

MQSeries Integrator for AIX V2 Performance Report

Version 1

September 8, 2000

Tim Dunn

Dan Jones

MQSeries Performance and Test
IBM UK Laboratories
Hursley Park
Winchester
Hampshire
SO21 2JN

Property of IBM

Take Note!

Before using this report be sure to read the general information under "Notices".

First Edition, September 2000

This edition applies to Version 1 of MQSeries Integrator for AIX - V2 Performance Report and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation** 2000. All rights reserved.
Note to U.S. Government Users -- Documentation related to restricted rights -- Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule contract with IBM Corp.

Notices

This report is intended to help the reader understand the performance characteristics and perform capacity planning for MQSeries Integrator for AIX - V2.0.1. The information is not intended as the specification of any programming interfaces that are provided by MQSeries or MQSeries Integrator for AIX - V2.0.1.

References in this report to IBM products or programs do not imply that IBM intends to make these available in all countries in which IBM operates.

Information contained in this report has not been submitted to any formal IBM test and is distributed "asis". The use of this information and the implementation of any of the techniques is the responsibility of the customer. Much depends on the ability of the customer to evaluate these data and project the results to their operational environment.

The performance data contained in this report was measured in a controlled environment and results obtained in other environments may vary significantly.

Trademarks and service marks

The following terms, used in this publication, are trademarks of the IBM Corporation in the United States or other countries or both:

- IBM
- AIX
- Netfinity
- MQSeries
- MQSeries Integrator
- DB2

The following terms are trademarks of other companies:

- Windows NT, Visual Studio Microsoft Corporation
- NEONFormatter, NEONRules, NEON

Other company, product, and service names may be trademarks or service marks of others.

Summary of Amendments

Date	Changes
8 September 2000	Initial Release.

Contents

1.0 CONCEPTS	1
1.1 Major Components	1
1.1.1 The Configuration Manager	1
1.1.2 The Brokers	2
1.1.3 The Control Center	2
1.1.4 The User Name Server	3
1.2 Message Flows	3
1.3 Messages and Message Sets	4
1.4 Message Parsers	5
2.0 BROKER THROUGHPUT MEASUREMENTS	7
2.1 MQInput/MQOutput Throughput	9
2.2 Compute Node Throughput	10
2.2.1 Simple Compute Node	11
2.2.2 Complex Compute Node	12
2.2.3 Multiple Complex Compute Nodes	13
2.2.4 Very Complex Compute Node	14
2.3 Database Node Throughput	15
2.4 Filter Node Throughput	17
2.5 Publication Node Throughput	18
2.6 What Is The Cost Of Converting Messages To Different Formats ?	20
2.7 How Well Do NEON Rules and NEON Formatter Run Within MQSI V2?	21
2.8 What Is The Effect of Increasing The Number of Execution Groups?	23
2.9 What is The Effect of Making a Message Flow Transactional?	25
2.10 What Is The Effect of Using coordinatedTransaction=yes on a Message Flow?	26
2.11 What Effect Does an Increasing Number of Subscribers Have on Publish/Subscribe Throughput?	28
2.12 What Is The Effect of Using Content Based Publish/Subscribe?	30
2.13 What Is The Effect of Examining an Increasing Number of Fields in a Message?	31
2.14 What Is The Effect of Using User Trace?	32
2.15 What Is The Effect on Message Throughput of Using Faster Disks?	32
3.0 CAPACITY PLANNING	33
3.1 Throughput	33
3.2 Scaling Message Throughput	34

3.3 Memory	36
3.4 Recommended Minimum Configurations	37
4.0 PERFORMANCE RECOMMENDATIONS	38
4.1 Understand Recovery Requirements	38
4.2 Optimize Queue Manager	38
4.3 Configuration Considerations	39
4.4 Maximizing Throughput	39
5.0 GLOSSARY	41
6.0 APPENDIX A - MEASUREMENT HARDWARE AND SOFTWARE	43
7.0 APPENDIX B - MEASUREMENT DATA	45
7.1 MQInput/MQOutput Throughput Results	45
7.2 Compute Node Throughput Results	46
7.2.1 Simple Compute Node	46
7.2.2 Complex Compute Node	46
7.2.3 Multiple Complex Compute Nodes	47
7.2.4 Very Complex Compute Node	47
7.3 Database Node Throughput Results	48
7.4 Filter Node Throughput Results	48
7.5 Publication Node Throughput Results	49
7.6 Converting Messages Between Formats	50
7.7 NEONRules and NEONFormatter Within MQSI V2	53
7.8 The Effect of Increasing The Number Of Execution Groups	54
7.9 The Effect of Making a Message Flow Transactional	55
7.10 The Effect of using coordinatedTransaction=yes	56
7.11 The Effect of Increasing the Number of Subscribers	57
7.12 Content vs Topic Based PubSub	60
7.13 The Effect of an Increasing Number of Message Fields	61
7.14 The Effect of userTrace options	62
8.0 APPENDIX C - COMPLEX COMPUTE NODE	63
8.1 Complex Compute Node	63
8.2 Multiple Complex Compute Node	63
8.3 Very Complex Compute Node	63

1.0 CONCEPTS

MQSeries Integrator for AIX - V2.0.1 (MQSI V2) is IBM's message broker product, addressing the needs of business and application integration. Business integration is the coordination of all of a company's processes. Application integration is the coordination of its applications. This process of integration involves the bringing together of the data and processes within an organization to maximize the sharing of data and applications in order to cope with changing organization structure (merger, acquisition etc.) and increase the effectiveness of the organization.

A key requirement of such business and application integration is that applications are able to communicate with each other without having to make code changes. MQSI V2 makes the required integration easier through the services that it provides. These services are:

- Route a message to several destinations, using rules that act on the contents of one or more of the fields in the message or message header.
- Transform a message, so that applications using different formats can exchange messages in their own formats.
- Store and retrieve a message, or part of a message, in a database.
- Modify the contents of a message (for example, by adding data extracted from a database).
- Publish a message to make it available to other applications. Other applications can choose to receive publications that relate to specific topics, have specific content, or both.
- Extend the capabilities of rules and formats defined in MQSeries Integrator V1.

The above services are based on the messaging transport services provided by the MQSeries Messaging products.

1.1 Major Components

The major components of MQSI V2 are:

- The Configuration Manager
- The Brokers
- The Control Center.
- The User Name Server.

The Configuration Manager and Control Center must run on a machine running the Windows NT operating system. The User Name Server can be run on AIX or Windows NT if required.

1.1.1 The Configuration Manager

The Configuration Manager is the main component of the MQSI environment. The components and resources managed by the Configuration Manager constitute the broker domain. The Configuration Manager servers three main functions:

- It maintains configuration details in the configuration repository. This is a set of database tables that provide a central record of the broker domain components.
- It manages the initialization and deployment of brokers and message processing operations in response to actions initiated through the Control Center. It communicates with other components in the broker domain using MQSeries transport services.

- It checks the authority of defined user Ids to initiate those actions.

There is a single Configuration Manager to manage a broker domain. The Configuration Manager provides a service to other components in the broker domain providing them with configuration updates in response to actions taken by the user of the Control Center.

1.1.2 The Brokers

The broker is a named resource that hosts and controls the business processes which are defined as message flows. Applications send new messages to the message flow and receive processed messages from the message flow, using MQSeries queues and connections.

Any number of brokers can be created within a broker domain. It is possible to create more than one broker on any one physical system if desired, but there must be a unique queue manager for each broker.

Within each broker it is possible to define execution groups which are responsible for running the message flows. An execution group is implemented as an operating system process. Within an execution group it is possible to define additional threads that will also perform the processing of the message flow. These are known as additional instances.

When creating message flows that provide a publish/subscribe service it is possible to connect a number of brokers in a collective using the Control Center. A collective contains a number of brokers that are all physically interconnected. All the broker queue managers must be connected by pairs of MQSeries channels.

A collective optimizes the publish/subscribe of messages in the broker domain by reducing the number of clients per broker, without increasing the hops taken by any message by more than one. In this way collectives are more efficient than a hierarchy.

It is possible to connect collectives to other collectives and to other individual brokers. When collectives are connected to a standalone broker only one broker in each collective must provide the connection.

Messages published to any one broker are propagated to all connected brokers (whether or not they are in a collective) to which an application has subscribed to the messages topic or content.

1.1.3 The Control Center

The Control Center interfaces with the Configuration Manager to allow the user to configure and control the broker domain. The Control Center and Configuration Manager exchange messages (using MQSeries) to provide the information requested and to make updates to the broker domain configuration.

It is possible to install and invoke any number of Control Center instances. The Control Center can be installed on the same physical system as the Configuration Manager, or any other system that can connect to the Configuration Manager.

The Control Center is structured as a number of views on the configuration and message repositories. The message repository contains all message definitions which have been created or imported through the Control Center. The configuration repository contains configuration information pertaining to all other resources within the broker domain; brokers, collectives, message processing nodes, message flows, topics and subscriptions.

The Control Center can be used to

- Develop, modify, assign and deploy message flows.
- Develop, modify, assign and deploy message sets.

- Define the broker domain topology and create collectives.
- Control topic security of messages by topic.
- View status information.

1.1.4 The User Name Server

The User Name Server monitors the underlying security subsystem provided by the operating system and provides information about the valid principals (users and groups of users) in the system. The User Name Server shares this information with the brokers and Configuration Manager and updates it at frequent intervals. The information can be used to control access to topic-based messages produced by the publish/subscribe service. Topic-based security gives the ability to control the authority of applications, identified by the user ID under which they are executing, to publish on topics, to subscribe to topics and to request persistent delivery of messages on topics.

1.2 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 supplies a number of predefined message processing node types, known as IBM primitives. These provide basic functions including input, output, filter (on message data content), and compute (manipulate message content: for example, add data from a database).

A message flow and the message processing nodes it contains describe 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 message flows to be defined and reused in other message flows when required.

When the message flow creation is complete, it can be assigned for execution to one or more brokers. The message flow must be operationally complete. That is it must contain at least MQInput node. Most message flows will also contain at least one MQOutput or one Publication node, although this is not required.

A message flow is transactional: it is possible to define 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. It is possible to define a message flow to work outside of a unit of work if this transactional support is not required.

When a message flow is deployed to a broker, the broker automatically starts an instance of the message flow for each input node (one or more). This is the default behaviour. Each instance retrieves a message from the input node, and runs in parallel with other instances that retrieve a message from other input nodes.

In order to further increase the throughput of the message flow, it is possible to 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. It is possible to set properties of the input node to exercise control over the order in which messages are processed.

It is possible to increase message flow throughput by assigning more than one copy of the message flow to the same broker. 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.

The broker provides the run-time environment for a set of deployed message flows: this environment is called an execution group. An execution group provides an isolated environment, because each is started as a separate operating system process.

One execution group, the default execution group, is set up for use whenever a broker is created. By setting up additional execution groups, it is possible to isolate message flows that handle sensitive data such as payroll records or security information, from other non-sensitive message flows.

Within an execution group the assigned message flows run in different thread pools. The size of the thread pool that is assigned for each message flow is set by specifying the number of additional instances of each message flow.

The broker guarantees operating isolation of each execution group, thus guaranteeing data integrity between execution groups and improving robustness of message flows.

1.3 Messages and Message Sets

In MQSI messages are always in one of two broad categories:

- Predefined. The content of a predefined message is described by the message template.
- Self-defining. The content of a self-defining message is described by the message itself.

The message definition process is managed by the Message Repository Manager (MRM) component of the Control Center.

Message definitions are created or modified using the Control Center, the MRM stores them in the message repository.

Predefined Messages

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

The logical structure defines the contents of the message using a tree structure that identifies each field and its relation to other fields. The applications sending and receiving messages like this understand the format and type of each field. For example they might use a C structure that shows AccountNumber is an eight byte field, AccountName is a 20 byte character field and AccountBalance is an 8 byte character field.

The physical structure, also known as a wire format, is a string of bytes. Without the logical structure the physical structure has no intrinsic meaning.

The physical structure of each element in a message is further defined by its Custom Wire Format(CWF) characteristics. These give the physical format (for example COBOL packed decimal), the length, whether the field is signed and so on.

Message Templates

A message template is made up of four values contained within the header information:

1. *Message Domain* which identifies the message parser that will interpret the bit-stream of the message. By default, MQSI supports the values MRM, XML, BLOB and NEON. MRM is the MRM-enabled parser and is used for all messages whose definitions have been created in, or imported to, the message repository. XML is for self defining messages only. BLOB is used for

messages whose format is not understood, in which case they are treated as a bit stream. NEON is for NEON messages only.

2. *Message set* which identifies the grouping of messages within the message domain, as it has been defined. Typically a message set contains a number of related messages that provide the definitions required for a specific business task or application suite. The message set is similar in concept to the application group in NEON.
3. *Message type* which identifies the logical structure of the data in the message. For example the number and location of character strings and their relationships.
4. *Message format* which identifies the physical representation of the message (its wire format). The MRM-enabled parser supports the following wire formats:

XML - the message is identified as an XML document that complies with a Document Type Descriptor (DTD) that can be generated for a message by the MRM. This option does not apply to self-defining messages, which have the domain XML, rather than MRM.

