

# **WebSphere MQ for Solaris V7 - Performance Evaluations**

Version 1.1

December 2008

Peter Toghil , Craig Stirling .

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.

**Second Edition, December 2008.**

This edition applies to *WebSphere MQ for Solaris V7* (and to all subsequent releases and modifications until otherwise indicated in new editions).

© Copyright International Business Machines Corporation 2008. 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*. 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.

**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

This SupportPac is designed for people who:

- Will be designing and implementing solutions using WebSphere MQ for Solaris.
- Want to understand the performance limits of WebSphere MQ for Solaris V7.
- Want to understand what actions may be taken to tune WebSphere MQ for Solaris.

The reader should have a general awareness of the Solaris operating system and of MQSeries in order to make best use of this SupportPac. Readers should read the section '**How this document is arranged**' on **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?

Please direct any comments of this nature to **WMQPG@uk.ibm.com**.

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 2K (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 were used as the Drivers

# How this document is arranged

## **Performance Headlines**

### **Pages: 2-19**

Chapter 2 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**

Chapter 3 contains performance measurements for *large messages*. This includes *MQI response times* of 50byte to 2Mb messages. It also includes *20K*, *200k* and *2M* messages using the same scenarios as for the “Performance Headlines”.

## **Application Bindings**

### **Pages: 42-47**

Chapter 4 contains performance measurements for *'trusted, normal, and isolated'* server applications, using the same three scenarios as for the “Performance Headlines”.

## **Short & Long Sessions**

### **Pages: 48-49**

Chapter 5 contains

## **Performance and Capacity Limits**

### **Pages: 50-52**

Chapter 6 contains

## **Tuning Recommendations**

### **Pages: 53-57**

Chapter 7 contains ....

## **Measurement Environment**

### **Pages: 58-59**

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

## **Compiler Glossary**

### **Pages: 60-60**

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

# CONTENTS

<b>1</b>	<b>Overview</b>	<b>1</b>
<b>2</b>	<b>Performance Headlines</b>	<b>2</b>
2.1	Local Queue Manager Test Scenario	2
2.1.1	Nonpersistent Messages – Local Queue Manager	3
2.1.2	Nonpersistent 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	Nonpersistent Messages – Client Channels	7
2.2.2	Nonpersistent Messages – Non Trusted Client Channels	8
2.2.3	Persistent Messages – Client Channels	9
2.2.4	Client Channels	11
2.3	Distributed Queuing Test Scenario	13
2.3.1	Nonpersistent 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
<b>3</b>	<b>Large Messages</b>	<b>20</b>
3.1	MQI Response Times: 50bytes to 100Mb – Local Queue Manager	20
3.1.1	50bytes to 32KB	20
3.1.2	32KB to 2Mb	22
3.1.3	2Mb to 100Mb	23
3.2	20K 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
<b>4</b>	<b>Application Bindings</b>	<b>42</b>
4.1	Local Queue Manager	42
4.1.1	Nonpersistent Messages	42
4.1.2	Persistent Messages	43
4.2	Client Channels	44
4.2.1	Nonpersistent Messages	44
4.2.2	Persistent Messages	45
4.3	Distributed Queuing	46
4.3.1	Nonpersistent Messages	46
4.3.2	Persistent Messages	47
<b>5</b>	<b>Short &amp; Long Sessions</b>	<b>48</b>
<b>6</b>	<b>Performance and Capacity Limits</b>	<b>50</b>
6.1	Client channels – capacity measurements	50
6.2	Distributed queuing – capacity measurements	51
<b>7</b>	<b>Tuning Recommendations</b>	<b>53</b>
7.1	Tuning the Queue Manager	53
7.1.1	Queue Disk, Log Disk, and Message Persistence	53
7.1.2	Log Buffer Size, Log File Size, and Number of Log Extents	53
7.1.3	Channels: Process or Thread, Standard or Fastpath?	55
7.1.4	Multiplexed clients	55
7.2	Applications: Design and Configuration	55
7.2.1	Standard (Shared or Isolated) or Fastpath?	55
7.2.2	Parallelism, Batching, and Triggering	55
7.3	Virtual Memory, Real Memory, & Paging	56
7.3.1	Queue Manager	56
7.3.2	Channels	56
7.3.3	Client Channels	56

7.3.4	Server – Server Channels .....	57
7.3.5	Reply Queue.....	57
7.3.6	BufferLength .....	57
7.3.7	MQIBINDTYPE .....	57
<b>8</b>	<b>Measurement Environment .....</b>	<b>58</b>
8.1	Workload description .....	58
8.1.1	MQI response time tool.....	58
8.1.2	Test scenario workload.....	58
8.2	Hardware .....	59
8.3	Software .....	59
<b>9</b>	<b>Compiler Glossary .....</b>	<b>60</b>



