WebSphere Message Broker V6.0

For z/OS

Performance report

Version 1.2

January, 2006

Tim Dunn

Kevin Braithwaite

Rich Bicheno

WebSphere Message Broker Development
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".

Third Edition, January 2006.
This edition applies to *WebSphere Message Broker V6 for z/OS* and to all subsequent releases and modifications until otherwise indicated in new editions.
© Copyright International Business Machines Corporation 2005. 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 for Architects, Systems Programmers, Analysts and Programmers wanting to understand the performance characteristics of **WebSphere Message Broker V6 for z/OS**. The information is not intended as the specification of any programming interfaces that are provided by WebSphere MQ or WebSphere Message Broker V6 for z/OS. It is assumed that the reader is familiar with the concepts and operation of WebSphere Message Broker V6.

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
- WebSphere MQ
- WebSphere Message Broker
- DB2

The following terms are trademarks of other companies:

- Windows 2000, Windows XP, Microsoft Corporation

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

Summary of Amendments

Date	Changes
1st December 2005	Initial Release
8th December 2005	Add comment about Java XPath performance and make some editing changes
11th January 2006	Correct the results for the Large Messaging usage scenario

Table of Contents

Table of Contents 5
Introduction 7

Part I 9
Release Highlights 10
Improvements over WebSphere Business Integration Message Broker V5 10
Improvements over WebSphere MQ Integrator V2.1 12
Use Case Outline 13
Additional Information 14

Part II 16
Routing and Transformation Processing Profiles 17
Minimal Processing 19
Message Parsing and Writing 19
Parsing a Message in the MRM Domain 19
Writing a Message in the MRM Domain 25
Parsing a Message in the XML Domain 28
Writing a Message in the XML Domain 29
External Resources 30
Accessing a Database from a Message Flow 30
Calling External Procedures 32
Routing and Transformation Logic 34
Using ESQL 34
Using Java 42
Using XMLT 47
Publish Subscribe 48
Scaling Message Throughput 49
Overheads 51
Using Accounting and Statistics 51
Using Trace 51
Resource Requirements 52
Recommended Minimum Specification 52
Memory Use 52
zAAP Utilization 54
Tuning 56
Message Broker 56
WebSphere MQ 57
TCP/IP 57
Database 58
Miscellaneous 58
Additional Tuning Information 58
Conclusion 59

Appendix A - Measurement Environment 60
Server Machine 60
Client Machines 60
Network Configuration 61
Appendix B - Evaluation Method 62
Point to Point testing 62
Message Generation and Consumption 62
Machine Configuration 63
Publish Subscribe testing 64
Message Generation and Consumption 64
Machine Configuration 65
Reported Message Rates 66
Appendix C - Test Messages 67
Input Message 67
Output Message 68
Appendix D - Use Case Descriptions 71

Aggregation 71
Coordinated Request/Reply 72
Data Warehouse 74
Large Messaging 75
Message Routing 77
SWIFT Message Parse 78
XML Transformation 79
Feedback 81

Introduction

The purpose of this report is to illustrate the key processing characteristics of WebSphere Message Broker. This has been done by measuring the message throughput which is possible for a number of different types of message processing, covering multiple message formats, types and sizes.

This report consists of two parts. These meet different requirements:

1. **Part I** contains the release highlights and some background information to help understand the context of the results. It shows:

    a. The improvement in performance with WebSphere Message Broker V6 when compared with WebSphere Business Integration Message Broker V5.

    b. The level of message throughput that is achievable when using WebSphere Message Broker in different ways. These tests use **multiple copies of the message flow and utilise as much of the server machine as possible** to illustrate the maximum message rate which can be sustained for the individual types of processing.

    The information in this part is presented at a high level and is intended to help you quickly understand WebSphere Message Broker throughput capabilities.

2. **Part II** contains measurement data for a wide variety of tests which examine the processing costs of individual functions **using a single copy of the message flow**. This information is provided for those who wish to understand the processing costs of different components within WebSphere Message Broker such as the differences in CPU cost between Fixed Length Tagged Delimited Strings and All Elements Delimited Tagged Delimited Strings. This information is intended for the more experienced WebSphere Message Broker user who is familiar with the product concepts and functions. **As these tests run a single copy of the message flow. They do not utilise the whole of the server machine and do not therefore represent the maximum message throughput which is achievable.**

There are a number of changes from previous performance reports. The most significant are:

1. Re-engineered tests to better reflect the processing costs which are encountered when processing messages with a WebSphere Message Broker message flow. The previous tests are deprecated and do not appear in this report.

2. Measurement of a selection of product samples which are available with WebSphere Message Broker V6. This is done for two reasons: Firstly it makes it easier for you to understand the volume of messages which can be processed for a variety of common use cases. Previous reports focused on the use of individual nodes which made it difficult to visualise particular applications. Secondly by using samples it is easy for you to take exactly the same message flows and message sets and run them in your own environment. You can then compare the results obtained in your environment against those published in this report. This can be very useful in validating that a broker environment is well configured.

3. More extensive analysis of product function, including incremental test cases.

4. Larger range of message sizes including a greater range of persistent message sizes.

5. A change in layout to separate the overview of message processing capability from the detailed data which shows the costs of using individual functions.

The performance measurements focus on the throughput capabilities of the broker using different message formats and processing node types. The aim of the measurements is to help you understand how many messages a second can be processed in different situations as well as helping you to understand the relative costs of the different node types and approaches to message processing.

You should not attempt to make any direct comparisons of the test results in this report with what may appear to be similar tests in previous performance reports. This is because the contents of the test messages are significantly different as is the processing in the tests. It is not meaningful to make such comparisons.

Some optimisations to the test environment and procedures have been implemented to minimise the effect of logging for example and to ensure that messages do not build up on output queues (which has a detrimental effect on message throughput). These are detailed in the section Summary of Tuning Information.

In many of the tests the business logic used is minimal so the results presented represent the best throughput that can be achieved for that node type. This should be borne in mind when performing sizing for WebSphere Message Broker.

Part I

This part contains an overview of the improvements in performance which were obtained with WebSphere Message Broker V6 when compared with WebSphere Business Integration Message Broker V5.

It contains the following sections:

- *Release Highlights* which outlines the main differences in performance when using WebSphere Message Broker V6 compared with WebSphere Business Integration Message Broker V5.

- *Additional Information* which provides links to other sources of information about WebSphere Message Broker and related products.

Release Highlights

Improvements over WebSphere Business Integration Message Broker V5

Improving Message Broker runtime performance has been a specific focus with WebSphere Message Broker V6 and as a result there are many improvements in the level of performance when compared with WebSphere Business Integration Message Broker V5. The improvements come from two sources - updates to existing function and provision of new function.

All key areas of Message Broker runtime function have been investigated and improvements made to improve performance. The main areas of focus were:

- The parsing and streaming of messages
- The cost of processing ESQL
- Message aggregation
- Message Broker infrastructure
- The calling of Java and database procedures

The improvements in these areas can be obtained by upgrading to WebSphere Message Broker V6. **No code or message model changes are required to benefit from the improvements**.

Further improvements are available if you take advantage of new functions such as

- The support for shared variables in ESQL which provides the capability to build an in-memory cache. This allows an in memory table to be built and accessed within message flows for example. The function can remove the need to access a database for read only routing or data validation. Previously a message flow had to issue a read against a database for each message flow invocation. The Message Routing sample shipped with the product provides an illustration of such processing.
- The MQGET node which makes it possible to use WebSphere MQ queues as an intermediate data store for communication between request and reply message flows for example. The Coordinated Request Reply sample provides an illustration of how such processing can be implemented. Previously a database had to be used to store the intermediate data.
- The extended and improved DATETIME functions which make it possible to perform complex date and time formatting operations using WebSphere Message Broker provided function. Previously a user had to write functions in ESQL or Java to perform such processing.

In addition the ability to more easily code Java processing for a message flow using the Java Compute node makes it much easier to take advantage of the IBM zSeries Application Assist Processor (zAAP). Whilst it was possible to do this previously through the Java plug-in node and Publication node the addition of the Java Compute node has made it much easier to code mainstream message flow processing using Java. Utilization of the zAAP is covered in greater depth in the section zAAP Utilization.

The Table below shows the results of running a series of use cases in WebSphere Business Integration Message Broker V5 and WebSphere Message Broker V6. The use cases are briefly described at the end of this section and more fully in Appendix D – Use Case Descriptions. The use cases are largely taken from the samples gallery of WebSphere Message Broker V6.

Use Case	**Message Size**	**V6 Msgs/sec**	**Improvement Ratio (V6/V5)**	**Note**
Aggregation	8K	180	**3.27**	**1**
Coordinated Request/Reply	1K	1150	**2.12**	**2**
Data Warehouse	1K	1500	**1.07**	**3**
Large Messaging	16K	530	**1.12**	**4**
Message Routing	1K	3540	**2.27**	**5**
SWIFT Message Parse	7K	200	**5.00**	**6**

Throughput Comparison for Use Cases.

Notes:

1. As Aggregation in WebSphere Message Broker V6 is now based on the use of WebSphere MQ queues and not a database the I/O bottleneck which was caused by database logging has been removed. This combined with a reduced CPU cost per message has made it possible to increase message throughput in V6 when compared with V5.

2. Use of a WebSphere MQ queue for intermediate data storage in the WebSphere Message Broker V6 edition of the message flow removed an I/O bottleneck (due to database logging requirements). This combined with a reduced cost CPU cost per message allowed a higher CPU utilization and message rate to be obtained.

3. There was no change in message throughput in this use case.

4. There was no change in message throughput in this use case.

5. By using a routing table which was held in shared variables the CPU cost per message was reduced. In WebSphere Business Integration Message Broker V5 a database read was required for every invocation of the message flow. By using shared variables this could be removed. The reduced CPU cost per message allowed a higher message rate to be achieved.

6. The rewrite of the MRM TDS parser has significantly reduced the CPU cost per message of parsing a TDS message. As a result it is possible to achieve a significantly higher message rate in WebSphere Message Broker V6 when compared with WebSphere Business Integration Message Broker V5.

Each of the use cases was implemented in WebSphere Business Integration Message Broker V5 and WebSphere Message Broker V6 using the same hardware and prerequisite software.

The Aggregation, Data Warehouse, Large Messaging and SWIFT Message Parse use cases contained no new code in the message flows. Exactly the same message flow was run on V5 and V6.

The results in the table above were obtained by running sufficient copies of each message flow to utilise as much of the available CPU as possible.

These results show that there are significant increases in the level of message throughput that is achievable with WebSphere Message Broker V6 when compared with WebSphere

Business Integration Message Broker V5. Those use cases showing the best gains where the SWIFT Message Parse, Aggregation, Message Routing and Coordinated Request/Reply, all of which doubled message throughput or better.

Most of the gains in message throughput for the use cases came from improvements to existing broker function. The CPU cost of many aspects of the broker has been significantly reduced and as such message throughput can increase for a given amount of CPU power.

Improvements over WebSphere MQ Integrator V2.1

In this report WebSphere Message Broker V6 performance has been compared with that of WebSphere Business Integration Message Broker V5.

No direct comparison with WebSphere MQ Integrator V2.1 was undertaken. The reader is reminded that there was a 20% improvement in message throughput in WebSphere Message Broker V5 when compared with WebSphere MQ Integrator V2.1.

When estimating the benefits of migrating from WebSphere MQ Integrator V2.1 to WebSphere Message Broker V6 this improvement in performance should be included.

Use Case Outline

This section contains a brief outline of the tests used to obtain the results presented in the table above. For more detail on individual test cases see the section Appendix D - Use Case Descriptions.

- **Aggregation**
  This represents the type of processing that is required when travel is booked and arrangements for a flight, hotel, car and money must be made. Requests to four different applications are made and the replies consolidated into a single reply. This test performs the processing required to split an incoming XML message and perform a four message aggregation using the Aggregation nodes which are supplied with WebSphere Message Broker.

- **Coordinated Request Reply**
  This performs the processing needed to enable two applications with different message formats to communicate with each other. One application has a message format of self-defining XML and the other uses Custom Wire Format (CWF) messages. The request and reply processing for a particular request must be coordinated so that data from the original request is restored to the reply message.

- **Data Warehouse**
  This demonstrates a scenario in which a message flow is used to perform the archiving of data, such as sales data, into a database. The data is stored for later analysis by another message flow or application.

- **Large Messaging**
  This is based on the scenario of end-of-day processing of sales data. Messages representing sales for the day are batched together for transmission to the IT center. On receipt at the IT center the batched messages are split back out into their constituent parts for subsequent processing.

- **Message Routing**
  This shows how a message flow can be used to route messages to different WebSphere MQ queues based on data stored in a database table. This is a commonly used scenario which is applicable to many different industries and applications.

- **SWIFT Message Parse**
  This demonstrates the use of WebSphere Message Broker to read and parse a SWIFT MT543 message for subsequent processing

Additional Information

This section contains links to information about WebSphere Message Broker and associated products.

The Web Resources section in the development toolkit of WebSphere Message Broker V6 contains links to many additional pieces of information on topics such as Education, Technical Resources and SupportPacs. The Web resources section can be accessed by selecting `Web Resources` from the Help drop down on the development toolkit menu bar.

For additional suggestions consider the following:

- See the announcement letters for
  - IBM WebSphere Message Broker V6 which is available at http://www.ibm.com/software/integration/wbimessagebroker/v6
  - IBM WebSphere Message Broker V6 for z/OS which is available at http://www.ibm.com/software/integration/wbimessagebroker/v6/zos.html
- For more information about the zSeries Application Assist Processor (zAAP) see http://www.redbooks.ibm.com/abstracts/sg246386.html
- IBM WebSphere MQ SupportPacs provide you with a wide range of downloadable code and documentation that complements the WebSphere MQ family of products. Additional performance reports are also available. These are available at http://www.ibm.com/software/integration/support/supportpacs.
- For more information about WebSphere Message Broker V6, go to the WebSphere Message Broker Web site. Product documentation is also available. This is available at http:/www.ibm.com/software/integration/wbimessagebroker.
- For more information about WebSphere MQ V6, go to the WebSphere MQ Web site. Product documentation is also available. This is available at http://www.ibm.com/software/integration/wmq.
- For more information about business integration software from IBM go to WebSphere Business Integration Web site. This is available at http://www.ibm.com/software/info1/websphere/index.jsp?tab=products/businessint.
- Get the latest WebSphere Message Broker technical resources at the WebSphere Business Integration zone. This is available at http://www.ibm.com/developerworks/websphere/zones/businessintegration.
- In order to obtain the maximum message rate for your implementation it is important that you understand the current best practices for WebSphere Message Broker. These practices cover the architecture of message flow processing, the coding of message flows as well as the configuration and tuning of the message broker and associated components. Such information can be found in the Business Integration Zone of WebSphere Developer Domain. A suggested starting place is the article at http://www.ibm.com/developerworks/websphere/library/techarticles/0403_dunn/0403_dunn.html which highlights the information available and where it may be found.

