

WebSphere MQ Windows 64 bit v7.1

Performance Evaluations

Version 1.2

February 2012

Fred Preston, Craig Stirling , Peter Toghil.

WebSphere MQ Performance

IBM UK Laboratories

Hursley Park

Winchester

Hampshire

SO21 2JN

Property of IBM

Please take Note!

Before using this report, please be sure to read the paragraphs on “disclaimers”, “warranty and liability exclusion”, “errors and omissions”, and the other general information paragraphs in the “Notices” section below.

First Edition, December 2011.

This edition applies to *WebSphere MQ for Windows 64bit v7.1* (and to all subsequent releases and modifications until otherwise indicated in new editions).

© Copyright International Business Machines Corporation 2011. 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

DISCLAIMERS

The performance data contained in this report were measured in a controlled environment. Results obtained in other environments may vary significantly.

You should not assume that the information contained in this report has been submitted to any formal testing by IBM.

Any use of this information and implementation of any of the techniques are the responsibility of the licensed user. Much depends on the ability of the licensed user to evaluate the data and to project the results into their own operational environment.

WARRANTY AND LIABILITY EXCLUSION

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore this statement may not apply to you.

In Germany and Austria, notwithstanding the above exclusions, IBM's warranty and liability are governed only by the respective terms applicable for Germany and Austria in the corresponding IBM program license agreement(s).

ERRORS AND OMISSIONS

The information set forth in this report could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; any such change will be incorporated in new editions of the information. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this information at any time and without notice.

INTENDED AUDIENCE

This report is intended for architects, systems programmers, analysts and programmers

wanting to understand the performance characteristics of *WebSphere MQ for Windows 64 bit v7.1*. The information is not intended as the specification of any programming interface that is provided by WebSphere. It is assumed that the reader is familiar with the concepts and operation of WebSphere MQ v7.1.

LOCAL AVAILABILITY

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. Consult your local IBM representative for information on the products and services currently available in your area.

ALTERNATIVE PRODUCTS AND SERVICES

Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

USE OF INFORMATION PROVIDED BY YOU

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

TRADEMARKS AND SERVICE MARKS

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

- IBM
- WebSphere
- DB2

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

EXPORT REGULATIONS

You agree to comply with all applicable export and import laws and regulations.

Preface

Target audience

The SupportPac was designed for people who:

- Will be designing and implementing solutions using WebSphere MQ for Windows 64 bit.
- Want to understand the performance limits of WebSphere MQ for Windows 64 bit v7.1.
- Want to understand what actions may be taken to tune WebSphere MQ for Windows 64 bit.

The reader should have a general awareness of the Windows operating system and of MQSeries in order to make best use of this SupportPac. Readers should read the section '**How this document is arranged**'—**Page VI** to familiarise themselves with where specific information can be found for later reference.

The contents of this SupportPac

This SupportPac includes:

- Release highlights performance charts.
- Performance measurements with figures and tables to present the performance capabilities of WebSphere MQ local queue manager, client channel, and distributed queuing scenarios.
- Interpretation of the results and implications on designing or sizing of the WebSphere MQ local queue manager, client channel, and distributed queuing configurations.

Feedback on this SupportPac

We welcome constructive feedback on this report.

- Does it provide the sort of information you want?
- Do you feel something important is missing?
- Is there too much technical detail, or not enough?
- Could the material be presented in a more useful manner?

Specific queries about performance problems on your WebSphere MQ system should be directed to your local IBM Representative or Support Centre.

Introduction

The three scenarios used in this report to generate the performance data are:

- Local queue manager scenario.
- Client channel scenario.
- Distributed queuing scenario.

Unless otherwise specified, the standard message sized used for all the measurements in this report is 2KB (2,048 bytes).

A xSeries 3850 box containing two dual core 3.33GHz Intel Xeon CPUs and 4GB of RAM was used as the Device under test.

A xSeries 3850 box containing two dual core 3.33GHz Intel Xeon CPUs and 4GB of RAM was used as the Driver.

How this document is arranged

Pages: 1-17

The first section contains the performance *headlines* for each of the three scenarios, with MQI applications connected to:

- A local queue manager.
- A remote queue manager over MQI-client channels.
- A local queue manager, driving throughput between the local and remote queue manager over server channel pairs.

The headline tests show:

- The maximum message throughput achieved with an increasing number of MQI applications.
- The maximum number of MQI-clients connected to a queue manager.
- The maximum number of server channel pairs between two queue managers, for a fixed think time between messages until the response time exceeds one second.

Large Messages

Pages: 20-41

The second section contains performance measurements for *large messages*. This includes *MQI response times* of 50B to 2MB messages. It also includes *20KB, 200KB and 2MB* messages using the same scenarios as for the 2KB messages”.

Application Bindings

Page: 42-47

The third section contains performance measurements for *'trusted, shared, and isolated'* server applications, using the same three scenarios as for the 2KB messages.

Performance and Capacity Measurements

Pages:49

Tuning Recommendations

Pages: 51

Windows 64-bit tuning recommendations

Measurement Environment

Pages: 55 56

A summary of the way in which the workload is used in each test scenario is given in the “ headlines” section. This includes a more detailed description of the workload, hardware and software specifications.

Windows : Microsoft Windows Server 2003 Enterprise x64 Edition

MQSeries: Version 6.0.2.11, Version 7.0, Version 7.0.1.6, Version 7.1

Compiler: Microsoft visual C compiler

Glossary

Page: 56

A short glossary of the terms used in the tables throughout this document.

CONTENTS

1	Overview	1
2	Performance Headlines	2
2.1	Local Queue Manager Test Scenario	2
2.1.1	Non-persistent Messages – Local Queue Manager	3
2.1.2	Non-persistent Messages – Non-trusted – Local Queue Manager	4
2.1.3	Persistent Messages – Local Queue Manager	5
2.2	Client Channels Test Scenario	6
2.2.1	Non-persistent Messages – Client Channels	7
2.2.2	Non-persistent Messages – Non-Trusted Client Channels	8
2.2.3	Persistent Messages – Client Channels	9
2.2.4	Client Channels	10
2.2.5	SSL	11
2.3	Distributed Queuing Test Scenario	13
2.3.1	Non-persistent Messages – Server Channels	14
2.3.2	Non-Persistent non-Trusted – Server Channels	15
2.3.3	Persistent Messages – Server Channels	16
2.3.4	Server Channels	17
2.3.5	SSL	18
3	Large Messages	20
3.1	MQI Response Times: 50B to 100MB – Local Queue Manager	20
3.1.1	50B to 32KB	20
3.1.2	32KB to 2MB	22
3.1.3	2MB to 100MB	23
3.2	20KB Messages	24
3.2.1	Local Queue Manager	24
3.2.2	Client Channel	26
3.2.3	Distributed Queuing	28
3.3	200K Messages	30
3.3.1	Local Queue Manager	30
3.3.2	Client Channel	32
3.3.3	Distributed Queuing	34
3.4	2MB Messages	36
3.4.1	Local Queue Manager	36
3.4.2	Client Channel	38
3.4.3	Distributed Queuing	40
4	Application Bindings	42
4.1	Local Queue Manager	42
4.1.1	Non-persistent Messages	42
4.1.2	Persistent Messages	43
4.2	Client Channels	44
4.2.1	Non-persistent Messages	44
4.2.2	Persistent Messages	45
4.3	Distributed Queuing	46
4.3.1	Non-persistent Messages	46
4.3.2	Persistent Messages	47
5	Short & Long Sessions	48
6	Performance and Capacity Limits	49
6.1	Client channels – capacity measurements	49
6.2	Distributed queuing – capacity measurements	49
7	Tuning Recommendations	51
7.1	Tuning the Queue Manager	51
7.1.1	Queue Disk, Log Disk, and Message Persistence	51
7.1.2	Log Buffer Size, Log File Size, and Number of Log Extents	51
7.1.3	Channels: Process or Thread, Standard or Fastpath?	53
7.2	Applications: Design and Configuration	53
7.2.1	Standard (Shared or Isolated) or Fastpath?	53
7.2.2	Parallelism, Batching, and Triggering	53
7.3	Virtual Memory, Real Memory, & Paging	54
7.3.1	BufferLength	54
7.3.2	MQIBINDTYPE	54
8	Measurement Environment	55

8.1	Workload description	55
8.1.1	MQI response time tool.....	55
8.1.2	Test scenario workload.....	55
8.2	Hardware	56
8.3	Software	56
9	Glossary	57

TABLES

Table 1 – Performance headline, non-persistent messages and local queue manager	3
Table 2 – Performance headline, non-persistent, non-trusted messages and local queue manager	4
Table 3 – Performance headline, persistent messages and local queue manager	5
Table 4 – Performance headline, non-persistent messages and client channels	7
Table 5 – Performance headline, non-persistent messages and client channels	8
Table 6 – Performance headline, persistent messages and client channels.....	9
Table 7 – 1 round trip per driving application per second, client channels	11
Table 8 – Ordered relative SSL Client cipher performance	12
Table 9 – Performance headline, non-persistent messages and server channels	14
Table 10 – Performance headline, non-persistent, non trusted messages and server channels.....	15
Table 11 – Performance headline, persistent messages and server channels.....	16
Table 12 – 1 round trip per driving application per second, client channels	18
Table 13 – Ordered relative SSL DQ cipher performance	19
Table 14 – 20KB non-persistent messages, local queue manager	24
Table 15 – 20KB persistent messages, local queue manager	25
Table 16 – 20KB non-persistent messages, client channels	26
Table 17 – 20KB persistent messages, client channels.....	27
Table 18 – 20KB non-persistent messages, client channels	28
Table 19 – 20KB persistent messages, client channels.....	29
Table 20 – 200KB non-persistent messages, local queue manager	30
Table 21 – 200KB persistent messages, local queue manager	31
Table 22 – 200KB non-persistent messages, client channels	32
Table 23 – 200KB persistent messages, client channels.....	33
Table 24 – 200KB non-persistent messages, distributed queuing	34
Table 25 – 200KB persistent messages, distributed queuing	35
Table 26 – 2MB non-persistent messages, local queue manager	36
Table 27 – 2MB persistent messages, local queue manager.....	37
Table 28 – 2MB non-persistent messages, client channels.....	38
Table 29 – 2MB persistent messages, client channels	39
Table 30 – 2MB non-persistent messages, distributed queuing	40
Table 31 – 2MB persistent messages, distributed queuing	41
Table 32 – Application binding, non-persistent messages, local queue manager.....	42
Table 33 – Application binding, persistent messages, local queue manager	43
Table 34 – Application binding, non-persistent messages, client channels	44
Table 35 – Application binding, persistent messages, client channels	45
Table 36 – Application binding, non-persistent messages, distributed queuing.....	46
Table 37 – Application binding, persistent messages, distributed queuing	47
Table 38 – Capacity measurements, client channels	49
Table 39 – Capacity measurements, server channels	50

FIGURES

Figure 1 – Connections into a local queue manager 2

Figure 2 – Performance headline, non-persistent messages and local queue manager. 3

Figure 3 Performance headline, non-persistent, non-trusted messages and local queue manager. 4

Figure 4 – Performance headline, persistent messages and local queue manager 5

Figure 5 – MQI-client channels into a remote queue manager 6

Figure 6 – Performance headline, non-persistent messages and client channels 7

Figure 7 – Performance headline, non-persistent messages with non-trusted client channels 8

Figure 8 – Performance headline, persistent messages and client channels 9

Figure 9 – 1 round trip per driving application per second, client channels and non-persistent messages
..... 10

Figure 10 – 1 round trip per driving application per second, client channels, persistent messages 10

Figure 11 – Client non-persistent message rates with various SSL ciphers 11

Figure 12 – Server channels between two queue managers 13

Figure 13 – Performance headline, non-persistent messages and server channels 14

Figure 14 – Performance headline, non-persistent, not trusted messages and server channels 15

Figure 15 – Performance headline, persistent messages and server channels 16

Figure 16 – 1 round trip per driving application per second, server channel, non-persistent messages . 17

Figure 17 – 1 round trip per driving application per second, server channel, persistent messages 17

Figure 18 – DQ non-persistent message rates with various SSL ciphers 18

Figure 19 –The effect of non-persistent message size on MQI response time (50B - 32KB) 20

Figure 20 –The effect of persistent message size on MQI response time (50B - 32KB)..... 21

Figure 21 –The effect of non-persistent message size on MQI response time (32KB – 2MB) 22

Figure 22 –The effect of persistent message size on MQI response time (32KB – 2MB) 22

Figure 23 –The effect of non-persistent message size on MQI response time (2MB – 100MB) 23

Figure 24 –The effect of persistent message size on MQI response time (2MB – 100MB)..... 23

Figure 25 – 20KB non-persistent messages, local queue manager 24

Figure 26 – 20KB persistent messages, local queue manager 25

Figure 27 – 20KB non-persistent messages, client channels 26

Figure 28 – 20KB persistent messages, client channels 27

Figure 29 – 20KB non-persistent messages, distributed queuing 28

Figure 30 – 20KB persistent messages, distributed queuing 29

Figure 31 – 200KB non-persistent messages, local queue manager 30

Figure 32 – 200KB persistent messages, local queue manager 31

Figure 33 – 200KB non-persistent messages, client channels 32

Figure 34 – 200KB persistent messages, client channels 33

Figure 35 – 200KB non-persistent messages, distributed queuing 34

Figure 36 – 200KB persistent messages, distributed queuing 35

Figure 37 – 2MB non-persistent messages, local queue manager 36

Figure 38 – 2MB persistent messages, local queue manager 37

Figure 39 – 2MB non-persistent messages, client channels 38

Figure 40 – 2MB persistent messages, client channels 39

Figure 41 – 2MB non-persistent messages, distributed queuing 40

Figure 42 - 2MB persistent messages, distributed queuing 41

Figure 43 – Application binding, non-persistent messages, local queue manager 42

Figure 44 – Application binding, persistent messages, local queue manager 43

Figure 45 – Application binding, non-persistent messages, client channels 44

Figure 46 – Application binding, persistent messages, client channels 45

Figure 47 – Application binding, non-persistent messages, distributed queuing 46

Figure 48 – Application binding, persistent messages, distributed queuing 47

1 Overview

WebSphere MQ v7.1 on Windows64 has improved performance in almost every area. For 2KB messages, almost every test shows improvements over earlier versions of WMQ. Most notable improvements have been in persistent messaging.

