

WebSphere MQ Solaris v7.1

Performance Evaluations

Version 1.2

February 2012

Fred Preston, Craig Stirling and Peter Toghill.

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 Solaris 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 Solaris 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 v7.1 for Solaris.
- Want to understand the performance limits of WebSphere MQ v7.1 for Solaris.
- Want to understand what actions may be taken to tune WebSphere MQ v7.1 for Solaris.

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 Sun T2000 server 8 core 1.4GHz with 64GB of RAM was used as the Device under test.

Two Sun V490s where used as the Drivers

How this document is arranged

Pages: 1-13

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: 18-39

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

Application Bindings

Page: 40-45

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 Limits

Pages: 48

Tuning Recommendations

Pages: 50

Tuning recommendations for Solaris

Measurement Environment

Pages: 55

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.

Glossary

Page: 57

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
2.3	Distributed Queuing Test Scenario.....	12
2.3.1	Non-persistent Messages – Server Channels.....	13
2.3.2	Non-Persistent non-Trusted – Server Channels.....	14
2.3.3	Persistent Messages – Server Channels.....	15
2.3.4	Server Channels.....	16
3	Large Messages	18
3.1	MQI Response Times: 50B to 100MB – Local Queue Manager.....	18
3.1.1	50B to 32KB.....	18
3.1.2	32KB to 2MB.....	20
3.1.3	2MB to 100MB.....	21
3.2	20KB Messages.....	22
3.2.1	Local Queue Manager.....	22
3.2.2	Client Channel.....	24
3.2.3	Distributed Queuing.....	26
3.3	200K Messages.....	28
3.3.1	Local Queue Manager.....	28
3.3.2	Client Channel.....	30
3.3.3	Distributed Queuing.....	32
3.4	2MB Messages.....	34
3.4.1	Local Queue Manager.....	34
3.4.2	Client Channel.....	36
3.4.3	Distributed Queuing.....	38
4	Application Bindings	40
4.1	Local Queue Manager.....	40
4.1.1	Non-persistent Messages.....	40
4.1.2	Persistent Messages.....	41
4.2	Client Channels.....	42
4.2.1	Non-persistent Messages.....	42
4.2.2	Persistent Messages.....	43
4.3	Distributed Queuing.....	44
4.3.1	Non-persistent Messages.....	44
4.3.2	Persistent Messages.....	45
5	Short & Long Sessions	46
6	Performance and Capacity Limits	48
6.1	Client channels – capacity measurements.....	48
6.2	Distributed queuing – capacity measurements.....	48
7	Tuning Recommendations	50
7.1	Tuning the Queue Manager.....	50
7.1.1	Queue Disk, Log Disk, and Message Persistence.....	50
7.1.2	Log Buffer Size, Log File Size, and Number of Log Extents.....	50
7.1.3	Channels: Process or Thread, Standard or Fastpath?.....	52
7.2	Applications: Design and Configuration.....	52
7.2.1	Standard (Shared or Isolated) or Fastpath?.....	52
7.2.2	Parallelism, Batching, and Triggering.....	52
7.3	Tuning the Operating System (Solaris).....	53
7.4	Virtual Memory, Real Memory, & Paging.....	53
7.4.1	BufferLength.....	53
7.4.2	MQIBINDTYPE.....	53
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 – Performance headline, non-persistent messages and server channels	13
Table 9 – Performance headline, non-persistent, non trusted messages and server channels.....	14
Table 10 – Performance headline, persistent messages and server channels.....	15
Table 11 – 1 round trip per driving application per second, client channels	17
Table 12 – 20KB non-persistent messages, local queue manager	22
Table 13 – 20KB persistent messages, local queue manager	23
Table 14 – 20KB non-persistent messages, client channels	24
Table 15 – 20KB persistent messages, client channels.....	25
Table 16 – 20KB non-persistent messages, client channels	26
Table 17 – 20KB persistent messages, client channels.....	27
Table 18 – 200KB non-persistent messages, local queue manager	28
Table 19 – 200KB persistent messages, local queue manager	29
Table 20 – 200KB non-persistent messages, client channels	30
Table 21 – 200KB persistent messages, client channels.....	31
Table 22 – 200KB non-persistent messages, distributed queuing	32
Table 23 – 200KB persistent messages, distributed queuing	33
Table 24 – 2MB non-persistent messages, local queue manager	34
Table 25 – 2MB persistent messages, local queue manager.....	35
Table 26 – 2MB non-persistent messages, client channels.....	36
Table 27 – 2MB persistent messages, client channels.....	37
Table 28 – 2MB non-persistent messages, distributed queuing	38
Table 29 – 2MB persistent messages, distributed queuing.....	39
Table 30 – Application binding, non-persistent messages, local queue manager.....	40
Table 31 – Application binding, persistent messages, local queue manager	41
Table 32 – Application binding, non-persistent messages, client channels	42
Table 33 – Application binding, persistent messages, client channels	43
Table 34 – Application binding, non-persistent messages, distributed queuing.....	44
Table 35 – Application binding, persistent messages, distributed queuing	45
Table 34 – Capacity measurements, client channels	48
Table 36 – Capacity measurements, server channels	49

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 – Server channels between two queue managers 12

Figure 12 – Performance headline, non-persistent messages and server channels 13

Figure 13 – Performance headline, non-persistent, not trusted messages and server channels 14

Figure 14 – Performance headline, persistent messages and server channels 15

Figure 15 – 1 round trip per driving application per second, server channel, non-persistent messages . 16

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

Figure 18 –The effect of non-persistent message size on MQI response time (50B - 32K) 18

Figure 19 –The effect of persistent message size on MQI response time (50byte - 32K)..... 18

Figure 20 –The effect of non-persistent message size on MQI response time (32KB – 2MB) 20

Figure 21 –The effect of persistent message size on MQI response time (32KB – 2MB) 20

Figure 22 –The effect of non-persistent message size on MQI response time (2MB – 100MB) 21

Figure 23 –The effect of persistent message size on MQI response time (2MB – 100MB)..... 21

Figure 24 – 20KB non-persistent messages, local queue manager..... 22

Figure 25 – 20KB persistent messages, local queue manager 23

Figure 26 – 20KB non-persistent messages, client channels 24

Figure 27 – 20KB persistent messages, client channels 25

Figure 28 – 20KB non-persistent messages, distributed queuing 26

Figure 29 – 20KB persistent messages, distributed queuing 27

Figure 30 – 200KB non-persistent messages, local queue manager..... 28

Figure 31 – 200KB persistent messages, local queue manager 29

Figure 32 – 200KB non-persistent messages, client channels 30

Figure 33 – 200KB persistent messages, client channels 31

Figure 34 – 200KB non-persistent messages, distributed queuing 32

Figure 35 – 200KB persistent messages, distributed queuing 33

Figure 36 – 2MB non-persistent messages, local queue manager 34

Figure 37 – 2MB persistent messages, local queue manager 35

Figure 38 – 2MB non-persistent messages, client channels 36

Figure 39 – 2MB persistent messages, client channels..... 37

Figure 40 – 2MB non-persistent messages, distributed queuing 38

Figure 41 - 2MB persistent messages, distributed queuing 39

Figure 42 – Application binding, non-persistent messages, local queue manager 40

Figure 43 – Application binding, persistent messages, local queue manager..... 41

Figure 44 – Application binding, non-persistent messages, client channels..... 42

Figure 45 – Application binding, persistent messages, client channels 43

Figure 46 – Application binding, non-persistent messages, distributed queuing 44

Figure 47 – Application binding, persistent messages, distributed queuing 45

Figure 48 – Short sessions, client channels 46

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 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:-

- 6% better than v6.0.2.11
- 14% better than v7.0
- 15% better than v7.0.1.6

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

- 36% better than v6.0.2.11
- 30% better than v7.0
- 38% better 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 shows 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 41% when comparing version 7.1 to 7.0.1.6 and 42% 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	12	22015	0.0006	63%
WMQv7.0	11	20600	0.00062	56%
WMQv7.0.1.6	13	22375	0.00064	66%
WMQv7.1	20	31487	0.00069	95%

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 shows that the peak throughput of non-persistent, non-trusted messages (shared bindings - MQIBINDTYPE=NORMAL) has increased by 21% when comparing version 7.1 to 7.0.1.6 and 22% 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	14	19154	0.00085	81%
WMQv7.0	14	18713	0.00085	77%
WMQv7.0.1.6	13	19345	0.00077	77%
WMQv7.1	19	23339	0.00093	96%

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 55% when comparing version 7.1 to 7.0.1.6 and 60% 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	16	2151	0.0086	17%
WMQv7.0	18	2944	0.0075	60%
WMQv7.0.1.6	22	2214	0.011	17%
WMQv7.1	18	3441	0.0063	23%

