/message sent/received?	<i>[Enter # of events/messages] Events/messages sent</i> <i>[Enter # of events/messages] Events/messages received</i>
2. For each event/message, list the number of unique transactions?	
3. Does the events/messages, share transactions – i.e. an Account trax?	<i>[Enter # of events/messages] Enter number of events/messages that share transactions</i>

12.3.1.2 Transaction Details - Transformation/Formatting

1. What event/message cleansing is performed today (i.e., removal of headers/trailers)?	
2. Where is the transformation Logic performed?	<input type="checkbox"/> Source Application <input type="checkbox"/> Target Application <input type="checkbox"/> “EAI Bus” – The transformation logic for the events/messages from the source to target format reside on the “EAI Bus” <input type="checkbox"/> Other (please specify)

<p>3. Fore each event/message, what business logic, transformation or formatting rules are applied (i.e. converting from one date format to another MMDDYY -> YYYY-MM-DD)</p>	<p>[Enter # of transformation of this complexity] No transformation required [Enter # of transformation of this complexity] Simple - Mostly pass through data (straight copies of fields), Very little reformatting, semantic or syntactic rules, Fewer than 50 elements [Enter # of transformation of this complexity] Medium - Some logic and rules based transformation, Format changes, Simple IF or CASE logic, Limited error handling, Very simple DART or CORBA integration, Simple EDI transactions, Between 25 and 75 elements [Enter # of transformation of this complexity] Complex - Multiple Rules, Complex Logic, Moderate to extensive error handling, Moderate to complex DART or CORBA integration complex logic, Moderate to complex EDI interactions, Between 50 and 150 elements [Enter # of transformation of this complexity] Super Complex - Joins of data from multiple sources, Complex Logic, Extensive data validation, Extensive error handling, Complex DART or CORBA integration, Complex EDI interactions, More than 100 elements</p>
<p>4. What transactionality is necessary (i.e. does processing depend on the data from multiple events/messages.</p>	
<p>5. Is there an acknowledgement once the transaction is completed?</p>	<p>[Enter #] Events/messages require acknowledgements</p>
<p>6. What is the current error handling procedures when the event data is incomplete?</p>	
<p>7. Does the event/message need to be persisted?</p>	

12.4 Detail System and Interface Overview

The following sections capture technical and operational information about the system and its interfaces in order to build the core services of the Integration Architecture. This section includes information about the detail system information, and detail interface information.

12.5 Detail System Information

This section outlines physical locations of environments, application platform information, key system contacts, system context diagram, and data store information.

12.5.1 Environments

List all of the application's environments and physical locations.

Environment	Location
Development	
Test	
Production	
Other	

12.5.2 Application Platform Information

Complete a separate platform section for each environment listed in the previous section (12.5.1 Environments) used by the application.

Platform 1 – (e.g. Development)	
Logical Platform Name or Description	
Hardware Platform	
Operating System	
OS Version	
OS Patches	
DBMS	
Transmission Protocol	
IP Address	
Node Name	
MQSeries Intercommunication Port	<i>MQSeries uses IP Port 1414 as a default for intercommunication. Is this port available? If not, please specify an available port number for MQSeries intercommunication.</i>
System Access	<i>Does the EAI Core development team have connectivity to the system from their location? Is this system only accessible through the EDNet network?</i>
List all programming languages used	
Compiler tool(s) and version(s)	
Source or Target Platform	<input type="checkbox"/> Source <input type="checkbox"/> Target <input type="checkbox"/> Both
Additional platform information:	
Security Access Requirements	
Contracting Organization	
Developer Location	

12.5.3 Contacts

Title/Role	Name	Contact Number	E-mail
SFA System Owner			
System Security Officer			
Systems Administrator			
EAI Team Liaison			
Technical SME			
Functional SME (developer)			
Contact for Development Access			
Contact for Production Access			

12.5.4 System Context Diagram

Provide a high-level context diagram showing this system with all the major interfaces. The diagram should also depict the system's inputs and outputs.