CWF - denotes legacy data structures used in common programming languages (C or COBOL). Data structures for CWF messages are typically imported into the message repository.

The message format value is only valid for predefined messages and not self-defining messages.

Self-defining messages

Self-defining messages use the XML standard to structure their content. They can be used in any message flow, and are supported by all message flow nodes.

Self-defining messages do not have to be defined to the Control Center, nor do they have to be assigned to brokers to ensure that they can be interpreted.

Self defining messages are said to use generic XML.

When a message is processed in a message flow, its format must be determined first so that the correct parser is used. The message characteristics are identified by the input node of a message flow in one of two ways:

- For messages with an MQRFH or MQRFH2 architected header, the input node checks the value in the message header.
- For messages which do not have an MQRFH or MQRFH2 header, the input node uses the default message template, defined as a property of the input node, to determine how the message must be parsed.

1.4 Message Parsers

MQSI can handle any message template for which a suitable parser is available. The parsers interact with the message templates stored in the message dictionaries. The range of messages supported can be extended by creating your own message parsers.

Message parsers are provided for :

- Predefined XML. Such messages have an MQRFH or MQRFH2 header.
- The standard MQSeries headers: MQCIH, MQCFH, MQDLH, MQIIH, MQMD, MQMDE, MQRFH, MQRFH2, MQRMH, QSAPH and MQWIH.
- Record-orientated C and COBOL language structures.
- Self-defining (generic XML) messages.

- Messages whose formats are defined in the NEON dictionary (These messages are defined using the NEON interface not the Control Center).

If no parser can be identified for a message, MQSI treats it as a binary object that passes, of necessity, unaltered through any message flow. However such a message can be stored in a database, be routed according to topic, and have headers added or removed.

MQSI V2 provides a function that allows messages to be transformed from one format to another.

2.0 BROKER THROUGHPUT MEASUREMENTS

In order to understand the processing characteristics of MQSI V2 a number of performance measurements have been taken using multiple aspects of the product. The test cases used have been deliberately made trivial in order to be able to report the cost of using MQSI V2, rather than to report the cost of running a particular application. It is very difficult to accurately represent what might be considered a typical application since the business logic is always enterprise specific.

The effect of the queue manager has been minimized where possible. This has meant predominantly using non persistent messages as well as having a compensating program to ensure that the queue manager queue cache did not overflow to disk. Where these two actions were not taken the throughput possible with MQSI V2 would have been constrained by the necessary I/O processing of the queue manager with which the MQSI V2 broker was associated.

The performance measurements have focused on the throughput capabilities of the broker using different processing node types. The aim of the measurements was to be able to answer questions such how many messages a second can be processed with each of the node types, how much CPU does it cost to do that, and finally, how much memory is required.

In the throughput measurements the following node types have been measured:

- MQInput and MQOutput
- Compute
- Database
- Filter
- NEONFormatter
- Publication

These nodes give a cross section of the possible node types and should be sufficient to cover most basic types of message transformation and distribution. Some node types have been measured in more than one configuration in order to investigate the various configuration effects, such as running multiple execution groups. All the nodes measured used minimal processing where it was possible (apart from the investigation into complex node processing so the results presented represent the best throughput that can be achieved for that node type. This should be borne in mind when performing capacity planning.

Minimum specification machines are recommended for each of the MQSI V2 components.

All measurements were conducted in the same measurement environment. This is described in Section 6.0 APPENDIX A - MEASUREMENT HARDWARE AND SOFTWARE.

All measurements used a broker which ran as a standard MQSeries application. The MQSeries listener process was also run as STANDARD.

All measurements were driven by a multithreaded MQSeries Client program written in C. There were 45 threads used for each measurement, with all threads sending the same message format and content. The Message Queue Interface programming interface was used to write and read messages.

There was no error processing or error conditions in the measurements. All messages were successfully passed from one node to another through the out or true terminal. No messages were passed through the failure terminal of a node.

The DB2 instance used with the broker was a default configuration with no tuning.

The message rates reported are the number of messages passing between the MQSeries multithreaded client and the MQSeries queue manager to which the input data is written and from which the reply data is read. Thus a client thread will issue an MQPUT to write an input message for a message flow and issue an MQGET to read the output message of the message flow. This results in a count of two messages. This method of reporting makes this report consistent with MQSeries performance reports. Those who wish to only deal with an input message rate should halve the message rates reported.

For an MQInput and MQOutput node it is possible to define transaction support for a node. Possible values are yes, no and automatic with the default value being yes for an MQInput node and automatic for an MQOutput node.

- A value of yes means that the message flow will take place under transaction control. Any derived messages subsequently sent by an MQOutput node in the same instance of the message flow will be sent transactionally unless the MQOutput node has explicitly overridden the use of transaction control.
- A value of no means that the message flow is not under transaction control. Any derived messages subsequently sent by an MQOutput node in the flow will be sent non-transactionally, unless the MQOutput node has specified that the message should be put as part of a transaction.
- A value of automatic means that the message flow will be under transaction control if the incoming message is marked as persistent, otherwise it will not. Any derived messages subsequently sent by an MQOutput node will be sent under transaction control or not, as determined by the persistence on the incoming message, unless the MQOutput node has specifically overridden the use of transaction control.

The use of transaction control means that message processing takes place within an MQSeries unit of work. This involves additional CPU and I/O processing by MQSeries because the unit of work is recoverable. The result is inevitably a reduction in message throughput for both persistent and non persistent messages.

In order to show optimal performance of MQSI V2 all the throughput measurements in this document used a value of automatic for the transaction parameter on the MQInput node unless otherwise specified. The MQOutput node used its default value.

2.1 MQInput/MQOutput Throughput

A message flow consisting of a single MQInput and MQOutput node represents a very simple message flow. Measuring the throughput achievable with such a message flow shows the maximum message rate that can be achieved using MQSI V2 to move messages between MQSeries queues.

A single message flow was defined, consisting of an MQInput node and MQOutput node. The transaction mode for the MQInput and MQOutput nodes was set to automatic.

Figure 1 below shows the results which were obtained as a result of running the message flow with varying message sizes and persistence. There was a single instance and single execution group running the message flow.




Figure 1: MQInput/MQOutput Throughput Results

With a 1K non persistent message it was possible to process approximately 1500 msgs/second. Increasing the message size to 4K had an effect on the throughput rate achieved (1100 msgs/second).

Increasing the message size to 16K and beyond had a significant effect on the maximum message throughput which could be achieved. This decrease in throughput is as a result of the additional volume of data that must be managed both within MQSI V2 and the associated MQSeries queue manager.

As the input message was non persistent there was no transactional control. We are, therefore observing the maximum rate at which MQSI V2 is able to transfer messages from the input queue to the output queue for a single execution group. Adding additional execution groups allows greater throughput to be achieved.

The use of persistent messages had a significant effect on the maximum message throughput rate which could be achieved. For a 1K persistent message the message rate was 52 msgs/second.

When persistent messages are used there are two additional effects that dominate the maximum message throughput rate achievable:

1. Any messages read from or written to an MQSeries queue now take place under MQSeries transaction control
2. The MQSeries queue manager must make the message persistent, which involves a synchronous write to the MQSeries log and a write to the file system for the message. The write to the file system may involve an I/O or not. This is dependent on the file system and is not forced by MQSeries.

As a result of the additional disk I/O required the message rate becomes dominated by I/O processing and is no longer CPU bound. The greatest message rate that is achievable is totally dependent on the speed of the I/O device on which the MQSeries log is located.

The detailed measurement data for the MQInput/MQOutput throughput measurements is available in **Section 7.1 - MQInput/MQOutput Throughput Results**.

2.2 Compute Node Throughput

A compute node provides the capability to derive an output message from an input message and also optionally include user specified processing as well as data values from an external relational database. The compute node has the potential to vary from simple to complex in its processing. The degree of complexity specified has a direct bearing on the message throughput rates which can be achieved using nodes of that type. A series of measurements were taken using varying numbers of compute nodes as well as varying levels of user specified processing in order to illustrate these effects.

Each test case consisted of an MQInput and MQOutput node with varying numbers of compute nodes in between. The level of complexity in the compute nodes was also varied. The following cases were measured:

- A simple compute node which copied the input message to an output message. The purpose of this measurement was to show the message throughput which is achievable when copying a message and modifying a single field. This represents the simplest form of compute node.
- A single complex compute node which contained user specified ESQL processing as well as the copying of the input message to an output message. The purpose of this measurement was to show the effect that additional CPU bound processing has on message throughput.
- Multiple complex compute nodes which consisted of five of the complex compute nodes connected in sequence. The purpose of this measurement was to establish the cost of using multiple complex compute nodes.
- A single very complex compute node which consisted of five times the processing of the single complex compute node. The purpose of this measurement was to illustrate the benefit which can be obtained by combining processing within a single compute node.

In these measurements the input message consisted of a single field. The transaction mode on the MQInput and MQOutput nodes was set to automatic.

2.2.1 Simple Compute Node

Figure 2 below shows the results which were obtained as a result of running the simple compute message flow with varying message sizes and persistence. There was a single instance and single execution group running the message flow.




Figure 2: Simple Compute Node Throughput Results

With a 1K non persistent message it was possible to process approximately 440 msgs/second. The message throughput rate declined with size, reflecting the increased volume of data and additional processing required to deal with the messages.

With 1K persistent messages it was possible to process approximately 55 msgs/second. This reduced message rate, when compared with 1K non persistent messages is as a result of the additional logging within the MQSeries manager.

The detailed measurement data for the Simple Compute Node throughput measurements is available in **Section 7.2 - Compute Node Throughput Results**.

2.2.2 Complex Compute Node

Figure 3 below shows the results which were obtained as a result of running a complex message flow with varying message sizes and persistence. See Appendix C for a description of this complex flow. There was a single instance and single execution group running the message flow.




Figure 3: Complex Compute Node Throughput Results

With a 4K non persistent message it was possible to process approximately 75 msgs/second. The message throughput rate declined with size, reflecting the increased volume of data and additional processing required to deal with the messages. The lower message rate achieved with this compute node compared with the simple compute node case above reflects the increased processing which was added to the compute node.

With 4K persistent messages it was possible to process approximately 34 msgs/second. This reduced message rate, when compared with 4K non persistent messages is as a result of the additional logging within the MQSeries manager.

The detailed measurement data for the Compute Node throughput measurements is available in **Section 7.2 - Compute Node Throughput Results**.

2.2.3 Multiple Complex Compute Nodes

Figure 4 below shows the results which were obtained as a result of running five of the above complex nodes daisy chained together for varying message sizes and persistence. See Appendix C for a description of this complex flow. There was a single instance and single execution group running the message flow.




Figure 4: Multiple Complex Compute Node Throughput Results

With a 4K non persistent message it was possible to process approximately 24 msgs/second. The message throughput rate declined with size, reflecting the increased volume of data and additional processing required to deal with the messages.

With 4K persistent messages it was possible to process approximately 18 msgs/second. This reduced message rate, when compared with 4K non persistent messages is as a result of the additional logging within the MQSeries manager.

The detailed measurement data for the Compute Node throughput measurements is available in **Section 7.2 - Compute Node Throughput Results**.

2.2.4 Very Complex Compute Node

Figure 5 below shows the results which were obtained as a result of running a very complex message flow with varying message sizes and persistence. Appendix C has a description of a very complex flow. Briefly, a very complex flow is defined as the ESQL of the complex flow repeated 5 times in the same node. There was a single instance and single execution group running the message flow.