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

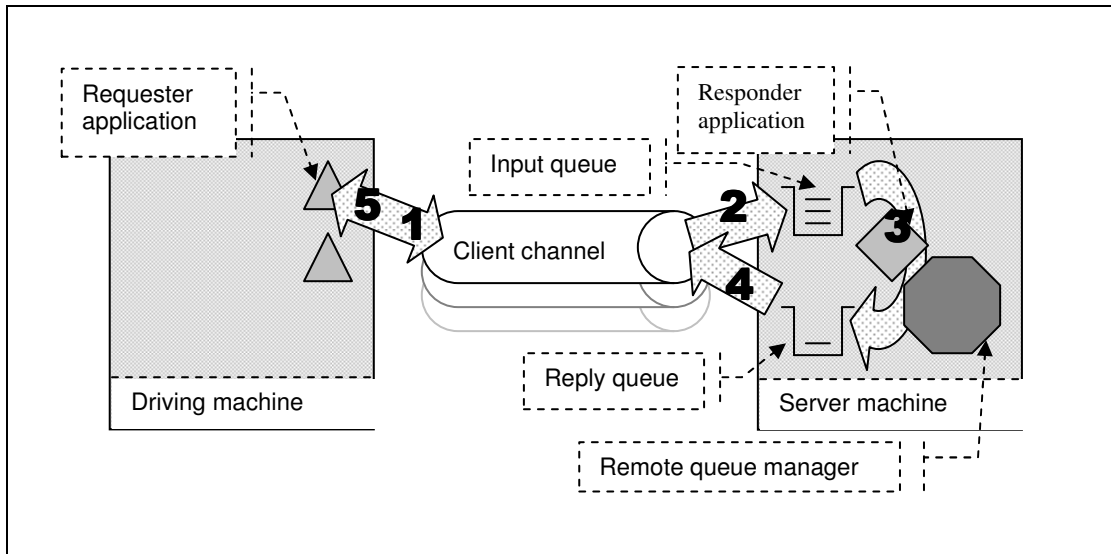


Figure 5 – MQI-client channels into a remote queue manager

- 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 shows 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 28% when comparing version 7.1 to 7.0.1.6 but has decreased by 12% 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	40	14319	0.003	64%
WMQv7.0	24	10069	0.0025	48%
WMQv7.0.1.6	26	9933	0.0028	52%
WMQv7.1	40	12699	0.0033	64%

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

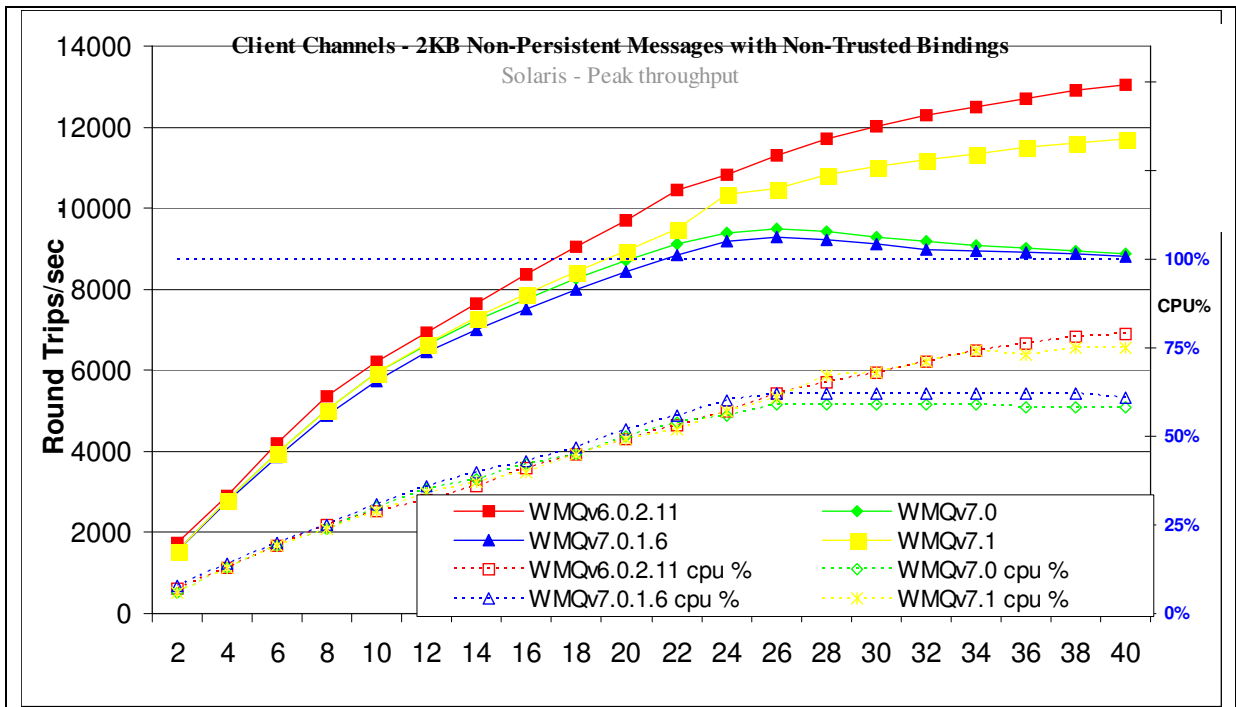


Figure 7 – Performance headline, non-persistent messages with non-trusted client channels

Figure 7 and Table 5 shows that the peak throughput of non-persistent, non-trusted messages (shared bindings - MQIBINDTYPE=NORMAL) has increased by 26% when comparing version 7.1 to 7.0.1.6 but has decreased by 10% 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	40	13056	0.0033	79%
WMQv7.0	26	9489	0.0029	59%
WMQv7.0.1.6	26	9280	0.003	62%
WMQv7.1	40	11723	0.0037	75%

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 throughput of persistent messages has increased by 36% when comparing version 7.1 to 7.0.1.6 and by 30% 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	80	2537	0.033	33%
WMQv7.0	120	2049	0.068	82%
WMQv7.0.1.6	104	2429	0.047	35%
WMQv7.1	88	3301	0.028	39%

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 15% when comparing version 7.1 to 7.0.1.6 and by 6% when comparing version 7.1 to 6.0.2.11. It also shows that the throughput of persistent messages has increased by 14% when comparing version 7.1 to 7.0.1.6 and by 14% 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	4100	4049	0.99	82%
WMQv7.0	3800	3789	0.74	81%
WMQv7.0.1.6	3800	3777	0.32	20%
WMQv7.1	4400	4359	0.18	25%

Test Name: Client Channels - R3600 Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	1700	1688	0.92	71%
WMQv7.0	1550	1532	0.98	72%
WMQv7.0.1.6	1750	1726	0.93	22%
WMQv7.1	2000	1977	0.19	22%

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

2.3 Distributed Queuing Test Scenario




Figure 11 – 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 12, Figure 13 and Figure 14 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 11 on the previous page) and WebSphere MQ (versions 7.1, 7.0.1.6, 7.0 and 6.0.2.11).

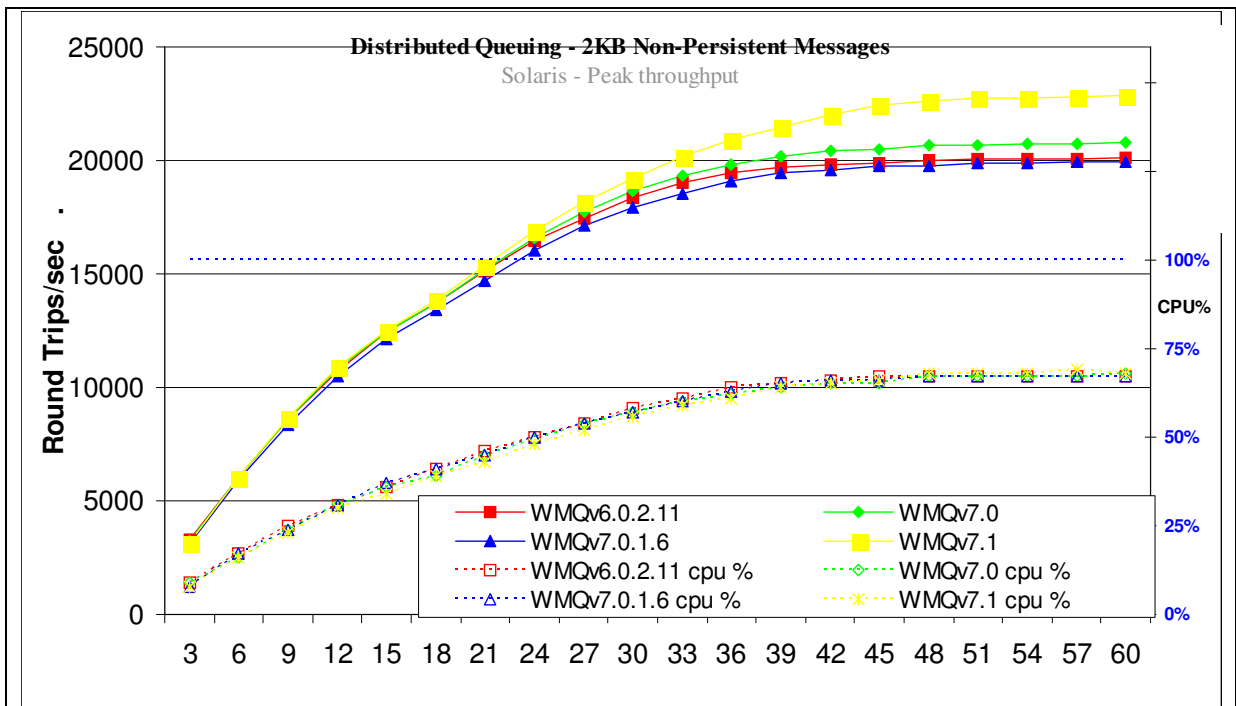


Figure 12 – Performance headline, non-persistent messages and server channels

Figure 12 and Table 8 shows that the peak throughput of persistent messages has increased by 15% when comparing version 7.1 to 7.0.1.6 and by 14% 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	60	20116	0.0032	67%
WMQv7.0	60	20791	0.0031	68%
WMQv7.0.1.6	60	19945	0.0032	67%
WMQv7.1	60	22841	0.0029	68%