Using the area under the graph to show performance differences, the following generalisations (to the nearest 5%) can be made about 2KB message sizes when comparing WMQv7.1 with previous releases;-

For client, distributed queuing and local **non-persistent** messages WMQv7.1 is;-

- 5% faster than v6.0.2.11
- 15% faster than v7.0
- 10% faster than v7.0.1.6

For client, distributed queuing and local **persistent** messages WMQv7.1 is;-

- 90% faster than v6.0.2.11
- 100% faster than v7.0
- 110% faster than v7.0.1.6

2 Performance Headlines

The measurements for the local queue manager scenario are for processing messages with no *think-time*. For the client channel scenario and distributed queuing scenario, there are also measurements for *rated* messaging.

No '*think-time*' is when the driving applications do not wait after getting a reply message before submitting subsequent request messages—this is also referred to as '*tight-loop*'.

The rated messaging tests used one round trip per driving application per second. In the client channel test scenarios, each driving application using a dedicated MQI-client channel, in the distributed queuing test scenarios, one or more applications submit messages over a fixed number of server channels.

All tests stop automatically after the response time exceeds 1 second.

2.1 Local Queue Manager Test Scenario




Figure 1 – Connections into a local queue manager

- 1) The Requester application puts a message to the common input queue on the local queue manager, and holds on to the message identifier returned in the message descriptor. The Requester application then waits indefinitely for a reply to arrive on the common reply queue.
- 2) The Responder application gets messages from the common input queue and places a reply to the common reply queue. The queue manager copies over the message identifier from the request message to the correlation identifier of the reply message.
- 3) The Requester application gets a reply from the common reply queue using the message identifier held from when the request message was put to the common input queue, as the correlation identifier in the message descriptor.

Non-persistent and persistent messages were used in the local queue manager tests, with a message size of 2KB. The effect of message throughput with larger messages sizes is investigated in the "*Large Messages*" section.

Application Bindings of the Responder program are 'Shared' and the Requester program is normally 'Trusted' except in the 'non-trusted' scenario where both programs use 'Shared' bindings.

2.1.1 Non-persistent Messages – Local Queue Manager

Figure 2, Figure 3 and Figure 4 show the non-persistent, non-persistent non-trusted and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario (see Figure 1 on the previous page) for different production levels of WebSphere MQ (versions 7.1, 7.0.1.6, 7.0 and 6.0.2.11).




Figure 2 – Performance headline, non-persistent messages and local queue manager.

Figure 2 and Table 1 show that the peak throughput of non-persistent messages has increased by 5% when comparing version 7.1 to 7.0.1.6 and 18% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 2KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	5	16052	0.00034	99%
WMQv7.0	4	17590	0.00027	99%
WMQv7.0.1.6	4	18078	0.00026	99%
WMQv7.1	4	19029	0.00027	99%

Table 1 – Performance headline, non-persistent messages and local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.1.2 Non-persistent Messages – Non-trusted – Local Queue Manager




Figure 3 Performance headline, non-persistent, non-trusted messages and local queue manager.

Figure 3 and Table 2 show that the peak throughput of non-persistent, non-trusted messages (shared bindings - MQIBINDTYPE=NORMAL) has increased by 5% when comparing version 7.1 to 7.0.1.6 and 12% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 2KB Non-Persistent Messages with Non-Trusted Bindings	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	7	12718	0.00067	99%
WMQv7.0	4	13638	0.00035	99%
WMQv7.0.1.6	4	13657	0.00043	99%
WMQv7.1	4	14299	0.00029	99%

Table 2 – Performance headline, non-persistent, non-trusted messages and local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.1.3 Persistent Messages – Local Queue Manager




Figure 4 – Performance headline, persistent messages and local queue manager

Figure 4 and Table 3 shows that the peak throughput of persistent messages has increased by 61% when comparing version 7.1 to 7.0.1.6 and 54% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 2KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	88	4015	0.027	81%
WMQv7.0	88	4052	0.026	78%
WMQv7.0.1.6	84	3835	0.026	76%
WMQv7.1	60	6174	0.013	95%

Table 3 – Performance headline, persistent messages and local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.2 Client Channels Test Scenario

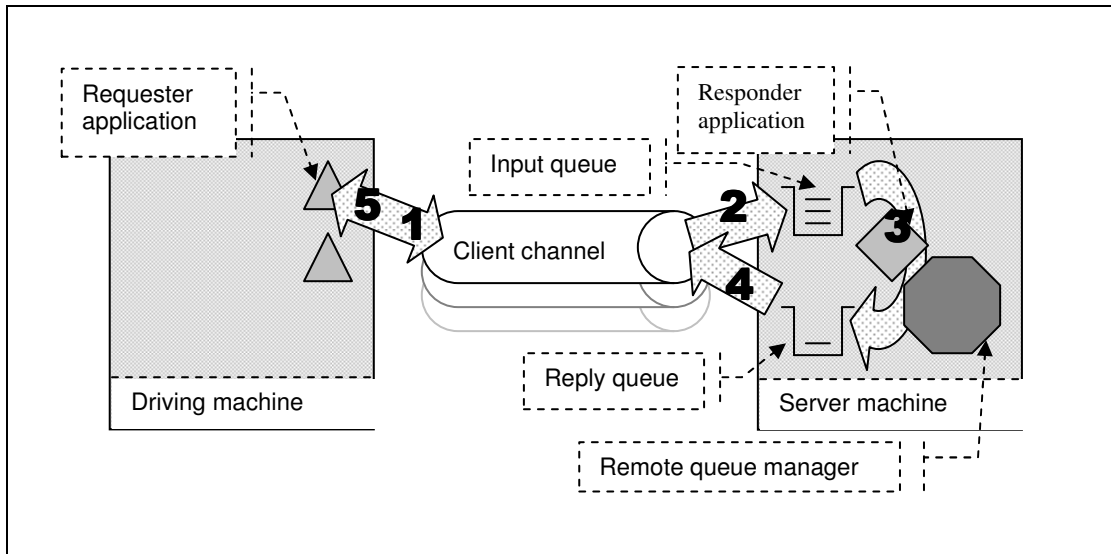


Figure 5 – MQI-client channels into a remote queue manager

- 1, 2) The Requester application puts a request message (over a client channel), to the common input queue, and holds on to the message identifier returned in the message descriptor. The Requester application then waits indefinitely for a reply to arrive on the common reply queue.
- 3) The Responder application gets messages from the common input queue and places a reply to the common reply queue. The queue manager copies over the message identifier from the request message to the correlation identifier of the reply message.
- 4, 5) The Requester application gets the reply message (over the client channel), from the common reply queue. The Requester application uses the message identifier held from when the request message was put to the common input queue, as the correlation identifier in the message descriptor.

Non-persistent and persistent messages were used in the client channel tests, with a message size of 2KB. The effect of message throughput with larger messages sizes is investigated in the “*Large Messages*” section.

Application Bindings of the Responder program are ‘Shared’ and the Client Channel is set to ‘MQIBindType = FASTPATH’ except in the ‘non-trusted’ scenario where ‘MQIBindType = STANDARD’ is used.

Version 7 onwards will multiplex multiple clients from the same process over one TCP socket. We have standardized all client measurements to use SHARECNV(1) since we have various tests that have between 1 and 100 clients per process and we are interested in results when all the clients come from different computers. Further information in section 7.1.4

2.2.1 Non-persistent Messages – Client Channels

Figure 6, Figure 7 and Figure 8 show the non-persistent, non-persistent non-trusted and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario (see Figure 5 on the previous page) for different production levels of WebSphere MQ (versions 7.1, 7.0.1.6, 7.0 and 6.0.2.11).




Figure 6 – Performance headline, non-persistent messages and client channels

Figure 6 and Table 4 show that the peak throughput of non-persistent messages has increased by 20% when comparing version 7.1 to 7.0.1.6 but has decreased by 5% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 2KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	9	8172	0.0013	96%
WMQv7.0	9	6727	0.0016	93%
WMQv7.0.1.6	11	6489	0.0023	94%
WMQv7.1	11	7752	0.0018	99%

Table 4 – Performance headline, non-persistent messages and client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.2.2 Non-persistent Messages – Non-Trusted Client Channels




Figure 7 – Performance headline, non-persistent messages with non-trusted client channels

Figure 7 and Table 5 shows that the peak throughput of non-persistent, non-trusted messages (shared bindings - MQIBINDTYPE=NORMAL) has increased by 12% when comparing version 7.1 to 7.0.1.6 but has decreased by 8% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 2KB Non-Persistent Messages with Non-Trusted Bindings	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	9	7162	0.0013	96%
WMQv7.0	10	6059	0.0019	95%
WMQv7.0.1.6	10	5835	0.0023	95%
WMQv7.1	15	6555	0.0028	99%

Table 5 – Performance headline, non-persistent messages and client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.2.3 Persistent Messages – Client Channels




Figure 8 – Performance headline, persistent messages and client channels

Figure 8 and Table 6 shows that the peak throughput of persistent messages has increased by 107% when comparing version 7.1 to 7.0.1.6 and by 48% when comparing version 7.1 to 6.0.2.11

Test Name: Client Channels - 2KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	112	2229	0.065	58%
WMQv7.0	96	1665	0.071	47%
WMQv7.0.1.6	120	1595	0.1	49%
WMQv7.1	32	3298	0.011	94%

Table 6 – Performance headline, persistent messages and client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.2.4 Client Channels

For the following client channel measurements, the messaging rate used is 1 round trip per second per MQI-client channel, i.e. a request message outbound over the client channel and a reply message inbound over the channel per second.




Figure 9 – 1 round trip per driving application per second, client channels and non-persistent messages

Note: Messaging in these tests is 1 round trip per driving application per second.




Figure 10 – 1 round trip per driving application per second, client channels, persistent messages

Figure 9, Figure 10 and Table 7 show that the peak throughput of non-persistent messages has increased by 9% when comparing version 7.1 to 7.0.1.6 but has decreased by 17% when comparing version 7.1 to 6.0.2.11. It also shows that the throughput of persistent messages has increased by 43% when comparing version 7.1 to 7.0.1.6 and by 18% when comparing version 7.1 to 6.0.2.11

Test Name: Client Channels - R3600 Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	3800	3549	0.31	5%
WMQv7.0	3000	2750	0.31	78%
WMQv7.0.1.6	2950	2694	0.42	80%
WMQv7.1	3050	2947	0.43	89%

Test Name: Client Channels - R3600 Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	1580	1469	0.49	79%
WMQv7.0	1360	1306	0.66	75%
WMQv7.0.1.6	1160	1125	0.67	68%
WMQv7.1	1780	1684	0.49	98%

Table 7 – 1 round trip per driving application per second, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.2.5 SSL

The following diagram shows how throughput varies depending on the cipher selected. The top line (using the standard 2KB message CLNP test) is not encrypted. The other lines show a selection of the available ciphers.