- Once you have WebSphere Message Broker V6 installed it is worthwhile ensuring that the broker, broker queue manager and z/Os image on which the components run are not constrained. To help you understand the capacity of the system on which the broker is running you are recommended to download and run SupportPac IP13: WebSphere Business Integration Broker – Sniff Test and Performance on z/OS. This can be found at this link: http://www.ibm.com/support/docview.wss?rs=203&uid=swg24006892&loc=en_US&cs=utf-8&lang=en.

- For information about capacity planning and tuning for WebSphere MQ for z/OS use SupportPac MP16 which is available at http://www.ibm.com/support/docview.wss?rs=203&uid=swg24007421&loc=en_US&cs=utf-8&lang=en.

- For information about the performance of WebSphere MQ for z/OS V6 use SupportPac MP1E which is available at http://www-1.ibm.com/support/docview.wss?rs=203&uid=swg24009932&loc=en_US&cs=utf-8&lang=en

Part II

This part contains the description and results of a series of tests which have been run in order to identify the processing costs of the different functions which are provided with WebSphere Message Broker.

It contains the following sections:

- *Routing and Transformation Processing Profiles* which describes the tests and shows the results obtained when a **single** copy of the message flow was run.

- *Resource Requirements* which provides a recommended minimum specification machine on which to install the product as well as some guidance on virtual memory use for execution groups running a variety of message flows.

- *zAAP Utilization* which provides an indication of the level of processing off-load that is available when using a zAAP.

- *Tuning* which describes the changes made to the default settings for WebSphere Message Broker V6 and WebSphere MQ in order to obtain the results detailed in this report.

- *Conclusion* which summarises the report.

Routing and Transformation Processing Profiles

This section contains the results of a series of micro tests which illustrate the costs of performing different types of processing using WebSphere Message Broker such as message parsing, message streaming, use of Filter nodes etc. These tests are not intended to represent applications. They are an illustration of the processing costs of specific functions.

The test results were all run using the same methodology. This was to run a single copy of the message flow (unless specified otherwise) to maximum CPU utilization and to observe the message rate obtained. From this a CPU cost per message was calculated. This is presented in the results table for each measurement.

When comparing the costs of different functions it is recommended to compare them on the basis of CPU cost per message rather than message rate.

There are many comparisons which can be made using the data in this section which will give some insight into the relative costs of different implementations such as what is the relative cost of ESQL and XSLT to process the same message.

The data in this section will allow you to make a comparison on the basis of CPU costs. Other factors such as the potential for code re-use and the operational considerations of using a particular technology are not discussed.

Messages Used in Processing
For the majority of tests the message content was common. Different formats (in XML, CWF, TDS) of a common input message were used. The output message varied dependent on the test case. The messages are described in the section Appendix C – Test Messages.

For the Publish Subscribe tests a 1K JMS Bytes message was used. This was a sequence of random data. In these tests the message content was not of interest.

Results Presentation
Each of the tests are described below and accompanied by a table of data which has a format such as this:

Msg Size	Persistent	Message Rate (Msgs/sec)	% CPU Busy	CPU ms/msg
1K	No			
4K	No			
16K	No			
64K	No			
256K	No			
1K	Yes			
4K	Yes			
16K	Yes			
64K	Yes			
256K	Yes			

The data in the columns is as follows:

Msg Size records the approximate size of the message used as input to the test. This is the size of the message payload and does not include the size of any message header. For the Message Repository Manager (MRM) tests which use CWF and TDS message formats the input message will be smaller. This is due to the differences in the way in which data is formatted. In these cases the input message will still contain the same amount of information but it will be the CWF or TDS representation of the generic XML representation of the same

data. Most test cases used messages of 1K, 4K, 16K, 64K and 256K. In some cases a more limited range of message sizes was run where the test was not suitable for the whole range of message sizes.

Persistent: Indicates whether the messages used in the test were persistent or not

Message Rate: The number of round trips or message flow invocations per second

% CPU Busy: System busy CPU percentage on the server machine. This includes the CPU used by all processes (WebSphere Message Broker, WebSphere MQ queue manager, database manager etc) on the system under test. The rate is expressed as a percentage utilization of all processors on the machine.

CPU ms/msg: Overall CPU cost per message, expressed as CPU milliseconds per message. The value is obtained using the calculation:

Total CPU used in the measurement slot / Number of messages processed in the measurement slot

This cost includes WebSphere Message Broker, WebSphere MQ, DB2, operating system costs etc. The CPU ms/msg figures reported are specific to the machine on which they were obtained and if projections of message processing capacity are to be made for other machines a suitable adjustment must be made in the costs to allow for differences in the capacity of the two systems.

Response Times
Response time data for the message flow execution is not reported. The tests are configured to maximise message throughput and minimise CPU costs. As such tests always have a number of messages waiting on the input node of the message flow so that there is a message ready to be processed immediately after processing of the current message has completed. This means that the processing of each message involves queuing time at the input node. Because of this it is not meaningful to report message processing times as observed by the client as it will not reflect the true execution time in the message flow.

It is possible to estimate the elapsed time within a message flow in milliseconds from the results of these tests by dividing 1000 (representing the number of milliseconds in 1 second) by the message rate for the test.

For example let us suppose that a test achieved a message rate of 2000 per second. The message flow average execution time is 1000 / 2000 = 0.5ms. For a message rate of 200 per second the average execution time is 1000/200 = 5ms.

These times are an estimate of the execution time in the message flow and as such represent the elapsed time between the message being read from the input queue and the result being placed on the output queue.

If messages are generated or consumed by remote clients an allowance needs to be made for network delays.

The test descriptions and results follow.

Minimal Processing

The test in this section illustrates some of the simplest processing which can be performed with WebSphere Message Broker. As such it illustrates the smallest processing cost that you could expect for a message flow. This is not typical of the majority of implementations of Message Broker though. The data is provided for reference purposes only to help you understand the maximum rate that could be expected for one copy of the message flow.

Typically the processing within a message flow involves message parsing, processing logic and message serialisation. Under these circumstances the CPU processing costs can increase significantly and as such the message rate obtained for given amount of CPU will be lower than for the very simple type of flow presented in this section.

Setting of the MQ Message Headers

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers for the outgoing message are created using ESQL. To minimise processing costs only the CodedCharSetId and Encoding fields in the MQMD header are set. The message body is ignored and therefore not used in the output message.

This test identifies the cost of setting the message header only and creating an output message with no payload.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1885	0.99
1k	Yes	465	1.79

Message Parsing and Writing

The tests in this section illustrate the cost of parsing input messages and writing output messages for different message formats.

Parsing a Message in the MRM Domain

The tests in this section illustrate the CPU processing costs of parsing different message formats in the MRM domain.

Parsing a Tagged Delimited String, All Elements Delimited Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing an All Elements Delimited, Tagged Delimited String input message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	757.60	1.88
4k	No	360.00	3.34
16k	No	115.60	9.34
64k	No	31.12	33.25
256k	No	7.90	129.30
1k	Yes	332.20	2.72
4k	Yes	228.80	4.31
16k	Yes	98.00	10.43
64k	Yes	29.88	34.41
256k	Yes	7.84	130.58

Parsing a Tagged Delimited String, Fixed Length Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a Fixed Length, Tagged Delimited String input message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	773.40	1.86
4k	No	394.00	8.34
16k	No	131.20	17.14
64k	No	35.80	54.97
256k	No	9.16	196.16
1k	Yes	344.20	11.53
4k	Yes	234.80	12.97
16k	Yes	106.60	25.50
64k	Yes	34.14	61.29
256k	Yes	9.00	230.15

Parsing a Tagged Delimited String, Tagged Delimited Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a Tagged Delimited String, Tagged Delimited input message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	611.60	2.19
4k	No	251.80	4.58
16k	No	75.00	14.23
64k	No	19.70	52.77
256k	No	5.00	205.25
1k	Yes	283.00	3.02
4k	Yes	179.80	5.47
16k	Yes	68.60	15.21
64k	Yes	19.04	53.90
256k	Yes	4.95	208.46

Parsing a Tagged Delimited String, Tagged Fixed Length Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a Tagged Fixed Length, Tagged Delimited String input message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	608.20	2.21
4k	No	254.20	4.59
16k	No	76.80	13.96
64k	No	20.14	51.69
256k	No	5.08	203.79
1k	Yes	277.40	3.10
4k	Yes	177.40	5.47
16k	Yes	69.20	15.18
64k	Yes	19.58	53.19
256k	Yes	5.04	206.36

Parsing an MRM XML Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing an MRM XML input message.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	667.20	2.05
4k	No	329.80	3.67
16k	No	110.00	10.06
64k	No	29.56	36.57
256k	No	7.50	142.26
1k	Yes	302.80	2.90
4k	Yes	198.00	4.66
16k	Yes	89.20	11.31
64k	Yes	27.98	37.81
256k	Yes	7.30	144.73

Parsing a Custom Wire Format Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a Custom Wire Format input message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	988.00	1.55
4k	No	405.80	3.02
16K	No	96.00	12.98
64k	No	24.40	42.44
256k	No	6.08	167.77
1k	Yes	364.20	2.40
4k	Yes	240.80	3.96
16k	Yes	84.40	12.10
64k	Yes	23.66	43.25
256k	Yes	6.00	169.88

Parsing a Comma Separated Value Input Message using Data Patterns

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the entire incoming message is copied to the outgoing message. In addition the format of the outgoing message is set to XML. This causes a full parse of the incoming message using the Tagged Delimited String Parser and a full write of the outgoing message using the Generic XML Writer.

This test identifies the cost of converting an incoming Comma Separated Value input message using the Data Pattern function with the Tagged Delimited String Parser, to an outgoing Generic XML Message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	296.20	4.02
1k	Yes	202.60	4.72

Parsing a SWIFT 543 Input Message using the Tagged Delimited String Parser

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a SWIFT MT543 message using the Tagged Delimited String format. A single implementation of this message was used which was approximately 7K in size.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
7k	No	51.40	20.33
7k	Yes	48.28	21.05

Parsing and Writing a SWIFT 543 Input Message using the Tagged Delimited String Parser

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the Envelope within the incoming SWIFT Message is copied over to the outgoing message. This causes a full parse of the incoming message and a full serialisation of the outgoing message.

This test identifies the cost of parsing a SWIFT MT543 message and serializing it again using the Tagged Delimited String format. A single implementation of this message was used which was approximately 7K in size.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
7k	No	26.84	38.32
7k	Yes	24.86	39.15

Parsing a JMS SOAP message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

The incoming JMS SOAP message is parsed. The output message consists of a message header only and no payload.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	637.20	2.25
1k	Yes	303.80	2.90

Parsing a JMS SOAP Message with Attachments

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

The incoming MIME JMS message is parsed and the soap envelope extracted and parsed using the MIME Parser within the MRM Domain. The output message consists of a message header only and no payload.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	113.40	10.50
1k	Yes	58.80	12.64

Parsing and Writing a JMS SOAP Message and Modifying a Field

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

The incoming JMS SOAP message is parsed. One field is modified and the resulting message is written to the output queue.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	207.60	5.57
1k	Yes	145.40	6.28

Writing a Message in the MRM Domain

The tests in this section illustrate the CPU processing costs of creating an output message with different formats in the MRM domain.   This is the processing associated with taking a message tree in OutputRoot and flattening it to create a bitstream which is the output message.

Writing a Tagged Delimited String, All Elements Delimited Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition the incoming Generic XML message is converted to an All Elements Delimited, Tagged Delimited String outgoing message. This causes a full parse of the incoming message payload which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML message and writing out an All Elements Delimited, Tagged Delimited String output message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	525.60	2.46
4k	No	228.40	5.06
16k	No	69.80	15.43
64k	No	18.44	56.68
256k	No	4.65	223.70
1k	Yes	258.80	3.38
4k	Yes	143.60	6.12
16k	Yes	58.00	16.53
64k	Yes	17.14	58.44
256k	Yes	4.50	227.66

Writing a Tagged Delimited String, Fixed Length Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition the incoming Generic XML message is converted to a Fixed Length, Tagged Delimited String outgoing message. This causes a full parse of the incoming message payload which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML message and writing out a Fixed Length, Tagged Delimited String output message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	531.60	2.44
4k	No	228.60	5.05
16k	No	69.60	15.41
64k	No	18.50	56.74
256k	No	4.66	223.99
1k	Yes	237.60	3.43
4k	Yes	132.40	6.00
16k	Yes	53.80	16.84
64k	Yes	16.78	58.26
256k	Yes	4.42	227.74

Writing a Tagged Delimited String, Tagged Fixed Length Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition the incoming Generic XML message is converted to a Tagged Fixed Length, Tagged Delimited String outgoing message. This causes a full parse of the incoming message payload which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML message and writing out a Tagged Fixed Length, Tagged Delimited String output message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	526.40	2.46
4k	No	226.40	5.08
16k	No	69.20	15.60
64k	No	18.22	57.58
256k	No	4.61	226.98
1k	Yes	251.00	3.40
4k	Yes	140.60	6.30
16k	Yes	50.96	16.87
64k	Yes	16.34	59.46
256k	Yes	4.30	231.62

Writing a Tagged Delimited String, Tagged Delimited Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition the incoming Generic XML message is converted to a Tagged Delimited String, Tagged Delimited outgoing message. This causes a full parse of the incoming message payload which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML message and writing out a Tagged Delimited String output message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	533.60	2.44
4k	No	226.80	5.06
16k	No	70.00	15.38
64k	No	18.34	57.32
256k	No	4.62	227.12
1k	Yes	256.40	3.39
4k	Yes	151.20	6.17
16k	Yes	54.20	16.89
64k	Yes	16.48	59.35
256k	Yes	4.37	230.56

Writing an MRM XML Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition the incoming Generic XML message is converted to an MRM XML outgoing message. This causes a full parse of the incoming message payload which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML message and writing out an MRM XML output message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	440.40	2.84
4k	No	166.20	6.71
16k	No	47.62	22.09
64k	No	12.34	83.74
256k	No	3.10	331.65
1k	Yes	238.80	3.76
4k	Yes	115.00	7.88
16k	Yes	40.88	23.43
64k	Yes	11.46	85.84
256k	Yes	2.95	336.75

Writing a Custom Wire Format Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition the incoming Generic XML message is converted to a Custom Wire Format outgoing message. This causes a full parse of the incoming message payload which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML message and writing out a Custom Wire Format output message.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	616.60	2.21
4k	No	262.60	4.50
16k	No	80.20	13.56
64k	No	21.24	49.59
256k	No	5.34	196.07
1k	Yes	284.60	3.09
4k	Yes	164.60	5.62
16k	Yes	66.20	14.71
64k	Yes	19.60	51.22
256k	Yes	5.14	200.63

Parsing a Message in the XML Domain

The tests in this section illustrate the CPU processing costs of parsing different message formats in the XML domain.