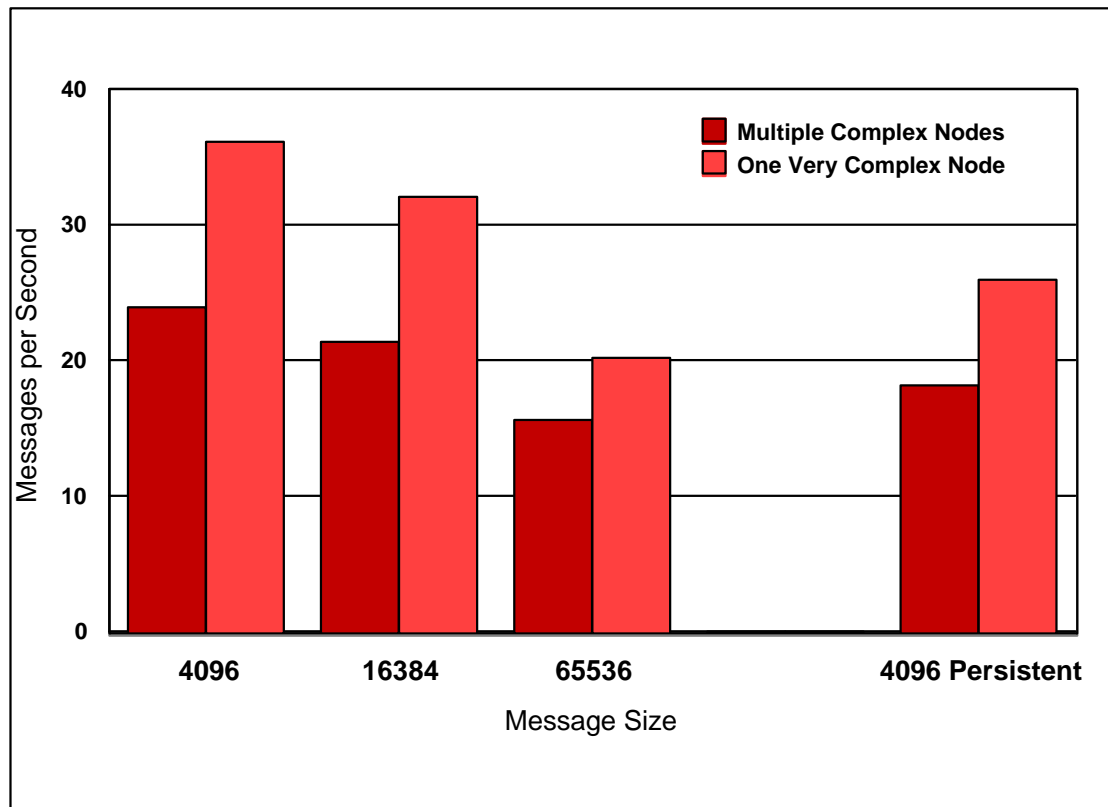


Figure 5: Very Complex Compute Node VS Multiple Complex Compute Node Throughput Results

With a 4K non persistent message it was possible to process approximately 36 msgs/second. The message throughput rate declined with size, reflecting the increased volume of data and additional processing required to deal with the messages. With 4K persistent messages it was possible to process approximately 26 msgs/second. This reduced message rate, when compared with 4K non persistent messages is as a result of the additional logging within the MQSeries manager.

For comparison purposes Figure 5 also shows the message throughput rates which were achieved for the multiple complex compute node case detailed in **Section 2.2.3 - Multiple Complex Compute Nodes**.

For 4K non persistent messages there was a 50% improvement in message throughput as a result of using a single compute node for the processing, rather than using 5 nodes. For performance reasons it is clearly better to have one node which does the work of several less complex nodes. This performance improvement has to be offset against the management and support of more complex nodes.

The detailed measurement data for the Very Complex Compute Node throughput measurements is available in **Section 7.2 - Compute Node Throughput Results**.

2.3 Database Node Throughput

A database node allows a database transaction in the form of an SQL expression to be applied to a specified ODBC data source. The statement to be applied and the data source are specified on the database node definition.

In order for the database transaction to be part of a global unit of work which incorporates the processing of the message within the same transaction, the MQSeries queue manager associated with the MQSI V2 broker must have a suitable X/Open XA interface connection with the database. The MQSeries queue manager associated with the MQSI V2 broker acts as the transaction coordinator. The database is a resource manager.

Without establishing such an XA connection it would not be possible for the database manager to commit (or backout) the database transaction at the same time as the message processing updates in the same unit work. This would mean data was an inconsistent state.

In this simple test to illustrate the effect of using a database node an XA connection was configured between the MQSeries queue manager and the database, DB2 in this case and a message flows consisting of an MQInput node, a database node and an MQOutput node was defined.

The message flow consisted of an insert/delete for a row in a table of a database. The transaction mode value on the MQInput node was set a value of automatic. The coordinatedTransaction value for the message flow was set to yes. The effect of doing this is to specify that the message flow should be a globally coordinated unit of work. At present the Globally Coordinated Unit of Work is only supported with DB2.

The maximum possible message throughput rates were determined for a single instance and single execution group running the message flow. **Figure 6** below shows the results which were obtained for varying message size and persistence.

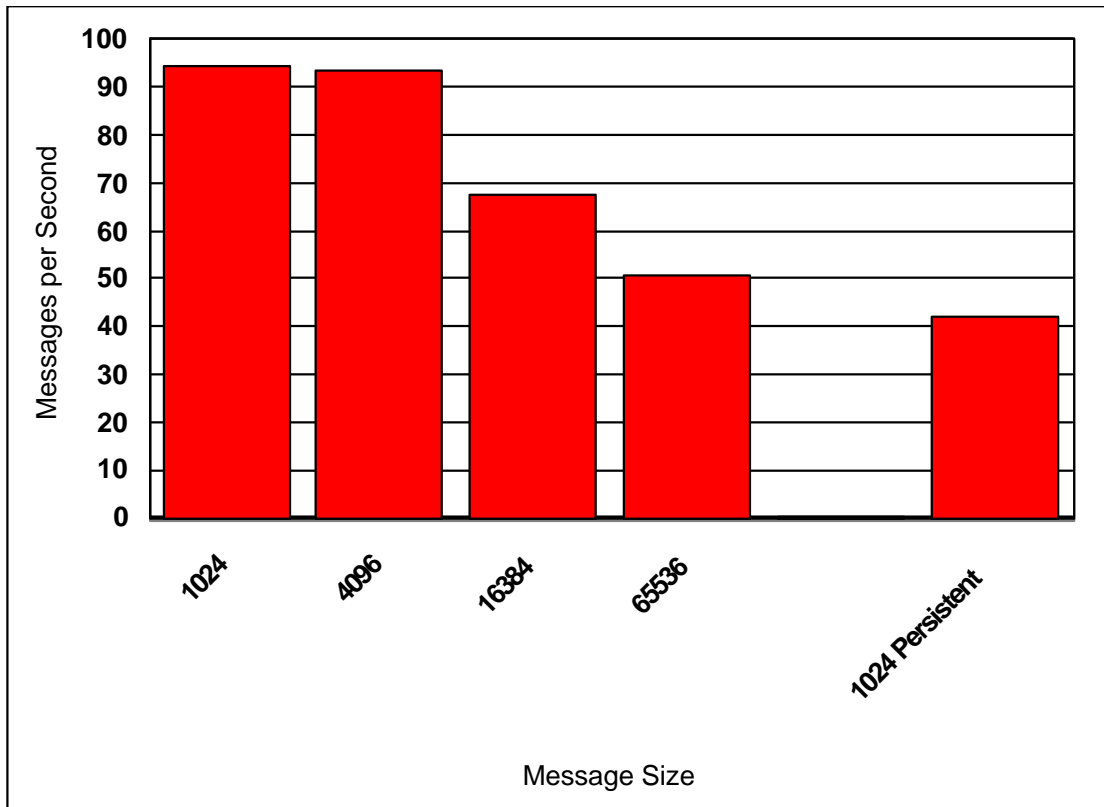


Figure 6: Database Insert/Delete Throughput Results

With 1K non persistent messages it was possible to achieve a message throughput rate of 95 msgs/second. This is 47 database insert and deletes per second. The rate of insert/delete activity reduced with message size as expected.

With 1K persistent messages it was possible to achieve a message throughput rate of 42 msgs/second. This is 21 insert and deletes per second. This lower rate is due to the increased volume of I/O processing to both the MQSeries queue manager log and the DB2 log.

The detailed measurement data for the Database node measurements is available in **Section 7.3 - Database Node Throughput Results**.

2.4 Filter Node Throughput

A Filter node evaluates an ESQL expression against the content of the input message. Based on the result of the expression evaluation the message is propagated to the true terminal if the expression evaluates to true. It is propagated to the false terminal if the expression evaluates to false.

A message flow consisting of an MQInput node, a Filter node, and an MQOutput node was defined. The Filter node processing involved selecting a message on the basis of the contents of a tag value. The input message consisted of an MQRFH2 message with two tags specified following the header. The transaction mode on the MQInput and MQOutput nodes was set to automatic.

Figure 7 below shows the results which were obtained as a result of running the message flow with varying message sizes and persistence. There was a single instance and single execution group running the message flow.




Figure 7: Filter Node Throughput Results

With 1K non-persistent messages it was possible to run approximately 1080 msg/s. The difference between this rate and the 1500 achieved with an MQInput/MQOutput node pair for 1K non-persistent messages represents the overhead of using a Filter node. The cost of the Filter node will vary with the complexity of the filter expression and the number of fields in the input message.

With 1K persistent messages the throughput was 54 msg/s. The reduction in throughput is as a result of using persistent messages which involves additional logging within the MQSeries manager as well as the fact that the message is processed under MQSeries transaction control.

The detailed measurement data for the Filter Node Throughput measurements is available in **Section 7.4 - Filter Node Throughput Results**.

2.5 Publication Node Throughput

A publication node may be used within a message flow to represent a point from which messages are "published" that is, a point from which messages are transmitted to a set of subscribers who have registered interest in a particular set of messages.

A message flow consisting of an MQInput node and a Publication node was defined. The transaction mode on the MQInput node was set to automatic. The measurement used topic routing. The messages used the MQRFH2 format header.

In the throughput measurements each client thread performed the role of publisher and subscriber. Firstly, an MQPUT was issued to publish a message on the given topic. Secondly, the client thread issued an MQGET to receive the published message. At this point the client thread was a subscriber.

With this mode of operation there were:

- Two messages between the MQ client and queue manager
- One published message
- One subscriber

Figure 8 below shows the results which were obtained as a result of running the message flow with varying message sizes and persistence. There was a single instance and single execution group running the message flow. The rates shown are the rate at which messages are being received by the subscriber.




Figure 8: Publication Node Throughput Results

With 1K non-persistent messages it was possible to run approximately 640 msgs/second. This represents a publication and subscription rate of 320 messages per second.

As the message size increased, the rate at which messages were published decreased. This is as expected.

With 1K persistent messages the throughput was approximately 65 msgs/second. The reduction in throughput is as a result of using persistent messages which involves additional logging within the MQSeries manager as well as the fact that the message is processed under MQSeries transaction control.

The detailed measurement data for the Publication Node Throughput measurements is available in **Section 7.5 - Publication Node Throughput Results**.

In a separate measurement the overhead of using Access Control for topic publications was measured. The measurement consisted of a 1K non persistent message size and 50 subscribers. The overhead was an 18% reduction in message throughput.

2.6 What Is The Cost Of Converting Messages To Different Formats ?

MQSI V2 provides the capability to process messages of different formats as well as the ability to convert messages between formats. Throughput measurements were taken to show the effect of using MQSI V2 to convert messages between MRM XML, Generic XML and CWF formats, where MRM XML refers to the predefined XML used within the MRM, Generic XML refers to self-defining XML and CWF denotes a legacy data structure such as a C structure or COBOL copybook.

The same message type was used for each of the conversions. This was a 4096 byte non persistent message containing 30 input fields, with 10 fields consisting of a short string (12 characters), 10 fields a floating point number, and 10 integer fields.

The format conversion was achieved using a Compute node with suitable ESQL statements. The input messages contained an MQRFH2 header in which the message type was set. The output format was specified in the Compute node processing. Each message format was converted to MRM XML, Generic XML and CWF and the message throughput achieved was measured. The results are presented in **Table 1** below. (Figures are given as an Input/Output message rate.)

Table 1: Message Rates when Converting Between Different Formats.

Conversion	TO	Generic XML	CWF	MRM XML
FROM				
Generic XML		164.00	109.00	107.00
CWF		78.70	47.80	47.80
MRM XML		94.80	66.40	66.80

Even when the output message is set to have the same format as the input message there are still significant costs in processing messages because the messages must be parsed, deconstructed by MQSI, and reconstructed into the required output format. However, the cost of converting between two formats occurs on a once per message flow basis and not in each node.

Because of the cost of conversion between different message formats it is recommended that is used only where necessary.

The detailed measurement data for cost of message conversion is available in **Section 7.6 - Converting Messages Between Formats Throughput Results.**

2.7 How Well Do NEON Rules and NEON Formatter Run Within MQSI V2?

MQSeries Integrator V1 consisted of NEONRules and NEONFormatter components. It is possible to continue to use those components, unchanged, in MQSI V2.

Throughput measurements were taken in order to show the relative performance of using NEONRules and NEONFormatter within MQSI V1 and MQSI V2 node. Two test cases were used. They were both run within MQSI V1 and MQSI V2.

In the first case an input message with 50 fields was copied to an output message with 1 NEON rule present. The purpose of this measurement was to show the cost of reformatting a message using NEONFormatter. In the second test case the 50 field message was again copied but there were 100 rules defined, only 1 of which was executed - all are evaluated though. The purpose of this measurement was to show the effect of having an increasing number of rules to evaluate.

In both test cases three copies of the NEONRules engine were run. A NEONRules node was used when running with MQSI V2. The same MQSeries input and output queues were used for all 3 copies of the NEONRules engine. When NEONRules was run within MQSI V2 it was run with transaction mode set to yes so as to be consistent with the transactional support provided by NEONRules in MQSI V1.

Figure 9 below shows the results which were obtained as a result of running with 1 rule defined when using the NEONRules engine within MQSI V1 and within MQSI V2.