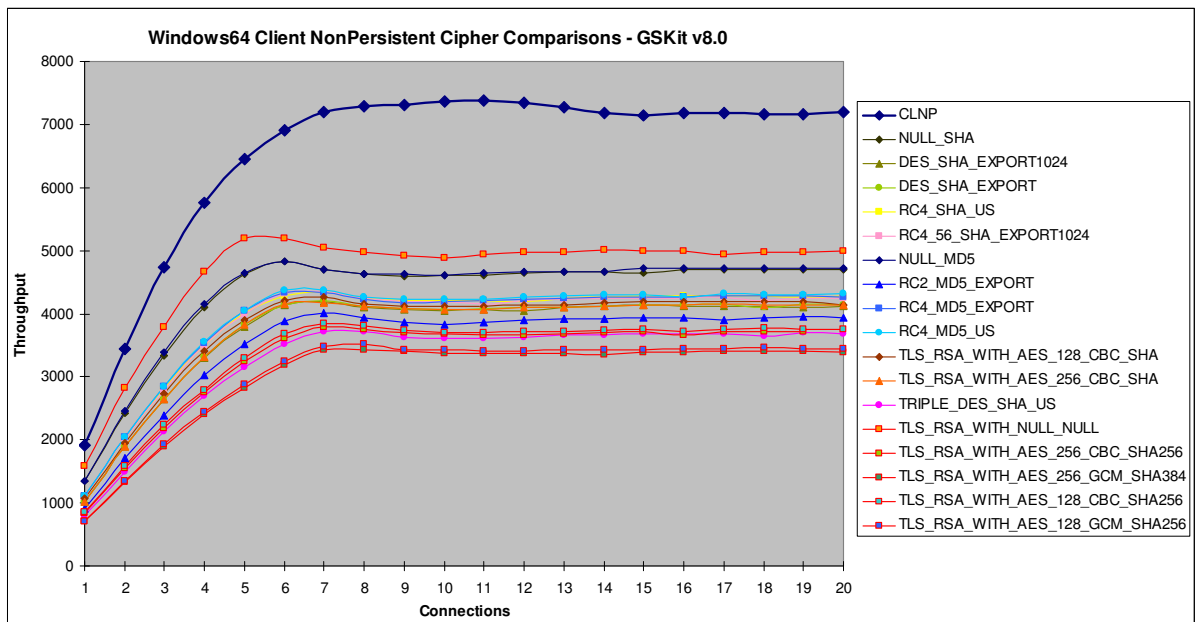


Figure 11 – Client non-persistent message rates with various SSL ciphers

The area under the curve values for each cipher are used to compile an ordered table of these ciphers and can be used as a guide to expected performance degradations that can be expected when these ciphers are used. The unencrypted value is used to rate the other values as a percentage. Thus for the SSL cipher *TLS_RSA_WITH_NULL_NULL*, the expected message rate would be approximately 71% of the unencrypted rate.

Windows64 - SSL Cipher Comparison	CLNP
Unencrypted	100%
TLS_RSA_WITH_NULL_NULL	71%
NULL_MD5	66%
NULL_SHA	66%

Windows64 - SSL Cipher Comparison	CLNP
RC4_MD5_US	60%
RC4_MD5_EXPORT	59%
RC4_SHA_US	59%
RC4_56_SHA_EXPORT1024	59%
TLS_RSA_WITH_AES_128_CBC_SHA	58%
DES_SHA_EXPORT	57%
TLS_RSA_WITH_AES_256_CBC_SHA	57%
DES_SHA_EXPORT1024	57%
RC2_MD5_EXPORT	54%
TLS_RSA_WITH_AES_128_CBC_SHA256	51%
TLS_RSA_WITH_AES_256_CBC_SHA256	51%
TRIPLE_DES_SHA_US	50%
TLS_RSA_WITH_AES_128_GCM_SHA256	47%
TLS_RSA_WITH_AES_256_GCM_SHA384	46%

Table 8 – Ordered relative SSL Client cipher performance

2.3 Distributed Queuing Test Scenario

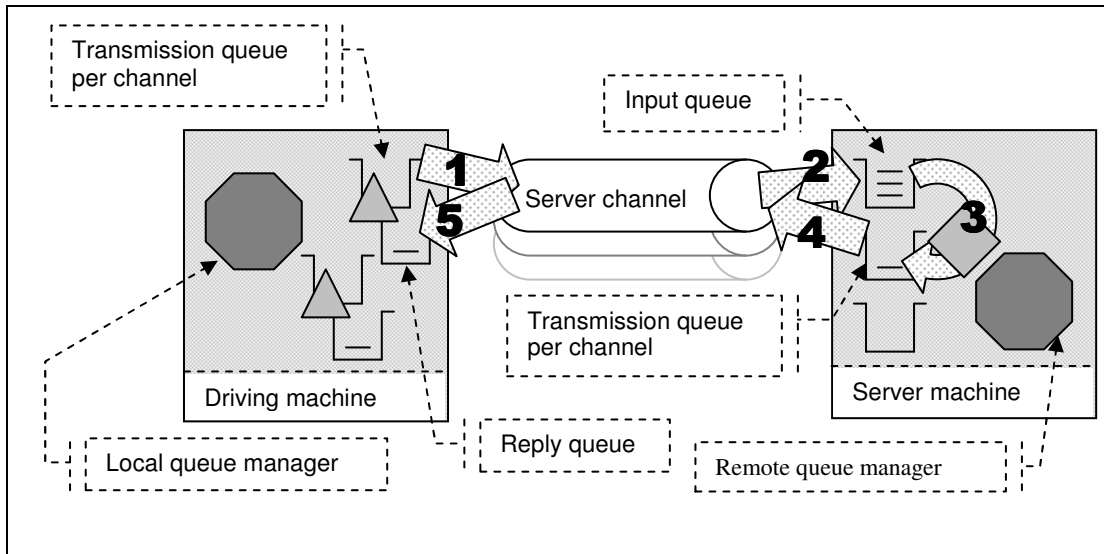


Figure 12 – Server channels between two queue managers

- 1) The Requester application puts a message to a local definition of a remote queue located on the server machine, and holds on to the message identifier returned in the message descriptor. The Requester application then waits indefinitely for a reply to arrive on a local queue.
- 2) The message channel agent takes messages off the channel and places them on the common input queue on the server machine.
- 3) The Responder application gets messages from the common input queue, and places a reply to the queue name extracted from the messages descriptor (the name of a local definition of a remote queue located on the driving machine). The queue manager copies over the message identifier from the request message to the correlation identifier of the reply message.
- 4) The message channel agent takes messages off the transmission queue and sends them over the channel to the driving machine.
- 5) The Requester application gets a reply from a local queue. The Requester application uses the message identifier held from when the request message was put to the local definition of the remote queue, as the correlation identifier in the message descriptor

Non-persistent and persistent messages were used in the distributed queuing tests, with a message size of 2KB. The effect of message throughput with larger messages sizes is investigated in the “*Large Messages*” section.

Application Bindings of the Responder program are ‘Shared’ , the Requester program is normally ‘Trusted’ , and the channels specified as ‘MQIBindType = FASTPATH’ except in the ‘non-trusted’ scenario where both programs use ‘shared’ bindings and the channels are specified as ‘MQIBindType = STANDARD’.

2.3.1 Non-persistent Messages – Server Channels

Figure 13, Figure 14 and Figure 15 show the non-persistent, non-persistent non-trusted and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario (see Figure 12 on the previous page) and WebSphere MQ (versions 7.1, 7.0.1.6, 7.0 and 6.0.2.11).

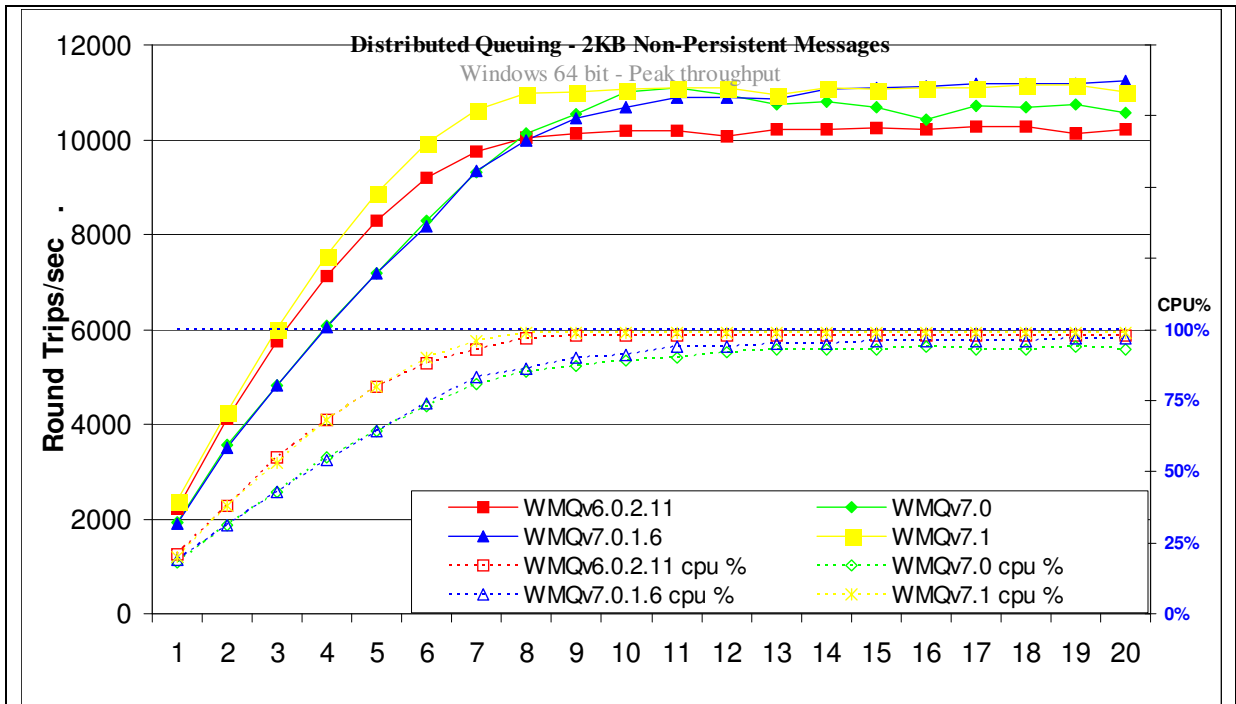


Figure 13 – Performance headline, non-persistent messages and server channels

Figure 13 and Table 9 shows that the peak throughput of persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 and has increased by 9% when comparing version 7.1 to 6.0.2.11

Test Name: Distributed Queuing - 2KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	18	10275	0.0018	98%
WMQv7.0	11	11104	0.0014	90%
WMQv7.0.1.6	20	11228	0.0025	97%
WMQv7.1	18	11148	0.0015	99%

Table 9 – Performance headline, non-persistent messages and server channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.3.2 Non-Persistent non-Trusted – Server Channels




Figure 14 – Performance headline, non-persistent, not trusted messages and server channels

Figure 14 and Table 10 shows that the peak throughput Table 10 of non-persistent, non-trusted messages has decreased by 8% when comparing version 7.1 to 7.0.1.6 but has increased by 3% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 2KB Non-Persistent Messages with Non-Trusted Bindings	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	30	8405	0.0045	94%
WMQv7.0	29	8742	0.0043	95%
WMQv7.0.1.6	30	9341	0.0041	97%
WMQv7.1	26	8656	0.0037	98%

Table 10 – Performance headline, non-persistent, non trusted messages and server channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.3.3 Persistent Messages – Server Channels




Figure 15 – Performance headline, persistent messages and server channels

Figure 15 and Table 11 shows that the peak throughput of persistent messages has increased by 108% when comparing version 7.1 to 7.0.1.6 and by 103% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 2KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	300	2424	0.15	31%
WMQv7.0	300	2596	0.13	30%
WMQv7.0.1.6	300	2489	0.18	37%
WMQv7.1	160	5184	0.041	86%

Table 11 – Performance headline, persistent messages and server channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.3.4 Server Channels

For the following distributed queuing measurements, the messaging rate used is 1 round trip per driving application per second, i.e. a request message outbound over the sender channel, and a reply message inbound over the receiver channel per second. Note that there are a fixed number of 4 server channel pairs for the non-persistent messaging tests, and 2 pairs for the persistent message tests.




Figure 16 – 1 round trip per driving application per second, server channel, non-persistent messages

Note: Messaging in these tests is 1 round trip per driving application per second.