Parsing a Generic XML Input Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a Generic XML input message. As the message body is ignored there are no writing costs.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1277.80	1.32
4k	No	734.40	1.99
16k	No	274.80	4.56
64k	No	78.80	14.90
256k	No	20.26	56.88
1k	Yes	409.20	2.18
4k	Yes	300.80	2.83
16k	Yes	159.80	5.63
64k	Yes	61.00	16.45
256k	Yes	18.00	59.70

Parsing a Generic XML Input Message Containing XML Entities

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a variable is declared and set to the last element in the incoming message. This causes a full parse of the incoming message including the tags and entities. The output message consists of a message header only and no payload.

This test identifies the cost of parsing a Generic XML input message containing many XML Entities to see what the effect of having entities present in the message is.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1619.40	1.11
64k	No	95.80	12.94

Writing a Message in the XML Domain

The test in this section illustrates the CPU processing costs of using the XML domain to create an output message. This is the processing associated with taking a message tree in OutputRoot and flattening it to create a bitstream which is the output message.

Writing a Generic XML Output Message

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the entire Message from the incoming message is copied over to the outgoing message. In addition the last element in the incoming message is modified. This causes a full parse of the incoming message which is then written as the payload of the output message.

This test identifies the cost of parsing a Generic XML input message and writing a Generic XML output message.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1018.60	1.54
4k	No	490.80	2.65
16k	No	162.40	7.13
64k	No	44.08	25.11
256k	No	11.20	98.40
1k	Yes	330.00	2.43
4k	Yes	208.40	3.75
16k	Yes	98.00	8.53
64k	Yes	32.82	27.40
256k	Yes	9.10	103.92

External Resources

The tests in this section illustrate the processing cost of accessing resources such as a database or external procedure.

Accessing a Database from a Message Flow

The tests in this section illustrate the processing cost of performing operations on a DB2 database.

Reading from a Database

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition a SELECT is performed to obtain a piece of data from the Database. This data is used to validate an element in the input message.

This test identifies the cost of performing a Database SELECT.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	588.40	2.28
4k	No	572.40	2.41
16k	No	513.20	2.86
64k	No	181.00	4.41
256k	No	45.58	10.06
1k	Yes	173.00	3.47
4k	Yes	169.00	3.63
16k	Yes	141.60	4.22
64k	Yes	96.80	6.00
256k	Yes	39.00	13.66

Inserting into a Database

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition an INSERT is performed to populate the database with a piece of data. This data is obtained from an element in the input message.

This test identifies the cost of performing a Database INSERT.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	294.40	2.28
4k	No	292.00	2.42
16k	No	278.00	2.84
64k	No	180.20	4.58
256k	No	45.74	10.03
1k	Yes	118.80	3.62
4k	Yes	115.40	3.75
16k	Yes	100.80	4.32
64k	Yes	76.20	6.14
256k	Yes	37.32	13.58

Updating a row in a Database

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. In addition an UPDATE is performed to update a piece of data in the database with a new value. This value is obtained from an element in the input message.

This test identifies the cost of performing a Database UPDATE.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	546.20	2.38
4k	No	547.40	2.48
16k	No	511.00	3.03
64k	No	181.20	4.71
256k	No	45.06	10.79
1k	Yes	170.60	3.92
4k	Yes	162.40	4.03
16k	Yes	137.20	4.20
64k	Yes	91.80	5.99
256k	Yes	35.32	13.75

Calling External Procedures

The tests in this section illustrate the processing cost of invoking an external procedure such as a Java class or database stored procedure with different parameters.

Calling an External Java Procedure with no Parameters

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. Two thousand identical calls are made to an external Java procedure. The procedure receives zero input parameters and passes back zero parameters returning immediately.

This test identifies the cost of calling a Java procedure with zero parameters.

The results of running this test are given in the table below. The CPU ms/msg figure has been adjusted to report a per procedure invocation cost by dividing the CPU cost obtained from the test results by 2000.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	88.200	0.006
1k	Yes	79.000	0.006

Calling an External Java Procedure with One Integer Input Parameter

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. Two thousand identical calls are made to an external Java procedure. The procedure receives one Integer parameter and passes back zero parameters returning immediately.

This test identifies the cost of calling a Java procedure with one Integer parameter.

The results of running this test are given in the table below. The CPU ms/msg figure has been adjusted to report a per procedure invocation cost by dividing the CPU cost obtained from the test results by 2000.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	47.360	0.011
1k	Yes	44.160	0.011

Calling an External Java Procedure with Twenty Integer Input Parameters

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. Two thousand identical calls are made to an external Java procedure. The procedure receives twenty parameters all of which are integers and passes back zero parameters returning immediately.

This test identifies the cost of calling a Java procedure with twenty parameters which are integers.

The results of running this test are given in the table below. The CPU ms/msg figure has been adjusted to report a per procedure invocation cost by dividing the CPU cost obtained from the test results by 2000.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	5.620	0.091
1k	Yes	5.600	0.090

Routing and Transformation Logic

The tests in this section illustrate the processing cost of simple routing and transformation logic using a variety of routing and transformation technologies (ESQL, JavaCompute node, XML Transformation). A number of the tests are performed for each of the technologies thus allowing a simple comparison of CPU processing costs to be made. In other cases a comparison is only made within a technology such as looking at the efficiency of different parsers whilst using ESQL.

These tests are not a definitive statement of the relative processing costs of the different technologies. They are provided for illustrative purposes only. Message processing performance will be affected by the complexity of the messages and processing to be performed on the messages.

Using ESQL

The tests in this section illustrate the processing costs of using ESQL for different routing and transformation operations.

Filter an Incoming Message based on the First Element in the Message using the XML Parser

This test consists of MQ Input Node -> Filter Node -> MQ Output Node.

Within the filter node the first element of the incoming message is examined. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the start of a message using the XML parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1980.00	0.98
4k	No	1789.00	1.03
16k	No	646.60	1.66
64k	No	169.60	3.32
256k	No	38.86	10.26
1k	Yes	435.60	1.84
4k	Yes	332.60	2.16
16k	Yes	193.60	2.87
64k	Yes	80.40	5.36
256k	Yes	20.92	15.57

Filter an Incoming Message Based on the Last Element in the Message using the XML Parser

This test consists of MQ Input Node -> Filter Node -> MQ Output Node.

Within the filter node the last element of the incoming message is examined. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the end of a message using the XML parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1557.00	1.15
4k	No	823.40	1.83
16k	No	287.20	4.43
64k	No	78.80	15.03
256k	No	20.16	57.56
1k	Yes	406.40	2.03
4k	Yes	264.00	2.83
16k	Yes	129.20	5.74
64k	Yes	48.22	17.19
256k	Yes	13.94	64.41

Filter an Incoming Message Based on the First Element in the Message using the XMLNSC Parser

This test consists of MQ Input Node -> Filter Node -> MQ Output Node.

Within the filter node the first element of the incoming message is examined. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the start of a message using the XMLNSC parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1888.00	0.96
4k	No	1647.40	1.09
16k	No	642.20	1.60
64k	No	168.60	3.26
256k	No	38.66	9.85
1k	Yes	421.20	1.84
4k	Yes	321.60	2.15
16k	Yes	193.40	2.84
64k	Yes	80.00	5.26
256k	Yes	20.66	16.28

Filter an Incoming Message Based on the Last Element in the Message using the XMLNSC Parser

This test consists of MQ Input Node -> Filter Node -> MQ Output Node.

Within the filter node the last element of the incoming message is examined. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the end of a message using the XMLNSC parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1534.20	1.16
4k	No	808.60	1.86
16k	No	278.80	4.49
64k	No	77.00	15.31
256k	No	19.56	58.99
1k	Yes	405.60	2.03
4k	Yes	266.00	2.83
16k	Yes	128.80	5.73
64k	Yes	47.68	17.16
256k	Yes	13.52	65.32

Computation on an Input Message using the XML Parser

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node ESQL is used to calculate the total of all items and prices within a repeating structure which is in the input message. The totals along with a copy of the input message are written in the outgoing message.

This test identifies the cost of using ESQL to perform computation and message parsing using the XML parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	822.40	1.77
4k	No	328.40	3.71
16k	No	96.60	11.36
64k	No	24.74	42.83
256k	No	5.92	176.91
1k	Yes	324.20	2.63
4k	Yes	175.40	4.75
16k	Yes	69.60	12.69
64k	Yes	21.22	44.92
256k	Yes	5.34	182.51

Computation on an Input Message using the XMLNSC Parser

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node ESQL is used to calculate the total of all items and prices within a repeating structure which is in the input message. The totals along with a copy of the input message are written out in the outgoing message.

This test identifies the cost of using ESQL to perform computation and message parsing using the XMLNSC parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	772.40	1.85
4k	No	309.60	3.94
16k	No	91.00	12.18
64k	No	23.38	45.80
256k	No	5.60	189.7
1k	Yes	275.80	2.57
4k	Yes	157.60	4.85
16k	Yes	63.20	13.59
64k	Yes	19.78	47.32
256k	Yes	5.00	194.73

Manipulation of an Input Message using the XML Parser

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node ESQL is written to significantly change the structure of the incoming message. The new structure is written as the output message

This test identifies the cost of using ESQL to perform message manipulation and message parsing using the XML parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	749.80	1.88
4k	No	280.80	4.24
16k	No	79.00	13.74
64k	No	19.38	54.06
256k	No	3.99	259.87
1k	Yes	280.40	2.75
4k	Yes	166.40	5.32
16k	Yes	61.40	14.94
64k	Yes	17.22	56.37
256k	Yes	3.73	267.09

Manipulation of an Input Message using the XMLNSC Parser

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node ESQL is written to significantly change the structure of the incoming message. The new structure is written as the output message

This identifies the cost of using ESQL to perform message manipulation and message parsing using the XMLNSC parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	759.20	1.86
4k	No	288.00	4.15
16k	No	81.60	13.35
64k	No	20.08	52.59
256k	No	4.13	252.74
1k	Yes	316.80	2.70
4k	Yes	170.40	5.21
16k	Yes	62.60	14.58
64k	Yes	17.94	54.32
256k	Yes	3.87	256.74

Manipulation of an Input Message using the EVAL Function on all Lines of ESQL in One Invocation

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node ESQL is written to significantly change the structure of the incoming message. The new structure is written as the output message The ESQL processing is run within a single EVAL statement.

This test identifies the cost of using the EVAL function to run a large amount of ESQL

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	109.80	9.73
4k	No	87.60	11.93
16k	No	48.32	21.49
64k	No	16.44	62.40
256k	No	3.74	271.79
1k	Yes	93.40	10.64
4k	Yes	74.20	13.05
16k	Yes	41.76	22.48
64k	Yes	15.32	63.92
256k	Yes	3.62	275.29

Manipulation of an Input Message using the EVAL Function on Each Line of ESQL

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node ESQL is written to significantly change the structure of the incoming message. The new structure is written as the output message. Each line of ESQL is run individually in an EVAL statement

This test identifies the cost of using the EVAL function on many lines of ESQL.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	58.80	17.77
4k	No	15.34	65.97
16k	No	3.87	260.36
64k	No	0.97	1043.11
256k	No	0.24	4226.37
1k	Yes	50.58	18.95
4k	Yes	14.82	67.22
16k	Yes	3.83	261.82
64k	Yes	0.96	1043.56
256k	Yes	0.23	4268.34

Manipulation of an Input Message Using the SELECT Function

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node an ESQL SELECT function is written to significantly change the structure of the incoming message. The new structure is written as the output message.

This test identifies the cost of using an ESQL SELECT function to perform message manipulation.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	782.80	1.82
4k	No	297.40	4.00
16k	No	86.40	12.57
64k	No	22.34	47.20
256k	No	5.60	187.12
1k	Yes	306.60	2.63
4k	Yes	175.60	5.05
16k	Yes	65.60	13.96
64k	Yes	19.78	49.84
256k	Yes	5.18	191.96

Manipulation of an Input Message using the ROW Function

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node the ESQL SELECT function within an ESQL ROW function is written to significantly change the structure of the incoming message. The new structure is written as the output message

This test identifies the cost of using an ESQL ROW function to perform message manipulation.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	756.60	1.83
4k	No	297.80	4.01
16k	No	85.20	12.72
64k	No	22.16	47.54
256k	No	5.58	187.06
1k	Yes	319.40	2.65
4k	Yes	176.20	5.08
16k	Yes	65.60	14.04
64k	Yes	19.66	49.59
256k	Yes	5.14	193.36

Manipulation of an Input Message using the ITEM Function

This test consists of MQ Input node -> Compute Node -> MQ Output Node.