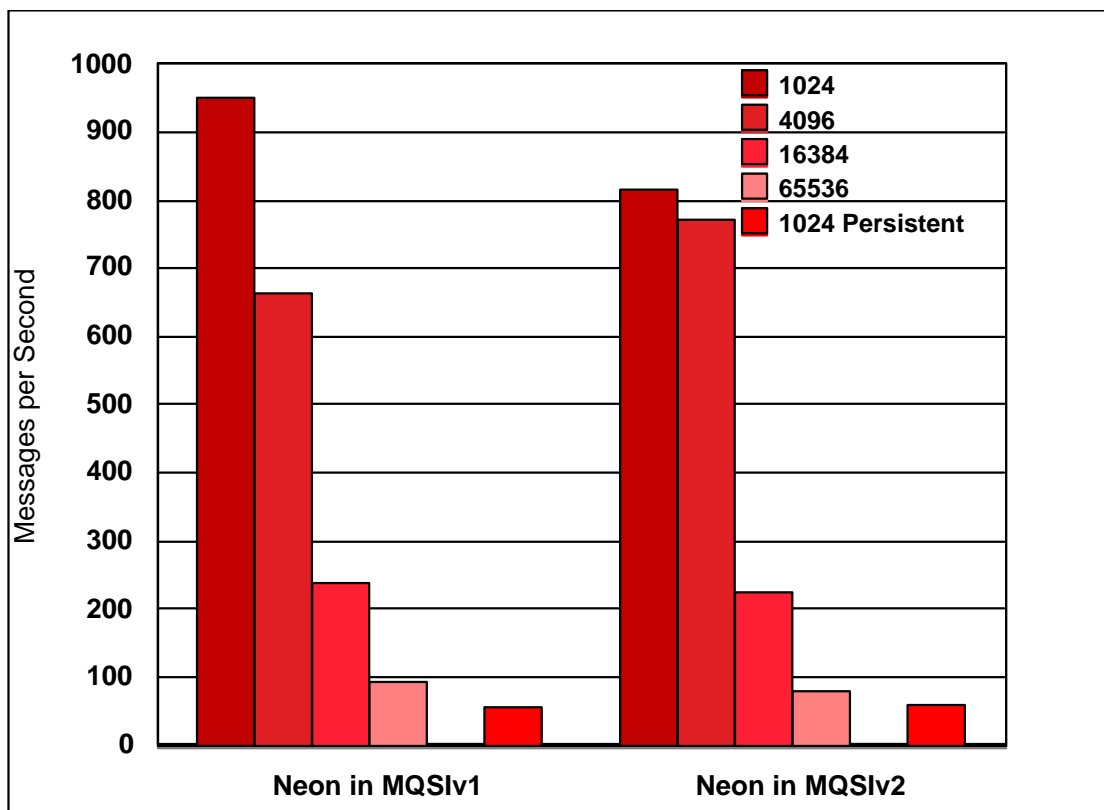


Figure 9: NEONRules Engine With One Rule Throughput Result

With 1K persistent messages the throughput is again almost equivalent when running NEONFormatter within MQSI V2 compared with running NEONFormatter by itself.

Figure 10 below shows the results which were obtained as a result of running with 100 rules defined when using the NEONRules engine within MQSI V1 and within MQSI V2.

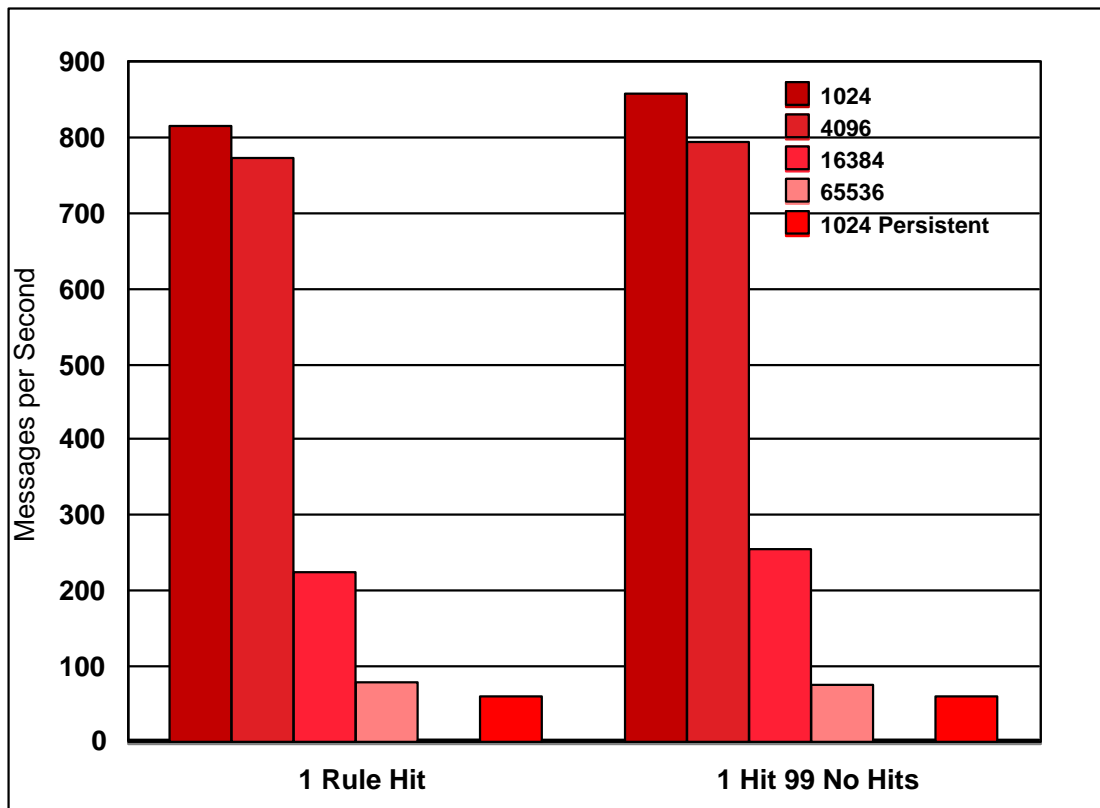


Figure 10: NEONRules Engine With One Hundred Rules Throughput Results

There was no significant difference between measuring one rule against one hundred rules. The slight difference in the charts can be attributed to experimental inaccuracies such as network traffic etc.

The detailed measurement data showing the effect of running the NEONRules engine within MQSI V2 is available in **Section 7.7 - NEONRules and NEONFormatter Within MQSI V2**.

2.8 What Is The Effect of Increasing The Number of Execution Groups?

A message flow runs within an execution group. An execution group is implemented as an operating system process. It is possible to have multiple execution groups running the same message flow. Reasons for doing this would be to achieve greater message throughput than is possible with a single execution group, and to provide some level of redundancy should an execution group process fail.

Figure 11 below shows the results which were obtained as a result of running one, two, and four execution groups for a complex compute message flow for varying message size and persistence. The message flow is described in Section 8.1 Complex Compute Node. The transaction mode values on the MQInput and MQOutput node were set to the value of automatic. The same input and output queues were used for all measurements.




Figure 11: Additional Execution Group Throughput Results

Figure 11 shows that greater message throughput can be achieved by using additional execution groups. With non persistent 1K messages and two execution groups it was possible to achieve twice the throughput that was achieved for a single execution group. When using four execution groups it was possible to achieve almost 4 times the throughput that was achieved with a single execution group.

Whilst non persistent messages to scale well the persistent messages show some improvement from one to two execution groups and then a limit is reached. This is due to contention for the MQSeries log disk.

The benefits of running multiple execution groups for non persistent messages in this case was significant. There was almost perfect scaling of message processing. This is principally because of the nature of the message flow which was used for the measurements. There was a significant amount of ESQL processing in the node. This meant that the level of queue access as a proportion of

all processing was low and so the potential for conflicts on queue access was low and consequently multiple execution groups were able to achieve greater throughput.

The lesser benefit gained from additional execution groups for persistent messages is due to the constraint of the MQSeries Queue Manager Log. By locating this log on a faster disk it is expected that multiple execution groups would show a more significant increase in message rate.

The detailed measurement data showing the effect of adding execution groups is available in **Section 7.8 - The Effect of Increasing The Number Of Execution Groups.**

2.9 What is The Effect of Making a Message Flow Transactional?

Making a message flow transactional (as opposed to making an individual node transactional) means that the unit of work is recoverable, but it does result in an additional overhead as work must now take place under transactional control. This involves the locking of data and logging of data images.

The purpose of these measurements was to illustrate the overhead of making a message flow transactional. A simple message flow was created consisting of a single MQInput and MQOutput node. The maximum message throughput rate was measured when the message flow had a transaction mode value of automatic and then with a value of yes.

Figure 12 below shows the results which were obtained as a result of running the message flow with varying message sizes and persistence.

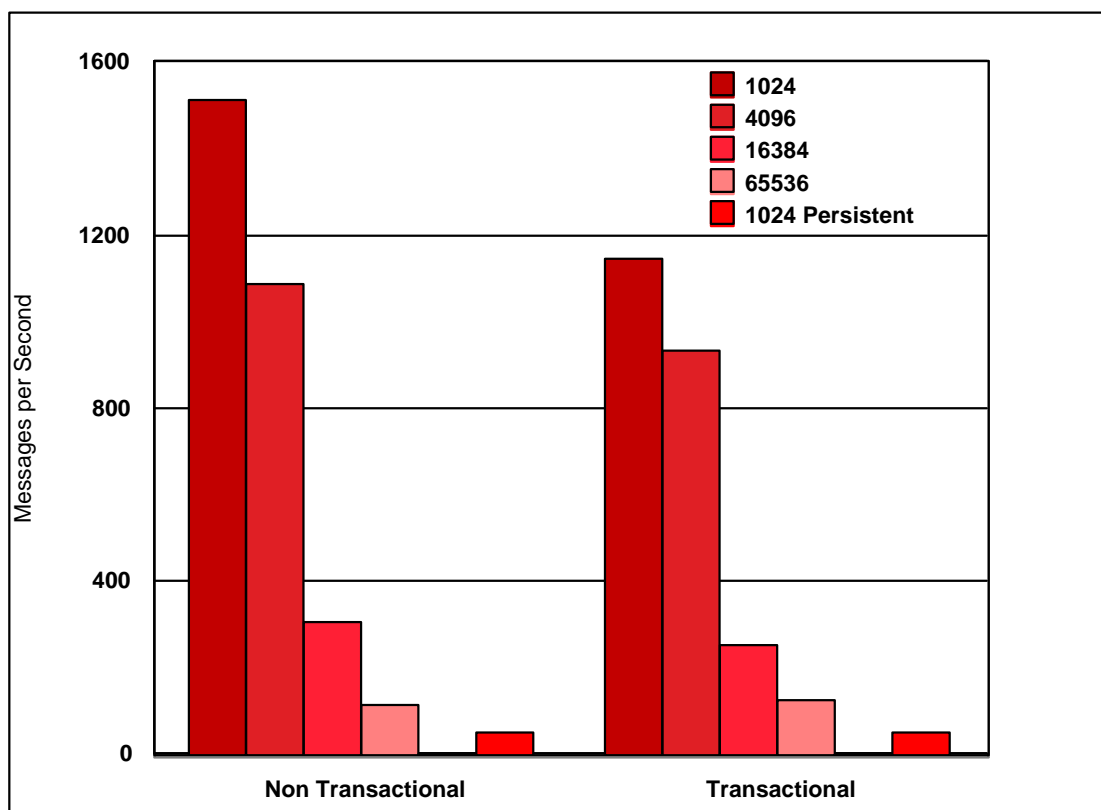


Figure 12: Making a Message Flow Transactional Throughput Results

Making non-persistent messages transactional had a significant effect on message throughput. The reduction in throughput is as a result of the additional CPU and I/O processing that must take place. The overhead of making the message flow transactional was most significant with the smaller message sizes.

For the persistent messages there is no difference in throughput as persistent messages would proceed under transaction control anyway with a transaction mode of automatic.

The detailed measurement data for showing the effect of making a node transactional is available in **Section 7.9 - The Effect of Making a Message Flow Transactional**.

2.10 What Is The Effect of Using `coordinatedTransaction=yes` on a Message Flow?

Specifying a value of `yes` for the `coordinatedTransaction` parameter on a message flow means that all updates performed within the message flow will take place as a global unit of work. Any database updates which are in the message flow will be committed atomically with the message processing. In order to ensure that the global unit of work is coordinated correctly a suitable XA connection must be configured between the MQSeries queue manager and the external relational data manager.

If data is to be updated in an MQSeries message and a relational database, and recovery is required, this configuration must be used.

Measurements were taken to illustrate the costs associated with using the `coordinatedTransaction` parameter on a message flow. The message flow consisted of an MQInput, database and MQOutput node. The database node performed an update of a row of a relational database. The purpose of this measurement was to illustrate the cost of an XA coordinated transaction.

Figure 13 below shows two sets of results. The first is for the case when `coordinatedTransaction` was set to `no` on the message flow and transaction mode was set to `automatic` on the MQInput and MQOutput nodes. The second case shows the results which were obtained when an XA connection was configured and a value of `yes` was specified for `coordinatedTransaction` on the message flow.




Figure 13: Database Update with XA and `coordinatedTransaction=yes`.