Figure 17 – 1 round trip per driving application per second, server channel, persistent messages

Figure 16, Figure 17 and Table 12 shows that the peak throughput of non-persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 or version 7.1 to 6.0.2.11 and for persistent messages has increased by 27% when comparing version 7.1 to 7.0.1.6 and by 15% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - R3600 Non-	Apps	Round	Response time	CPU
---	------	-------	---------------	-----

Persistent Messages		Trips/Sec	(s)	
WMQv6.0.2.11	5000	4995	0.019	45%
WMQv7.0	5000	4997	0.015	40%
WMQv7.0.1.6	5250	5000	0.013	44%
WMQv7.1	5000	4990	0.0028	42%

Test Name: Distributed Queuing - R3600 Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	3050	2802	0.96	40%
WMQv7.0	2700	2496	0.94	32%
WMQv7.0.1.6	2450	2286	0.94	37%
WMQv7.1	2650	2649	0.034	44%

Table 12 – 1 round trip per driving application per second, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

2.3.5 SSL

The following diagram shows how throughput varies depending on the cipher selected. The top line (using the standard 2KB message DQNP test) is not encrypted. The other lines show a selection of the available ciphers.

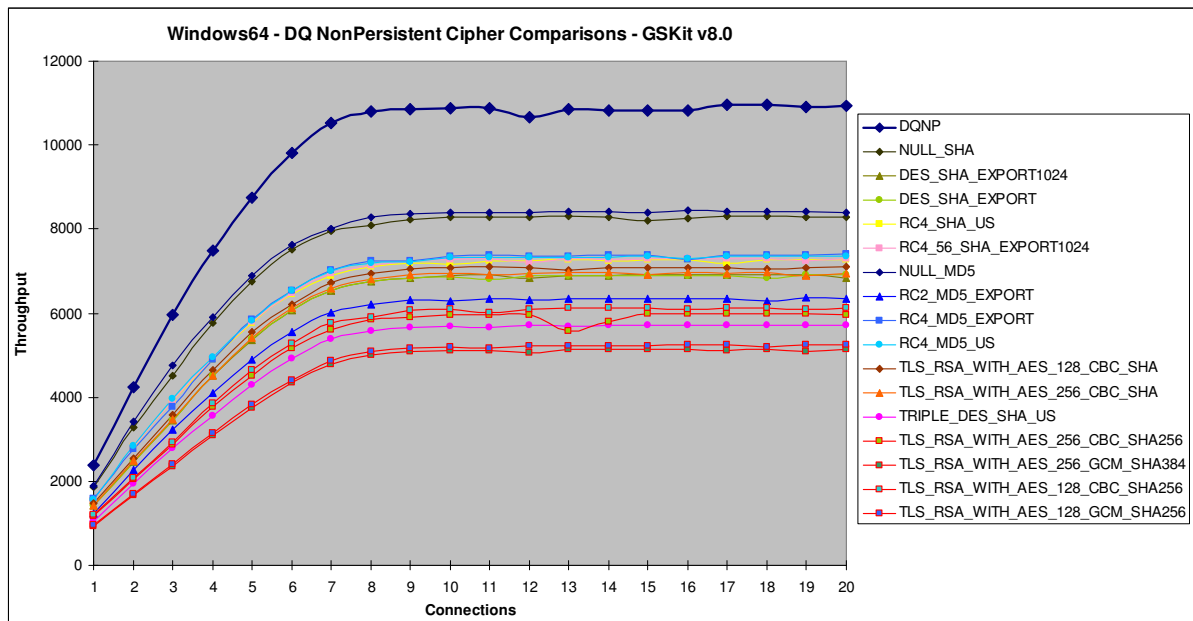


Figure 18 – DQ non-persistent message rates with various SSL ciphers

The area under the curve values for each cipher are used to compile an ordered table of these ciphers and can be used as a guide to expected performance degradations that can be expected when these ciphers are used. The unencrypted value is used to rate the other values as a percentage. Thus for the SSL cipher *NULL_MD5*, the expected message rate would be approximately 78% of the unencrypted rate.

Windows64 - SSL Cipher Comparison	DQNP
Unencrypted	100%
NULL_MD5	78%
NULL_SHA	76%
RC4_MD5_EXPORT	67%
RC4_MD5_US	67%
RC4_56_SHA_EXPORT1024	67%
RC4_SHA_US	66%
TLS_RSA_WITH_AES_128_CBC_SHA	64%

Windows64 - SSL Cipher Comparison	DQNP
TLS_RSA_WITH_NULL_NULL	64%
TLS_RSA_WITH_AES_256_CBC_SHA	63%
DES_SHA_EXPORT1024	63%
DES_SHA_EXPORT	63%
RC2_MD5_EXPORT	57%
TLS_RSA_WITH_AES_128_CBC_SHA256	55%
TLS_RSA_WITH_AES_256_CBC_SHA256	54%
TRIPLE_DES_SHA_US	51%
TLS_RSA_WITH_AES_128_GCM_SHA256	47%
TLS_RSA_WITH_AES_256_GCM_SHA384	46%

Table 13 – Ordered relative SSL DQ cipher performance

3 Large Messages

3.1 MQI Response Times: 50B to 100MB – Local Queue Manager

3.1.1 50B to 32KB

Figure 19 shows the response time for MQPut/MQGet pairs for non-persistent message sizes between 50B and 32KB.




Figure 19 –The effect of non-persistent message size on MQI response time (50B - 32KB)

Figure 20 shows the response time for MQPut/MQGet pairs for persistent message sizes between 50B and 32KB.




Figure 20 –The effect of persistent message size on MQI response time (50B - 32KB)

3.1.2 32KB to 2MB

Figure 21 show that the response time for MQPut/MQGet pairs has improved for all non-persistent message sizes between 32KB and 2MB.




Figure 21 –The effect of non-persistent message size on MQI response time (32KB – 2MB)

Figure 22 shows the response time for MQPut/MQGet pairs for persistent message sizes between 32KB and 2MB.




Figure 22 –The effect of persistent message size on MQI response time (32KB – 2MB)

3.1.3 2MB to 100MB

Figure 23 Response time for MQPut/MQGet pairs for NP message between 2MB and 100MB.




Figure 23 –The effect of non-persistent message size on MQI response time (2MB – 100MB)

Figure 24 Shows the response time for MQPut/MQGet pairs for persistent message sizes between 2MB and 100MB.




Figure 24 –The effect of persistent message size on MQI response time (2MB – 100MB)

3.2 20KB Messages

3.2.1 Local Queue Manager

Figure 25 and Figure 26 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

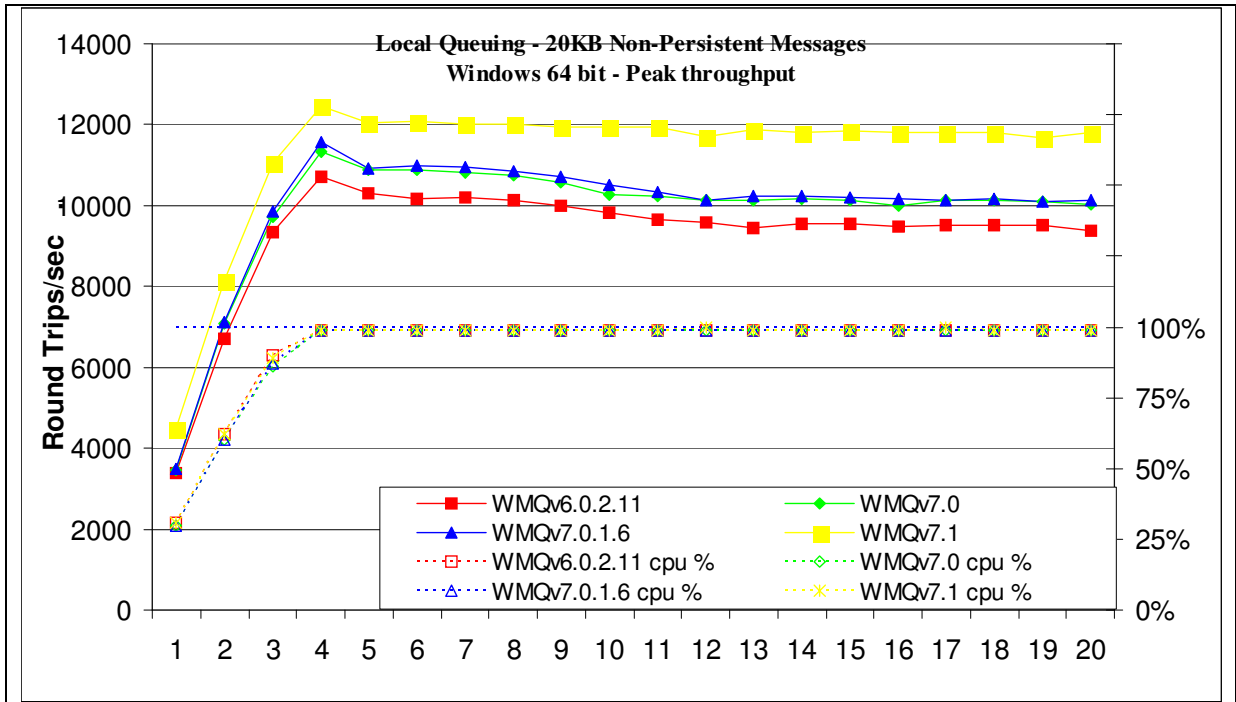


Figure 25 – 20KB non-persistent messages, local queue manager

Figure 25 and Table 14 shows that the peak throughput of non-persistent messages has increased by 8% when comparing version 7.1 to 7.0.1.6 and by 17% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 20KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	4	10710	0.0012	99%
WMQv7.0	4	11320	0.00051	99%
WMQv7.0.1.6	4	11584	0.00045	99%
WMQv7.1	4	12449	0.0004	99%

Table 14 – 20KB non-persistent messages, local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.2.1.1 Persistent Messages

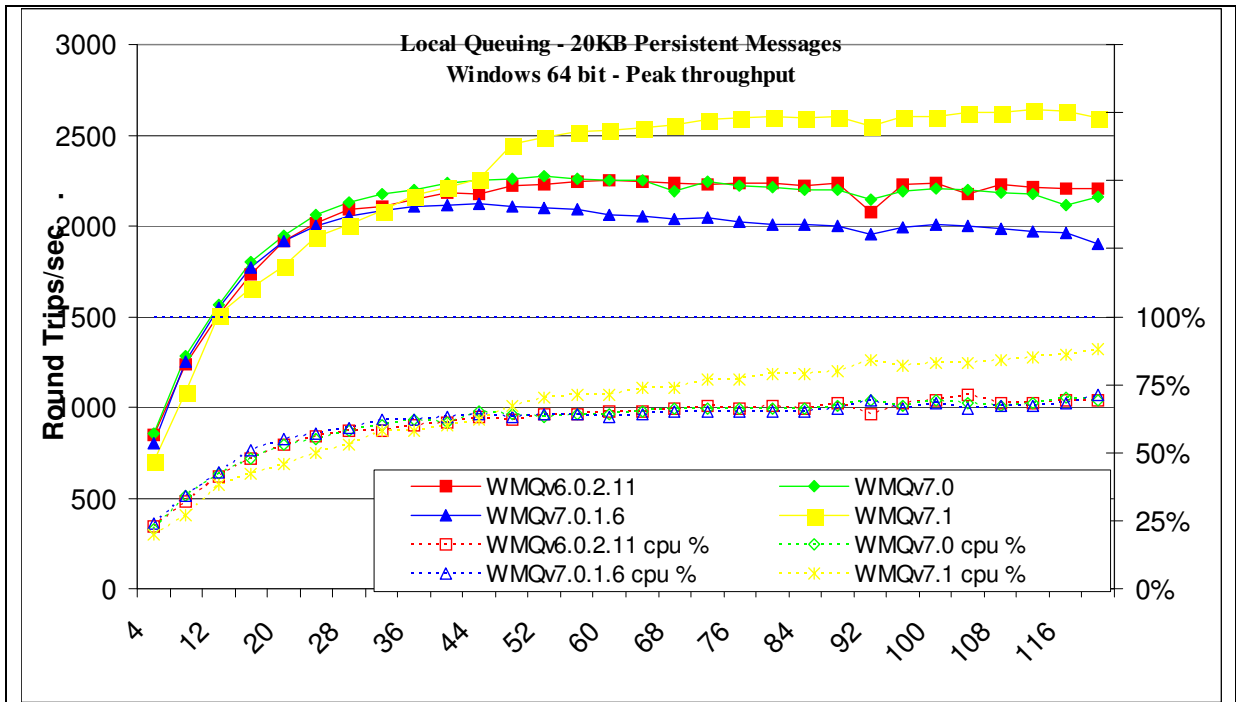


Figure 26 – 20KB persistent messages, local queue manager

Figure 26 and Table 15 shows that the peak throughput of persistent messages has increased by 24% when comparing version 7.1 to 7.0.1.6 and 17% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 20KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	60	2248	0.32	65%
WMQv7.0	52	2277	0.027	63%
WMQv7.0.1.6	44	2120	0.024	64%
WMQv7.1	112	2641	0.057	85%

Table 15 – 20KB persistent messages, local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.2.2 Client Channel

Figure 27 and Figure 28 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario.