Within the compute node an ESQL SELECT function using the ITEM clause is written to significantly change the structure of the incoming message. The new structure is written as the output message.

This test identifies the cost of using an ESQL ITEM function to perform message manipulation.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	784.20	1.82
4k	No	300.00	3.98
16k	No	86.40	12.56
64k	No	22.36	47.23
256k	No	5.64	186.06
1k	Yes	321.00	2.67
4k	Yes	177.80	5.03
16k	Yes	66.00	13.97
64k	Yes	19.72	49.82
256k	Yes	5.18	192.82

Calling an Internal ESQL Procedure with No Parameters

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. Two thousand identical calls are made to an internal ESQL procedure. The procedure receives zero input parameters and passes back zero parameters returning immediately.

This test identifies the cost of calling an ESQL procedure with zero parameters.

The results of running this test are given in the table below. The CPU ms/msg figure has been adjusted to report a per procedure invocation cost by dividing the CPU cost obtained from the test results by 2000.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	295.600	0.002
1k	Yes	208.000	0.004

Calling an Internal ESQL Procedure with One Integer Input Parameter

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message. Two thousand identical calls are made to an internal ESQL procedure. The procedure receives one integer parameter and passes back zero parameters returning immediately.

This test identifies the cost of calling an ESQL procedure with one Integer parameter.

The results of running this test are given in the table below. The CPU ms/msg figure has been adjusted to report a per procedure invocation cost by dividing the CPU cost obtained from the test results by 2000.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	191.400	0.003
1k	Yes	145.600	0.003

Calling an Internal ESQL Stored Procedure with Twenty Integer Input Parameters

This test consists of MQ Input Node -> Compute Node -> MQ Output Node.

Within the compute node the message headers from the incoming message are copied over to the outgoing message Two thousand identical calls are made to an internal ESQL procedure. The procedure receives twenty parameters all of which are integers and passes back zero parameters returning immediately.

This test identifies the cost of calling an ESQL procedure with twenty parameters which are integers.

The results of running this test are given in the table below. The CPU ms/msg figure has been adjusted to report a per procedure invocation cost by dividing the CPU cost obtained from the test results by 2000.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	48.040	0.011
1k	Yes	44.260	0.011

Using Java

The tests in this section illustrate the processing costs of using the JavaCompute node for different routing and transformation operations.

WebSphere Message Broker development is aware of a performance problem when using XPath and No XPath with the JavaCompute node which results in higher CPU usage than would be expected in some circumstances. This occurs in the test Manipulation of an Input Message using the Java Compute Nodes XPath Capability for example. The problem is being investigated and a fix will be shipped in due course. Contact IBM service for the latest status on this problem.

Until this problem is fixed you should not make performance comparisons between the JavaCompute test results and other techniques for computation or transformation.

Filter an Incoming Message Based on the First Element in the Message using the Java Compute Nodes XPath Capability

This test consists of MQ Input Node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node the first element of the incoming message is examined using the XPath capability. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the start of a message using the Java Compute Node XPath capability.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1828.20	0.99
4k	No	1598.60	1.12
16k	No	638.80	1.61
64k	No	168.40	3.20
256k	No	38.80	9.88
1k	Yes	401.00	1.86
4k	Yes	307.20	2.13
16k	Yes	171.20	2.84
64k	Yes	61.80	5.26
256k	Yes	16.34	15.66

Filter an Incoming Message Based on the Last Element in the Message using the Java Compute Nodes XPath Capability

This test consists of MQ Input Node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node the last element of the incoming message is examined using the XPath capability. The result is always set to be true and thus the message is propagated to the MQOutput node

This test identifies the cost of filtering on an element at the end of a message using the Java Compute Node XPath capability.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1501.60	1.17
4k	No	810.20	1.84
16k	No	285.00	4.43
64k	No	79.40	14.81
256k	No	20.42	56.61
1k	Yes	377.00	2.08
4k	Yes	251.60	2.85
16k	Yes	119.40	5.72
64k	Yes	45.06	17.07
256k	Yes	13.04	63.15

Filter an Incoming Message Based on the First Element in the Message using the Java Compute Nodes No XPath Capability

This test consists of MQ Input Node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node the first element of the incoming message is examined using the No XPath capability. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the start of a message using the Java Compute Node No XPath capability.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1864.40	1.03
4k	No	1622.20	1.11
16k	No	639.60	1.59
64k	No	168.00	3.15
256k	No	38.78	9.76
1k	Yes	425.00	1.83
4k	Yes	320.80	2.12
16k	Yes	195.80	2.83
64k	Yes	75.60	5.35
256k	Yes	19.92	15.43

Filter an Incoming Message Based on the Last Element in the Message using the Java Compute Nodes No XPath Capability

This test consists of MQ Input Node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node the last element of the incoming message is examined using the No XPath capability. The result is always set to be true and thus the message is propagated to the MQOutput node.

This test identifies the cost of filtering on an element at the end of a message using the Java Compute Node No XPath capability.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	1517.00	1.18
4k	No	813.80	1.85
16k	No	285.00	4.45
64k	No	79.20	14.90
256k	No	20.30	57.15
1k	Yes	406.80	2.05
4k	Yes	268.40	2.88
16k	Yes	130.60	5.77
64k	Yes	48.74	17.16
256k	Yes	13.98	63.28

Computation on an Input Message using the Java Compute Nodes XPath Capability

This test consists of MQ Input node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node Java code is used to calculate the total of all items and prices within a repeating structure which is in the input message. The totals along with a copy of the input message are written in the outgoing message.

This test identifies the cost of using Java to perform computation and message parsing using the XML parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	448.00	2.86
4k	No	234.40	5.04
16k	No	77.00	14.32
64k	No	20.54	52.20
256k	No	5.18	206.46
1k	Yes	238.20	3.75
4k	Yes	142.00	6.10
16k	Yes	56.00	16.09
64k	Yes	17.58	55.30
256k	Yes	4.69	212.07

Manipulation of an Input Message using the Java Compute Nodes XPath Capability

This test consists of MQ Input node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node Java code, utilising the XPath capability is used to significantly change the structure of the incoming message. The new structure is written as the output message.

This test identifies the cost of using Java code and XPath to perform message manipulation.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	116.60	10.15
4k	No	100.60	11.34
16k	No	39.78	26.89
64k	No	12.30	85.78
256k	No	3.25	323.86
1k	Yes	83.60	12.05
4k	Yes	73.80	13.54
16k	Yes	34.08	28.46
64k	Yes	11.60	86.99
256k	Yes	3.10	328.29

Manipulation of an Input Message using the Java Compute Nodes No XPath Capability

This test consists of MQ Input node -> Java Compute Node -> MQ Output Node.

Within the Java Compute Node Java code, utilising the No XPath capability is used to significantly change the structure of the incoming message. The new structure is written as the output message.

This test identifies the cost of using Java code and No XPath to perform message manipulation.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	561.60	2.36
4k	No	182.60	6.19
16k	No	49.48	21.43
64k	No	12.64	82.48
256k	No	3.16	328.28
1k	Yes	267.80	3.22
4k	Yes	127.20	7.27
16k	Yes	42.60	22.68
64k	Yes	11.80	84.20
256k	Yes	3.01	333.04

Using XMLT

The tests in this section illustrate the processing costs of using an XML Transformation node to perform a computation and manipulation of an input message.

Computation on an Input Message

This test consists of MQ Input node -> XMLT Node -> MQ Output Node.

Within the XMLT Node a compiled stylesheet is used to calculate the total of all items and prices within a repeating structure which is in the input message. The totals along with a copy of the input message are written in the outgoing message.

This test identifies the cost of using an XSL stylesheet to perform computation and message parsing using the XML parser.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	305.00	3.92
4k	No	129.60	8.52
16k	No	38.70	27.23
64k	No	10.16	102.28
256k	No	3.40	309.57
1k	Yes	200.00	4.78
4k	Yes	97.40	9.61
16k	Yes	33.76	28.58
64k	Yes	9.44	104.41
256k	Yes	3.16	313.86

Manipulation of an Input Message

This test consists of MQ Input node -> XMLT Node -> MQ Output Node.

Within the XMLT Node a compiled stylesheet is used to significantly change the structure of the incoming message. The new structure is written as the output message.

This test identifies the cost of using an XSL stylesheet to perform message manipulation.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	328.00	3.68
4k	No	145.60	7.63
16k	No	45.00	23.86
64k	No	12.08	87.20
256k	No	4.35	244.01
1k	Yes	210.40	4.60
4k	Yes	105.20	8.51
16k	Yes	38.62	24.92
64k	Yes	11.24	89.27
256k	Yes	4.07	249.20

Publish Subscribe

The tests in this section illustrate the processing costs of using the publish/subscribe functions within WebSphere Message Broker with different message protocols and varying numbers of subscribers.

Topic Based Publish/Subscribe using Non Persistent MQ Messages

This test consists of MQInput node -> Publication node.

A publisher publishes a message on a single topic. The test is run repeatedly with varying numbers of subscribers (1, 10, 100 and 1000). All subscribers are registered to receive messages on the single topic.

Non persistent MQ messages 1K in size are used by the publisher.
This test identifies the cost of using the Publication node for a single publisher, varying subscribers, single topic and a single copy of the message flow when using non persistent MQ messages.

The results of running this test are given in the table below.

Publishers	Topics	Subscribers	Msg Size	Persistent	Message Rate (Msgs/sec)	% CPU Busy	CPU ms/msg
1	1	1	1024	No	2650.00	71.00	0.54
1	1	10	1024	No	5225.00	100.00	0.38
1	1	100	1024	No	4444.00	99.00	0.45
1	1	1000	1024	No	1801.80	80.00	0.89

Topic Based Publish/Subscribe using MQ Real-time Messages

This test consists of a Real-time OptimizedFlow Node.

A publisher publishes a message on a single topic. The test is run with a single subscriber which is registered to receive messages on the single topic.

Both the publisher and subscriber use the WebSphere MQ Real-time transport to publish and subscribe to messages.

This test identifies the cost of using the Publication node for a single publisher, subscriber, topic and a single copy of the message flow when using the WebSphere MQ Real-time transport.

The results of running this test are given in the table below.

Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1k	No	31600.00	0.04

Scaling Message Throughput

The tests in this section show the effect of using two different approaches to increase message throughput for a message flow. These are the use of additional instances and assigning one copy of the message flow to each of multiple execution groups.

Using Additional Instances

This test consists of running the Large Messaging sample with a varying number of instances of the message flow in a single execution group.

The purpose of this is to see how effective the use of additional instances is in increasing message throughput and achieving higher system CPU utilization. The benefits observed in any given situation will depend on the processing requirements of the message flow. CPU bound message flows will have different scaling characteristics from those which are I/O bound for example.

The results of running this test are given in the table below.

Instances	Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1	1k	No	659.2	2.15
2	1k	No	1230.8	2.23
4	1k	No	1992.2	2.38
8	1k	No	1591.4	2.67

The results in the table show that when running with 2 instances of the message flow there is an increase in message throughput. It was possible to achieve 1.87 times that achieved when with running one instance.

When running with 4 instances of the message flow message throughput continued to increase. It was possible to achieve 3.02 times that of one instance.

When using 8 instances of the message flow declined when compared with 4 instances. It was possible to achieve 2.41 times that of one instance.

This drop in throughput for 8 instances when compared with 4 instances is being investigated and a fix will be shipped in due course. Contact IBM service for the latest status on this problem.

These figures illustrate the scaling behaviour for one, CPU bound, workload. In an I/O bound workload the use of a higher number of instances is more likely to be effective. As each case can be different you are recommended to determine the optimum number of instances to use for each message flow individually through experimentation with a varying number of instances.

From these measurements we can see that use of additional instances can be used as a mechanism for increasing message throughput. Whilst not realising the full potential because of the scaling problem it was still possible to increase message throughput from 659 messages/second to a peak of 1992 messages/second.

Using Multiple Execution Groups

This test consists of running the Large Messaging sample with a single instance of the message flow in a varying number of execution groups.

The purpose of this is to see how effective the use of multiple execution groups is in increasing message throughput and achieving higher system CPU utilization. The benefits observed in any given situation will depend on the processing requirements of the message

flow. CPU bound message flows will have different scaling characteristics from those which are I/O bound for example.

The results of running this test with are given in the table below.

Execution Groups	Msg Size	Persistent	Message Rate (Msgs/sec)	CPU ms/msg
1	1k	No	659.2	2.15
2	1k	No	1278.4	2.18
4	1k	No	2434.8	2.24
8	1k	No	3215.6	2.43

The results in the table show that when running with 2 execution groups there is an increase in message throughput. It was possible to achieve 1.93 times that achieved when with running one instance.

When running with 4 instances of the message flow message throughput continued to increase. It was possible to achieve 3.69 times that of one instance.

When using 8 copies of the message flow message throughput continued to increase. It was possible to achieve 4.88 times that of one instance.

These figures illustrate the scaling behaviour for one, CPU bound, workload. In an I/O bound workload the use of a higher number of instances is more likely to be effective. As each case can be different you are recommended to determine the optimum number of copies to use for each message flow individually through experimentation.

From these measurements we can see that use of execution groups is an effective mechanism for increasing message throughput and allowing a machine to be fully utilised.

Overheads

The tests in this section indicate the processing costs of using Accounting and Statistics and Trace on a message flow.

Using Accounting and Statistics

This test consists of running a single copy of the Large Messaging sample with basic thread level and advanced node level accounting activated.

Using a 1K message size there was a 7.7% reduction in message throughput. This is a CPU overhead and reflects the additional cost of processing needed to collect the data.

Using a lower level of reporting would have resulted in a lower overhead.

Using Trace

This test consists of running a single copy of the Large Messaging sample whilst taking a user trace of type normal at the same time.

Using a 1K message size there was a 17.3% reduction in message throughput. This reflects the CPU and I/O overhead of writing user trace.

With debug trace the overhead will be even higher as debug trace is more extensive.

You are strongly recommended not to use WebSphere Message Broker trace in a production system. You are also strongly recommended not to have any trace nodes in the main processing paths of message flows. Even if trace is not active a penalty is still incurred to evaluate the expression specified in the Trace Node.