# TABLES

Table 1 – Performance headline, nonpersistent messages, local queue manager .....	3
Table 2 – Performance headline, non-persistent messages, non-trusted bindings, local queue manager .	4
Table 3 – Performance headline, persistent messages, local queue manager .....	5
Table 4 – Performance headline, nonpersistent messages, client channels .....	7
Table 5 – Performance headline, non-persistent messages, non-trusted client channels .....	8
Table 6 – Performance headline, persistent messages, client channels .....	10
Table 7 – 1 round trip per driving application per second, client channels .....	12
Table 8 – Performance headline, nonpersistent messages, server channels .....	14
Table 9 – Performance headline, nonpersistent messages, server channels .....	15
Table 10 – Performance headline, persistent messages, server channels .....	16
Table 11 – 1 round trip per driving application per second, client channels .....	18
Table 12 – effect of SSL Cipher Specifications on Throughput with 40 Connections .....	19
Table 13 – 20K nonpersistent messages, local queue manager .....	24
Table 14 – 20K persistent messages, local queue manager .....	25
Table 15 – 20K nonpersistent messages, client channels .....	26
Table 16 – 20K persistent messages, client channels .....	27
Table 17 – 20K nonpersistent messages, client channels .....	28
Table 18 – 20K persistent messages, client channels .....	29
Table 19 – 200K nonpersistent messages, local queue manager .....	30
Table 20 – 200K persistent messages, local queue manager .....	31
Table 21 – 200K nonpersistent messages, client channels .....	32
Table 22 – 200K persistent messages, client channels .....	33
Table 23 – 200K nonpersistent messages, distributed queuing .....	34
Table 24 – 200K persistent messages, distributed queuing .....	35
Table 25 – 2M nonpersistent messages, local queue manager .....	36
Table 26 – 2M persistent messages, local queue manager .....	37
Table 27 – 2M nonpersistent messages, client channels.....	38
Table 28 – 2M persistent messages, client channels .....	39
Table 29 – 2M nonpersistent messages, distributed queuing .....	40
Table 30 – 2M persistent messages, distributed queuing .....	41
Table 31 – Application binding, nonpersistent messages, local queue manager .....	42
Table 32 – Application binding, persistent messages, local queue manager .....	43
Table 33 – Application binding, nonpersistent messages, client channels .....	44
Table 34 – Application binding, persistent messages, client channels .....	45
Table 35 – Application binding, nonpersistent messages, distributed queuing .....	46
Table 36 – Application binding, persistent messages, distributed queuing .....	47
Table 37 – Short sessions, client channels.....	49
Table 38 - Client Capacity, throughput .....	50
Table 39 – Client capacity, memory utilisation.....	51
Table 40 – Capacity measurements, server channels .....	51
Table 41 – DQ capacity, memory utilisation.....	51

## FIGURES

Figure 1 – Connections into a local queue manager .....	2
Figure 2 – Performance headline, nonpersistent messages, local queue manager .....	3
Figure 3 – Performance headline, non-persistent messages, non-trusted bindings, local queue manager .....	4
Figure 4 – Performance headline, persistent messages, local queue manager .....	5
Figure 5 – MQI-client channels into a remote queue manager .....	6
Figure 6 – Performance headline, nonpersistent messages, client channels .....	7
Figure 7 – Performance headline, nonpersistent messages, non-trusted client channels .....	8
Figure 8 – Performance headline, persistent messages, client channels .....	9
Figure 9 – Performance headline, persistent messages, client channels, optimized .....	9
Figure 10 – 1 round trip per driving application per second, client channels, nonpersistent messages ..	11
Figure 11 – 1 round trip per driving application per second, client channels, persistent messages .....	11
Figure 12 – Server channels between two queue managers .....	13
Figure 13 – Performance headline, nonpersistent messages, server channels .....	14
Figure 14 – Performance headline, nonpersistent messages, non-trusted, server channels .....	15
Figure 15 – Performance headline, persistent messages, server channels .....	16
Figure 16 – 1 round trip per driving application per second, server channel, nonpersistent messages ..	17
Figure 17 – 1 round trip per driving application per second, server channel, persistent messages .....	17
Figure 18 – effect of SSL Cipher Specifications on Throughput .....	18
Figure 19 –The effect of nonpersistent message size on MQI response time (50byte - 32K) .....	20
Figure 20 –The effect of persistent message size on MQI response time (50byte - 32K) .....	21
Figure 21 –The effect of nonpersistent message size on MQI response time (32K – 2Mb) .....	22
Figure 22 –The effect of persistent message size on MQI response time (32K – 2Mb) .....	22
Figure 23 –The effect of nonpersistent 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 – 20K nonpersistent messages, local queue manager .....	24
Figure 26 – 20K persistent messages, local queue manager .....	25
Figure 27 – 20K nonpersistent messages, client channels .....	26
Figure 28 – 20K persistent messages, client channels .....	27
Figure 29 – 20K nonpersistent messages, distributed queuing .....	28
Figure 30 – 20K persistent messages, distributed queuing .....	29
Figure 31 – 200K nonpersistent messages, local queue manager .....	30
Figure 32 – 200K persistent messages, local queue manager .....	31
Figure 33 – 200K nonpersistent messages, client channels .....	32
Figure 34 – 200K persistent messages, client channels .....	33
Figure 35 – 200K nonpersistent messages, distributed queuing .....	34
Figure 36 – 200K persistent messages, distributed queuing .....	35
Figure 37 – 2M nonpersistent messages, local queue manager .....	36
Figure 38 – 2M persistent messages, local queue manager .....	37
Figure 39 – 2M nonpersistent messages, client channels .....	38
Figure 40 – 2M persistent messages, client channels .....	39
Figure 41 – 2M nonpersistent messages, distributed queuing .....	40
Figure 42 – 2M persistent messages, distributed queuing .....	41
Figure 43 – Application binding, nonpersistent messages, local queue manager .....	42
Figure 44 – Application binding, persistent messages, local queue manager .....	43
Figure 45 – Application binding, nonpersistent messages, client channels .....	44
Figure 46 – Application binding, persistent messages, client channels .....	45
Figure 47 – Application binding, nonpersistent messages, distributed queuing .....	46
Figure 48 – Application binding, persistent messages, distributed queuing .....	47
Figure 49 – Short sessions, client channels .....	48