3.2.2.1 Non-persistent Messages




Figure 27 – 20KB non-persistent messages, client channels

Figure 27 and Table 16 shows that the peak throughput of non-persistent messages has increased by 30% when comparing version 7.1 to 7.0.1.6 but has not changed when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 20KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	18	3800	0.016	96%
WMQv7.0	10	3064	0.0037	94%
WMQv7.0.1.6	9	3031	0.0037	95%
WMQv7.1	20	3937	0.0056	99%

Table 16 – 20KB non-persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.2.2.2 Persistent Messages

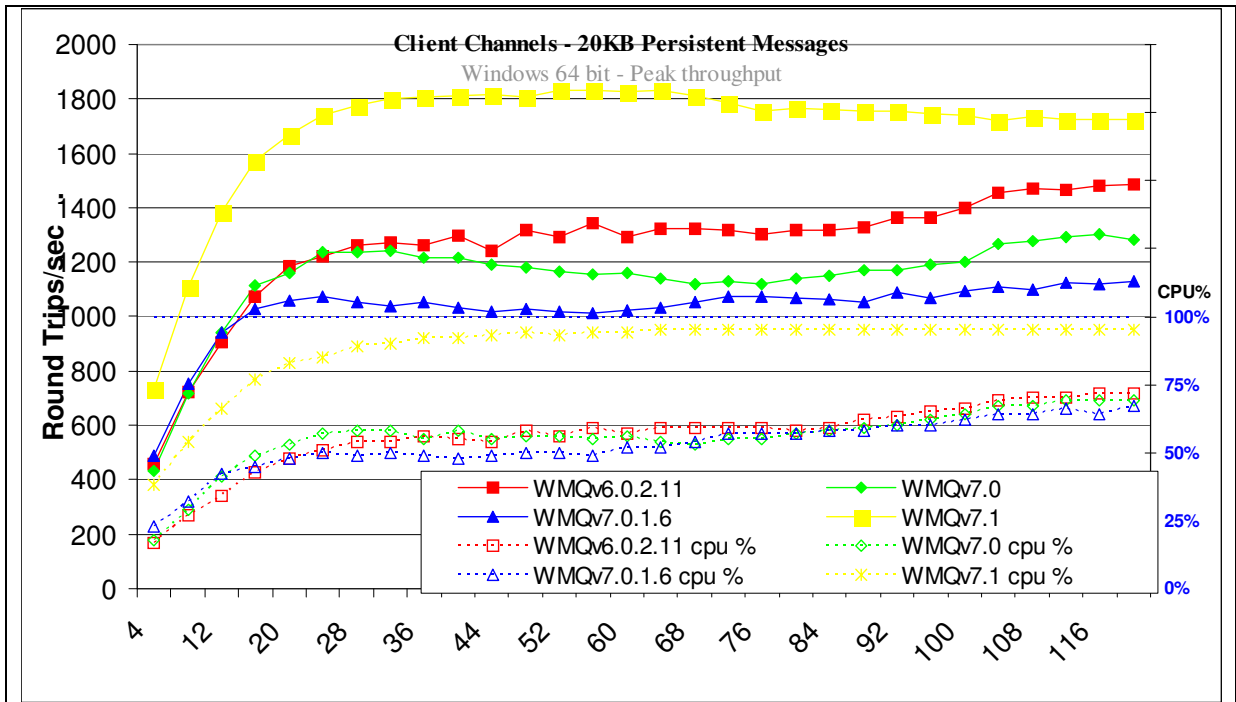


Figure 28 – 20KB persistent messages, client channels

Figure 28 and Table 17 shows that the peak throughput of persistent messages has increased by 62% when comparing version 7.1 to 7.0.1.6 and by 23% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 20KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	120	1486	0.17	72%
WMQv7.0	116	1302	0.1	69%
WMQv7.0.1.6	120	1132	0.13	67%
WMQv7.1	64	1832	0.041	95%

Table 17 – 20KB persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.2.3 Distributed Queuing

Figure 29 and Figure 30 shows the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario

3.2.3.1 Non-persistent Messages




Figure 29 – 20KB non-persistent messages, distributed queuing

Figure 29 and Table 18 shows that the peak throughput of non-persistent messages has increased by 6% when comparing version 7.1 to 7.0.1.6 and by 4% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 20KB Non Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	30	4415	0.017	98%
WMQv7.0	13	4368	0.0034	96%
WMQv7.0.1.6	12	4527	0.0026	98%
WMQv7.1	11	4781	0.0031	98%

Table 18 – 20KB non-persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.2.3.2 Persistent Messages




Figure 30 – 20KB persistent messages, distributed queuing

Figure 30 and Table 19 shows that the peak throughput of persistent messages has increased by 43% when comparing version 7.1 to 7.0.1.6 and by 23% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 20KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	470	1790	0.55	59%
WMQv7.0	490	1598	0.36	56%
WMQv7.0.1.6	470	1546	0.34	56%
WMQv7.1	340	2217	0.19	82%

Table 19 – 20KB persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.3 200K Messages

3.3.1 Local Queue Manager

Figure 31 and Figure 32 shows the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

3.3.1.1 Non-persistent Messages




Figure 31 – 200KB non-persistent messages, local queue manager

Figure 31 and Table 20 shows that the peak throughput of non-persistent messages has increased by 117% when comparing version 7.1 to 7.0.1.6 and by 116% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 200KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	8	1590	0.037	84%
WMQv7.0	9	1609	0.0062	82%
WMQv7.0.1.6	8	1582	0.0058	81%
WMQv7.1	4	3432	0.0012	99%

Table 20 – 200KB non-persistent messages, local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.3.1.2 Persistent Messages




Figure 32 – 200KB persistent messages, local queue manager

Figure 32 and Table 21 shows that the peak throughput of persistent messages has increased by 22% when comparing version 7.1 to 7.0.1.6 and by 18% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 200KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	68	233	0.00015	35%
WMQv7.0	68	226	0.35	35%
WMQv7.0.1.6	68	225	0.35	36%
WMQv7.1	76	274	0.32	43%

Table 21 – 200KB persistent messages, local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.3.2 Client Channel

Figure 33 and Figure 34 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario.

3.3.2.1 Non-persistent Messages




Figure 33 – 200KB non-persistent messages, client channels

Figure 33 and Table 22 shows that the peak throughput of non-persistent messages has increased by 43% when comparing version 7.1 to 7.0.1.6 and by 15% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 200KB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	12	481	0.013	93%
WMQv7.0	9	427	0.024	90%
WMQv7.0.1.6	8	433	0.021	89%
WMQv7.1	9	618	0.016	98%

Table 22 – 200KB non-persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.3.2.2 Persistent Messages




Figure 34 – 200KB persistent messages, client channels

Figure 34 and Table 23 shows that the peak throughput of persistent messages has increased by 31% when comparing version 7.1 to 7.0.1.6 and by 27% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 200KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	112	195	0.0014	60%
WMQv7.0	100	190	0.64	59%
WMQv7.0.1.6	100	181	0.68	55%
WMQv7.1	36	237	0.17	66%

Table 23 – 200KB persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.3.3 Distributed Queuing

Figure 35 and Figure 36 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario

3.3.3.1 Non-persistent Messages




Figure 35 – 200KB non-persistent messages, distributed queuing

Figure 35 and Table 24 shows that the peak throughput of non-persistent messages has increased by 13% when comparing version 7.1 to 7.0.1.6 and by 1000% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 200KB Non Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	20	58	0.018	4%
WMQv7.0	13	506	0.03	93%
WMQv7.0.1.6	20	528	0.046	94%
WMQv7.1	9	599	0.017	92%

Table 24 – 200KB non-persistent messages, distributed queuing

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.3.3.2 Persistent Messages




Figure 36 – 200KB persistent messages, distributed queuing

Figure 36 and Table 25 shows that the peak throughput of persistent messages has increased by 62% when comparing version 7.1 to 7.0.1.6 and by 600% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 200KB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	32	25	0.41	5%
WMQv7.0	44	128	0.41	30%
WMQv7.0.1.6	56	146	0.48	33%
WMQv7.1	112	237	0.55	56%

Table 25 – 200KB persistent messages, distributed queuing

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.4 2MB Messages

3.4.1 Local Queue Manager

Figure 37 and Figure 38 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

3.4.1.1 Non-persistent Messages




Figure 37 – 2MB non-persistent messages, local queue manager

Figure 37 and Table 26 shows that the peak throughput of non-persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 and is improved by 4% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 2MB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	4	70	0.0004	91%
WMQv7.0	4	74	0.054	92%
WMQv7.0.1.6	4	73	0.054	92%
WMQv7.1	4	72	0.055	94%

Table 26 – 2MB non-persistent messages, local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.4.1.2 Persistent Messages




Figure 38 – 2MB persistent messages, local queue manager

Figure 38 and Table 27 shows that the peak throughput of persistent messages has increased by 8% when comparing version 7.1 to 7.0.1.6 and by 8% when comparing version 7.1 to 6.0.2.11.

Test Name: Local Queuing - 2MB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	28	13	0.94	24%
WMQv7.0	32	13	2.8	25%
WMQv7.0.1.6	28	13	2.5	24%
WMQv7.1	12	14	0.94	26%

Table 27 – 2MB persistent messages, local queue manager

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.4.2 Client Channel

Figure 39 and Figure 40 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario.

3.4.2.1 Non-persistent Messages




Figure 39 – 2MB non-persistent messages, client channels

Figure 39 and Table 28 shows that the peak throughput of non-persistent messages has increased by 86% when comparing version 7.1 to 7.0.1.6 and by 11% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 2MB Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	7	35	0.0056	93%
WMQv7.0	10	22	0.53	71%
WMQv7.0.1.6	8	21	0.43	69%
WMQv7.1	6	39	0.17	97%

Table 28 – 2MB non-persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.4.2.2 Persistent Messages




Figure 40 – 2MB persistent messages, client channels

Figure 40 and Table 29 shows that the peak throughput of persistent messages has increased by 27% when comparing version 7.1 to 7.0.1.6 and by 27% when comparing version 7.1 to 6.0.2.11.

Test Name: Client Channels - 2MB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	32	11	0.041	36%
WMQv7.0	28	12	2.8	38%
WMQv7.0.1.6	28	11	2.9	34%
WMQv7.1	28	14	2.2	44%

Table 29 – 2MB persistent messages, client channels

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.4.3 Distributed Queuing

Figure 41 and Figure 42 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario.

3.4.3.1 Non-persistent Messages




Figure 41 – 2MB non-persistent messages, distributed queuing

Figure 41 and Table 30 shows that the peak throughput of non-persistent messages has increased by 4% when comparing version 7.1 to 7.0.1.6 and by 200% when comparing version 7.1 to 6.0.2.11..

Test Name: Distributed Queuing - 2MB Non Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	10	8	0.54	17%
WMQv7.0	12	26	0.54	60%
WMQv7.0.1.6	11	23	0.56	57%
WMQv7.1	6	24	0.28	57%

Table 30 – 2MB non-persistent messages, distributed queuing

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

3.4.3.2 Persistent Messages




Figure 42 - 2MB persistent messages, distributed queuing

Figure 42 and Table 31 shows that the peak throughput of persistent messages has increased by 75% when comparing version 7.1 to 7.0.1.6 and by 1300% when comparing version 7.1 to 6.0.2.11.

Test Name: Distributed Queuing - 2MB Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	8	1	2.5	3%
WMQv7.0	4	8	0.45	24%
WMQv7.0.1.6	28	8	3.8	25%
WMQv7.1	32	14	2.5	41%

Table 31 – 2MB persistent messages, distributed queuing

Note: The numbers in the table above show the peak number of round trips per second, the number of driving applications used, the response time and the server CPU at that time

4 Application Bindings

This report analyzes the message rate between a Requester (Driver) application and a Responder (Server) application. This chapter looks at the effect of various combinations of application bindings for Requester and Responder programs.