Figure 13 shows that a message rate of approximately 230 msgs/second was achieved for 1K non-persistent messages and 59 msgs/second for persistent messages when `coordinatedTransaction` was set to `no` on the message flow, transaction mode was set to `automatic` on the MQInput and MQOutput nodes and there was no XA connection configured.

With an XA connection configured and `coordinatedTransaction` set to `yes` on the message flow a rate of 100 msgs/second was achieved for non persistent messages and 46 msgs/second for persistent messages.

The difference in message rates represents the overhead of coordinating the updates as a single unit of work. In this case the cost is 57% for 1K non persistent tending to zero as the message size increases. This can be attributed to the extra data MQ is transporting. For persistent 1K messages the difference in message rate is much smaller. This can be attributed to the fact that MQ is already logging the messages as they are persistent. The database processing in these measurements was simple. In practice it would typically be more involved, in which case the overhead imposed as a result of using `coordinatedTransaction=yes` would be less as a percentage of the total cost.

The detailed measurement data showing the effect of using `coordinatedTransaction` set to the value of `yes` is available in **Section 7.10 - The Effect of using `coordinatedTransaction=yes`**.

2.11 What Effect Does an Increasing Number of Subscribers Have on Publish/Subscribe Throughput?

As an increasing number of subscribers register an interest in receiving published messages on a given topic, so the broker must undertake additional processing to maintain a list of currently subscriptions and write a message to each subscribers queue when a message is published.

In order to illustrate the effect of coping with an additional number of subscribers for a given topic a series of measurements were taken with 1, 10, 30, 50, 70, 100 and 1000 subscribers. Messages of varying size and persistence were published to a single topic. The results obtained are presented in **Figure 14** below. The X axis shows the number of subscribers. The Y axis shows the number of seconds taken to process a message. It is derived from the reciprocal of the message rate.




Figure 14: Varying Number of Subscribers.

From the graph it is possible to see that the processing required to deliver messages to the subscribers rises with the increasing number of subscribers. This makes sense since with each additional subscriber there is an additional MQSeries queue to write a message to.

The cost of publishing persistent messages is significantly higher as the processing is dominated by the necessary I/O processing. This is reflected in the steeper gradient of the 1K persistent message measurements.

Before examining the measurement data for varying number of subscribers it is important to understand the way in which the measurement was taken.

For each subscriber that registered to receive publications the published message was written to a queue for that subscriber. With 30 subscribers for example, a single message was written to each of 30 queues. In the measurement environment there was a background program consuming all but one of the published messages. Taking the example of 10 subscribers, 9 of the published messages were consumed by this program. The remaining message was read by the client program emulating the subscriber. In this situation a message count of 2 was registered for the purposes of reporting message rates, although the MQSI V2 broker had written multiple messages. It is because of this that the reported message rate declines with an increasing number of subscribers, although the level of work performed by the broker is obviously much greater with an increasing number of subscribers.

The 1000 subscribers case has not been shown graphically since there is not adequate space to show all the measurements meaningfully. However, the results show the straight line trend continuing for 1000 subscribers as it does for 1 through 100. This can be seen from the detailed measurement data held in the appendix. For 1K non persistent messages the rates measured for 100 and 1000 subscribers were 27.6 and 2.6 messages per second respectively. It can be seen that the addition of extra subscribers corresponds to a scalar cost to the message rate with 1000 subscribers giving 1/10th of the rate achievable with 100.

From some additional measurements which were taken it has been shown that the performance obtained when using publication nodes is dependent on the number of subscriptions which receive messages as a result of a message being published rather than the total number of registered subscribers. For example consider the case where there are 1000 subscribers registered to receive messages on two topics, with 990 registered for Topic 1 and 10 for Topic 2. When a message is published on Topic 2 it will only be the cost of publishing the message to the 10 subscribers registered for Topic 2 that will be seen. The fact that there are 1000 subscribers in total has no noticeable effect on performance.

All measurements in this section with more than 20 subscribers used the `mqsichangeproperties` command to increase the size of the cache used to hold open queue descriptors. The default value is 30. If the number of open queue descriptors has to increase by 1 beyond the cache size, a queue must be closed before another can be opened and the published message delivered. This can have a significant effect on the rate at which messages can be published. This sequence of closing one queue and opening another will occur each time a message is published unless the cache size is increased. It is recommended to increase the size of the cache to exceed the number of registered subscribers.

The `mqsichangeproperties` command used to increase the cache for the 1000 subscriber measurements was as follows.

```
mqsichangeproperties CSIM -e pubsub1 -o ComIbmMQConnectionManager  
-n queueCacheMaxSize -v 1050.
```

Where CSIM was the name of the MQSI Broker and pubsub1 was the name of the execution group running the message flow. This command was issued on the machine on which the Broker was running.

In one measurement there was a 60% increase in message throughput as a result of increasing the cache size.

The command must be issued for each execution group containing a publication node with more than 30 subscribers.

The detailed measurement data showing the effect of an increasing number of subscribers is available in **Section 7.11 - The Effect of Increasing the Number of Subscribers**.

2.12 What Is The Effect of Using Content Based Publish/Subscribe?

With MQSI V2 it is possible for a subscriber to register to receive publications based on the contents of a message. This is known as content based routing. This is more complex than topic based routing (in which a subscriber receives all messages on the topic to which they have subscribed) because the message contents must be examined.

In order to determine the cost of using content based routing a measurement was taken to examine the message rates which could be achieved for both content and topic based routing when delivering messages of the same structure and size.

Figure 15 below shows the message throughput rates which were achieved for content and topic based routing with varying message size and persistence. The topic and content based routing used MQRFH2 messages since this format must be used for content based routing.

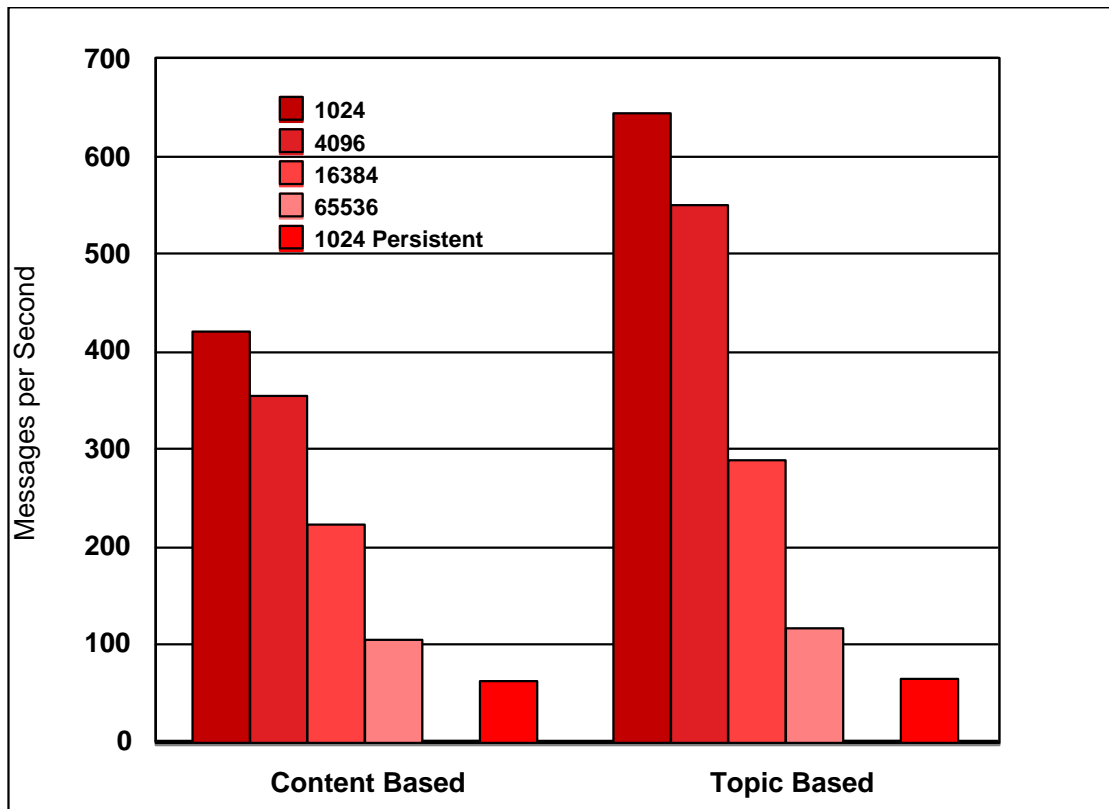


Figure 15: Content vs Topic Based PubSub

Figure 15 shows that it is possible to achieve greater message throughput using topic based routing when compared with content based routing. This is as expected since topic based routing is able to publish a message without regard to the message contents. Content based must parse the message contents, then determine which of the subscribers is to receive the message before finally publishing to those subscribers.

The detailed measurement data showing the comparison of content and topic based PubSub is available in **Section 7.12 - Content vs Topic Based PubSub**.

2.13 What Is The Effect of Examining an Increasing Number of Fields in a Message?

When message manipulation is field related, as opposed to copying the entire message for example, then the number of fields makes a significant difference to the processing that is required to process the message.

In order to illustrate the effect of processing an increasing number of fields in a message, message throughput was measured for an XML message with a single field and compared with the results for a 50 field XML message when being processed by a compute node.

Figure 16 below shows the message throughput rates obtained with varying message size and type.

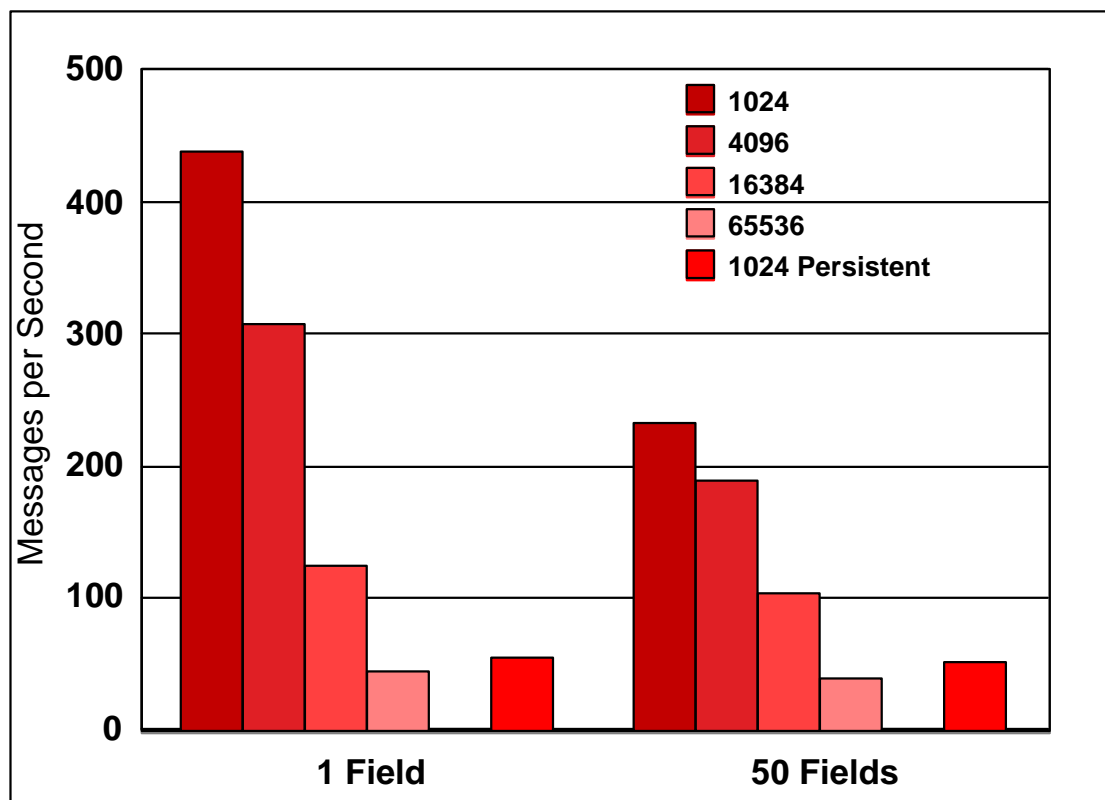


Figure 16: The Effect of an Increasing Number of Fields.

With non persistent messages for all message sizes it was possible to process a greater throughput of messages where there was a single field in the input message. The number of messages processed with persistent messages was very similar. This is because message processing is dominated by MQSeries log I/O.

A recommendation from this measurement is to minimize message size and number of fields where possible.

The detailed measurement data showing the effect of an increasing number of fields in a message is available in **Section 7.13 - The Effect of an Increasing Number of Message Fields**.

2.14 What Is The Effect of Using User Trace?

On occasions it is necessary to trace the course of processing within MQSI V2. There are 2 levels of trace available, service trace and user trace. There are also two modes of trace, safe and fast. With safe mode trace entries are always written to disk. This mode involves a greater amount of I/O activity to the trace file in order to ensure that all entries are written. With fast mode, data is buffered in an internal buffer and subsequently written to disk.

From measurements taken with MQSI for AIX - V2.0.1 there is an overhead of upto 97% when using user trace with a level of normal and mode of safe. The overhead is most significant for the smaller non persistent message sizes. The overhead is high because of the use of the safe mode, which immediately forces writes of trace data to disk.