# 1 Overview

WebSphere MQ V7 on Solaris has similar performance characteristics to the V6 product. The comparisons in this report show that throughput is similar overall (for Local, Client and Distributed Queuing) when the Clients are running in V6 compatibility mode (see section “Multiplexed clients” on page 55.). The default enhanced client support that provides Heartbeating, enhanced reliability, and multiplexing degrades Client benchmarks by 8%.

There are new functions in V7 that provide enhanced performance to applications that are able to use them and they include Asynchronous Puts, Read-ahead, Properties, and selectors but they are not covered in this document.

## 2 Performance Headlines

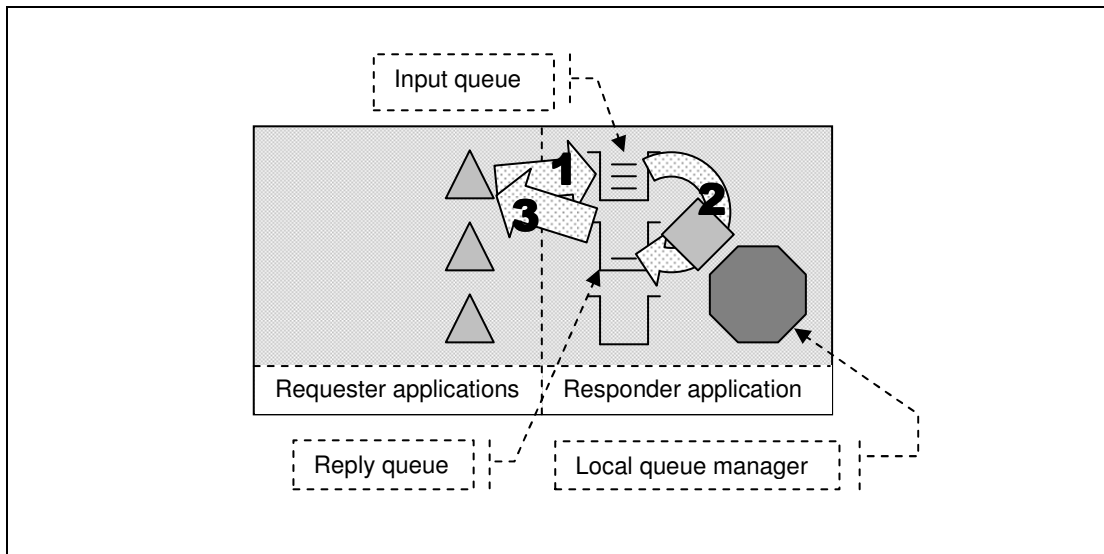
The measurements for the local queue manager scenario are for processing messages with no *think-time*. For the client channel scenario and distributed queuing scenario, there are also measurements for *rated* messaging.

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 are automatically stopped 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.

Nonpersistent and persistent messages were used in the local queue manager tests, with a message size of 2K. 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 Nonpersistent Messages – Local Queue Manager

Figure 2, Figure 3 and Figure 4 show the nonpersistent 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), and WebSphere MQ V7 compared to Version 6.