	Requester	Responder
Normal	Trusted	Shared
Isolated	Isolated	Isolated
Trusted	Trusted	Trusted
Non Trusted	Shared	Shared

4.1 Local Queue Manager

Figure 43 and Figure 44 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

4.1.1 Non-persistent Messages




Figure 43 – Application binding, non-persistent messages, local queue manager

Figure 43 and Table 32 show that the throughput of non-persistent messages when comparing Normal, Isolated, Trusted and Shared bindings.

Test Name: WMQv7.1 - Local Queuing - Application Bindings with Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	4	19029	0.00027	99%
Isolated	3	8238	0.00053	98%
Trusted	4	30023	0.00015	99%
Non-Trusted	4	14299	0.00029	99%

Table 32 – Application binding, non-persistent messages, local queue manager

4.1.2 Persistent Messages

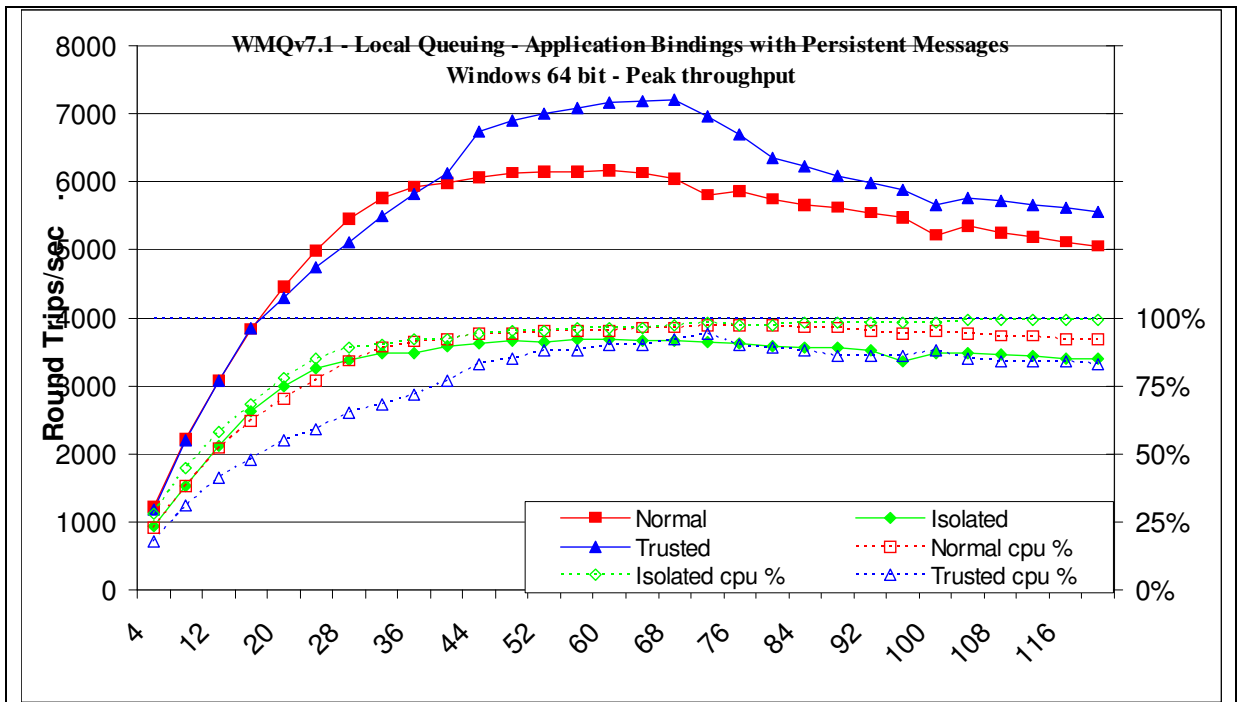


Figure 44 – Application binding, persistent messages, local queue manager

Figure 44 and Table 33 show that the throughput of persistent messages when comparing Normal, Isolated and Trusted bindings.

Test Name: WMQv7.1 - Local Queuing - Application Bindings with Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	60	6174	0.013	95%
Isolated	60	3688	0.021	96%
Trusted	68	7216	0.013	92%

Table 33 – Application binding, persistent messages, local queue manager

4.2 Client Channels

Figure 45 and Figure 46 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario.

4.2.1 Non-persistent Messages




Figure 45 – Application binding, non-persistent messages, client channels

Figure 45 and Table 34 show that the peak throughput of non-persistent messages when comparing Normal, Isolated and Trusted bindings.

Test Name: WMQv7.1 - Client Channels - Application Bindings with Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	11	7752	0.0018	99%
Isolated	16	6274	0.0026	99%
Trusted	10	8599	0.0014	97%

Table 34 – Application binding, non-persistent messages, client channels

4.2.2 Persistent Messages

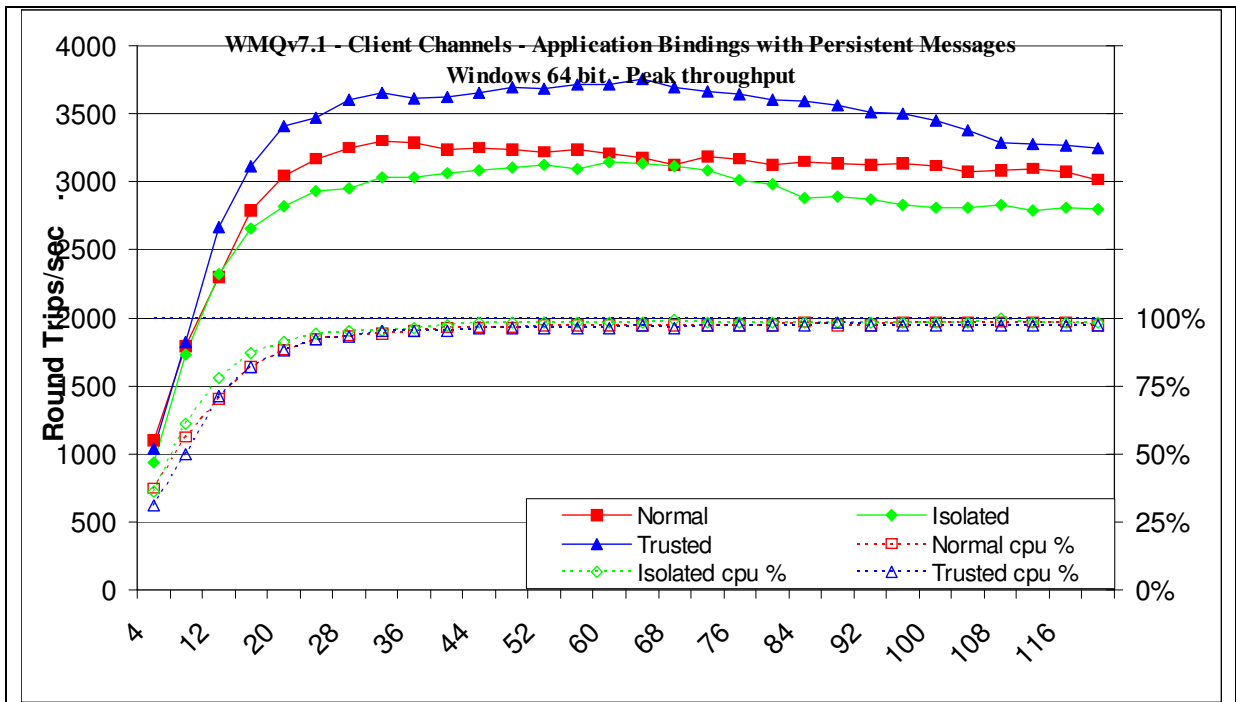


Figure 46 – Application binding, persistent messages, client channels

Figure 46 and Table 35 show that the peak throughput of non-persistent messages when comparing Normal, Isolated and Trusted bindings.

Test Name: WMQv7.1 - Client Channels - Application Bindings with Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	32	3298	0.011	94%
Isolated	60	3148	0.023	98%
Trusted	64	3755	0.018	97%

Table 35 – Application binding, persistent messages, client channels

4.3 Distributed Queuing

Figure 46 and Figure 47 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario.

4.3.1 Non-persistent Messages

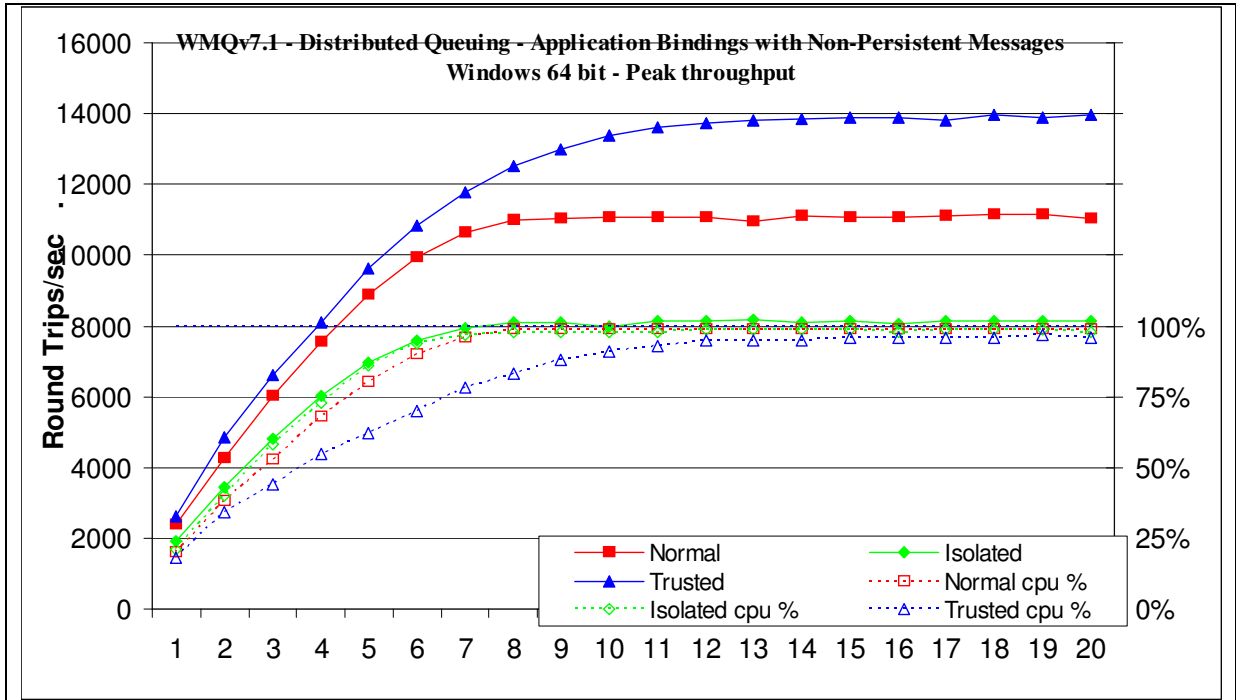


Figure 47 – Application binding, non-persistent messages, distributed queuing

Figure 47 and Table 36 show that the peak throughput of non-persistent messages when comparing Normal, Isolated and Trusted bindings.

Test Name: WMQv7.1 - Distributed Queuing - Application Bindings with Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	18	11148	0.0015	99%
Isolated	13	8161	0.0019	99%
Trusted	20	13981	0.0014	96%

Table 36 – Application binding, non-persistent messages, distributed queuing

Note: The large bold numbers in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput.