Resource Requirements

This section details the recommended minimum specification of a machine on which to install the development toolkit and Message Broker runtime. It also illustrates virtual memory use for message flows.

Recommended Minimum Specification

The recommended minimum specification machine to install and run the development toolkit is:

- Any Intel Pentium III (or higher) processor-based IBM PC or compatible with 700 or more MHZ processor speed. This is the minimum supported level. For improved performance use a 2 GHz or faster processor.
- Up to 6.1 GB of disk space
  - 4.5 GB disk plus 1.5 GB temporary space for WebSphere Message Broker
  - 105 MB for ODBC drivers for Cloudscape
- 512MB memory. This is the minimum requirement though and 1GB is recommended.

The recommended minimum specification machine to install and run the broker runtime is:

- Any zSeries processor which meets the required minimum specification as detailed in the announcement letter. For improved performance use a faster processor. For production a multi-processor machine is recommended.
- Up to 915 MB disk space
  - 315 MB disk plus 300 MB temporary space for WebSphere Message Broker
  - 300 MB for DB2 Enterprise Server compact version (assuming DB2 as the Message Broker database)
- 2 GB real memory available to the z/OS image. This is the minimum requirement though and more is recommended for busy systems. How much will depend on the complexity of the messages and message flows. For an active development and test environment running multiple brokers and other work 4GB is recommended. Production systems should be sized using experience gained with development and test systems.

These are recommended minimum specifications which are suitable to enable the processing of simple messages with simple message transformation or routing. Situations requiring more intensive processing are likely to need greater resources.

For more guidance on the support processors and configuration requirements see the WebSphere Message Broker *Managing Your Installation* manual.

Memory Use

The amount of virtual and real memory used by a message flow running within an execution group will vary, dependent on the complexity of the message flow, the style of processing within the message flow and the size of the messages being processed. This is a complex subject and a detailed discussion is beyond the scope of this document. However to assist with planning the memory used for a variety of tests is reported.

The figures in the table below record the amount of real memory used by an execution group for the message flow when it is initially deployed and after it has processed a number of messages and the size stabilised.

In each case a single copy of the message flow was deployed to a single execution group. Each use case was deployed to a new execution group. The figures are those reported by the `rmfwdm` command and are given to the nearest 1 MB.

Use Case	Real memory usage with message flow deployed but no messages processed	Real memory usage after processing messages
Empty Execution Group	209	N/A
Aggregation	210	214
Coordinated Request Reply	211	213
Data Warehouse		
Large Messaging	209	209
Message Routing	209	209
SWIFT Message parse	304	322
XMLT	209	351

Real Memory Use in MB for a Variety of Use Cases.

In taking these figures the minimum heap size of the WebSphere Message Broker Java Virtual Machine (JVM) was allowed to default to the value of 128MB.

zAAP Utilization

Conceptually a zAAP is very similar to a System Assist Processor (SAP); they are not able to execute an Initial Program Load and can only assist the general purpose Control Processors (CPs) in the execution of Java programming under control of the IBM JVM. However, unlike standard CPs, Internal Coupling Facility (ICFs), and Integrated Facility for Linux (IFLs), zAAPs can do nothing on their own. For this reason, IBM does not impose software charges on zAAP capacity. As such the ability to run processing on a zAAP helps reduce the total cost of ownership (TCO) of a software component. The more processing that is eligible to run on a zAAP the greater the benefit of having one.

zAAPs are designed to operate asynchronously with the general CPs to execute Java programming under control of the IBM Java Virtual Machine (JVM). This is an important point because zAAPs can only help execute Java applications that use the IBM JVM. The IBM JVM processing cycles can be executed on the configured zAAPs with no modifications to the Java applications.

For more information on the zAAP see the references listed in the section "Additional Information".

Within Message Broker the ability to utilise a zAAP comes from the execution of

- Business processing code written in Java using the Java node or a Java plug-in node
- XSLT processing
- The Publication node

Processing associated with message parsing and serialisation of messages using the standard broker parsers (XML, XMLNS, XMLNSC, MRM etc) or the execution of ESQL is not eligible to run on a zAAP as it is not written in Java. A customer parser written in Java would be eligible though.

The extent to which a particular message flow can benefit from the use of a zAAP will depend very much on the nature of the processing in it. To provide an illustration of the extent to which processing can utilise a zAAP the offload percentage for three of the tests is given. These tests all involve some Java processing but to differing extents as the results will show. The off-load is a measure of how much processing is eligible to run on a zAAP rather than a CP.

The percentage off-load figures given are those reported by the IBM JVM. For these particular tests a zAAP was not available on the image on which the test was run but the IBM JVM is still able to provide an estimate of what the off-load would be were a zAAP available.

To enable the reporting of the off-load information the command `Set IBM_JAVA_OPTIONS=-Xifa:project1` was coded in the Message Broker ENVFILE prior to Broker start.

The value of 1 specifies that the reports of processor utilization should be produced every 1 minute.

The results for three tests which are covered in the report are presented in the tables below. The figures presented for each test are typical values for the reported distribution of CPU processing by the IBM JVM while each test was executing in steady state for a period which was longer than the reporting interval.

Computation on an Input Message using the Java Compute Nodes XPath Capability

Message Size	Total Address Space CPU seconds	zAAP eligible CPU seconds	CP CPU seconds	% Off-load to zAAP
1K	60	2	58	3.33
4K	60	3	57	5.00
16K	60	3	57	5.00
64K	60	3	57	5.00
256K	60	3	57	5.00

Manipulation of an Input Message using the Java Compute Nodes XPath Capability

Message Size	Total Address Space CPU seconds	zAAP eligible CPU seconds	CP CPU seconds	% Off-load to zAAP
1K	60	1	59	1.67
4K	63	3	60	5.00
16K	60	4	56	6.67
64K	60	4	56	6.67
256K	60	4	56	6.67

Manipulation of an Input Message using an XSL stylesheet

Message Size	Total Address Space CPU seconds	zAAP eligible CPU seconds	CP CPU seconds	% Off-load to zAAP
1K	62	50	12	80.65
4K	63	54	6	85.71
16K	60	58	2	96.67
64K	60	59	1	98.33
256K	61.5	60.8	0.7	98.86

These tests have shown that there is significant potential to off-load processing onto a zAAP. The maximum off-load was 98% of processing in the message flow. The minimum was 3%.

Although all three tests involved some Java processing there was clearly a big difference in the amount of processing which could be off-loaded. The amount which can be off-loaded is a function of the amount of Java processing in the message flow and the proportion of that Java processing as a part of all processing in the message flow.

You should not take away from these results the belief that the JavaCompute node is not capable of achieving a high zAAP utilization. With a different test case, containing more extensive Java code implementing business rules the percentage off-load could be much larger. The test involving the XSL Stylesheet was largely implemented in Java and so achieved a very high off-load percentage.

Tuning

This section details the parameters which were reviewed or changed in the course of obtaining the measurement results.

The description of each parameter is brief as a detailed discussion of the effects of any changes are beyond the scope of this document.

Message Broker

The Message Broker used in the measurements was configured in the following ways for all tests:

1. Transactional support was used where appropriate. When processing persistent messages it was used, with non persistent messages it was not. The use of transaction control means that message processing takes place within a WebSphere MQ unit of work. This involves additional CPU and I/O processing by WebSphere MQ because the unit of work is recoverable. The result is inevitably a reduction in message throughput for persistent messages. By default the transaction parameter on the MQInput node was set to automatic. This is the recommended value to use for transaction mode unless there is a specific requirement to use a particular value since persistent messages will be processed within transactional control and non persistent messages will not.

Additional tuning was performed for the **publish subscribe** tests. This was as follows:

1. The heap size of the The WebSphere Message Broker Java Virtual Machine (JVM) (in which much of the publish subscribe code is executed) was set to 512MB. For the non Publish Subscribe tests the default value of 128MB was used.

2. The thread settings of the RealtimeOptimizedNode used a default value of 10 read and write threads. These were sufficient to cater for the test cases run in this report. However if more clients are used increasing these values could be beneficial.

3. Client Pinging - The ping protocol implements a “keep alive” protocol where the broker is periodically verifying that connected clients are alive. This process allows the broker to detect disconnected clients and maintain an updated subscription list. In situations where there are a large number of clients connected to a broker, this pinging process may account for a large proportion of the messages exchanged between the broker and clients and can impact the broker's message throughput. In such circumstances, the ping interval can be turned off or alternatively increased to reduce the amount of traffic generated by pinging. For the tests in the report the value was set to 0.

4. Client Queue Size - The broker employs a set of internal queues which are used to regulate the delivery of messages to subscribers. Note: these are not the queues used by the MQ Transport. The size of the queue specifies the number of bytes of data that the broker will store for one client. If this maximum is exceeded, the broker will take action which is determined by the value of the parameter "Client disconnection due to queue overflow". The default queue size is 100,000 bytes. Setting the value to zero allows the broker to grow the queue size as required. In this case, the queue size will only be limited by the available system memory. In the tests detailed in this report a value of 0 was used.

5. Client disconnection due to queue overflow - When the depth of the client queue exceeds the Client Queue Size value, the broker can choose between two courses of action. The default action is to disconnect the client; in this case, all the queued messages are lost. This is specified through a value of true for the parameter. The alternative course of action, specified with a value of false, is to keep the client connection alive but remove any excess messages from the client's queue. In the tests detailed in this report a value of false was used.

6. Maximum message size - If the broker receives a message that is bigger than the maximum message size value, it will disconnect the client that sent the message. This feature is useful for protecting the broker from applications sending excessively large messages. The default maximum message size is 100,000 bytes and this was adequate for the tests run in the report so the value was left unchanged.

7. Maximum number of client connections - The broker has the ability to limit the number of client connections that it will handle. This is useful in situations where the number of client applications that will connect to the broker is unknown. In such cases limiting the number of connections will allow the broker to maintain a particular level of service (this level will depend upon the particular environment in which the broker is being used). The default setting is unlimited. This was also the value used for the tests run in the report.

8. Interval statistics reporting was enabled and set to an interval of 10000 (10 seconds) so that the value of ClientBytesQueued could be monitored.

There were no error processing or error conditions in any of 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.

WebSphere MQ

The Message Broker queue manager on z/OS was allowed to default in its parameter settings.

The following changes were made to all queue managers running on Windows and Linux in the tests:

1. The value of DefaultQBufferSize was increased to a value of 1000000 for each queue used in the tests.

2. Given the use of persistent messages in the tests the following MQ log parameters were modified:
   - LogBufferPages was set to 0 allowing the value to go to its maximum
   - LogFileSize was set to 1024
   - LogType was set to circular
   - LogPrimaryFiles was set to 3
   - LogSecondaryFiles was set to 2

3. Circular logging was set for all WebSphere MQ queue managers used in the tests.

4. The Message Broker queue manager MQ listener and channels were run as trusted applications. In the queue manager qm.ini the value MQIBindType was set to FASTPATH in the channel stanza. The environment variable MQ_CONNECT_TYPE=FASTPATH was present in the environment in which the broker queue manager was started.

TCP/IP

No specific tuning was performed for TCP/IP. All machines used the operating system default values.

Database

The DB2 instance used with the message broker was a default configuration and the only tuning performed on the instance was placement of the database data and log files on different disks.

Miscellaneous

Although not implemented in all cases the following additional tuning changes are recommended

- Follow the recommendations described in SupportPac MP1E: WebSphere MQ for z/OS Performance and SupportPac MP16: Capacity planning and tuning for WebSphere MQ for z/OS. Links for both documents are given in the section Additional Information.

- Locate the log of any WebSphere MQ queue manager through which persistent messages pass on a dedicated disk.

- Locate the WebSphere MQ queue manager log on a very fast disk such as one with a non-volatile fast write cache. Such disks are consistently capable of I/O times of 1ms compared with a time of 6 ms for a 10,000 RPM SCSI disk. When using a disk with a fast write cache it is essential that it has a non-volatile capability as the log data is critical to the integrity of your queue manager.

- Note that there is no need to locate the WebSphere queue manager queue file on a fast disk. It is advisable to locate it on a dedicated disk in order to improve the efficiency of queue manager checkpoint processing.

- Locate the log of the Message Broker database on a dedicated disk.

- Locate the log of the Message Broker database on a very fast disk such as one with a non-volatile fast write cache.

- When performing BLOB inserts to a database locate the data portion of the database on a very fast disk such as one with a non-volatile fast write cache. BLOB I/O is not buffered by a database such as DB2 and is written to disk immediately.

- When using the aggregation node follow the message flow coding advice provided in **Supportpac IP05, WebSphere MQ Integrator V2.1 - Optimizing Use of Aggregation Nodes** which is available at http://www.ibm.com/software/integration/support/supportpacs/individual/supportpacs/ip05.pdf.

  NOTE: When using WebSphere Message Broker V6 there is no need to follow the recommendations in the document about Message Broker database configuration. This is because the aggregation mechanism is now based on the use of WebSphere MQ queues, rather than a database table as with previous versions of the Message Broker.

Additional Tuning Information

In order to obtain the maximum message rate for your implementation it is important that you understand the current best practices for WebSphere Message Broker. These practices cover the architecture of message flow processing, the coding of message flows as well as the configuration and tuning of the message broker and associated components.

Such information can be found in the Business Integration Zone of WebSphere Developer Domain. A suggested starting place is the article http://www.ibm.com/developerworks/websphere/library/techarticles/0403_dunn/0403_dunn.html which highlights the information available and where it may be found.

Conclusion

This report has detailed the key performance characteristics of the WebSphere Message Broker V6 runtime. The primary focus in the report has been on identifying the CPU costs of different functions. Additional information has been supplied on virtual memory requirements for the use cases.

Part I showed the level of message throughput that can be expected for a variety of common use cases. You have the ability to run these same tests in your own environment as the messages flows are shipped as product samples. Using machines with even faster processors it will be possible to achieve high message rates.