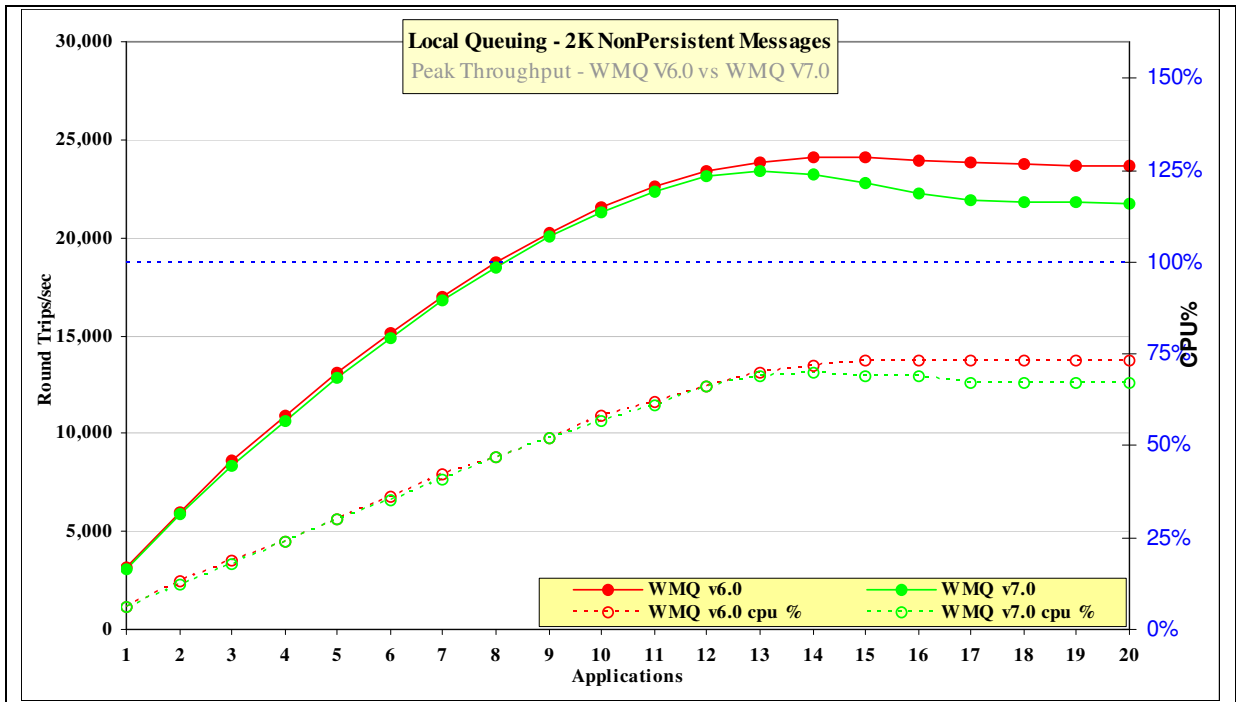


Figure 2 – Performance headline, nonpersistent messages, local queue manager

Figure 2 and Table 1 shows that the throughput of nonpersistent messages has degraded by 3% comparing Version 6 to Version 7

Test name: <b>local_np</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>14</b>	<b>24069</b>	<b>0.0006</b>	<b>72%</b>
WebSphere MQ V7	<b>13</b>	<b>23402</b>	<b>0.0007</b>	<b>69%</b>

Table 1 – Performance headline, nonpersistent messages, local queue manager

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.