Table 8 – 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 13 – Performance headline, non-persistent, not trusted messages and server channels

Figure 13 and Table 9 shows that the peak throughput Table 9 of non-persistent, non-trusted messages has increased by 18% when comparing version 7.1 to 7.0.1.6 and by 10% 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	40	12980	0.0035	56%
WMQv7.0	32	13462	0.0028	53%
WMQv7.0.1.6	36	12152	0.0033	51%
WMQv7.1	40	14306	0.0032	53%

Table 9 – 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

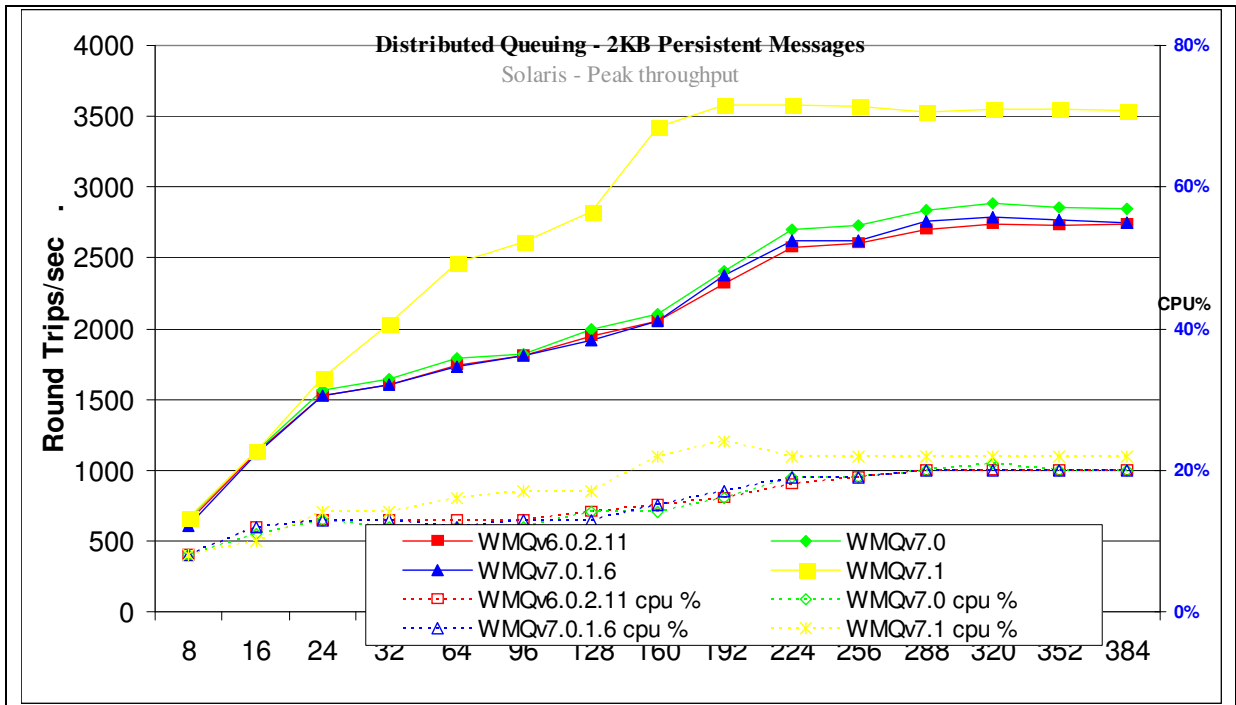


Figure 14 – Performance headline, persistent messages and server channels

Figure 14 and Table 10 shows that the peak throughput of persistent messages has increased by 29% when comparing version 7.1 to 7.0.1.6 and by 31% 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	384	2738	0.15	20%
WMQv7.0	320	2888	0.11	21%
WMQv7.0.1.6	320	2784	0.12	20%
WMQv7.1	192	3584	0.057	24%

Table 10 – 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 15 – 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 16 – 1 round trip per driving application per second, server channel, persistent messages

Figure 15, Figure 16 and Table 11 shows that the 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 40% 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 - R3600 Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	5400	4988	0.16	46%
WMQv7.0	5050	4979	0.093	48%
WMQv7.0.1.6	5100	4983	0.14	15%
WMQv7.1	4900	5007	0.012	14%

Test Name: Distributed Queuing - R3600 Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	2850	2822	0.17	60%
WMQv7.0	2850	2820	0.16	64%
WMQv7.0.1.6	2500	2477	0.17	16%
WMQv7.1	3550	3463	0.58	20%

Table 11 – 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

3 Large Messages

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

3.1.1 50B to 32KB

Figure 17 shows the response time for MQPut/MQGet for non-persistent messages between 50B and 32KB.




Figure 17 –The effect of non-persistent message size on MQI response time (50B - 32K)

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

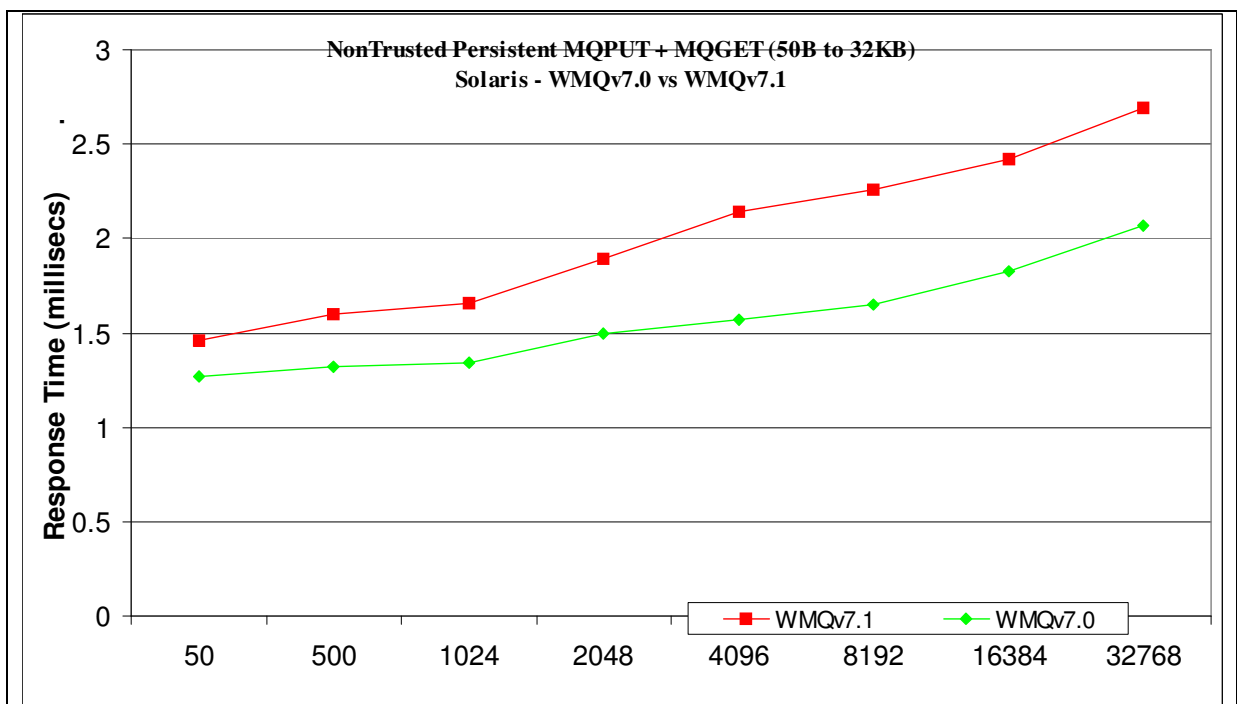


Figure 18 –The effect of persistent message size on MQI response time (50byte - 32K)

3.1.2 32KB to 2MB

Figure 19 show that the response time for MQPut/MQGet pairs has improved for all non-persistent message sizes between 32KB and 2MB.