2.15 What Is The Effect on Message Throughput of Using Faster Disks?

With persistent messages the limiting factor for message throughput is the speed at which the MQSeries queue manager log is able to operate at. The lower the seek time and latency of the disk on which the MQSeries log is located, the more beneficial the effect on message throughput.

From measurements taken with MQSI for Windows NT - V2 using SSA disks with a fast write cache defined for the disk it was possible to show that there was a three and four times increase in message throughput rate.

Although the measurements were taken on the Windows NT platform there is every expectation that the same results would be obtained on an AIX system with similarly fast disks. The benefit in persistent message throughput comes as a result of reduced I/O time.

The fast write cache was the key to the improvement in message rate. When running with the MQSeries queue manager log located on a SSA disk without a fast write cache defined the message throughput rate resulted in a similar message rate to that achieved with a SCSI disk.

Further details of the measurements can be obtained from version 1.1 of the MPI4 Support Pac.

3.0 CAPACITY PLANNING

This section gives general guidelines on capacity planning for MQSI V2. A tool to provide detailed capacity planning which takes into account the contents of a message flow will be provided in the near future. The remainder of this section provides general guidelines to assist with the implementation of MQSI V2.

When capacity planning for the introduction of a new software it is important to be able to establish two things:

- How much resource (CPU, disk, memory) is required to support the required message rate.
- IF it is possible to run the software at the expected message rate. Has the software been run at that rate before, will the deployed system be working within known limits ?

The way in which MQSI V2 can be used varies enormously. For this reason it is not possible to provide detailed guidance on the resources which would be required for all possible configurations of MQSI V2. It is possible, however to provide a series of guidelines to get an initial estimate of the required capacity. Once a prototype implementation has been developed, future resource requirements can be based on measurement and observation of the prototype and its successor implementations, which is the only meaningful exercise for you. The problem is invariably one of getting started. This section helps with that process.

In gauging the capabilities of MQSI V2 to process messages at the required rate examine the various throughput measurements detailed in Section 2.0 - BROKER THROUGHPUT MEASUREMENTS. This section details the highest rates at which MQSI V2 has been run.

Projecting measurement results to other machine types is difficult because performance depends on many factors, such as processor speed and instruction cache sizes. The only method currently available to gauge relative machine performance is to use published performance figures such as SPECint benchmark results which are released by manufacturers for their hardware.

3.1 Throughput

Key factors affecting the throughput rate that is achievable are:

- Use of non-persistent vs persistent messages.
- The types of MQSI V2 nodes being used.
- The amount of processing in the nodes, simple vs complex compute nodes for example.
- The number of MQSI V2 nodes being used.
- The number of times messages are written to an MQSeries queue from a node.
- The complexity of processing in the nodes, for example complex ESQL statements vs simple copying.

Of the above, those having most effect on message throughput are:

- The use of persistent messages. If persistent messages are used the maximum message rate will be governed by disk performance rather than by processor speed.
- The amount of user supplied (ESQL) processing in compute nodes.
- The average number of subscribers to match a topic for a Publication node.

Processing of non persistent messages is CPU intensive. It is therefore important to ensure that there is sufficient CPU available to process the message in order to maximize throughput. The measurements contained in this report are based on the use of POWER3-II processors (375MHz). Using faster or slower processors will have a corresponding effect.

Processing of persistent messages is I/O intensive. It is therefore important to ensure that the MQSeries queue manager log is located on a dedicated fast disk. When running persistent messages CPU utilization will be lower than for non persistent messages and you would not expect to be able to fully utilize processors in the same way in which it is possible to do so for non persistent messages.

3.2 Scaling Message Throughput

The addition of multiple execution groups provides the opportunity to increase message throughput. This is normally done because the message rate that is achievable with a single execution group is not sufficient for the planned message rates. In general the ability to scale message throughput depends on a number of factors:

- The availability of sufficient resources (CPU, disks, memory) to cope with the increased resource demands as a result of simply processing more messages.
- The ability to schedule multiple pieces of work in parallel.
- A minimum of contention between the parallel pieces of work.

If we look at each of these issues with respect to MQSI V2: resolving item 1 above depends on having sufficient hardware of sufficient speed (CPU and disk) available. This is a principally a planning issue; For item 2, MQSI V2 provides the ability to schedule multiple pieces of work in parallel by the use of multiple execution groups for example; The contention between pieces of work (item 3) depends on the message type (persistent vs non persistent messages), the amount of access to MQSeries queues (high level of access vs low level) and the nature of additional database processing where it is used (Insert/Update/Delete vs read only).

The level of contention in any implementation is largely determined by the nature of the application being implemented and minimizing the effects of contention is an essential part of the application architecture, design and implementation. This is a large subject and is not covered in any level of detail in this report. Consider the following cases as an illustration:

Case 1

In a very simple message flow where messages are simply being copied from one queue to another as fast as possible the potential to keep increasing message throughput by the addition of more execution groups is limited because of contention for access to the queue in MQSeries. This is even more of an issue with persistent message since queues are locked for longer and there must be I/O to the MQSeries log.

Measurements show that when using 1, 2 and 4 execution groups to perform copying of non persistent messages from one queue to another at the maximum rate possible that there can be some benefit (maybe an additional 30% in throughput) from using a second execution group but there after there will be reduced benefit and even a decrease because of the contention for queue access. With a higher number of execution groups context switching between processes can rise significantly and although the processors may become busier the level of useful work does not rise. One way around this contention would be to use multiple message flows each with its own pair of input and output queues.

It is also possible to show an increase in throughput with persistent messages. The maximum number of persistent messages that can be processed will be lower than for non persistent messages since there must be I/O to the MQSeries log. In order to maximize message throughput the MQSeries log I/O processing time must be reduced. This can be done in several ways. The most effective method

would be to use a solid state disk on which to locate the MQSeries queue manager log. Another possibility is to use a disk subsystem with a fast write nonvolatile cache.

A message flow that has a significant amount of processing over and above any queue processing has much greater potential for increasing message throughput by the addition of more execution groups for example.

The increased throughput rates are possible because firstly the level of queue access overall is much lower and secondly there is processing that each execution group can perform independently.

Case 2

The Complex Compute node described in Appendix C, Complex Compute Node is a good example of a message flow that scales well. This message flow was measured with 1, 2 and 4 execution groups. The results are presented in Section 2.8, What Is The Effect Of Increasing the Number of Multiple Execution Groups.

The results of the measurements show that it was possible to achieve almost perfect scaling in message throughput with non persistent messages as a result of adding additional execution groups. By using non persistent messages the processing is CPU bound. The level of queue access is reduced because of the ESQL processing on the message. As a result the contention for queue access seen in Case 1 is not an issue.

With persistent messages the increase in message throughput also increased with more execution groups but at a lower rate than for non persistent messages. The maximum rate achievable is determined by the rate at which the MQSeries queue manager log is able to operate at. This is in turn determined by the speed of I/O for the device on which the log is located.

These cases above show that there is a very wide range of scaling from little benefit to 100%. Given the availability of sufficient resources, the following will determine the scaling ratio that can be achieved:

- The type of message being processed (persistent vs non persistent). Persistent messages have a lower message processing rate limit than non persistent messages. Persistent messages are I/O bound whilst non persistent are CPU bound. Once the I/O limit is reached throughput cannot be increased.
- The level of contention for resources such as MQSeries queues or rows in a database. If there is a common input and common output queue for message processing there is a greater potential for conflict than if there are multiple input and output queues. Similarly if multiple message flows are attempting to update the same row in a database there will be conflict for access to that row. Bypassing these issues largely comes down to application architecture.
- The amount of additional CPU processing in the message flow, such as ESQL in a compute node.
- The message throughput rate that is currently being achieved with a single execution group. If a single execution group is able to fully utilize the I/O capacity of the MQSeries queue manager log or database manager log for example then it is unlikely that additional instances will allow any more throughput to be obtained because the message flow is already constrained by the rate at which log I/O can take place. Similarly, if a message flow is simply copying messages from one queue to another at a very high rate, it is unlikely that the addition of another copy of this message flow would achieve significantly more throughput as there would be contention over queue access.

In planning your configuration it is important to establish where the constraints in the configuration will be, since these will affect the ability to scale throughput. Determine whether the system will be CPU or I/O bound. If it is CPU, ensure that the fastest processors available are used. If the system is I/O bound allocate as many physical disks as possible to split log and data. Also use the fastest devices available (disks with fast write nonvolatile cache or solid state disk). If it is not possible to remove the

constraints consider using multiple brokers on the same machine or multiple brokers over multiple machines in order to achieve the required message rate.

3.3 Memory

In estimating memory requirements for MQSI V2 there are a number of components that need to be considered. These are:

- The Control Center. There are likely to be multiple Control Centers in use.
- The Configuration Manager. There is one Configuration Manager per MQSI V2 implementation.
- The Broker. There may be multiple brokers and within these multiple execution groups and therefore multiple operating system processes.
- MQSeries Queue Manager. There will be one queue manager per broker. There will also be another Queue Manager for the Configuration Manager and another if a User Name Server is required.
- Relational Database. A database is required to hold information on behalf of the Configuration Manager and Broker. Additional relational databases may be in use which hold business data. Throughout this report a DB2 database was used for all of these purposes.

For the Control Center an initial recommendation is to allow 100MB memory per Control Center. This would be for development use.

The Configuration Manager and its associated DB2 database and queue manager should have a minimum of 256MB of memory available in a development environment, but you are recommended to have more available.

The amount of memory required by a broker will depend on the way in which it is configured. A guideline is to allow 300MB for MQSI V2 and its dependent software(broker related components only, no Configuration Manager or Control Center), with an additional 25 MB per running execution group. This recommendation is based on an MQSeries queue manager configuration consisting of 45 SVRCONN channels and a small number of queue definitions (less than 25). If the number of MQSeries resources (channels, queues etc.) to be configured in a system is different you must make an allowance and amend the amount of memory required accordingly.

3.4 Recommended Minimum Configurations

This section contains recommendations on the type of hardware on which an MQSI V2 configuration should be based when running in production. These are only recommendations and are not a substitute for a formal planning and sizing exercise in which requirements are accurately determined.

For production use it is recommended that the components of MQSI V2 are allocated over multiple machines with the following purposes:

- One or more machines to support Control Center usage.
- One machine to support the Configuration Manager. This may also include one Control Center.
- One or more machines to support brokers.

By following this organization the brokers can run in a shielded environment as they process messages. It is important that this processing proceeds without competition for resources from other processes in order to ensure the smooth flow of messages through the enterprise.

A recommended machine specification for the Control Center is a fast uni-processor (Pentium III 500 Mhz processor) with 256MB memory.

A recommended machine specification for the Configuration Manager is a fast uni-processor (Pentium III 500Mhz processor) with 512MB or more of memory.

The specification of the broker machine is more difficult to determine since it requires knowledge of the expected message rate, the types of node which are to be used, and the level of transaction control which is to be used. A recommended minimum specification would be a 2 way processor with 512MB memory. The specification may need to be upgraded if message rates are very high or there are many execution groups. In such cases more detailed planning would be required. Prototyping and benchmarking should be considered in order to accurately determine resource requirements. The results produced will then be specific, and tailored to, the individual configuration being built.

If persistent messages are to be used, the use of solid state disks or disks with a nonvolatile fast write cache is recommended for the device on which the MQSeries queue log manager is located. Where the message rate is less than 50 msgs/second per second fast I/O will improve message response time only. Where the rate is greater than 50 msgs/second then there will be an improvement in message throughput.

A separate disk is also recommended for the MQSeries queue manager queue data. This disk need not have a fast write capability.

If business data is accessed from a relational database the database log and data should each be located on dedicated disks. Consider using a fast device for the database manager log.

4.0 PERFORMANCE RECOMMENDATIONS

This section contains a number of recommendations to assist in the planning and implementation of an efficient MQSI V2 configuration.

4.1 *Understand Recovery Requirements*

In designing message flows within MQSI V2 it is important that the subject of data recovery is approached from the top down rather than bottom up. If you do not consider the recovery needs as a whole, it is possible that more logging will be undertaken, than is actually required. This will lead to a drop in the throughput rate which is achievable with a message flow as the flow becomes I/O bound.

In designing the message flow it is important to establish whether the whole message flow is to be made a recoverable or whether only certain parts of it are. It is also important to establish whether external resource managers such as a database are required. Establish whether data updated in an external resource manager is to be committed within a global unit of work or not.

If a message flow is to involve data held in an external resource manager, i.e. not a queue manager, then consider using the `coordinatedTransaction` parameter in order to make all changes to external data within the scope of a single unit of work. In order for this mode to function correctly the MQSeries queue manager associated with the broker must have an XA connection to each of the external resource managers.

Think carefully about when deciding to use MQSeries transactional control on messageflows. Maybe this is something that is only required for persistent messages.

4.2 *Optimize Queue Manager*

The performance of the underlying queue manager for a broker plays a key role in the performance which can be obtained from using MQSI V2.

In order to improve overall performance with the queue manager consider minimizing message sizes and only use persistent messages where required.