### 2.1.2 Nonpersistent Messages – Non trusted – Local Queue Manager

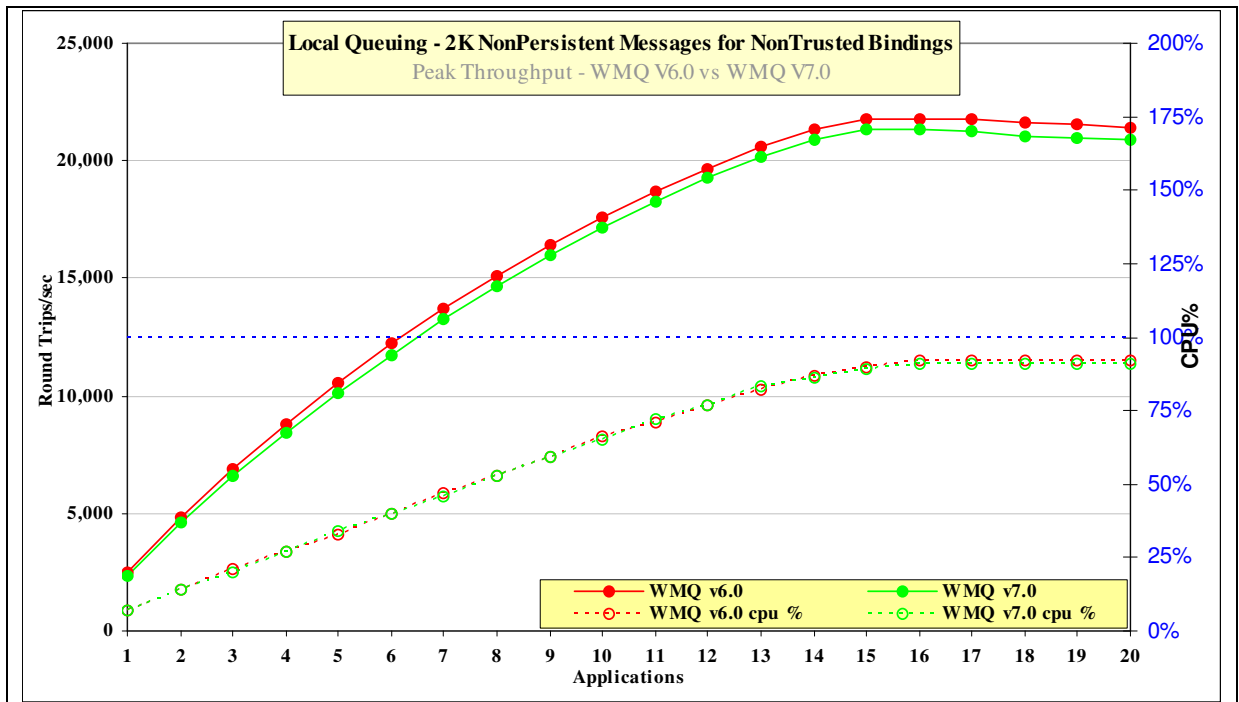


Figure 3 – Performance headline, non-persistent messages, non-trusted bindings, local queue manager

Test name: <b>local_np_nt</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	16	21806	0.0008	92%
WebSphere MQ V7	16	21344	0.0009	91%

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

Figure 3 and Table 2 show that the throughput of nonpersistent messages when the Requester and Responder both use Shared bindings has degraded by 3% comparing Version 6 to Version 7.

### 2.1.3 Persistent Messages – Local Queue Manager

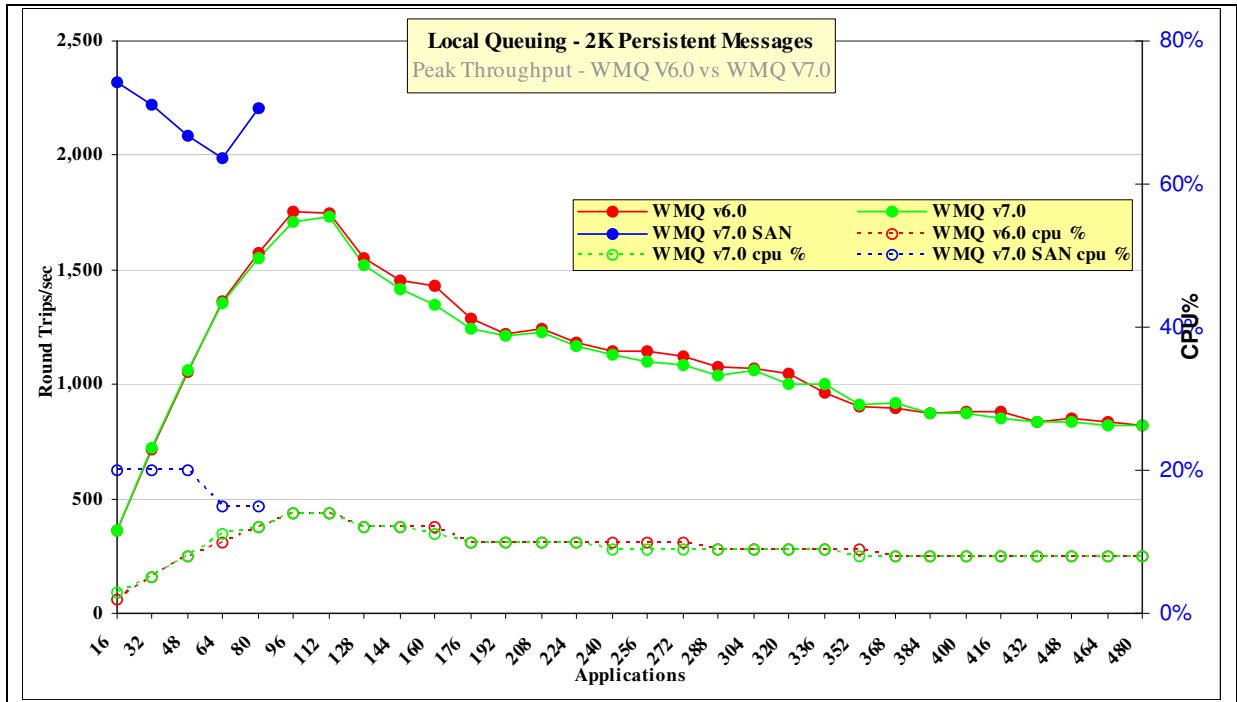


Figure 4 – Performance headline, persistent messages, local queue manager


Figure 4 and Table 3 show that the throughput of persistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>local_pm</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>96</b>	<b>1752</b>	<b>0.063</b>	<b>14%</b>
WebSphere MQ V7	<b>112</b> (96)	<b>1731</b> (1708)	<b>0.074</b> (0.063)	<b>14%</b> (14%)
WebSphere MQ V7 SAN	<b>16</b>	<b>2321</b>	<b>0.008</b>	<b>20%</b>

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

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. The numbers in brackets are included in the table to provide meaningful comparison between WebSphere MQ V6 and Version 7.