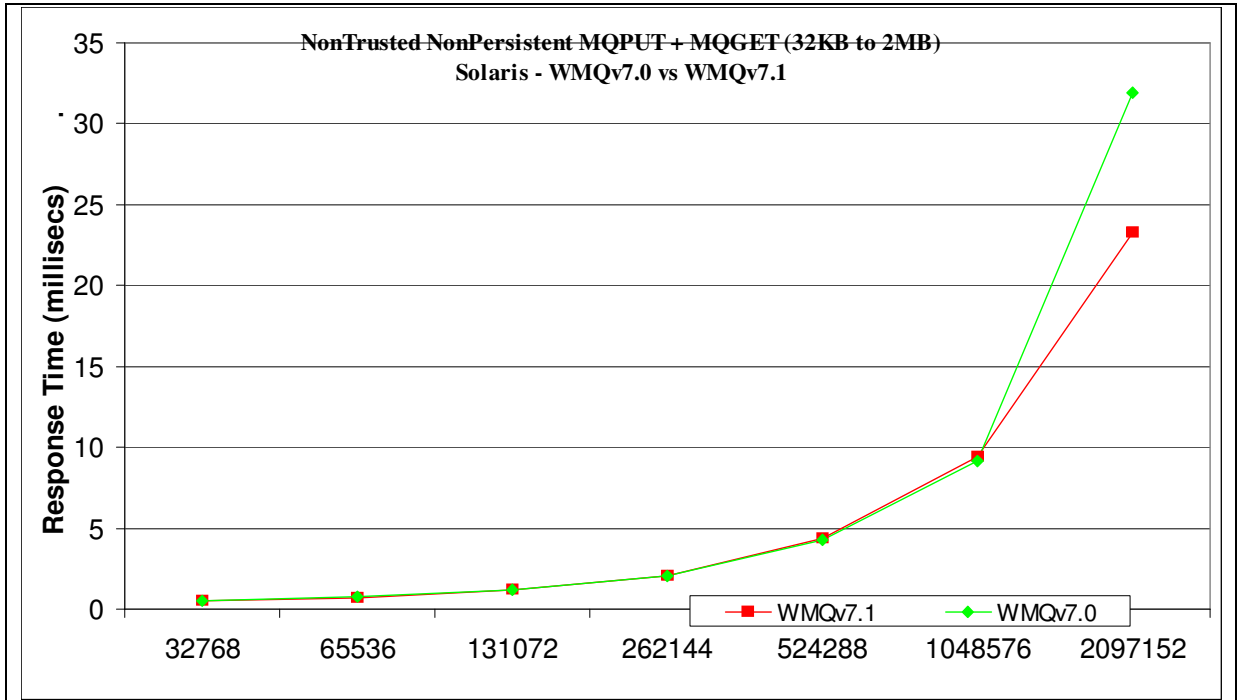


Figure 19 –The effect of non-persistent message size on MQI response time (32KB – 2MB)

Figure 20 show that the response for MQPut/MQGet pairs for persistent message sizes between 32KB and 2MB.




Figure 20 –The effect of persistent message size on MQI response time (32KB – 2MB)

3.1.3 2MB to 100MB

Figure 21 Response time for MQPut/MQGet pairs for NP message between 2MB and 100MB.

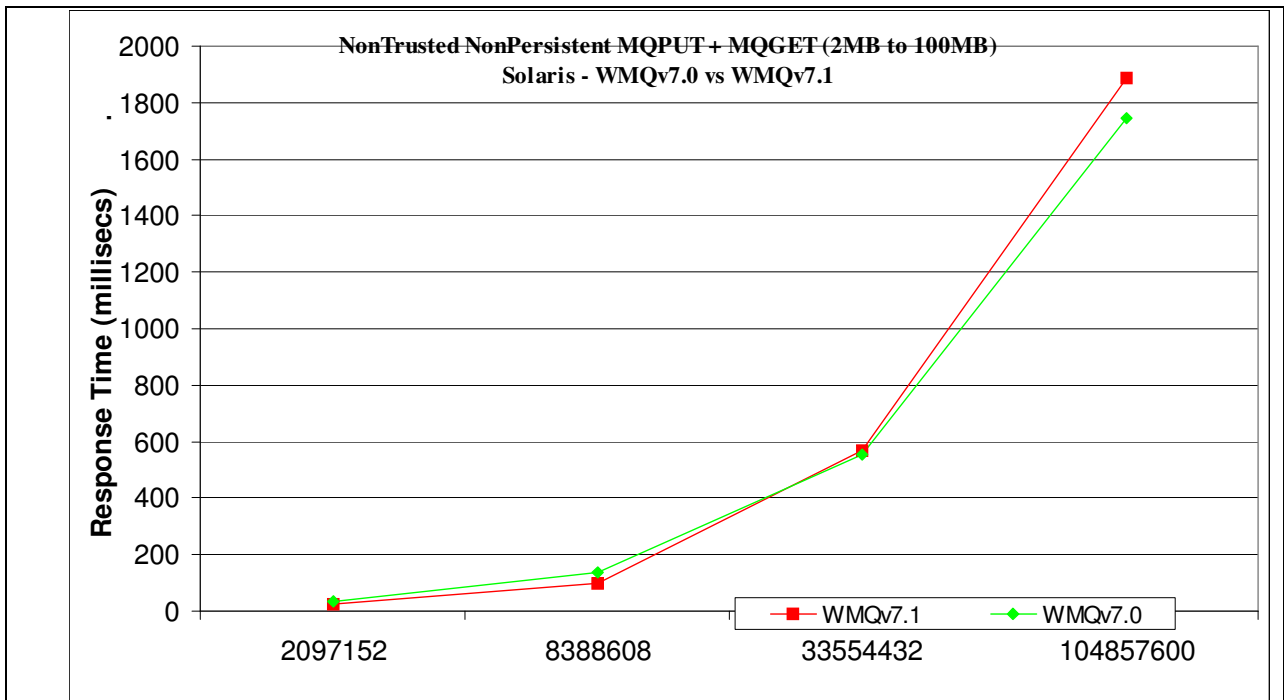


Figure 21 –The effect of non-persistent message size on MQI response time (2MB – 100MB)

Figure 22 The response for MQPut/MQGet pairs for persistent message sizes between 2MB and 100MB.

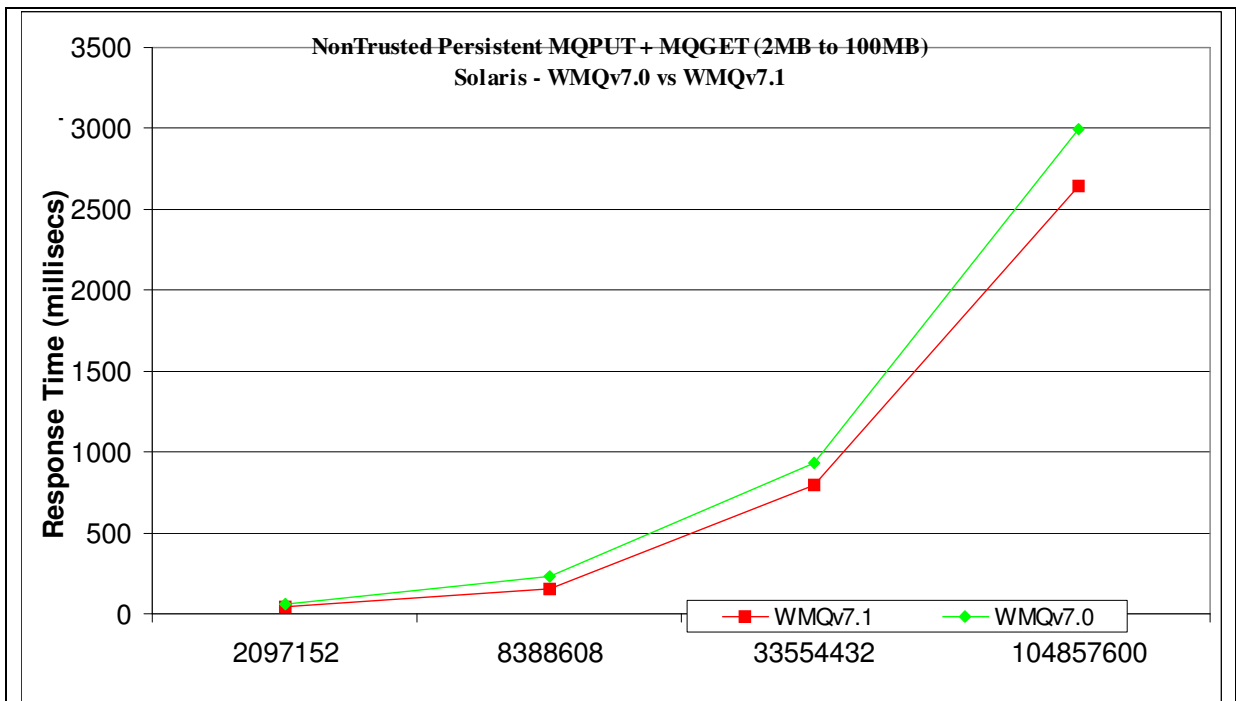


Figure 22 –The effect of persistent message size on MQI response time (2MB – 100MB)