With non persistent messages there is little that can be done to optimize queue manager performance other than ensuring that there is sufficient memory and CPU available.

With the persistent messages the limiting factor is the speed at which the MQSeries queue manager log operates. To minimize the amount of logging taking place, and to improve the efficiency where possible consider the following points:

- The MQSeries queue manager log and queue data should be configured on individual, dedicated, disks.
- Use the fastest disks available for the MQSeries log. Solid state disks give the best possible performance but are expensive. The next alternative is to use disks which have a non volatile fast write cache, such as SSA .
- Run with parallel applications rather than a single application. There is some benefit to be obtained from tailcoating on log I/O which occurs when there is more than one application running with the queue manager.

4.3 Configuration Considerations

Consider the following points when building an MQSI V2 configuration:

- You are recommended to use a separate database instance for each of the Configuration Manager and each Broker.
- You are recommended not to use the database instances for the Configuration Manager or Broker to hold business data.
- You are recommended to ensure that the database instance for the Configuration Manager is local to the machine on which the Configuration Manager is installed.
- You are recommended to ensure that the database instance for the broker is local to the machine on which the broker runs.
- You are recommended to use a local database for business data. Where such a database is remote from the broker machine, ensure that there is a fast, preferably dedicated, communications link between the broker machine and the database manager.
- Carefully examine default settings for nodes and messageflows, especially those related to recovery, to ensure that the values are those required. For example, the transaction mode parameter for an MQInput node will default to yes, meaning that the message flow will proceed under transaction control. This may not be what you required.
- When creating and deploying large messageflows increase the heap allocation of the Configuration Manager database. In DB2 this is the APP_CTL_HEAP_SIZE parameter. You should increase the value empirically.

4.4 Maximizing Throughput

In order to improve the message throughput for a message flow consider the following points:

- Achieve as much parallelism as possible. This can be achieved in MQSI V2 by running multiple execution groups or using additional instances as this will add an additional process or thread respectively which provides the potential to increase CPU utilization by using another processor. Where a message flow is I/O bound because of database I/O in a database node for example consider using an additional instance of the message flow. These approaches are only effective on a machine which has multiple processors. With a single processor machine it will not be possible to improve throughput in this way. However, it may still be necessary to configure multiple execution groups for other reasons though.
- In maximizing multiprocessor exploitation remember that there are processes associated with the MQSeries queue manager and any databases which may be used from within message flows. For example, the MQSeries queue manager listener process is capable of fully utilizing a processor dependent on the message rate. This was the case with the MQSI V2 performance measurements.
- When creating multiple execution groups be careful about the number of times messages are written back to an MQSeries queue as a means of communicating between execution groups since this is an expensive operation especially with persistent messages. It may be better to form one larger single execution group.
- Compute nodes are expensive in processing costs since they take a copy of the input message. Aim to minimize the number of compute nodes therefore. Dependent on the circumstances it may be possible to use a filter node instead of a compute node if message selection is required.
- When designing messages, make them as simple as possible. Large and more complex messages require more parsing. This consumes more CPU.

- Monitor the number of application handles in any databases containing business data (non MQSI V2 databases). In DB2 for example this is the MAX_APPLS parameter. Each database node will obtain an application handle. The node will retain this handle until the execution group terminates. If there are multiple message flows each with multiple database nodes then the DB2 default value of 40 can be quickly exceeded.
- Regularly monitor the performance of any database containing business data.
- Ensure that there are sufficient resources available to the database manager (CPU, memory, disk) so that it does not becoming a limiting factor in message throughput.

5.0 GLOSSARY

- **CPU Usage (%)**
% Processor busy time.
- **Get**
Get time in seconds as reported by the Requester application. (This is an MQGET with a wait for the reply message)
- **Message Size**
Size of the user portion of a message. Does not include the MQSeries header size.
- **Mbits/Sec**
Number of bits per second being transferred between the client and server.
- **Msgs/Sec**
Messages per second measured by the Server. Each request and each reply is a message.
- **Pages In per Second**
Number of pageins per second.
- **Pages On Free List**
Number of available pages of real memory to be used as cache.
- **Pages Out per Second**
Number of pageouts per second.
- **Persistence**
Indicates whether the message type was persistent (yes) or non persistent (no).
- **Put**
Put time in seconds as reported by the Requester application.
- **Response**
Response time in seconds as reported by the Requester application. (This is the round trip time from putting the request message on a queue to receiving the reply message. It should be the total of get and put).

- **MQRFH**

Rules and Formatting Header(MQRFH) used for Publish/Subscribe applications. Publishing and Subscribing applications send their messages to the broker in MQRFH format

- **MQRFH2**

MQSeries Rules and Formatting Header version 2 (MQRFH2) The principal use is to contain message format information and for Publish/Subscribe additional control information for both Publications and Subscriptions.

- **SSA**

Serial Storage Architecture (SSA). A serial storage system available from IBM.

6.0 APPENDIX A - MEASUREMENT HARDWARE AND SOFTWARE

All throughput measurements were taken on a single server machine driven by MQSeries clients running on a separate client machine connected by a 100Mb Ethernet link.

MQSeries Clients communicated with the MQ queue manager on the server machine using an MQI channel.

Server Machine

The server machine hardware consisted of

- An IBM 44P Model 270 with 4 * 375MHz POWER3-II Processors.
- Three 8.0 GB SCSI hard drives.
- 1GB RAM.
- 100Mb Ethernet Card.

The server machine software consisted of:

- IBM AIX 4.3.3.0
- MQSeries V5.1 and PTF 04.
- MQSeries Integrator for AIX V2.0
- DB2 for AIX V6.1.

Client Machine

The client machine hardware consisted of

- An IBM Netfinity 5500 with 2 * 450Mhz processors.
- Three x GB SCSI hard drives
- 1GB RAM.
- 100Mb Ethernet Card.

The client machine software consisted of:

- Microsoft Windows NT V4.0 with Service Pack 5.
- MQSeries V5.1.

Network Configuration

The client and server machines were connected on a full duplex 100Mb Ethernet LAN with a single hub.

Although the nominal data transfer rate is 200 Mb/second the practical limit proved to be around 130 Mb/second. This is only of relevance for the measurements with large message sizes.

7.0 APPENDIX B - MEASUREMENT DATA

This section contains the detailed results of the MQSI V2 performance measurements described in **Section 2.0 - Broker Throughput Measurements**. For an explanation of column headings see **Section 5.0 - Glossary**. The CPU utilization reported for the server machine, under the heading “rem CPU” in each table, is the system CPU utilization on the server machine. This includes **all** processes on the server machine. The figure therefore reflects the cost of the MQSeries queue manager processes, the MQ Listener, DB2 where appropriate etc. in addition to the CPU used by the MQSI V2 broker.

The default messages used in the performance measurements were not set to any particular type, i.e. MQRFH, MQRFH2 or XML. Where a particular type was used it is indicated.

7.1 MQInput/MQOutput Throughput Results

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	1511	0.063	0.061	0.002	49	14622	0	0	11.8
no	4096	1086	0.08	0.077	0.003	37	14532	0	0	33.9
no	16384	306	0.28	0.257	0.021	16	14171	0	0	38.3
no	65536	116	0.746	0.49	0.289	13	12282	0	0	58
yes	1024	52.2	1.7	1.642	0.058	4	14107	0	0	0.4

7.2 Compute Node Throughput Results

7.2.1 Simple Compute Node

This measurement used a message type of XML.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	438	0.202	0.201	0.001	30	14031	0	0	3.4
no	4096	308	0.289	0.287	0.001	30	13935	0	0	9.6
no	16384	125	0.698	0.689	0.004	28	13540	0	0	15.6
no	65536	43.7	2.049	2.033	0.015	28	11516	0	0	21.9
yes	1024	54.6	1.704	1.628	0.069	6	13544	0	0	0.4

7.2.2 Complex Compute Node

This measurement used a message type of XML.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	75.3	1.18	1.177	0.001	27	13227	0	0	2.4
no	16384	55.5	1.602	1.612	0.001	27	12895	0	0	6.9
no	65536	28.1	3.202	3.191	0.01	27	10825	0	0	14.1
yes	4096	34.6	2.528	2.492	0.036	15	13910	0	0	1.1

7.2.3 Multiple Complex Compute Nodes

This measurement used a message type of XML.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	23.9	3.797	3.797	0	26	12957	0	0	0.7
no	16384	21.4	4.093	4.093	0.001	26	12601	0	0	2.7
no	65536	15.7	5.72	5.719	0.001	26	10489	0	0	7.9
yes	4096	18.3	5.037	5.024	0.012	20	12408	0	0	0.6

7.2.4 Very Complex Compute Node

This measurement used a message type of XML.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	36.2	2.444	2.444	0	26	13129	0	0	1.1
no	16384	32.1	2.865	2.864	0	26	12796	0	0	4
no	65536	20.3	4.392	4.388	0.004	26	10759	0	0	10.2
yes	4096	26.1	3.46	3.437	0.021	19	13463	0	0	0.8

7.3 Database Node Throughput Results

Database Insert/Delete With XA Connection Configured using XML messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	94.4	0.974	0.973	0.001	8	110684	0	0	0.7
no	4096	93.6	0.92	0.918	0.001	9	111010	0	0	2.9
no	16384	67.7	1.281	1.276	0.004	9	110945	0	0	8.5
no	65536	50.5	1.819	1.801	0.019	16	109328	0	0	25.3
yes	1024	42	2.155	2.135	0.021	5	110663	0	0	0.3

7.4 Filter Node Throughput Results

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	1079	0.084	0.083	0.001	39	14431	0	0	8.4
no	4096	794	0.112	0.111	0.002	39	14339	0	0	24.8
no	16384	288	0.304	0.299	0.006	26	13971	0	0	36
no	65536	118	0.761	0.747	0.013	28	11953	0	0	59
yes	1024	54	1.637	1.588	0.053	4	13928	0	0	0.4

7.5 Publication Node Throughput Results

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	1024	645	0.136	0.135	0.001	34	37059	0	0	5
no	4096	550	0.139	0.137	0.002	33	37161	0	0	17.2
no	16384	289	0.291	0.283	0.009	21	37186	0	0	36.1
no	65536	117	0.76	0.664	0.105	16	35314	0	0	58.5
yes	1024	64.9	1.429	1.391	0.039	6	37689	0	0	0.5

7.6 Converting Messages Between Formats

Conversion of Generic XML to Generic XML

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	164	0.542	0.541	0.001	28	15104	0	0	5.1
no	16384	96.1	0.921	0.918	0.003	28	14761	0	0	12
no	65536	36.6	2.408	2.394	0.015	27	12667	0	0	18.3
yes	4096	42.4	2.117	2.027	0.083	9	15745	0	0	1.3

Conversion of Generic XML to CWF

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	109	0.815	0.814	0.002	27	14462	0	0	3.4
no	16384	72.3	1.181	1.178	0.002	27	14091	0	0	9
no	65536	31.2	2.847	2.834	0.013	27	11993	0	0	15.6
yes	4096	38.3	2.227	2.187	0.042	12	15072	0	0	1.2

Conversion of Generic XML to MRM XML.

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	107	0.84	0.839	0.002	27	14444	0	0	3.3
no	16384	70.9	1.312	1.306	0.003	27	14069	0	0	8.9
no	65536	31.2	2.837	2.825	0.012	27	11973	0	0	15.6
yes	4096	39.8	2.317	2.274	0.052	12	15051	0	0	1.2

Conversion of CWF to Generic XML

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	4096	78.7	1.129	1.128	0.001	26	14594	0	0	2.5
no	16384	77.7	1.132	1.131	0.001	27	14548	0	0	9.7
no	65536	77.7	1.195	1.194	0.001	26	14489	0	0	38.9
yes	4096	42.1	2.218	2.203	0.016	15	16047	0	0	1.3

Conversion of CWF to CWF

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	4096	47.8	1.82	1.819	0.001	26	14573	0	0	1.5
no	16384	48.5	1.825	1.824	0.001	26	14544	0	0	6.1
no	65536	47.6	1.818	1.818	0	26	14514	0	0	23.8
yes	4096	36.6	2.399	2.384	0.017	20	15377	0	0	1.1

Conversion of CWF to MRM XML

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	4096	47.8	1.842	1.841	0.001	26	14567	0	0	1.5
no	16384	47.7	1.867	1.866	0.001	26	14538	0	0	6
no	65536	47.8	1.837	1.835	0.001	26	14509	0	0	23.9
yes	4096	38.1	2.395	2.381	0.016	21	15378	0	0	1.2

Conversion of MRM XML to Generic XML

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	4096	94.8	0.962	0.96	0.002	26	14395	0	0	3
no	16384	64.3	1.366	1.364	0.002	28	14034	0	0	8
no	65536	29.7	3.041	3.028	0.013	27	11947	0	0	14.9
yes	4096	38.2	2.459	2.411	0.04	12	14791	0	0	1.2

Conversion of MRM XML to CWF

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	4096	66.4	1.305	1.3	0.001	26	14269	0	0	2.1
no	16384	51.2	1.744	1.744	0	26	13913	0	0	6.4
no	65536	25.8	3.389	3.379	0.01	27	11797	0	0	12.9
yes	4096	34.6	2.507	2.477	0.033	15	14480	0	0	1.1

Conversion of MRM XML to MRM XML

This measurement used MQRFH2 messages.