## 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.

Nonpersistent and persistent messages were used in the client channel tests, with a message size of 2K. 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 will multiplex multiple clients from the same process over one TCP socket. The version 6 behavior where each client had its own TCP socket can be set by specifying Sharecnv(0) on the client channel definition and is shown in the charts as ‘optimized’. Further information in section 7.1.4



### 2.2.1 Nonpersistent Messages – Client Channels

Figure 6, Figure 7, Figure 8 and Figure 9 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario (see Figure 5 on the previous page), and WebSphere MQ V7 compared to Version 6.

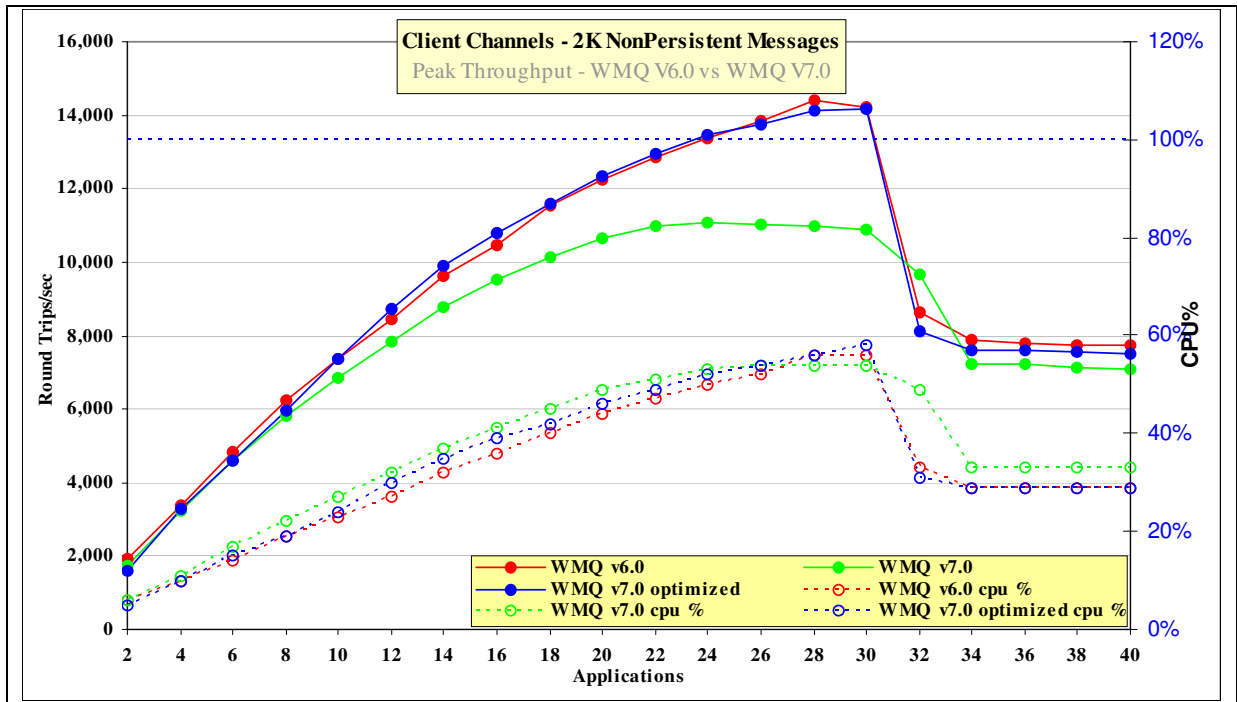


Figure 6 – Performance headline, nonpersistent messages, client channels

Figure 6 and Table 4 show that the throughput of nonpersistent messages is similar with optimised setup and degraded by 10% with default setup when comparing Version 6 to Version 7.

Test name: cInp	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>28</b>	<b>14395</b>	<b>0.002</b>	<b>56%</b>
WebSphere MQ V7 Optimised	<b>30</b>	<b>14161</b>	<b>0.002</b>	<b>58%</b>
WebSphere MQ V7	<b>24</b> (30)	<b>11069</b> (10895)	<b>0.002</b> (0.003)	<b>53%</b> (54%)

Table 4 – Performance headline, nonpersistent messages, client channels

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. The numbers in brackets are included in the table to provide meaningful comparison between WebSphere MQ V6 and Version 7.