3.2 20KB Messages

3.2.1 Local Queue Manager

Figure 23 and Figure 24 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

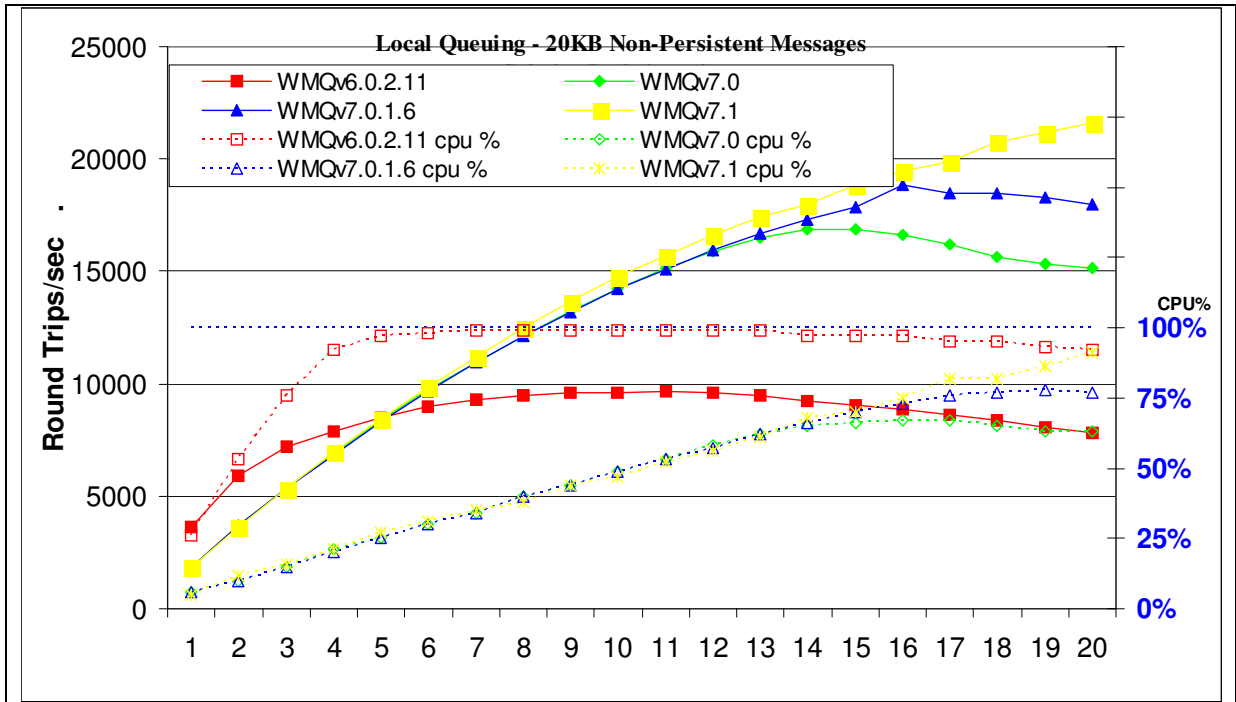


Figure 23 – 20KB non-persistent messages, local queue manager

Figure 23 and Table 12 shows that the peak throughput of non-persistent messages has increased by 15% when comparing version 7.1 to 7.0.1.6 and by 124% 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	11	9659	0.0014	99%
WMQv7.0	15	16872	0.001	66%
WMQv7.0.1.6	16	18839	0.00099	73%
WMQv7.1	20	21638	0.00092	91%

Table 12 – 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

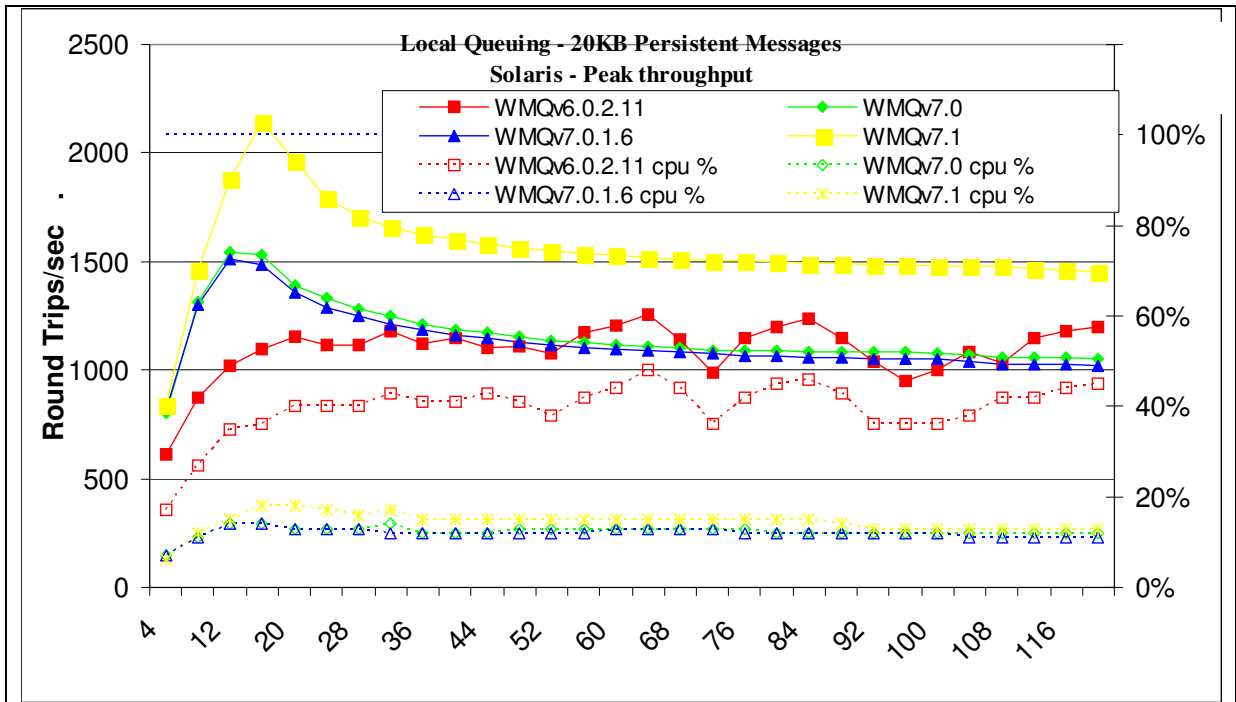


Figure 24 – 20KB persistent messages, local queue manager

Figure 24 and Table 13 shows that the peak throughput of persistent messages has increased by 41% when comparing version 7.1 to 7.0.1.6 and by 70% 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	64	1257	0.062	48%
WMQv7.0	12	1542	0.0089	14%
WMQv7.0.1.6	12	1514	0.0092	14%
WMQv7.1	16	2140	0.0086	18%

Table 13 – 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 25 and Figure 26 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 25 – 20KB non-persistent messages, client channels

Figure 26 and Table 14 shows that the peak throughput of non-persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 and has increased by 34% 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	11	2732	0.0047	68%
WMQv7.0	20	3433	0.0066	24%
WMQv7.0.1.6	20	3596	0.0058	24%
WMQv7.1	20	3655	0.0059	23%

Table 14 – 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 26 – 20KB persistent messages, client channels

Figure 27 and Table 15 shows that the peak throughput of persistent messages has increased by 40% when comparing version 7.1 to 7.0.1.6 and by 72% 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	112	1109	0.12	66%
WMQv7.0	20	1374	0.015	20%
WMQv7.0.1.6	20	1365	0.015	20%
WMQv7.1	20	1908	0.011	25%

Table 15 – 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 27 and Figure 28 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 27 – 20KB non-persistent messages, distributed queuing

Figure 27 and Table 16 shows that the peak throughput of non-persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 and has increased by 56% 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	8	3007	0.0031	66%
WMQv7.0	14	4622	0.0032	24%
WMQv7.0.1.6	15	4632	0.0038	24%
WMQv7.1	15	4699	0.0031	23%

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

3.2.3.2 Persistent Messages




Figure 28 – 20KB persistent messages, distributed queuing

Figure 29 and Table 17 shows that the peak throughput of persistent messages has increased by 35% when comparing version 7.1 to 7.0.1.6 and by 87% 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	32	754	0.054	32%
WMQv7.0	24	1085	0.024	11%
WMQv7.0.1.6	24	1043	0.024	11%
WMQv7.1	120	1408	0.092	14%