From the data supplied in Part I of the report it is possible to see that there have been some significant reductions in CPU costs and as a result increases in message throughput in WebSphere Message Broker V6 when compared with WebSphere Business Integration Message Broker V5. Most notably improvements in message throughput of

- 3.2 times for the Aggregation use case
- 2.1 times for the Coordinated Request/Reply use case
- 2.2 times for the Message Routing use case
- 5.0 times for the SWIFT Message Parse use case

No application changes are needed to obtain the majority of these performance improvements. They come as standard with version 6 and are available immediately on installation and migration of the message flows.

Some additional gains in performance are available by using new function such as shared variables.

The provision of the new JavaCompute node has significantly extended the ease with which it is possible to utilise a zAAP and so benefit from reduced software costs. Not only is it now possible to code all business processing in a message flow within Java but it is also much easier to code and deploy that code using the JavaCompute node.

In addition the significant reduction in CPU usage in the key areas of WebSphere Message Broker V6 means that is possible to reduce the CPU required to run a given workload or alternatively run at a high processing rate for a given amount of CPU.

Collectively these changes represent a major advance in performance and a lowering of the TCO of WebSphere Message Broker V6 on z/OS.

Appendix A - Measurement Environment

All throughput measurements where taken on a single server machine. The client type and machine on which they ran varied with the test. The details are given below.

Server Machine

The hardware consisted of

- An LPAR on an IBM zSeries 990 2084was used for all measurements.
    - For those measurements listed in Part I the LPAR was configured with 8 processors and 8GB memory. This is equivalent to an IBM zSeries 990 8 way machine, a 2084-308.
    - For those measurements listed in Part II the LPAR was configured with 2 processors and 8 GB memory. This is equivalent to an IBM zSeries 990 2 way machine, a 2084-302.
- Enterprise Storage Server (ESS) Model 800 disk storage system featuring redundant hardware, mirrored write caches, and RAID-5 and RAID-10 protection for the disks.
- 1 Gb Ethernet card

The software consisted of:

- z/OS Version 1 Release 6 including the required RRS (Resource Recovery Services) and OMVS Hierarchical File Subsystem (HFS)
- Java ™ Runtime Environment (JRE) 1.4.2
- WebSphere MQ for z/OS Version 6
- DB2 Universal Database for OS/390 and z/OS Version 8
- WebSphere Business Integration Message Broker for z/OS V6.0

Client Machines

A number of different client machines were used dependent on the tests being run. The different configurations are described below.

Point to Point Testing
The hardware consisted of:

- An IBM Netfinity 8500R with 4 * 900 MHz Pentium III Xeon processors
- Four 34 GB SCSI hard drives formatted with NTFS
- Two 8.5 GB SCSI hard drives formatted with NTFS
- 1 GB RAM
- 1 Gb Ethernet card

The software consisted of:

- Microsoft Windows 2000 with Service Pack 4
- WebSphere MQ V5.3

Publish Subscribe Testing
The hardware consisted of multiple machines of this specification:

- An IBM xSeries processor with 4 * 1800 MHz Pentium 4 Xeon processors
- 2 GB SCSI hard drives
- 4 GB RAM
- 1 Gb Ethernet card

The software on all client machines consisted of:

- Linux Red Hat Advanced Server Version 2.1
- IBM Java 1.4.2
- WebSphere MQ V5.3.6.

Network Configuration

The client and server machines were connected using a full duplex 1 Gigabit Ethernet LAN with a single hub.

Appendix B - Evaluation Method

This section outlines the software components that were used to produce the measurement results which are contained in this report.

Two different configurations were used in the generation and consumption of input and output messages. This is because different test cases required different types of input and output messages. The methods used were:

1. Point to Point Message Processing

2. Publish Subscribe Message Processing

These are described in the remainder of this section.

A series of parameter configuration changes were made to improve message throughput. These are discussed in the section Tuning.

Point to Point testing

This section describes how messages were generated and consumed for the point to point messaging tests, such as the Database Read tests or Filter an Incoming Message based on the First Element in the Message. The configuration of the software components is also discussed.

Message Generation and Consumption

A multi threaded WebSphere MQ Client program written in C was used to generate input messages for the test case being run and to consume the WebSphere MQ output messages.

The client program used the Message Queue Interface (MQI). Both persistent and non persistent messages were generated from this program.

Sufficient threads were run in the multi threaded client to ensure that there were always messages on the input queue waiting to be processed. This is important when measuring message throughput.

Any thread within the client program was able to retrieve any message which had been processed by a message flow. No use was made of the WebSphere MQ correlation identifiers to limit consumption of a message to the thread which created it.

Machine Configuration

The WebSphere MQ client program used to generate and consume messages for the message flows was run on a dedicated machine, the Client Machine. The Message Broker, its dedicated WebSphere MQ queue manager and broker database were all located on a dedicated machine, the Server Machine.

There was a single client machine.

Messages were transmitted from the client machine to the server machine over WebSphere MQ SVRCONN channels. The messages were received on the server machine through use of a WebSphere MQ queue manager listener process.

The database used for the database related test cases used the same database instance as the Message Broker.

Messages were transmitted from the client machines to the server machine using the WebSphere MQ transport.

The diagram below illustrates the major components in the measurement environment and their location.

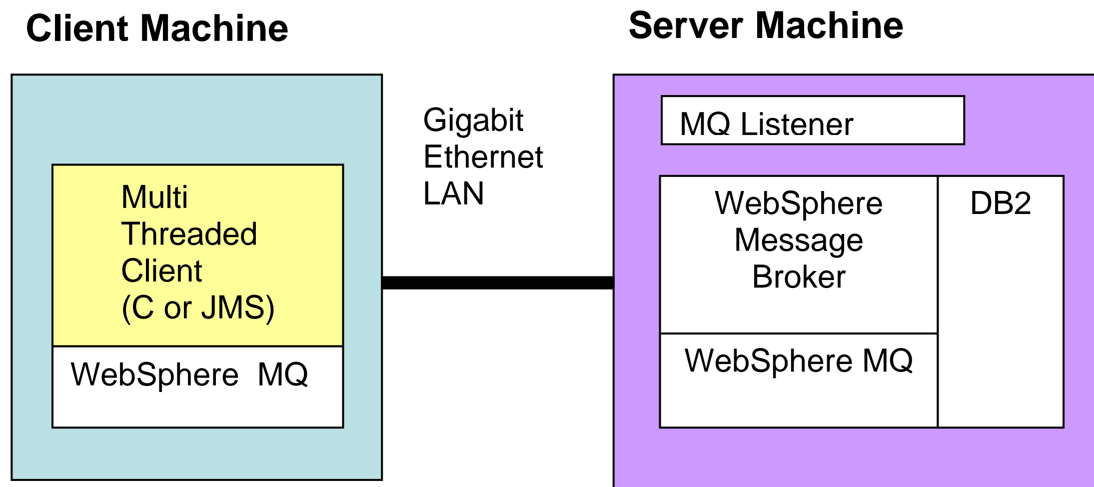

Both the client and server machine were configured with sufficient memory to ensure that no paging took place during the tests.

Publish Subscribe testing

This section describes how messages were generated and consumed for tests which used the publish subscribe message processing model.

Message Generation and Consumption

A multi threaded JMS client application was used to publish JMS messages and consume messages received by JMS Subscribers. The JMS client application used WebSphere Message Broker as the JMS Provider. The same client application was able to generate JMS messages using the WebSphere MQ and WebSphere MQ Real-time transports.

JMS bytes messages were used in the testing. The message content was not of interest in the tests only the topic under which it was published.

Messages were transmitted from the client machines to the server machine using either the WebSphere MQ or WebSphere Real-time transport depending on the test case.

When using the WebSphere MQ transport the publish rate was set to a high value, this publish rate was then throttled by the MQ acknowledgement protocol to a rate which was sustainable by the broker. The publisher acknowledgement interval was set to ensure messages were always available on the brokers input queue. Details of how to set the broker acknowledgement interval can be seen in the WebSphere MQ "Using Java" manual. Each subscriber was allocated its own temporary dynamic queue to store its messages on the broker.

The WebSphere MQ Real-time transport does not have a self throttling protocol like that of MQ. As a result the publish rate was manually adjusted until the optimum message throughout for the broker is found. The optimum level of message throughput was determined by monitoring the ClientBytesQueued value in broker statistics.

The value for ClientBytesQueued shows the number of bytes waiting to be delivered to subscriber clients. When the broker becomes overloaded it is unable to service this buffer fast enough and so the number of bytes that are queued increases. For a test to be successful the buffer size must not continually increase during the test run. Constant growth of this buffer indicates too high a publish rate. The point at which the buffer starts to fill is dependent on a combination of factors such as network bandwidth, system memory and client performance.

The subscriber, when using WebSphere MQ Real-time client, contains a message buffer to protect itself from message rate spikes. For these tests this was increased from the default to hold 3000 messages. For information on how to set the subscriber max buffer size see in the WebSphere MQ "Using Java" manual.

In all of the tests it was verified that all publications were delivered to subscribers without any loss of messages. As part of ensuring this all subscribers were started first before the publishing of messages commenced.

Queue depths and buffer sizes were monitored to ensure that the system was running in a stable manner and that there was no backlog of messages to be processed.

Publishers
The JMS Publisher sent non-persistent publications only. A non transacted JMS Session was used. The publisher produced publications at a constant rate, i.e. a fixed number of publications per second.

Subscribers
The JMS Subscribers were non durable and non transacted. Each subscriber had a separate TopicConnection and a TopicSession therefore each subscriber was associated with one physical TCP/IP connection (socket pair). A single topic was used for all tests and so all

subscribers were subscribed to the same topic. This meant that for every message published a copy was received by each subscriber.

Machine Configuration

The programs used to publish and subscribe messages for the tests were run on dedicated Client Machines which were separate from the machine on which the message broker was run.

The number of client machines used varied depending on the test case, for example for a 1 to 1 test (1 publisher and 1 subscriber) 2 client machines were used but for a 1 publisher to 1000 subscriber test the 1000 subscribers were spread over 10 Client Machines. The publisher used a different dedicated machine. The distribution of subscribers over multiple client machines was essential to avoid the client Machines becoming the bottleneck in the test case.

The Message Broker, its dedicated WebSphere MQ queue manager and broker database were all located on a dedicated machine, the Server Machine. The figure below shows the configuration of software components and machines.

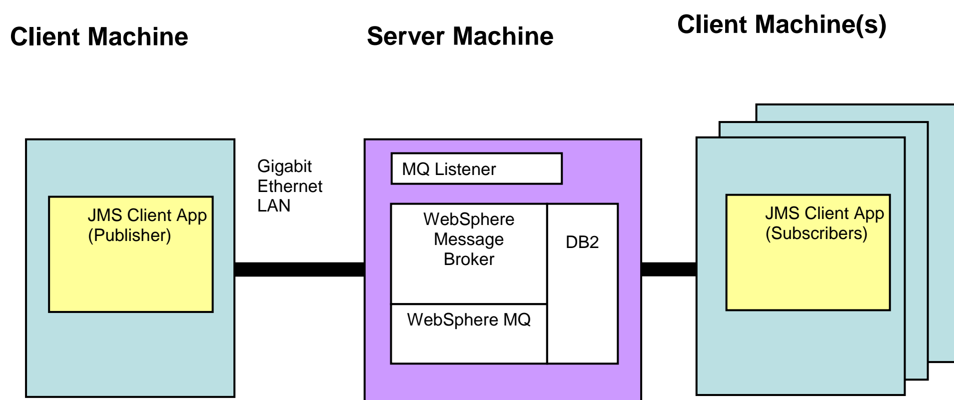

Both the client and server machine were configured with sufficient memory to ensure that no paging took place during the tests.

For those tests using WebSphere MQ as the transport a message flow consisting of an MQInput node wired to a Publication node was used.

For tests using WebSphere MQ Real-time as the transport a message flow consisting of a RealtimeOptimizedFlow node was used.

Reported Message Rates

For tests which did not involve publish subscribe the message rates reported are the number of invocations of the message flow per second.

For tests involving several message flows such as the message aggregation test the rate reported is the number of complete operations or aggregations per second. Fan-out and Fan-in processing is counted as one rather than separately.

For tests using publish subscribe the message rate reported is the total message rate. That is the number processed by all publishers and all subscribers. The total number of messages reported is calculated using the formula (number of subscribers +1) * publication rate.

For a configuration consisting of one publisher and 10 subscribers where the publication rate was 10 messages/second the total message rate is (10 +1) * 10 = 110 messages second.

The message rates quoted are an average taken over the measurement period. This starts once the system initialisation period has completed.

Appendix C - Test Messages

This section describes the input and output messages used for the tests detailed in this report.

The messages which are in this section have been formatted for this report and as such contain white space between tags. When used in measurements all such white space is removed.

Input Message

An input message of the type shown below was used for the non publish/subscribe tests in the report.

The publish/subscribe tests used a 1K JMS Bytes message.

The message shown below is in Generic XML format but it was also represented in a variety of other formats such as MRM XML, CWF and TDS where this was required in the test.

The different message sizes used in testing are achieved by repeating the content of the SaleList tag to give the required size. Larger messages thus result in more tags.
A Perl script ensures that the names and values in the tags are different as the SaleList structure is repeated. This is to stop a limited number of strings being used in very large messages which could lead to over optimistic results.

```
<Parent>
    <First>1</First>
    <SaleList>
        <Invoice>
            <Initial>K</Initial>
            <Initial>A</Initial>
            <Surname>Braithwaite</Surname>
            <Item>
                <Code>00</Code>
                <Code>01</Code>
                <Code>02</Code>
                <Description>Twister</Description>
                <Category>Games</Category>
                <Price>00.30</Price>
                <Quantity>01</Quantity>
            </Item>
            <Item>
                <Code>02</Code>
                <Code>03</Code>
                <Code>01</Code>
                <Description>The Times Newspaper</Description>
                <Category>Books and Media</Category>
                <Price>00.20</Price>
                <Quantity>01</Quantity>
            </Item>
            <Balance>00.50</Balance>
            <Currency>Sterling</Currency>
        </Invoice>
        <Invoice>
            <Initial>T</Initial>
            <Initial>J</Initial>
            <Surname>Dunnwin</Surname>
            <Item>
```

```
                <Code>04</Code>
                <Code>05</Code>
                <Code>01</Code>
                <Description>The Origin of Species</Description>
                <Category>Books and Media</Category>
                <Price>22.34</Price>
                <Quantity>02</Quantity>
            </Item>
            <Item>
                <Code>06</Code>
                <Code>07</Code>
                <Code>01</Code>
                <Description>Microscope</Description>
                <Category>Miscellaneous</Category>
                <Price>36.20</Price>
                <Quantity>01</Quantity>
            </Item>
            <Balance>81.84</Balance>
            <Currency>Euros</Currency>
        </Invoice>
    </SaleList>
    <Last>Test</Last>
</Parent>
```

Output Message

Two message types exist for the output messages dependent on the test case. These are the Compute and Transform messages.