### 2.2.2 Nonpersistent Messages – Non Trusted Client Channels

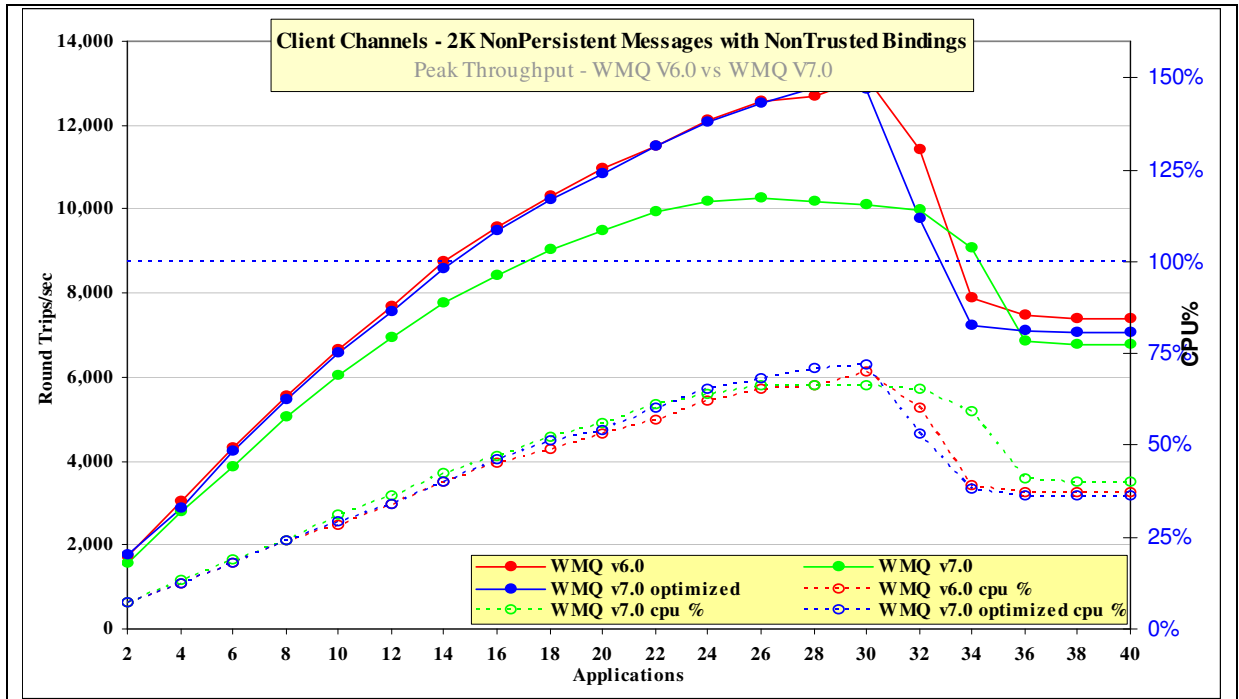


Figure 7 – Performance headline, nonpersistent messages, non-trusted client channels

Test name: Clnp_nt	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	30	13152	0.02	70%
WebSphere MQ V7 Optimised	28	12880	0.002	71%
WebSphere MQ V7	26	10247	0.003	66%

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

The throughput of nonpersistent messages when the channel has used the default MQIBINDTYPE=NORMAL has degraded by 3% with optimised setup and degraded by 11% with default setup when comparing Version 6 to Version 7.