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

3.3 200K Messages

3.3.1 Local Queue Manager

Figure 29 and Figure 30 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

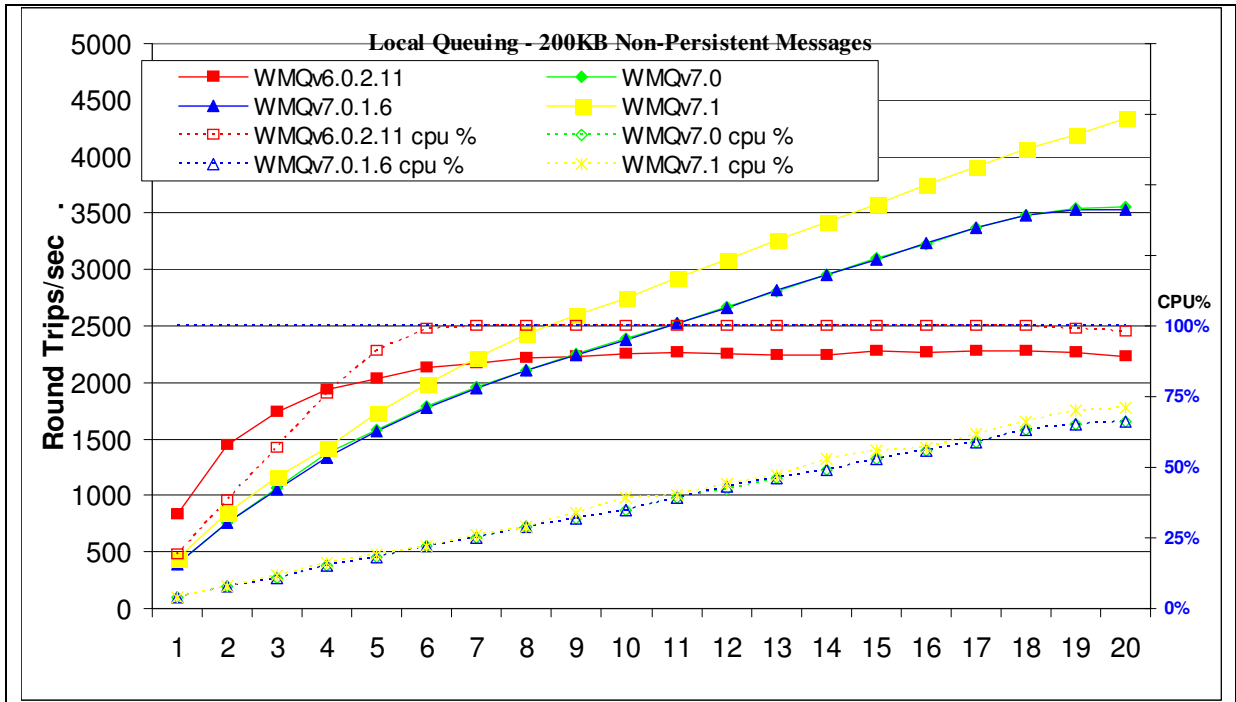


Figure 29 – 200KB non-persistent messages, local queue manager

Figure 30 and Table 18 shows that the peak throughput of non-persistent messages has increased by 23% when comparing version 7.1 to 7.0.1.6 and by 90% 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	17	2286	0.009	100%
WMQv7.0	20	3554	0.0067	66%
WMQv7.0.1.6	20	3535	0.0067	66%
WMQv7.1	20	4334	0.0041	71%

Table 18 – 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 30 – 200KB persistent messages, local queue manager

Figure 31 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 152% 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	88	147	0.7	16%
WMQv7.0	108	257	0.49	8%
WMQv7.0.1.6	60	259	0.27	8%
WMQv7.1	92	371	0.29	12%

Table 19 – 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 31 and Figure 32 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

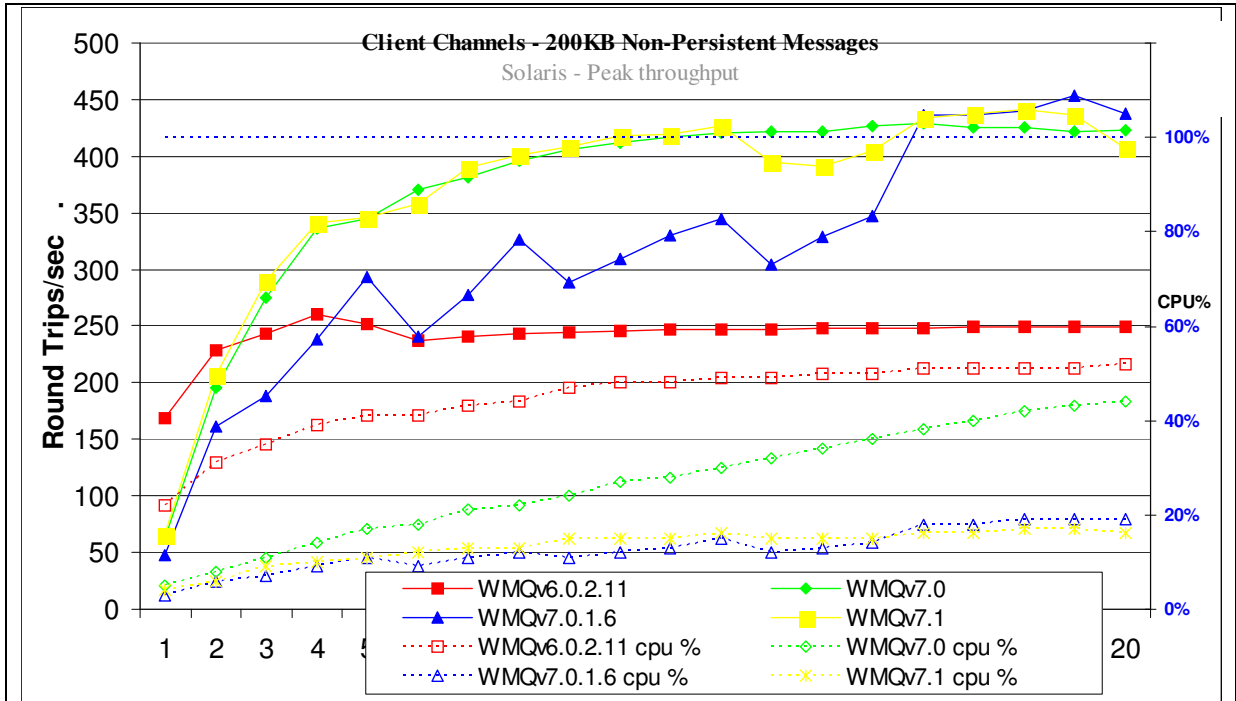


Figure 31 – 200KB non-persistent messages, client channels

Figure 32 and Table 20 shows that the peak throughput of non-persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 and has increased by 69% 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	4	261	0.017	39%
WMQv7.0	16	429	0.039	38%
WMQv7.0.1.6	19	454	0.046	19%
WMQv7.1	18	441	0.043	17%

Table 20 – 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 32 – 200KB persistent messages, client channels

Figure 33 and Table 21 shows that the peak throughput of persistent messages has increased by 49% when comparing version 7.1 to 7.0.1.6 and by 157% 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	108	137	0.97	42%
WMQv7.0	56	222	0.27	35%
WMQv7.0.1.6	24	237	0.1	15%
WMQv7.1	68	352	0.2	23%

Table 21 – 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 33 and Figure 34 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 33 – 200KB non-persistent messages, distributed queuing

Figure 34 and Table 22 shows that the peak throughput of non-persistent messages is unchanged when comparing version 7.1 to 7.0.1.6 but has increased by 25% 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	7	300	0.027	46%
WMQv7.0	20	389	0.055	25%
WMQv7.0.1.6	15	401	0.04	16%
WMQv7.1	7	376	0.019	12%

Table 22 – 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 34 – 200KB persistent messages, distributed queuing

Figure 35 and Table 23 shows that the peak throughput of persistent messages has increased by 74% when comparing version 7.1 to 7.0.1.6 and by over 200% 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	76	84	1.1	20%
WMQv7.0	20	158	0.13	8%
WMQv7.0.1.6	16	154	0.11	8%
WMQv7.1	120	268	0.48	13%

Table 23 – 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 35 and Figure 36 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 35 – 2MB non-persistent messages, local queue manager

Figure 36 and Table 24 shows that the peak throughput of non-persistent messages has increased by over 400% when comparing version 7.1 to 7.0.1.6 and by over 200% 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	12	164	0.086	68%
WMQv7.0	15	90	0.19	14%
WMQv7.0.1.6	20	88	0.27	14%
WMQv7.1	20	495	0.033	65%

Table 24 – 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 36 – 2MB persistent messages, local queue manager

Figure 37 and Table 25 shows that the peak throughput of persistent messages has increased by 28% when comparing version 7.1 to 7.0.1.6 and by 53% 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	16	15	1.2	14%
WMQv7.0	28	17	1.8	5%
WMQv7.0.1.6	32	18	2	5%
WMQv7.1	8	23	0.35	6%

Table 25 – 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 37 and Figure 38 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 37 – 2MB non-persistent messages, client channels

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 but has increased by 23% 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	4	26	0.17	42%
WMQv7.0	18	45	0.43	46%
WMQv7.0.1.6	20	46	0.47	21%
WMQv7.1	16	45	0.38	17%

Table 26 – 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 38 – 2MB persistent messages, client channels

Figure 39 and Table 27 shows that the peak throughput of persistent messages has increased by 64% when comparing version 7.1 to 7.0.1.6 and by 180% 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	16	13	1.4	32%
WMQv7.0	24	22	1.1	24%
WMQv7.0.1.6	28	22	1.3	14%
WMQv7.1	40	36	1.1	20%