4.3.2 Persistent Messages

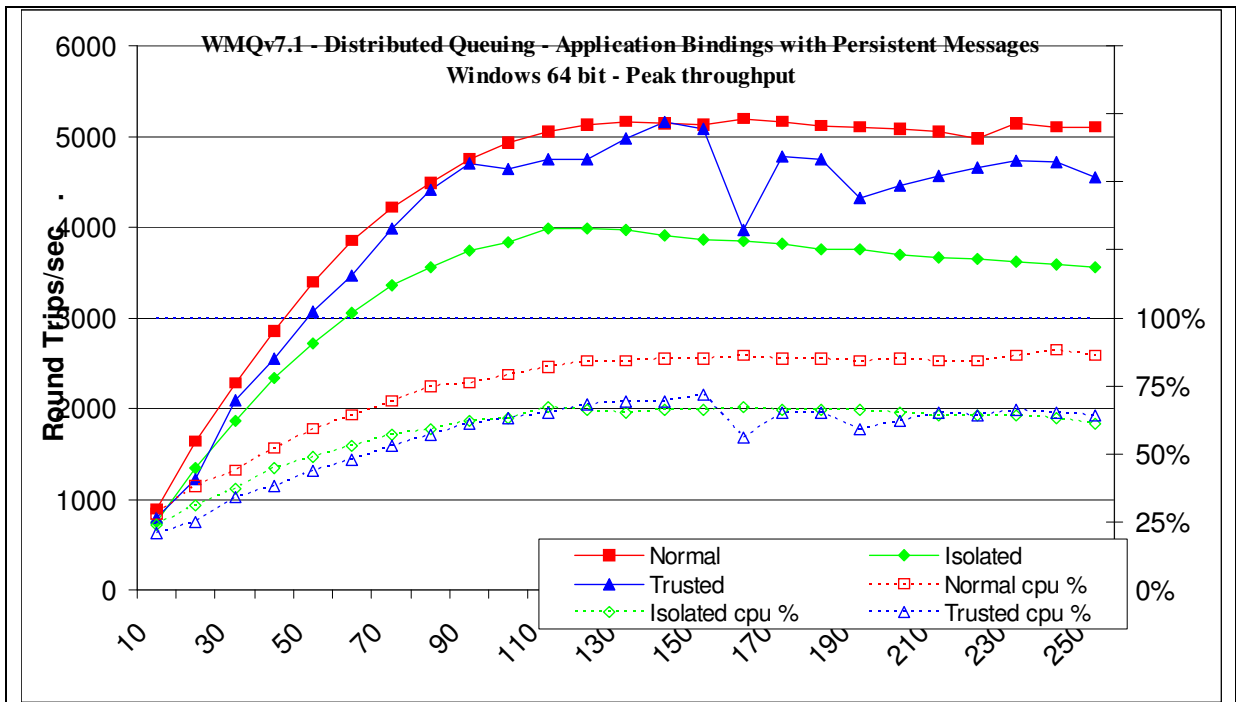


Figure 48 – Application binding, persistent messages, distributed queuing

Figure 48 and Table 37 show that the peak throughput of non-persistent messages when comparing Normal, Isolated and Trusted bindings.

Test Name: WMQv7.1 - Distributed Queuing - Application Bindings with Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	160	5184	0.041	86%
Isolated	110	3985	0.036	67%
Trusted	140	5165	0.035	69%

Table 37 – Application binding, persistent messages, distributed queuing

Note: The large bold numbers in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput.

5 Short & Long Sessions

NP 525 clients, 603 mws/sec, 0.368 resp, 79%cpu

P 325 clinets, 380 msg/sec, 0.806 resp, 76% cpu

6 Performance and Capacity Limits

6.1 Client channels – capacity measurements

The measurements in this section are intended to test the maximum number of client channels into a server queue managers with a messaging rate of 1 round trip per client channel per *minute* while additional connections are made. The maximum number of connected applications is likely to be determined by other criteria such as recovery time or manageability. Measurements are also made with smaller number of Client channels where the message insertion rate is increased until the system gets congested. This information is intended to be useful to the reader sizing a system with similar scenarios. These client measurements of V7.1 allocate a separate socket for each client (sharecnv=1 on svrcon channel).

Queue manager configuration for client channels capacity tests:

MaxChannels=50000 (100,000 for clnp_cmax). MQIBINDTYPE=FASTPATH

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
clnp	11	n/a*	7752	0.0018	99%
clnp_r3600	3050	3600	2947	0.43	89%
clnp_c6000	6000	1140	1880	0.943	91%
clnp max	10000	60 1033	166 1042	0.0016 0.084	39% >95%
cl_persist_c6000	6000	450	747	0.71	84%
clnp_cmax_no_correllid	6000	570	935	0.346	88%
cl_persist_cmax					

Table 38 – Capacity measurements, client channels

* There was no delay between the response to the previous message and the insertion of the next message with 11 clients.

The maximum message throughput is achieved when there are a small number of requester applications. The clnp_3600 measurement peaks when the queue of input messages waiting to be processed by the Server application builds up because the server application threads can no longer keep up with the demand. Although this ensures the server threads are always busy, the messages are being spilt from the Queue buffer to the file system and possibly to the disk. Each client uses a thread in the AMQRRMPA processes and the management of lots of threads and lots memory objects results in a larger CPU cost to handle each message.

Measurements normally use a Get by Correlation_Id from a common reply queue for all clients whereas the tests labelled ‘no_correllid’ have a separate reply queue per client. Each additional Client needs a thread in the AMQRRMPA process. Using a separate queue per client needs additional shared memory per client.

6.2 Distributed queuing – capacity measurements

The measurements in this section are intended to test the maximum number of server channel pairs between two queue managers with a messaging rate of 1 round trip per server channel per *minute* while applications are being attached. For the same number of server channel pairs, a faster message rate gives a higher total message throughput over each channel pair. This information is intended to be useful to the reader sizing a system with similar scenarios.

Queue manager and log configuration for distributed queuing capacity tests:

MaxChannels=20000, LogPrimaryFiles=12, LogFilePages=16384, LogBufferPages=512

Note: The large log capacity for this test is for writing the object definitions to the log disk (the transmission queue definitions for both sides of the server channel pair, and reply queue per receiver channel on the driving machine).

Test name:	Apps	Rate/app/hr	Round	Response time (s)	CPU
------------	------	-------------	-------	-------------------	-----

			Trips/sec		
dqnp	18	n/a*	11148	0.0015	99%
dqnp_r3600	5000	3600	4990	0.0028	42%
dqnp_q1000#	1000	23300	5224	0.161	97%
dqnp_qmax	2000	3000	1598	0.4614	88%
dq-persist-q1000	1000	3700	1014	0.064	78%
dq-persist_q2000	2000	1170	626	0.1129	84%

Table 39 – Capacity measurements, server channels

* *There was no delay between the response to the previous message and the insertion of the next message with 18 driving applications..*

This measurement used MQ V7.1.0.1

The dqnp and dqnp_r3600 both used a total of 4 pairs of Sender/Receiver pairs of channels between queue managers while the dq_qmax and dq_persist_q1000 used a pair of channels per application. The dqnp_q1000 shows the throughput experienced when 1000 queue managers are connected into a central hub.

7 Tuning Recommendations

7.1 Tuning the Queue Manager

This section highlights the tuning activities that are known to give performance benefits for WebSphere MQ V7.1; The reader should note that the following tuning recommendations **may not necessarily need** to be applied, especially if the message throughput and/or response time of the queue manager system already meets the required level. Some tuning recommendations that follow may degrade the performance of a previously balanced system if applied inappropriately. The reader should carefully monitor the results of tuning the queue manager to be satisfied that there have been no adverse effects.

Customers should test that any changes have not used excessive real resources in their environment and make only essential changes. For example, allocating several megabytes for multiple queues reduces the amount of shared and virtual memory available for other subsystems, as well as over committing real storage.

Note: The 'TuningParameters' stanza is not a documented external interface and maybe changed or be removed in future releases.

7.1.1 Queue Disk, Log Disk, and Message Persistence

Non-persistent messages are held in main memory, spilt to the file system as the queues become deep and lazily written to the Queue file. Persistent messages are synchronously written to the log by an MQCmit that are also periodically flushed to the Queue file.

To avoid potential queue and log I/O contention due to the queue manager simultaneously updating a queue file and log extent on the same disk, it is important that queues and logs are located on *separate* and *dedicated* physical devices. Multiple disks can be redirected to a Storage Area Network (SAN) but multiple high volume Queue managers can require different Logical Volumes to avoid congestion.

With the queue and log disks configured in this manner, careful consideration must still be given to message persistence: persistent messages should only be used if the message needs to survive a queue manager restart (forced by the administrator or as the result of a power failure, communications failure, or hardware failure). In guaranteeing the recoverability of persistent messages, the pathlength through the queue manager is three times longer than for a non-persistent message. This overhead does not include the additional time for the message to be written to the log, although this can be minimised by using cached disks or SAN.

7.1.1.1 Non-persistent and Persistent Queue Buffer

The default non-persistent queue buffer size is 64K per queue and the default persistent is 128K per queue for 32 bit Queue Managers and 128K /256K for 64 bit Queue Managers (AIX, Solaris, HPUX, Linux_64, z_Linux, and Windows64). They can all be increased to 1Mb using the TuningParameters stanza and the *DefaultQBufferSize* and *DefaultPQBufferSize* parameters. (For more details see SupportPac MP01: MQSeries – Tuning Queue Limits). Increasing the queue buffer provides the capability to absorb peaks in message throughput at the expense of real storage. Once these queue buffers are full, the additional message data is given to the file system that will eventually find its way to the disk. Defining queues using large non-persistent or persistent queue buffers can degrade performance if the system is short of real memory either because a large number of queues have already been defined with large buffers, or for other reasons -- e.g. large number of channels defined.

Note: The queue buffers are allocated in shared storage so consideration must be given to whether the agent process or application process has the memory addressability for all the required shared memory segments.

Queues can be defined with different values of *DefaultQBufferSize* and *DefaultPQBufferSize*. The value is taken from the TuningParameters stanza in use by the queue manager when the queue was defined. When the queue manager is restarted existing queues will keep their earlier definitions and new queues will be created with the current setting. When a queue is opened, resources are allocated according to the definition held on disk from when the queue was created.

7.1.2 Log Buffer Size, Log File Size, and Number of Log Extents

The Log component is often the bottleneck when processing persistent messages. Sufficient information is stored on the log to restart the queue manager after failure. Circular logging is sufficient to recover from application, software, or power failure while linear logging will also recover from media (or disk) failure. Log

records are written at each MQPut, MQGet, and MQCmit into the log buffer. This information is moved onto the log disk. Periodically the Checkpoint process will decide how many of these logfile extents are in the Active log and need to be kept online for recovery purposes. Those extents no longer in the active log are available for achieving when using Linear logging or available for reuse when using circular. There should be sufficient Primary logs to hold the Active log plus the new log extents used until the next checkpoint otherwise some Secondary logs are temporarily included in the log set and they have to be instantly formatted which is an unnecessary delay when using circular logging.

The log buffer is a circular piece of main memory where the log records are concatenated so that multiple log records can be written to the log file in a single I/O operation. The default values used for `LogBufferPages` and `LogFilePages` have been increased in V7 and are probably suitable for most installations. The default size of the log buffer is 512 pages with a maximum size of 4096 pages. To improve persistent message throughput of large messages (messages size > 1MB) the `LogBufferPages` could be increased to improve likelihood of messages only needing one I/O to get to the disk. Environments that process under 100 small (< 10KB messages) Persistent messages per second can reduce the memory footprint by using smaller values like 32 pages without impacting throughput. `LogFilePages` (i.e. `crtmqm -lf <LogFilePages>`) defines the size of one physical disk extent (default 4096 pages). The larger the disk extent, the longer the elapsed times between changing disk extents. It is better to have a smaller number of large extents but long running UOW can prevent Checkpointing efficiently freeing the disk extent for reuse. The largest size (maximum 65536 pages) will reduce the frequency of switching extents. The number of `LogPrimaryFiles` (i.e. `crtmqm -lp <LogPrimaryFiles>`) can be configured to a large number and the maximum number of Primary plus Secondary extents is 255 (Windows) and 511 (UNIX) but it is for functional reasons rather than performance that need more than 20 primary extents for Circular logging. Circular logging should be satisfied by Primary logs because Secondary logs are formatted each time they are reused. The Active log set is the number of extents that are identified by the Checkpoint process as being necessary to be kept online. As additional messages are processed, more space is taken by the active log. As UOWs complete, they enable the next Checkpoint process to free up extents that now become available for archiving with Linear logging. Some installation will use Linear logging and not archive the redundant logs because archiving impacts the run time performance of logging. They will periodically (daily or twice daily) use 'rcdmqmg' on the main queues thus moving the 'point of recovery' forward, compacting the queues, and freeing up log disk extents. The cumulative effect of this tuning will:

- Improve the throughput of persistent messages (enabling by default a possible 2MB of log records to be written from the log buffer to the log disk in a single write). Initial target - half to one second of log datastreaming into the Logbuffer.
- Reduce the frequency of log switching (permitting a greater amount of log data to be written into one extent). Initial target - LogFile extent hold at least 10 seconds of log datastreaming.
- Allow more time to prepare new linear logs or recycle old circular logs (especially important for long-running units of work).

Changes to the queue manager `LogBufferPages` stanza take effect at the next queue manager restart. The number of pages can be changed for all subsequent queue managers by changing the `LogBufferPages` parameter in the product default Log stanza.

It is unlikely that poor persistent message throughput will be attributed to a 2MB queue manager log but processing of large messages will be helped by these enhanced limits. It is possible to fill and empty the log buffer several times each second and reach a CPU limit writing data into the log buffer, before a log disk bandwidth limit is reached.

7.1.2.1 LogWriteIntegrity: SingleWrite or TripleWrite

The default value is TripleWrite. MQ writes log records using the TripleWrite method because it provides full write integrity where hardware that assures write integrity is not available.