[Insert diagram]

12.5.5 Data Sources

Provide details on the type and structure of data stores (e.g., databases, flat-files, VSAM, etc.) used to support this system.

Data Store	Version	Platform	Description

12.6 Detail Interface Information

Information regarding system interfaces is particularly important in determining the scope of requirements for integration between the EAI infrastructure and the application. Interfaces are defined as systems/applications outside the immediate domain of the subject system. Generally, interfaces support connectivity to other systems. Respondents are encouraged to provide as much detail and supporting documentation as possible. Graphical representation of interfaces is of particular value in defining these specifics.

12.6.1.1 Interface Flow Diagram

Provide a high-level diagram that depicts the flow of control between the system and its online and batch interfaces.

[Insert diagram]

12.6.1.2 Interface Transactions and Transaction Volume

The table below lists questions related to the application’s business and interface transactions (IT) transactions. Business transactions are transactions that support a particular business function (e.g., update address). IT transactions are application-level or system-level transactions that support business transactions (e.g., standardize the address according to USPS standards, update the HR database with the new address, update the organization directory with the new address, etc.). One business transaction can spawn more than one IT transactions.

1. Business Transactions Supported by Interface	List the transaction names or business objects:
2. Please define any seasonal specifics (Certain interfaces encounter volume fluctuations depending on the time of year and the various academic dependencies. Please provide details on whether or not this interface is subject to such dependencies.)	
3. Data Volume per Transaction (i.e. What are the maximum message lengths?)	
4. # of IT Transactions per Business Transaction	
5. What is the maximum number of IT transactions per hour that must be supported by the interface?	
6. What is the growth projection for the interface throughput over the next three (3) years?	

7. How many concurrent users (average and peak) will the application support?	Average: Peak:
8. What are the assumptions about key capacity metrics?	

12.6.1.3 System Availability

Please provide information on the system’s availability, as well as the business impact of lost availability. (Provide details on the period of time that the system and its interfaces must be available. Some of the system’s interfaces must be available on demand while others are based on scheduled availability.)

1. What are the system’s peak and off-peak times?	
2. What is the uptime percent (percentage of time during which the system is functioning and available) during peak hours?	
3. What is the uptime percent during non-peak hours?	
4. What is the maximum acceptable time to recover from a single interface connection failure?	
5. Does a reply timeout exist? If so, what is its duration?	
6. What is the maximum response time degradation acceptable under load?	
7. What is the average percentage of transactions that fail and require reprocessing?	
8. What is the acceptable window for any required real-time processing?	
9. What is the acceptable window for any required batch operations?	
10. Do other systems that provide required data meet the availability profile of this application?	

11. Does the user interface portion of the application system need to be operational during real-time updates?	
12. Does the user interface portion of the application system need to be operational during batch updates?	
13. Can a replicated version or subset of the database fulfill availability requirements when the production database is offline for real-time processing?	
14. Can a replicated version or subset of the database fulfill availability requirements when the production database is offline for batch processing?	
15. If your answer to question 11 is "YES", then what would be the required currency of the replicated data?	

12.6.1.4 System Security

Describe the system's security measures. Include information on user groups, and if the system is a mainframe include LPAR or region information. (Please provide specifics on the tools/products and procedures that are used to control online access to the application.)

1. Does the system operation environment use/require a single point for Identification, Authentication, and Access Control? (Please specify any software that is required.)	
2. Do users or applications require identification and authentication?	
3. Do any transactions or actions need to be authorized?	
4. Does the application maintain a history/audit	

log of messages sent and received?	
5. Does the application prevent information from being viewed by unauthorized parties?	
6. Does the application track inactivity?	
7. Where is security access defined and enforced?	
8. Is there a hierarchy (i.e. role-based groups) for organizing authorization so that access rights do not have to be individually assigned?	