Table 27 – 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 39 and Figure 40 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 39 – 2MB non-persistent messages, distributed queuing

Figure 40 and Table 28 shows that the peak throughput of non-persistent messages is unchanged 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: Distributed Queuing - 2MB Non Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
WMQv6.0.2.11	12	25	0.56	51%
WMQv7.0	16	38	0.45	27%
WMQv7.0.1.6	16	39	0.44	18%
WMQv7.1	15	37	0.44	14%

Table 28 – 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 40 - 2MB persistent messages, distributed queuing

Figure 41 and Table 29 shows that the throughput of persistent messages has increased by over 200% when comparing version 7.1 to 7.0.1.6 and by over 200% 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	12	8	1.6	19%
WMQv7.0	20	12	1.8	8%
WMQv7.0.1.6	8	10	0.86	6%
WMQv7.1	44	31	1.5	15%

Table 29 – 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	Non Trusted
Isolated	Isolated	Isolated
Trusted	Trusted	Trusted
Non Trusted	Shared	Shared

4.1 Local Queue Manager

Figure 41 and Figure 42 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

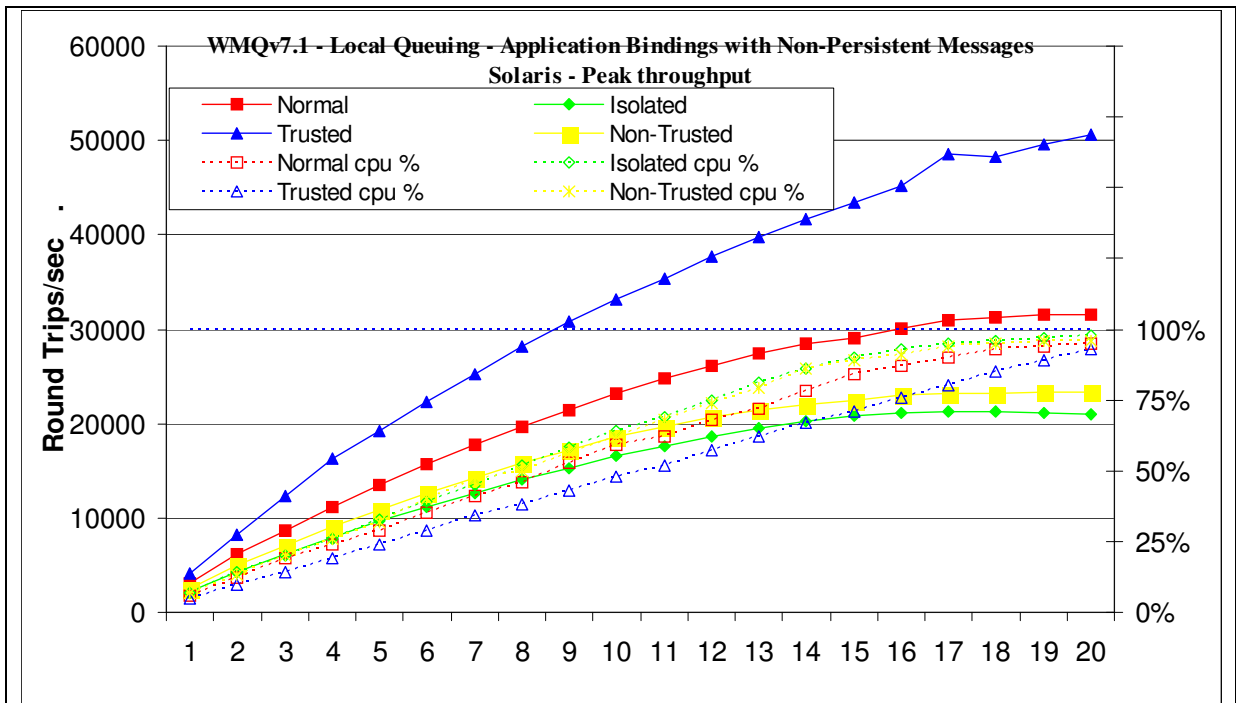


Figure 41 – Application binding, non-persistent messages, local queue manager

Figure 41 and Table 30 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	20	31487	0.00069	95%
Isolated	17	21205	0.00089	95%
Trusted	20	50547	0.00042	93%
Non-Trusted	19	23339	0.00093	96%

Table 30 – Application binding, non-persistent messages, local queue manager

4.1.2 Persistent Messages

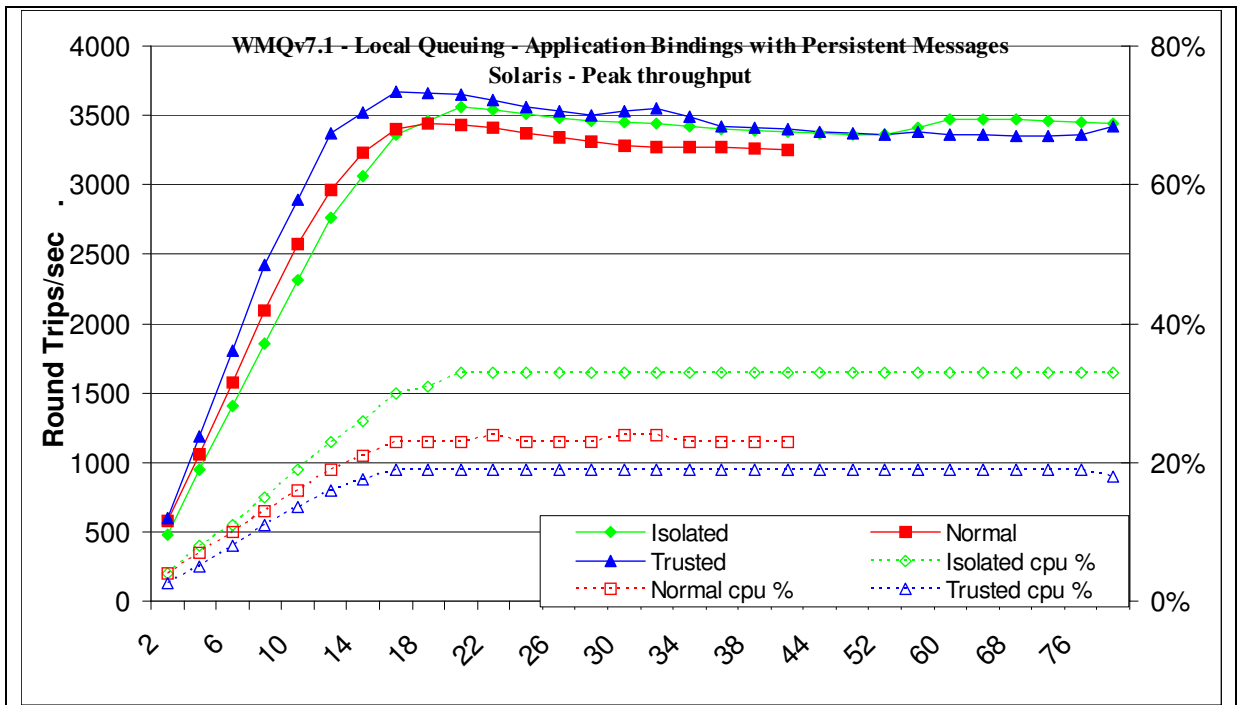


Figure 42 – Application binding, persistent messages, local queue manager

Figure 42 and Table 31 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
Isolated	20	3558	0.0066	33%
Normal	18	3441	0.0063	23%
Trusted	16	3672	0.0052	19%

Table 31 – Application binding, persistent messages, local queue manager

4.2 Client Channels

Figure 43 and Figure 44 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 43 – Application binding, non-persistent messages, client channels

Figure 43 and Table 32 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	40	12699	0.0033	64%
Isolated	40	12450	0.0034	67%
Trusted	20	10215	0.0021	33%

Table 32 – Application binding, non-persistent messages, client channels

4.2.2 Persistent Messages




Figure 44 – Application binding, persistent messages, client channels

Figure 44 and Table 33 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	88	3301	0.028	39%
Isolated	24	3299	0.0077	34%
Trusted	24	3468	0.0073	29%