Compute Message

For compute test cases the balance field for each invoice is validated and the currency is converted into sterling. So there is minor modification of the input message.

The message layout is shown below

```
<Parent>
    <First>1</First>
    <SaleList>
        <Invoice>
            <Initial>K</Initial>
            <Initial>A</Initial>
            <Surname>Braithwaite</Surname>
            <Item>
                <Code>00</Code>
                <Code>01</Code>
                <Code>02</Code>
                <Description>Twister</Description>
                <Category>Games</Category>
                <Price>00.30</Price>
                <Quantity>01</Quantity>
            </Item>
            <Item>
                <Code>02</Code>
                <Code>03</Code>
                <Code>01</Code>
                <Description>The Times Newspaper</Description>
                <Category>Books and Media</Category>
                <Price>00.20</Price>
                <Quantity>01</Quantity>
            </Item>
            <Balance>00.50</Balance>
```

```
                <Currency>Sterling</Currency>
            </Invoice>
            <Invoice>
                <Initial>T</Initial>
                <Initial>J</Initial>
                <Surname>Dunnwin</Surname>
                <Item>
                    <Code>04</Code>
                    <Code>05</Code>
                    <Code>01</Code>
                    <Description>The Origin of Species</Description>
                    <Category>Books and Media</Category>
                    <Price>22.34</Price>
                    <Quantity>02</Quantity>
                </Item>
                <Item>
                    <Code>06</Code>
                    <Code>07</Code>
                    <Code>01</Code>
                    <Description>Microscope</Description>
                    <Category>Miscellaneous</Category>
                    <Price>36.20</Price>
                    <Quantity>01</Quantity>
                </Item>
                <Balance>80.88</Balance>
                <Currency>Euros</Currency>
                </Invoice>
            <InvoicesTotal Currency="Sterling">57.116</InvoicesTotal>
        </SaleList>
        <Last>Test</Last>
</Parent>
```


Transform Message
For the transformation test the input message is modified and takes a different layout. For each invoice a statement is created for each customer within a SaleList.

The message layout is shown below.


```
<Parent>
   <SaleList>
        <Statement Type="Monthly" Style="Full">
            <Customer>
                <Initials>KA</Initials>
                <Name>Braithwaite</Name>
                <Balance>00.50</Balance>
            </Customer>
            <Purchases>
                <Article>
                    <Desc>Twister</Desc>
                    <Cost>4.8E-1</Cost>
                    <Qty>01</Qty>
                </Article>
                <Article>
                    <Desc>The Times Newspaper</Desc>
                    <Cost>3.2E-1</Cost>
                    <Qty>01</Qty>
                </Article>
            </Purchases>
            <Amount>8E-1</Amount>
        </Statement>
```

```
        <Statement Type="Monthly" Style="Full">
            <Customer>
                <Initials>TJ</Initials>
                <Name>Dunnwin</Name>
                <Balance>81.84</Balance>
            </Customer>
            <Purchases>
                <Article>
                    <Desc>The Origin of Species</Desc>
                    <Cost>3.5744E+1</Cost>
                    <Qty>02</Qty>
                </Article>
                <Article>
                    <Desc>Microscope</Desc>
                    <Cost>5.792E+1</Cost>
                    <Qty>01</Qty>
                </Article>
            </Purchases>
            <Amount>1.29408E+2</Amount>
        </Statement>
    </SaleList>
</Parent>
```

Appendix D - Use Case Descriptions

This section contains a description of the processing in each of the use cases which are used to characterise the performance of WebSphere Message Broker V6.

Aggregation

The Aggregation use case demonstrates a simple four-way aggregation operation, using the Aggregate Control, Request, and Reply nodes. It contains three message flows to implement a four-way aggregation: FanOut, RequestReplyApp, and FanIn. This is the type of processing that might be used to invoke four different applications to process a travel booking, one to organise each of the flight, hotel, car and money.

FanOut Message Flow
This is the flow that takes the incoming request message, generates four different request messages, sends them out on request/reply, and starts the tracking of the aggregation operation:

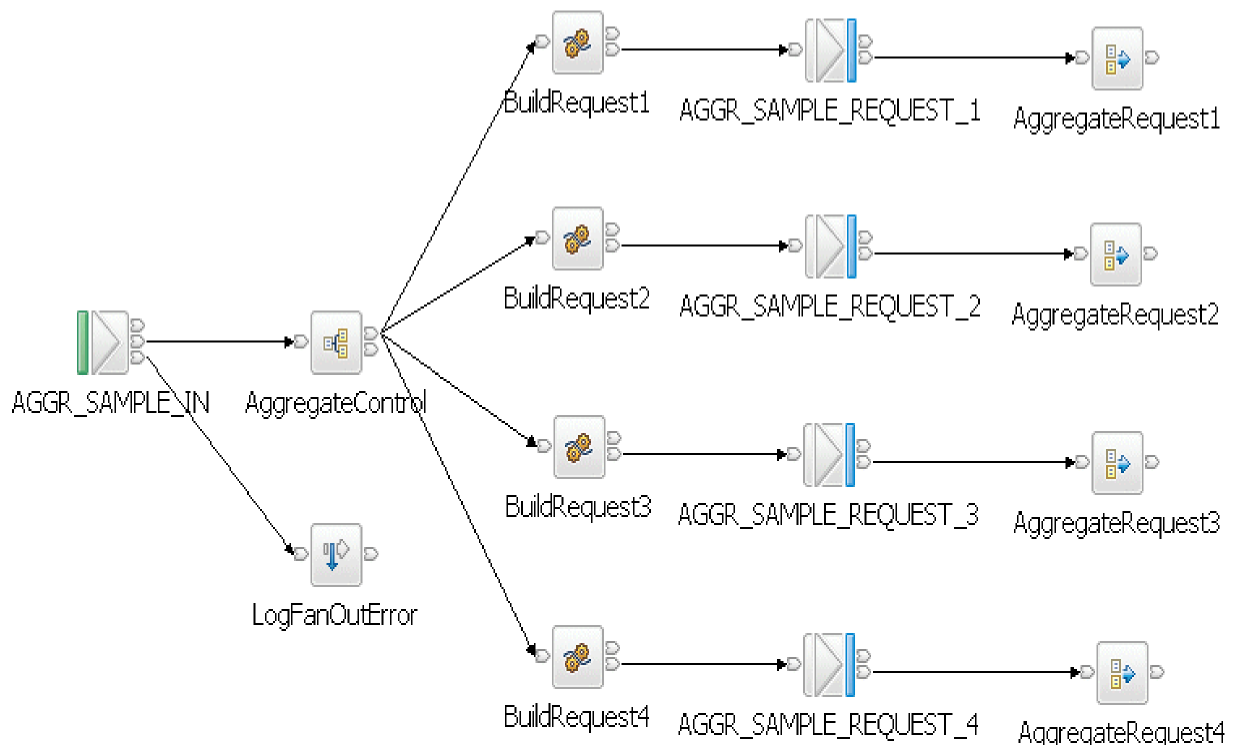

RequestReplyApp Message Flow
This message flow simulates the back-end service applications that would normally process the request messages from the aggregation operation. In a real system, these could be other message flows or existing applications. This message flow reads from the same queue that the MQOutput nodes in the FanOut flow write to, and it outputs to the queue that the input node which the FanIn flow reads from - it provides a messaging bridge between the two flows. The messages are put to their reply-to queue (as set by the MQOutput nodes in the FanOut flow).

FanIn Message Flow
This flow receives all the replies from the RequestReplyApp flow, and aggregates them into a single output message. The output message from the Aggregate Reply node cannot be output directly by an MQOutput node without some processing so a Compute node is added to process the data into a format where it can be written out to a queue.

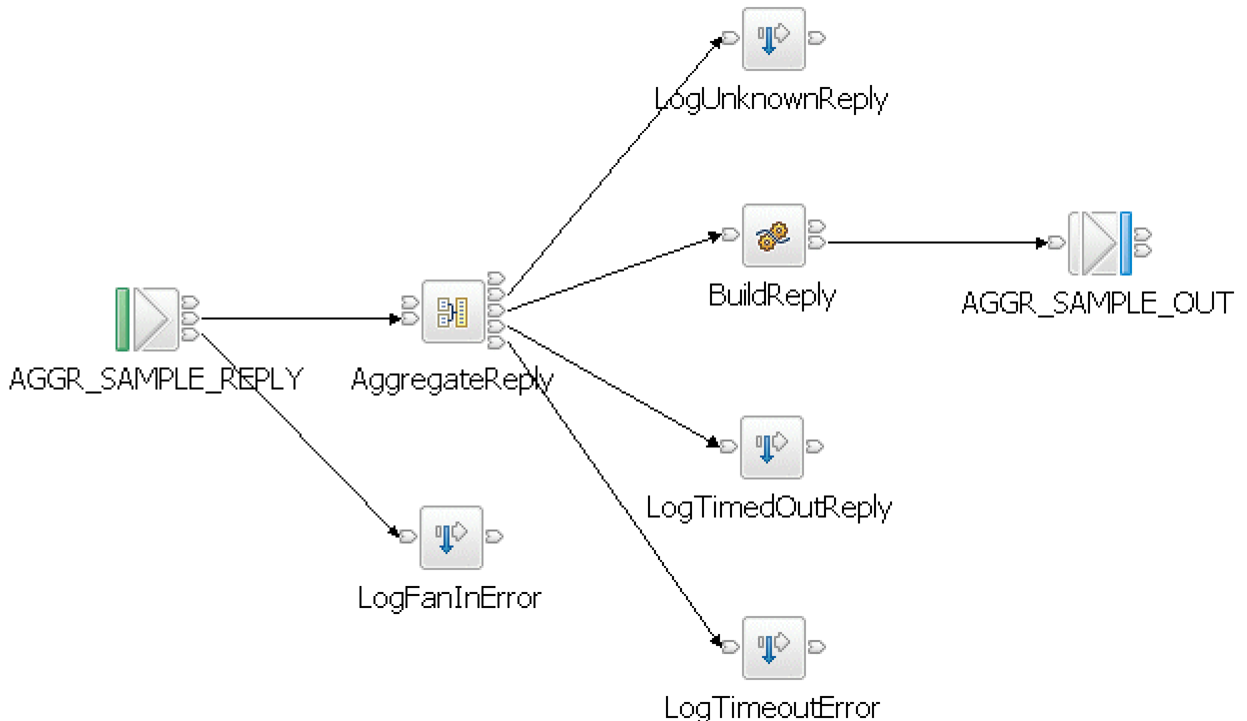

Further information about the Aggregation sample can be found in the Message Brokers section of the Technology samples category which is in the samples gallery of the WebSphere Message Broker development toolkit.

Coordinated Request/Reply

The coordinated request reply sample is based on the scenario of a contemporary and established application communicating through the use of WebSphere MQ messages in a request/reply processing pattern. The contemporary application uses self-defining XML messages and issues a request message. The established application uses Custom Wire Format (CWF) messages. It receives a request message, processes it and delivers a reply message. For the applications to successfully communicate, the message formats must be transformed for both the request and reply messages.

The processing in the sample consists of three message flows and one message set. The message flows are:

Request Message Flow
The request message flow performs the following processing:

- Reads a WebSphere MQ message containing an XML payload.
- Converts the message into the equivalent CWF format.
- Creates a WebSphere MQ message containing the transformed message.

- Saves the original ReplyToQ and ReplyToQMgr details in a separate WebSphere MQ message for subsequent retrieval by the Reply message flow.
- Sets the ReplyToQ and ReplyToQMgr details to be the input of the Reply message flow.
- Sends the message on to the Backend Reply message flow.

The Request message flow consists of the following nodes:

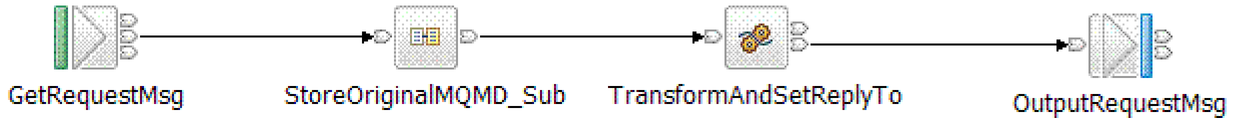

Backend Reply Message Flow

The backend reply message flows performs the following processing:

- Reads a WebSphere MQ message.
- Adds the time the message was modified to the payload of the message.
- Writes a WebSphere MQ message.

The Backend Reply message flow consists of the following nodes:

Reply Message Flow

The reply message flow performs the following processing:

1. Reads a WebSphere MQ message containing a message in CWF format.
2. Converts the message into the equivalent XML format.
3. Obtains the ReplyToQ and ReplyToQ Mgr of the original request message by reading the WebSphere MQ message which was used to store this information in the Request message flow. This is done by using the MQGET node.
4. Creates a WebSphere MQ message containing the transformed message and the retrieved ReplyToQ and ReplyToQMgr values.

The Reply message flow consists of the following nodes:

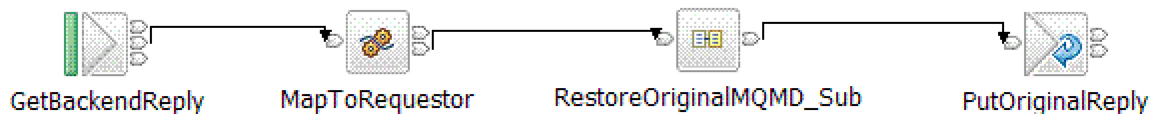

Further information about the Coordinated Request Reply sample can be found in the Message Brokers section of the Application samples category which is in the samples gallery of the WebSphere Message Broker development toolkit.