Persistence	Message Size	Msgs/Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/Sec
no	4096	66.8	1.298	1.292	0.001	26	14249	0	0	2.1
no	16384	52.1	1.664	1.664	0	26	13886	0	0	6.5
no	65536	25.9	3.33	3.321	0.01	27	11768	0	0	13
yes	4096	34.7	2.504	2.465	0.043	15	14499	0	0	1.1

7.7 NEONRules and NEONFormatter Within MQSI V2

NEONFormatter in MQSI V1 With 50 Fields Using MQRFH Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	951	0.107	0.105	0.003	54	1181	0	0	7.4
no	4096	663	0.193	0.185	0.007	41	1080	0	0	20.7
no	16384	238	0.398	0.373	0.026	19	716	0	0	29.8
no	65536	93.6	1.001	0.944	0.057	14	552	0	0	46.8
yes	1024	54.8	1.654	1.59	0.063	5	2482	0	0	0.4

NEONFormatter in MQSI V2 With 50 Fields using MQRFH Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	815	0.114	0.112	0.002	64	3450	0	0	6.4
no	4096	773	0.121	0.118	0.003	64	3318	0	0	24.2
no	16384	224	0.393	0.372	0.026	24	2926	0	0	28
no	65536	78.3	1.155	1.08	0.077	14	870	0	0	39.2
yes	1024	59.4	1.474	1.412	0.049	7	3087	0	0	0.5

NEONFormatter in MQSI V2 With 100 Rules defined and 1 Rule executed using MQRFH Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	905	0.096	0.094	0.002	74	26489	0	0	7.1
no	4096	798	0.101	0.098	0.003	66	26359	0	0	24.9
no	16384	231	0.367	0.341	0.024	24	25966	0	0	28.9
no	65536	72	1.219	1.13	0.074	11	23908	0	0	36
yes	1024	73	1.264	1.19	0.039	8	26241	0	0	0.6

7.8 The Effect of Increasing The Number Of Execution Groups

One Execution Group Running a Very Complex Compute Message Flow

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	36.2	2.444	2.444	0	26	13129	0	0	1.1
no	16384	32.1	2.865	2.864	0	26	12796	0	0	4
no	65536	20.3	4.392	4.388	0.004	26	10759	0	0	10.2
yes	4096	26.1	3.46	3.437	0.021	19	13463	0	0	0.8

Two Execution Groups Running a Very Complex Compute Message Flow

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	73.1	1.199	1.196	0.001	51	9062	0	0	2.3
no	16384	62.1	1.505	1.497	0.001	51	8661	0	0	7.8
no	65536	39.6	2.294	2.278	0.016	53	6404	0	0	19.8
yes	4096	37.4	2.513	2.464	0.048	28	9483	0	0	1.2

Four Execution Groups Running a Very Complex Compute Message Flow

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	4096	141	0.655	0.655	0.006	99	1535	0	0	4.4
no	16384	118	0.75	0.74	0.009	98	984	0	0	14.8
no	65536	66.6	1.307	1.261	0.044	88	586	0	0	33.3
yes	4096	36.2	2.516	2.398	0.094	27	3655	0	0	1.1

7.9 The Effect of Making a Message Flow Transactional

Non Transactional MQInput/MQOutput Message Flow (Transaction Mode set to automatic)

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	1511	0.063	0.061	0.002	49	14622	0	0	11.8
no	4096	1086	0.08	0.077	0.003	37	14532	0	0	33.9
no	16384	306	0.28	0.257	0.021	16	14171	0	0	38.3
no	65536	116	0.746	0.49	0.289	13	12282	0	0	58
yes	1024	52.2	1.7	1.642	0.058	4	14107	0	0	0.4

Transactional MQInput/MQOutput Message Flow (Transaction Mode set to yes)

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	1148	0.084	0.082	0.002	38	14563	0	0	9
no	4096	931	0.105	0.103	0.002	36	14469	0	0	29.1
no	16384	252	0.361	0.341	0.019	15	14110	0	0	31.5
no	65536	124	0.646	0.465	0.239	14	12213	0	0	62
yes	1024	52.3	1.697	1.646	0.052	4	13941	0	0	0.4

7.10 The Effect of using coordinatedTransaction=yes

coordinatedTransaction=no, Transaction Mode=automatic, no XA connection, Using XML Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	233	0.376	0.375	0.001	11	111119	0	0	1.8
no	4096	235	0.38	0.378	0.001	14	111026	0	0	7.3
no	16384	214	0.415	0.413	0.002	22	110647	0	0	26.8
no	65536	92.2	0.975	0.959	0.009	23	108714	0	0	46.1
yes	1024	58.5	1.532	1.504	0.026	5	110572	0	0	0.5

coordinatedTransaction=yes, XA connection, Using XML Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	93.8	0.95	0.949	0.002	7	110964	0	0	0.7
no	4096	94.3	0.962	0.961	0.002	8	111280	0	0	2.9
no	16384	68.8	1.263	1.256	0.003	9	111177	0	0	8.6
no	65536	51.6	1.692	1.673	0.018	15	109519	0	0	25.8
yes	1024	42.8	2.184	2.158	0.02	5	110897	0	0	0.3

7.11 The Effect of Increasing the Number of Subscribers

1 Subscriber Receiving MQRFH2 Type Published Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	645	0.136	0.135	0.001	34	37059	0	0	5
no	4096	550	0.139	0.137	0.002	33	37161	0	0	17.2
no	16384	289	0.291	0.283	0.009	21	37186	0	0	36.1
no	65536	117	0.76	0.664	0.105	16	35314	0	0	58.5
yes	1024	64.9	1.429	1.391	0.039	6	37689	0	0	0.5

10 Subscribers Receiving MQRFH2 Type Published Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	199	0.529	0.528	0.001	29	13155	0	0	1.6
no	4096	166	0.5	0.498	0.001	28	13039	0	0	5.2
no	16384	128	0.732	0.728	0.003	25	12671	0	0	16
no	65536	57.8	1.446	1.344	0.098	18	10752	0	0	28.9
yes	1024	16.7	5.329	5.284	0.047	5	13801	0	0	0.1

30 Subscribers Receiving MQRFH2 Type Published Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	86.4	0.998	0.997	0.001	30	12651	0	0	0.7
no	4096	77	1.122	1.115	0.001	29	12533	0	0	2.4
no	16384	62.5	1.4	1.394	0.002	29	12170	0	0	7.8
no	65536	29.6	2.882	2.864	0.017	20	10264	0	0	14.8
yes	1024	4.1	22.433	22.4	0.036	3	12167	0	0	0

50 Subscribers Receiving MQRFH2 Type Published Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	55.6	1.6	1.6	0	30	12345	0	0	0.4
no	4096	48.6	1.775	1.808	0	30	12228	0	0	1.5
no	16384	40.2	2.217	2.219	0	30	11871	0	0	5
no	65536	19.7	5.267	5.257	0.01	20	9974	0	0	9.9
yes	1024	2.5	34.04	34	0.039	3	11919	0	0	0

70 Subscribers Receiving MQRFH2 Type Published Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	40.9	2.146	2.146	0	28	12310	0	0	0.3
no	4096	37	2.306	2.308	0	28	12194	0	0	1.2
no	16384	31.3	2.91	2.909	0	29	11838	0	0	3.9
no	65536	9.5	9.467	9.459	0.008	13	9945	0	0	4.8
yes	1024	1.9	41.058	41.01	0.038	3	11963	0	0	0

100 Subscribers Receiving MQRFH2 Type Published Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	27.6	3.209	3.209	0	28	12283	0	0	0.2
no	4096	25.8	3.592	3.589	0	30	12163	0	0	0.8
no	16384	20.8	4.189	4.187	0	29	11810	0	0	2.6
no	65536	7.1	11.79	11.79	0.001	17	9932	0	0	3.6
yes	1024	1.4	41.845	41.81	0.03	4	11833	0	0	0

1000 Subscribers Receiving MQRFH2 Type Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	2.6	34.264	34.26	0	28	8210	0	0	0
no	4096	2.6	34.187	34.19	0	29	8092	0	0	0.1
no	16384	2.2	38.602	38.6	0	29	7730	0	0	0.3
no	65536	1.6	47.645	47.65	0	25	5932	0	0	0.8
yes	1024	1.2	77.284	77.28	0.014	20	5879	0	0	0

7.12 Content vs Topic Based PubSub

Topic Based Routing, Single Subscriber Receiving MQRFH2 Type Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	645	0.136	0.135	0.001	34	37059	0	0	5
no	4096	550	0.139	0.137	0.002	33	37161	0	0	17.2
no	16384	289	0.291	0.283	0.009	21	37186	0	0	36.1
no	65536	117	0.76	0.664	0.105	16	35314	0	0	58.5
yes	1024	64.9	1.429	1.391	0.039	6	37689	0	0	0.5

Content Based Routing, Single Subscriber Receiving MQRFH2 Type Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	421	0.203	0.203	0.001	31	36542	0	0	3.3
no	4096	354	0.261	0.26	0.002	30	36444	0	0	11.1
no	16384	223	0.437	0.434	0.004	29	36102	0	0	27.9
no	65536	106	0.822	0.814	0.008	31	34185	0	0	53
yes	1024	62.6	1.378	1.332	0.034	7	36184	0	0	0.5

7.13 The Effect of an Increasing Number of Message Fields

One Field in an Input Message to a Compute Node Using XML Type Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	438	0.202	0.201	0.001	30	14031	0	0	3.4
no	4096	308	0.289	0.287	0.001	30	13935	0	0	9.6
no	16384	125	0.698	0.689	0.004	28	13540	0	0	15.6
no	65536	43.7	2.049	2.033	0.015	28	11516	0	0	21.9
yes	1024	54.6	1.704	1.628	0.069	6	13544	0	0	0.4

50 Fields in an Input Message to a Compute Node Using XML Type Messages

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	233	0.384	0.383	0.001	28	13972	0	0	1.8
no	4096	189	0.473	0.471	0.001	28	13875	0	0	5.9
no	16384	103	0.915	0.912	0.003	28	13462	0	0	12.9
no	65536	38.9	2.282	2.266	0.016	27	11438	0	0	19.5
yes	1024	51.9	1.695	1.645	0.056	9	13423	0	0	0.4

7.14 The Effect of userTrace options

No Trace Active

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	1511	0.063	0.061	0.002	49	14622	0	0	11.8
no	4096	1086	0.08	0.077	0.003	37	14532	0	0	33.9
no	16384	306	0.28	0.257	0.021	16	14171	0	0	38.3
no	65536	116	0.746	0.49	0.289	13	12282	0	0	58
yes	1024	52.2	1.7	1.642	0.058	4	14107	0	0	0.4

UserTrace=normal, mode=safe

Persistence	Message Size	Msgs/ Sec	Response	Get	Put	CPU Usage (%)	Pages On Free List	Page In per Second	Page Out per Second	Mbits/ Sec
no	1024	40.6	2.2	2.2	0	25	13711	0	0	0.3
no	4096	40.1	2.193	2.199	0	25	13660	0	0	1.3
no	16384	39.1	2.338	2.338	0	25	13310	0	0	4.9
no	65536	38.9	2.362	2.345	0.016	26	11432	0	0	19.5
yes	1024	27.8	3.285	3.263	0.022	18	12430	0	0	0.2

8.0 APPENDIX C - COMPLEX COMPUTE NODE

This section contains details of the complex, multiple complex and very complex compute nodes.

8.1 *Complex Compute Node*

The ESQL statements used in the complex compute node are given below. The variable *i* has a maximum value of 20.

```
SET OutputRoot = InputRoot;
```

```
DECLARE i INTEGER;
```

```
SET i = 1;
```

```
WHILE i &lt;= CARDINALITY(OutputRoot.XML.CSIM.TestCase.Stack.ProcessingPath.Element[]) DO
```

```
    SET OutputRoot.XML.CSIM.TestCase.ProcessingPath.Component[i].Name =
```

```
        OutputRoot.XML.CSIM.TestCase.Stack.ProcessingPath.Element[i].COMPONENT;
```

```
    SET OutputRoot.XML.CSIM.TestCase.ProcessingPath.Component[i].Transport.(XML.attr)Type='A';
```

```
    SET OutputRoot.XML.CSIM.TestCase.ProcessingPath.Component[i].Transport.Queue =  
OutputRoot.XML.CSIM.TestCase.Stack.ProcessingPath.Element[i].QUEUE;
```

```
    SET i = i + 1;
```

```
END WHILE;
```

8.2 *Multiple Complex Compute Node*

The multiple complex compute nodes consisted of five identical complex compute nodes which were daisy chained. The logic within each of the complex compute nodes was the same as that for the complex compute node given in **Section 8.1 - Complex Compute Node**.

8.3 *Very Complex Compute Node*

The very complex compute node consisted of five repetitions of the logic for complex compute node (see **Section 8.1 - Complex Compute Node**) all contained within one compute node.

End of Document