Table 33 – Application binding, persistent messages, client channels

4.3 Distributed Queuing

Figure 44 and Figure 45 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

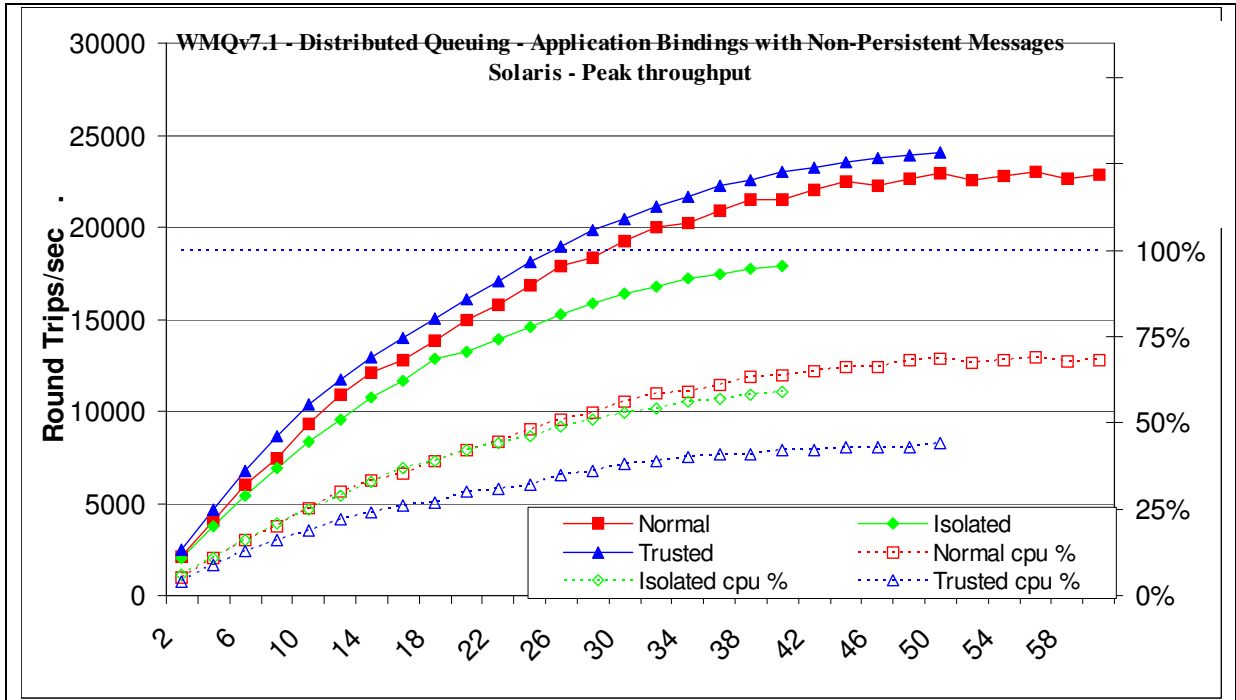


Figure 45 – Application binding, non-persistent messages, distributed queuing

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 - Distributed Queuing - Application Bindings with Non-Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	56	23004	0.0027	69%
Isolated	40	17892	0.0026	59%
Trusted	50	24074	0.0022	44%

Table 34 – 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 46 – Application binding, persistent messages, distributed queuing

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 - Distributed Queuing - Application Bindings with Persistent Messages	Apps	Round Trips/Sec	Response time (s)	CPU
Normal	200	3582	0.059	24%
Isolated	180	3852	0.051	24%
Trusted	170	3869	0.047	17%

Table 35 – 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

The previous chapters in this report only reported on steady state messaging that does not include any session setup and termination function. This chapter specifically bracket groups of five MQPut/MQGet pairs with MQConn/MQDisc and MQOpen/MQClose calls so a comparison of this overhead can be seen.

A short session is a term used to describe the behaviour of an MQI application as it processes a small number of messages using one or more queues and a queue manager. The measurements in this document use an MQI-client application and the following sequence:

- connects to the queue manager
- opens the common input queue, and common reply queue
- puts a request message to the common input queue
- gets the reply message from the common reply queue
- wait one second
- closes both queues
- disconnects from the queue manager



“Why measure short sessions?”

For each new connecting application or disconnecting application, the queue manager and Operating System must start a new process or thread and set up the new connection. As the number of connecting and disconnecting applications increases, the Operating System and queue manager are subjected to a higher load. While these requests are being serviced, the queue manager has less time available to process messages, so fewer driving applications can be reconnected to the queue manager per second before the response time exceeds one second.

This effect is greater than that of reducing the total messaging throughput of the queue manager by connecting thousands of MQI applications to the queue manager (refer to **Figure 47** for an illustration).




Figure 47 – Short sessions, client channels

Test Name: Client Channels Scenario	Apps	Round Trips/Sec	Response time (s)	CPU
Client Persistent	1750	1637	0.89	71%
Client Non-Persistent	2000	1796	0.93	58%

Note: Messaging in these tests is 1 round trip per driving application per second, i.e. 1 short session per driving application every 5 seconds

Note: The figures for non-persistent short sessions were generated with all message processing within sync-point control. All other non-persistent messages within this report were generated outside sync-point control.

The 'runmqtsr' has a much smaller overhead of connecting to and disconnecting from the queue manager because it only uses a single thread per connection rather than an entire process. INETD listener has a significantly smaller capacity because of the need to create a new process for every client.

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 . MQIBINDTYPE=FASTPATH

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
clnp	40	n/a*	12699	0.003	64%
clnp_r3600	4400	3600	4359	0.18	25%
clnp_c6000	6000	4080	6739	0.13	33%
		6960	7102	0.70	37%
clnp max	20000	690	3833	0.144	23%
	20000	840	4669	0.925	28%
	30000	190	1597	0.042	3%
cl_persist_c6000	6000	1060	3510	0.57	20%
clnp_cmax_no_correllid	15000	60	250	0.04	7%
	21000	60	357	0.6	11%
cl_persist_cmax	25000	120	616	0.059	20%

Table 36 – Capacity measurements, client channels

* *There was no delay between the response to the previous message and the insertion of the next message with 40 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 AMQRMPPA 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 Trips/sec	Response time (s)	CPU
dqnp	104	n/a*	36545	0.0033	54%
dqnp_r3600	3450	3600	3428	0.259	84%
dqnp_q1000	1000	72000	14546	0.0125	62%
dqnp_qmax	12000	1680	5554	0.2065	94%
dq-persist-qmax	4000	570	630	0.92	12%
	6000	320	524	0.672	12%
dq-persist_q1000	1000	9500	5260	0.029	23%

Table 37 – Capacity measurements, server channels

* *There was no delay between the response to the previous message and the insertion of the next message with 104 driving applications..*

The dqnp and dqnp_r3600 both used a total of 4 pairs of Sender/Receiver pairs of channels between queue managers while the dqnp_qmax and dq_persist_q4000 used a pair of channels per application. The dqnp_q1000 shows the reduced 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 archiving 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 the 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 data-streaming 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 data-streaming.
- 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 ('runmqslr', and for server channels an MCATYPE of 'THREAD') the threaded listener 'runmqslr' can now be used in all scenarios with client and server channels. Additional resource savings are available using the 'runmqslr' 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 managers 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 'runmqtsr' 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 Tuning the Operating System (Solaris)

For Solaris 9 and Solaris 10 these values are required (refer to info centre for implementation techniques). For more information on Solaris, use the following url
<http://publib.boulder.ibm.com/infocenter/wmqv7/v7r1/index.jsp>

Minimum Values

```
shmmax=8294967295
shmmni=1024
semmni=1024
semaem=16384
semvmx=32767
semmns=16384
semmsl=100
semopm=100
semmnu=16384
semume=256
shmseg=1024
lim_fd_cur=10000
lim_fd_max=10000
```

7.4 Virtual Memory, Real Memory, & Paging

7.4.1 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 than 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.

7.4.2 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

Sun T2000:	Server system (Device under test)
Model:	T201108C-64GA2G SFT2000
Processor:	1.4GHz
Architecture:	8 core
Memory (RAM):	64GB
Disk:	Internal disks – not used for measurements (Chapter 1-9)
SAN:	Fibre channel HBA PCI-X 4Gb FC Dual Port HBA (chapter 10) 2 SAN disks on DS8700 (5GB each, 1 queue, 1 log)
Network:	1Gb Ethernet Adapter
Sun V490:	Driver system (2 machines)
Model:	SUNW, Sun-Fire-V490
Processor:	SPARCv9 @ 1.5GHz
Architecture:	4 CPU
Memory (RAM):	32GB
Disk:	Internal disks – not used for measurements
SAN:	Fibre channel HBA PCI-X 4Gb FC Dual Port HBA (chapter 10) 2 SAN disks on DS8700 (5GB each, 1 queue, 1 log)
Network:	1Gb Ethernet Adapter

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 which is connected to the DCX at 4Gb.

8.3 Software

Operating system:	SunOS 5.10
MQSeries:	Version 6.0.2.11, Version 7.0, Version 7.0.1.6 and Version 7.1

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><i>local</i> => local queue manager test scenario</p> <p><i>cl</i> => client channel test scenario</p> <p><i>dq</i> => distributed queuing test scenario</p> <p><i>np</i> => non-persistent messages</p> <p><i>pm</i> => persistent messages</p> <p><i>r3600</i> => 1 round trip per driving application per second</p> <p><i>runmqslr</i> => channels using the 'runmqslr' listener (client channel test scenario, in addition to 'runmqchi' for distributed queuing test scenarios)</p> <p><i>c6000</i> => 6,000 client driving applications (i.e. 6,000 MQI-client connections)</p> <p><i>q1000</i> => 1,000 server channel pairs</p> <p><i>max</i> => maximum number of channels (or channel pairs)</p> <p><i>no_correl_id</i> => 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