Some hardware guarantees that, if a write operation writes a page and fails for any reason, a subsequent read of the same page into a buffer results in each byte in the buffer being either:

- The same as before the write, or
- The byte that should have been written in the write operation

On this type of hardware (for example, SSA write cache enabled), it is safe for the logger to write log records in a single write as the hardware assures full write integrity. This method provides the highest level of performance.

Queue manager workloads that have multiple streams asynchronously creating high volume log records will not benefit from 'SingleWrite' because the logger will not need to rewrite partial pages of the log file. Workloads that serialize on a small number of threads where the response time from an MQGet, MQPut, or MQCmit

inhibits the system throughput are likely to benefit from Singlewrite and could enhance throughput by 25%. Measurements in this report used LogWriteIntegrity=TripleWrite

7.1.3 Channels: Process or Thread, Standard or Fastpath?

Threaded channels are used for all the measurements in this report ('runmqtsr', and for server channels an MCATYPE of 'THREAD') the threaded listener 'runmqtsr' can now be used in all scenarios with client and server channels. Additional resource savings are available using the 'runmqtsr' listener rather than 'inetd', including a reduced requirement on: virtual memory, number of processes, file handles, and System V IPC.

Fastpath channels, and/or fastpath applications - see later paragraph for further discussion, can increase throughput for both non-persistent and persistent messaging. For persistent messages, the improvement is only for the path through the queue manager, and does not affect performance writing to the log disk.

Note: The reader should note that since the greater proportion of time for persistent messages is in the queue manager writing to the log disk, the performance improvement for fastpath channels is less apparent with persistent messages than with non-persistent messages.

7.2 Applications: Design and Configuration

7.2.1 Standard (Shared or Isolated) or Fastpath?

The reader should be aware of the issues associated with writing and using fastpath applications—described in the 'MQSeries Application Programming Guide'. Although it is recommended that customers use fastpath channels, it is not recommended to use fastpath applications. If the performance gain offered by running fastpath is not achievable by other means, it is essential that applications are rigorously tested running fastpath, and never forcibly terminated (i.e. the application should always disconnect from the queue manager). Fastpath channels are documented in the 'MQSeries Intercommunication Guide'.

7.2.2 Parallelism, Batching, and Triggering

An application should be designed wherever possible to have the capability to run *multiple instances* or *multiple threads* of execution. Although the capacity of a multi-processor (SMP) system can be fully utilised with a small number of applications using non-persistent messages, more applications are typically required if the workload is mainly using persistent messages. Processing messages inside syncpoint can help reduce the amount of time the queue manager takes to write a group of persistent messages to the log disk. The performance profile of a workload will also be subject to variability through cycles of low and heavy message volumes, therefore a degree of experimentation will be required to determine an optimum configuration.

Queue avoidance is a feature of the queue manager that allows messages to be passed directly from an 'MQPuter' to an 'MQGeter' without the message being placed on a queue. This feature only applies for processing messages outside of syncpoint. In addition to improving the performance of a workload with multiple parallel applications, the design should attempt to ensure that an application or application thread is always available to process messages on a queue (i.e. an 'MQGeter'), then messages outside of syncpoint do not need to ever be physically placed on a queue.

The reader should note that as more applications are processing messages on a single queue there is an increasing likelihood that queue avoidance will not be maintainable. The reasons for this have a cumulative and exponential effect, for example, when messages are being placed on a queue quicker than they can be removed. The first effect is that messages begin to fill the queue buffer—and MQGetters need to retrieve messages from the buffer rather than being received directly from an MQPuter. A secondary effect is that as messages are spilled from the buffer to the queue disk, the MQGetters must wait for the queue manager to retrieve the message from the queue disk rather than being retrieved from the queue buffer. While these problems can be addressed by configuring for more MQGetters (i.e processing threads in the server application), or using a larger queue buffer, it may not be possible to avoid a performance degradation.

Processing persistent messages inside syncpoint (i.e. in batches) can be more efficient than outside of syncpoint. As the number of messages in the batch increases, the average processing cost of each message decreases. For persistent messages the queue manager can write the entire batch of messages to the log disk in one go while outside of syncpoint control, the queue manager must wait for each message to be written to the log before returning control to the application.

Only one log record per queue can be written to the disk per log I/O when processing messages outside of syncpoint. This is not a bottleneck when there are a lot of different queues being processed. When there are a small number of queues being processed by a large number of parallel application threads, it is a bottleneck. By changing all the messages to be processed inside syncpoint, the bottleneck is removed because multiple log records per queue can share the same log I/O for messages processed within syncpoint.

A typical triggered application follows the performance profile of a short session. The 'runmqldr' has a much smaller overhead compared to inetd of connecting to and disconnecting from the queue manager because it does not have to create a new process. The programmatical implementation of triggering is still worth consideration with regard to programming a disconnect interval as an input parameter to the application program. This can provide the flexibility to make tuning adjustments in a production environment, if for instance, it is more efficient to remain connected to the queue manager between periods of message processing, or disconnect to free queue manager and Operating System resources.

7.3 Virtual Memory, Real Memory, & Paging

7.3.1 BufferLength

The AMQRMPPA process contains a thread per connected client. The BufferLength parameter of the MQGet is also used to allocate a long term piece of storage of this size in which the message is held before being retrieved by the client. If the size of the arriving messages cannot be predicted then the application should provide a buffer that can deal with 90% of the messages and redrive the MQGet after return code **2080 (X'0820') MQRC_TRUNCATED_MSG_FAILED** by providing a larger BUFFER for retrieving this particular message. There is a mechanism to gradually reduce the size of the storage in AMQRMPPA if the recent BufferLength size is significantly smaller than previous BufferLength.

7.3.2 MQIBINDTYPE

MQIBINDTYPE=FASTPATH will cause the channel to run 'Trusted' mode. Trusted applications do not use a thread in the Agent (AMQZLLA) process. This means there is no IPC between the Channel and Agent because the Agent does not exist in this connection. If the channel is run in STANDARD mode then any messages passed between the channel and agent will use IPCC memory (size = BufferSize with a maximum size of 1MB) that is dynamically obtained and only held for the lifetime of the MQGet. Standard channels each require an additional 80KB of memory. As the message rate increases, there will be more IPCC memory used in parallel.

The power of the machine used to process a workload needs to handle the peaks of troughs. Customers may specify a daily workload but this number cannot be divided by the number of seconds in a day to find the necessary system configuration. The peak hourly rate cannot be divided by 3600 because the peak rate per second will probably be 2-3 times higher. The system must process these peak loads without building up a backlog of queued work. It is important to prevent the queue depths increasing because they will occupy memory from the 'fre' pool or be spilled out to disk. Over commitment of real memory is handled by the page manager but sudden large jumps (storms) possibly due to queues becoming deep can cause the throughput to break down completely if the page manager chooses too much working set memory to be paged. Gradual over commitment enables the page manager to shuffle out those pages that are not part of the working set.

8 Measurement Environment

8.1 Workload description

8.1.1 MQI response time tool

The MQI tool exercises the local queue manager by measuring elapsed times of the 8 main MQSeries verbs: MQConn(x), MQDisc, MQOpen, MQClose, MQPut, MQGet, MQCmit, and MQBack. The following MQI calls are paired together inside a test application:

- MQConn(X) with MQDisc
- MQOpen with MQClose
- MQPut with MQGet
- MQCmit and MQBack with MQPut and MQGet

Note: MQClose elapsed time is only measured for an empty queue.

Note: Performance of MQCmit and MQBack is measured in conjunction with MQPut and MQGet, putting and getting messages inside a unit of work (i.e. inside syncpoint control). The unit of work is committed at the end of each batch. The number of messages per batch is a parameter of the test.

Note: This tool is not used to measure the performance of verbs: MQSet, MQInq, or MQBegin.

8.1.2 Test scenario workload

The MQI applications use 64 bit libraries for MQ

8.1.2.1 The driving application programs

The test scenario workload simulates many driving applications running on a single driving machine. This is not typical of a customer environment and is only used to facilitate test coordination. Driving applications were multi-threaded with each thread performing a sequence of MQI calls. The driving applications (Requesters) for Local and DQ tests used Trusted bindings. The number of threads in each application was adjusted according to whether the test was measuring a local queue manager, a client channel, or distributed queuing scenario. This was done to reduce storage overheads on the driving system.

Message rate: in all but the *rated* and *capacity limit* tests, message processing was performed in a *tight-loop*. In the *rated* tests a message rate of 1 round trip per driving application per *second* was used, and in the *capacity limit* tests a message rate of 1 round trip per channel per *minute* was used.

Non-persistent and persistent messages were used in all but the *capacity limit* tests.

Note: The driving applications gathered timing information for all MQI calls using a high-resolution timer.

8.1.2.2 The server application program

The server application was written as a multi-threaded program configured to use various threads for processing non-persistent messages and persistent messages. Each server thread performed the sequence of actions as outlined in the test scenario illustrations.

Non-persistent messaging was done outside of syncpoint control. Persistent messaging was done inside of syncpoint control. The average message throughput expressed as a number of round trips per second was calculated and reported by the server program.

8.1.2.3 Analysis techniques

In the overview section, the percentage throughput comparison used the area under the graph as an alternative method of interpreting the performance data. Elsewhere, the percentage throughput comparison used the peak throughputs found in the tables associated with the graphs. The area under the curve was favoured in this instance as it gives a much more general performance indicator than a single point.

NB: Locking improvements in WMQv7.1 have improved the right hand side of the graphs but came with path length costs that may affect the rate of growth on left hand side of the graph when there is only a small number of parallel applications.

8.2 Hardware

IBM x3850: Server system (Device under test)
Model: x3850 M2
Processor: 3.3GHz Intel xeon
Architecture: 4 CPU
Memory (RAM): 4GB
Disk: 2 SAN disks on DS8700 (5GB each, 1 queue, 1 log partition)
Network: 10Gb Ethernet

IBM x3850: Driver system
Model: x3850 M2
Processor: 3.3GHz Intel xeon
Architecture: 4 CPU
Memory (RAM): 4GB
Disk: 2 SAN disks on DS8700 (5GB each, 1 queue, 1 log partition)
Network: 10Gb Ethernet

The machines under test are connected to a SAN via a dedicated SVC. The SVC provides a transparent buffer between the server and SAN that will smooth any fluctuations in the response of the SAN due to external workloads. The server machines are connected via a fibre channel trunk to a 8Gb Brocade DCX director. The speed of each server is dictated by the server's HBA (typically 2Gb). 5GB generic LUNs are provisioned via SVC. The SVC is a 2145-8G4 that connects to the DCX at 4Gb. The SAN storage is provided by an IBM DS8700 that is connected to the DCX at 4Gb.

8.3 Software

Windows : Microsoft Windows Server 2003 Enterprise x64 Edition
MQSeries: Version 6.0.2.11, Version 7.0, Version 7.0.1.6, Version 7.1
Compiler: Microsoft visual C compiler

9 Glossary

Test name	<p>The name of the test.</p> <p><i>Note: The test names in some cases are rather long. This is done to provide a descriptive qualification of the test measurement to relate to the performance discussion in the sections throughout the document:</i></p> <p>local => local queue manager test scenario</p> <p>cl => client channel test scenario</p> <p>dq => distributed queuing test scenario</p> <p>np => non-persistent messages</p> <p>pm => persistent messages</p> <p>r3600 => 1 round trip per driving application per second</p> <p>runmqtsr => channels using the 'runmqtsr' listener (client channel test scenario, in addition to 'runmqchi' for distributed queuing test scenarios)</p> <p>c6000 => 6,000 client driving applications (i.e. 6,000 MQI-client connections)</p> <p>q1000 => 1,000 server channel pairs</p> <p>max => maximum number of channels (or channel pairs)</p> <p>no_correl_id => correlation identifier not used in the response messages (as each response is placed on a unique reply-to queue per driving application)</p>
Apps	The number of driving applications connected to the queue manager at the point where the performance measurement is given.
Rate/App/hr	The target message throughput rate of each driving application.
Round T/s	The average achieved message throughput rate of all the driving applications together, measured by the server application.
% (Round T/s)	<p>The percentage increase in the total message throughput rate.</p> <p><i>Note: The nature of the comparison is noted under each table where percentage improvements have been given.</i></p>
CPU	As reported by VMSTAT
Resp time (s)	The average response time each round trip, as measured and averaged by all the driving applications.
Swap	The total amount of swap area reservation for all processes in MB, unless otherwise specified as swap/app (i.e. swap area reservation per driving application).
FREE	Free memory as reported by IOSTAT