Data Warehouse

The Data Warehouse sample demonstrates a scenario in which a message flow is used to perform the archiving of data, such as sales data, into a database. The data is stored for later analysis by another message flow or application.

Because the sales data is analyzed at a later date, the storage of the messages has been organized in a way that makes it easy to select records for specified times. The date and time at which the WebSphere MQ message containing the sales record was written are stored as separate column values when the message is inserted into the database. The database table contains four columns:

- The message data - the payload of the WebSphere MQ message stored as a BLOB.
- The date on which the WebSphere MQ message was created.
- The time when the WebSphere MQ message was created.
- A time stamp created by the database to record the time when the record was inserted.

By storing the data in this way it is possible to retrieve records between specific periods of time, say between the hours of 9:00 a.m. to 12:00 p.m. or 12:01 p.m. and 5:00 p.m. which would allow a comparison of morning and afternoon sales to be made.

The data archiving is performed by the WarehouseData message flow. This is described below.

WarehouseData Message Flow

The WarehouseData message flow performs the following processing.

1. Reads a WebSphere MQ message containing an XML payload. The payload contains the data to be archived.
2. Converts a portion of the message tree to a BLOB ready for insertion into the database.
3. Inserts the message BLOB along with the date and time at which the WebSphere MQ message was written into a database.
4. Sends a WebSphere MQ confirmation message to signal successful insertion of the message into the database.

The WarehouseData message flow consists of the following nodes:

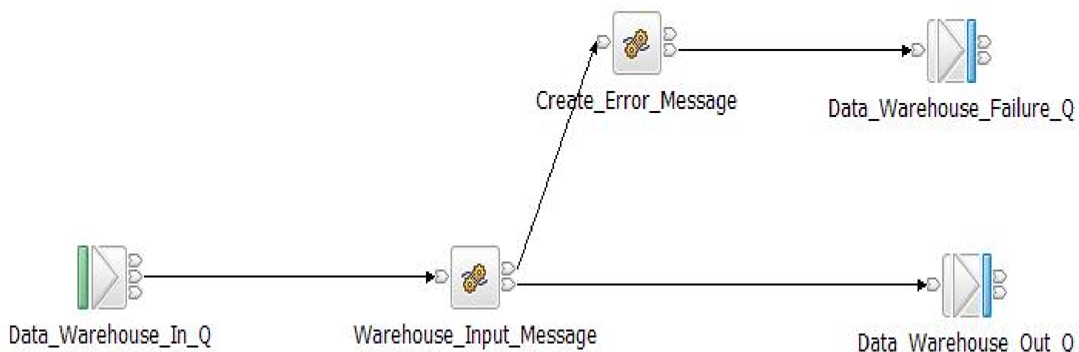

Further information about the Data Warehouse sample can be found in the Message Brokers section of the Application samples category which is in the samples gallery of the WebSphere Message Broker development toolkit.

Large Messaging

The Large Messaging sample is a sample based on the scenario of end-of-day processing of sales data. Messages recording the details of sales through the day are batched together in the store for transmission to the IT center. On receipt at the IT center the batched messages are split back out into their constituent parts for subsequent processing.

This splitting is achieved using a WebSphere Message Broker message flow. Each of the individual messages representing a sale has the same structure.

The input and output messages in this sample are implemented as self-defining XML messages for simplicity. Other message formats could easily be used.

Each input message consists of three parts:

- A header containing a count of the number of repetitions of the repeating SaleList structure that follows.
- The body that contains the repetitions of the repeating SaleList structure.
- The trailer that contains the time the message was processed.

The aim of the processing in this sample is to write each of the instances of the SaleList structure as a separate WebSphere MQ message while minimizing overall memory requirements.

The message flow implements a memory saving technique through the use of a mutable message tree.

The processing in the sample consists of one message flow. The processing it performs is described below.

Large Messaging Message Flow

The large messaging message flow performs the following processing:

1. Reads a WebSphere MQ message containing an XML payload under transactional control.
2. Formats a WebSphere MQ message for each instance of the SaleList structure.
3. Writes the WebSphere MQ messages to the output queue.
4. Produces a WebSphere MQ message to signal completion of the processing when the final element has been processed.

The Large Messaging message flow consists of the following nodes:

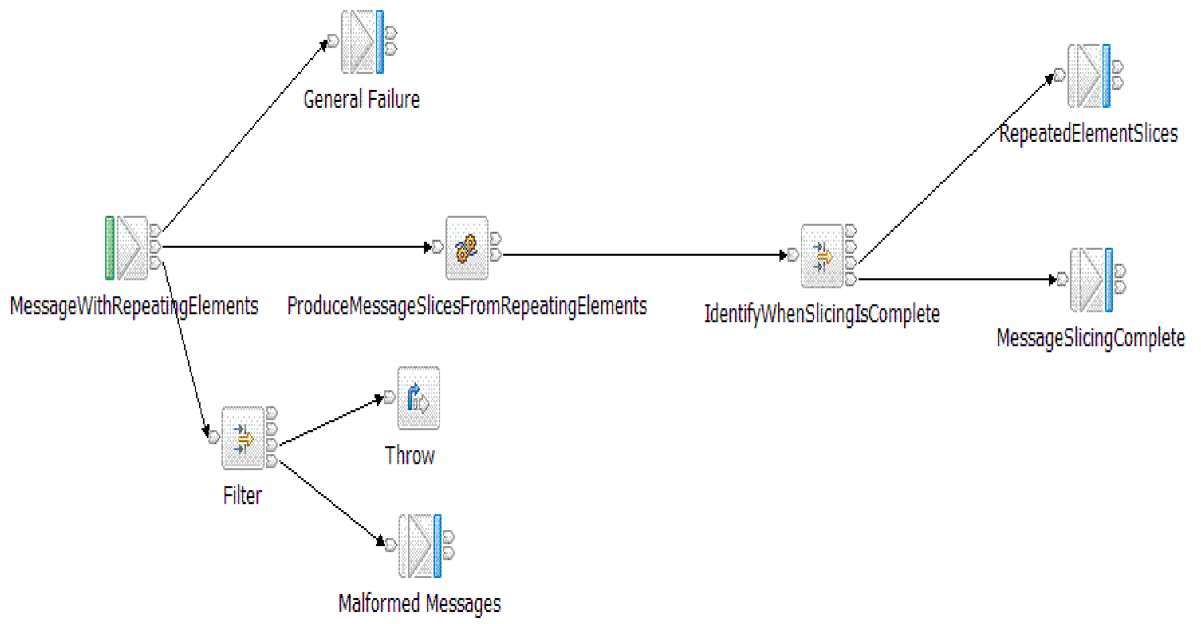

Further information about the Large Messaging sample can be found in the Message Brokers section of the Application samples category which is in the samples gallery of the WebSphere Message Broker development toolkit.

Message Routing

The message routing sample shows how a database table can be used to store routing information which a message flow can then use to route messages to WebSphere MQSeries queues. This uses function which is new in WebSphere Message Broker V6.

The message routing sample shows how to implement a routing table, using shared variables, to route messages in a message flow. Two versions of the message flow were used in these evaluations. One using a database was run as the WebSphere Business Integration Message Broker V5 test case and the second using the routing table implemented using shared variables was run as the WebSphere Message Broker V6 test case.

The processing in the message flows is described below:

Routing_using_database_table Message Flow

The message flow performs the following processing:

1. Reads a WebSphere MQ message containing an XML payload under transactional control.
2. Creates a destination list based on data in a database table and then routes the message to the entries in the destination list. Note this involves a read to the database for every message processed.
3. Produces a WebSphere MQ output message. The destination of the message is specified in the destination list.

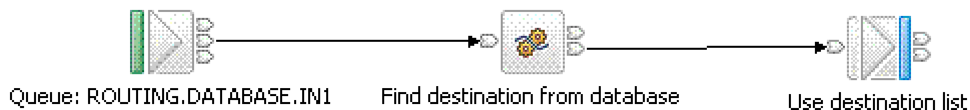

This version of the message flow was used for the WebSphere Business Integration Message Broker V5 measurements.

Routing_using_memory_cache Message Flow

The message flow performs the following processing:

1. Reads a WebSphere MQ message containing an XML payload under transactional control.
2. Creates a destination list based on data which is held in shared variables.
3. Produces a WebSphere MQ output message. The destination of the message is specified in the destination list.

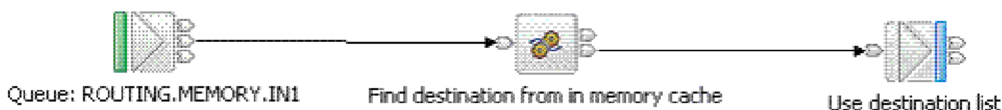

This version of the message flow was used for the WebSphere Message Broker V6 measurements.

Further information about the Message Routing sample can be found in the Message Brokers section of the Application samples category which is in the samples gallery of the WebSphere Message Broker development toolkit.

SWIFT Message Parse

The processing of SWIFT messages is a common requirement for financial institutions. The parsing of the messages is achieved using the MRM parser with a message format of Tagged Delimited String (TDS).

The processing in this test consists of a full parse of a SWFIT MT543 message format.

SWIFT Message Parsing Message Flow

The processing in the SWIFT Message Parse message flow consists of the following:

1. Reads a WebSphere MQ message containing a SWIFT MT543 message in tagged delimited string format.
2. Accesses the last element in the input message.
3. Produces a WebSphere MQ message to signal completion of the processing.

The SWIFT Message Parse processing consists of the following nodes:

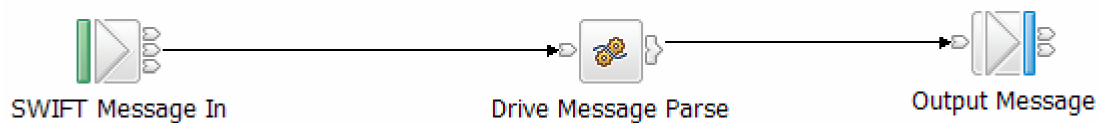

The output message is a minimal WebSphere MQ message.

XML Transformation

The XMLT sample shows how an XML message can be transformed into a different layout using an XMLTransformation node and a XSL stylesheet. This type of processing could be performed in any situation where the layout of a message needs to be changed to suit the requirements of different application.

This technology provides the ability to use an XSL stylesheet in a new way, using it as part of a message flow.

XSL Transformation Message Flow

The processing in the message flow consisted of the following:

1. Reads a WebSphere MQ message containing an XML payload.
2. Invoke the XSL Stylesheet transformation.
3. Write an MQOutput message containing the modified message.

The following figure shows the XSL Transformation message flow:

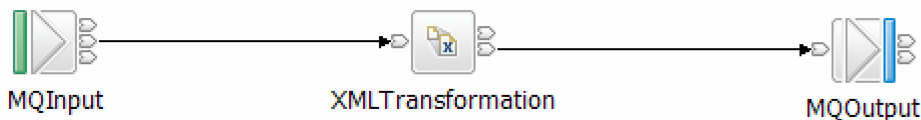

The XSL stylesheet used in the processing was as follows:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:template match="/">
<Parent>
<xsl:for-each select="/Parent/SaleList">
<SaleList>
<xsl:for-each select="Invoice">
<xsl:if test="not(contains(Surname,'Shop'))">
<Statement>
<xsl:attribute name="Type">Monthly</xsl:attribute>
<xsl:attribute name="Style">Full</xsl:attribute>
<Customer>
<Initials>
<xsl:for-each select="Initial">
<xsl:value-of select="."/>
</xsl:for-each>
</Initials>
<Name><xsl:value-of select="Surname"/></Name>
<Balance><xsl:value-of select="Balance"/></Balance>
</Customer>
<Purchases>
<xsl:for-each select="Item">
<Article>
<Desc><xsl:value-of select="Description"/></Desc>
<Cost><xsl:value-of select='format-number((number(Price)*1.6),"####.##")'/></Cost>
<Qty><xsl:value-of select="Quantity"/></Qty>
</Article>
</xsl:for-each>
</Purchases>
<Amount>
<xsl:attribute name="Currency">
<xsl:value-of select="Currency" />
```

```
</xsl:attribute>
<xsl:call-template name="sumSales">
<xsl:with-param name="list" select="Item"/>
</xsl:call-template>
</Amount>
</Statement>
</xsl:if>
</xsl:for-each>
</SaleList>
</xsl:for-each>
</Parent>
</xsl:template>
<xsl:template name="sumSales">
<xsl:param name="list" />
<xsl:param name="result"  select="0"/>
<xsl:choose>
<xsl:when test="$list">
<xsl:call-template name="sumSales">
<xsl:with-param name="list"
select="$list[position()!=1]"/>
<xsl:with-param name="result"
select="$result + number($list[1]/Price)*number($list[1]/Quantity)*1.6"/>
</xsl:call-template>
</xsl:when>
<xsl:otherwise>
<xsl:value-of select='format-number(number($result),"####.##")'/>
</xsl:otherwise>
</xsl:choose>
</xsl:template>
</xsl:stylesheet>
```

The message being transformed was the same format as that described in the Section Input Message.

Further information about this sample be found under the XMLT entry in the Message Brokers section of the Technology samples category which is in the samples gallery of the WebSphere Message Broker development toolkit.

Feedback

This report and other tools that are produced by the performance group are produced in order to help you understand the performance characteristics of WebSphere Message Broker and to assist you with sizing.

It is important that the reports and tools are effective in what they do and it is very useful to have feedback on the content and style of the information which is produced. Your comments, both positive and negative, are therefore welcome.

Your answers to the following questions are particularly interesting:

- What are your most common performance questions?
- Do the reports provide what is needed?
- Is there any other performance information which is required to help you do your job?
- Would you like to see any other aspects of WBIMB performance discussed?

Please supply feedback to Tim Dunn (dunnt@uk.ibm.com) or use the feedback facility on the SupportPac web page where you obtained this report.