### 2.2.3 Persistent Messages – Client Channels

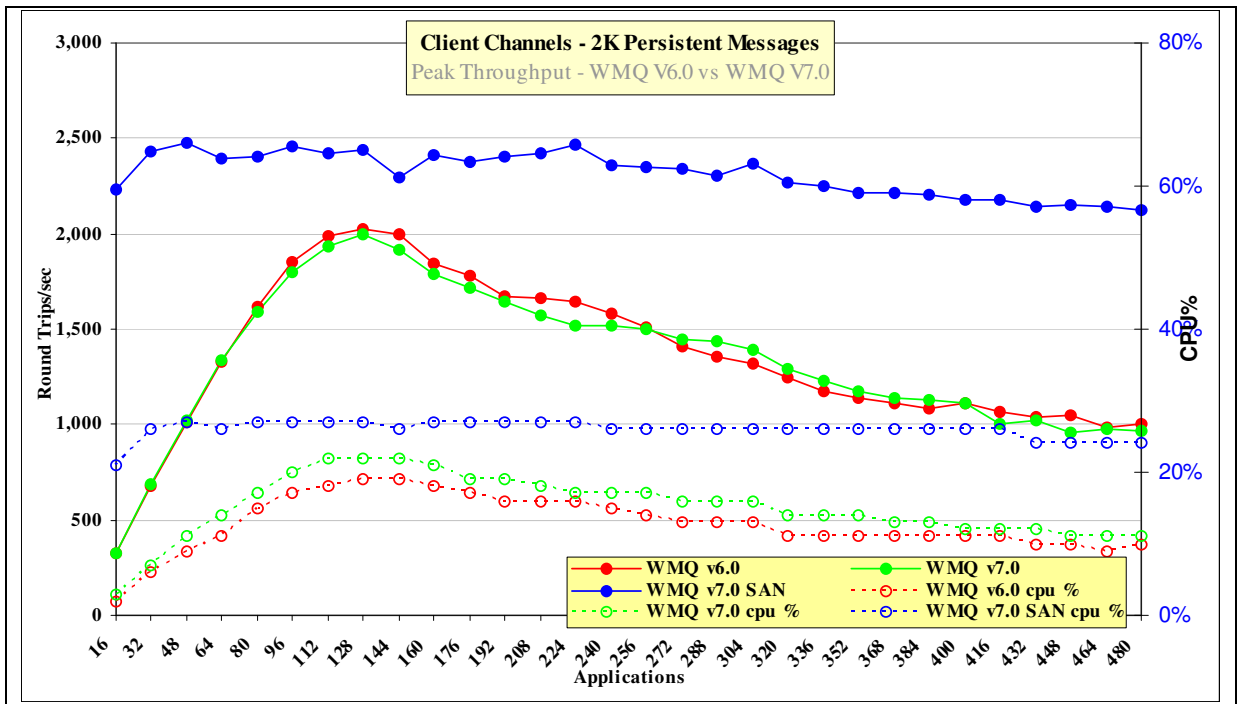


Figure 8 – Performance headline, persistent messages, client channels

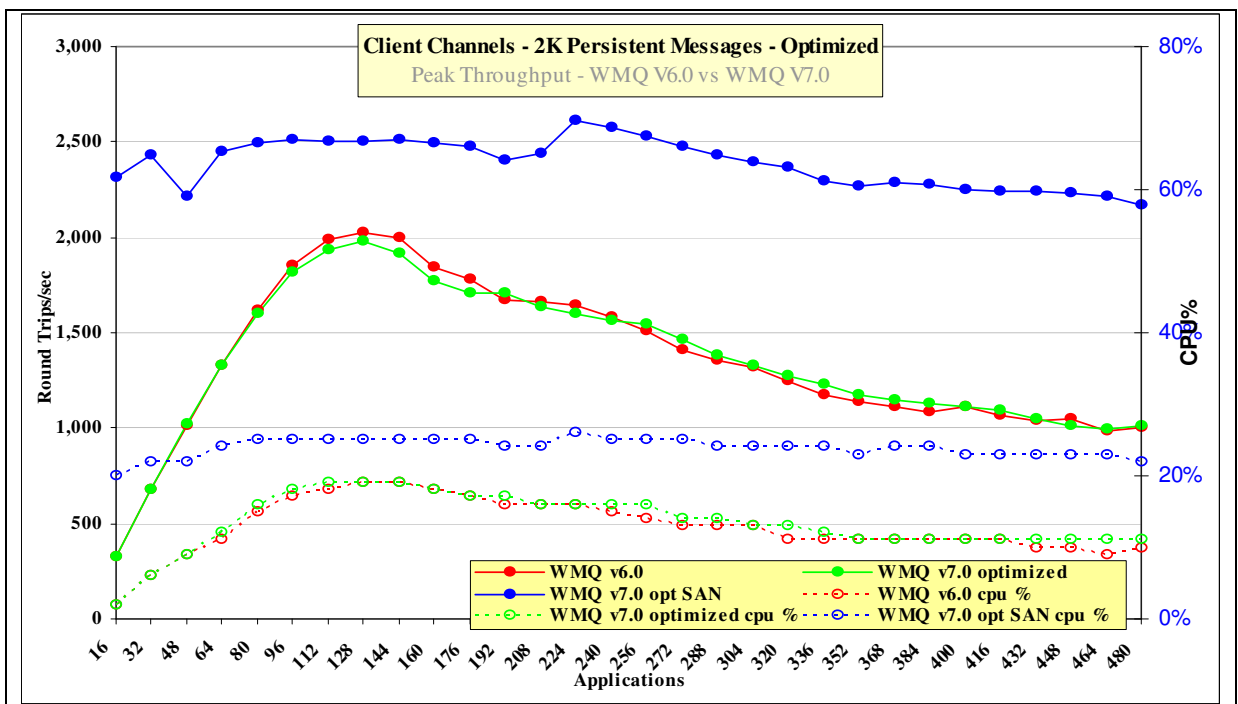


Figure 9 – Performance headline, persistent messages, client channels, optimized

Figure 8, Figure 9 and Table 6 show that the throughput of persistent messages is similar when using optimized setup and is similar using default setup when comparing Version 6 to Version 7.

Test name: <b>Clpm</b>	<b>Apps</b>	<b>Round Trips/sec</b>	<b>Response time (s)</b>	<b>CPU</b>
WebSphere MQ V6	<b>128</b>	<b>2026</b>	<b>0.068</b>	<b>19%</b>
WebSphere MQ V7 Optimised	<b>128</b>	<b>1978</b>	<b>0.067</b>	<b>19%</b>
WebSphere MQ V7 Optimised SAN	<b>128</b>	<b>2502</b>	<b>0.054</b>	<b>25%</b>
WebSphere MQ V7	<b>128</b>	<b>1994</b>	<b>0.067</b>	<b>22%</b>
WebSphere MQ V7 SAN	<b>128</b>	<b>2435</b>	<b>0.057</b>	<b>27%</b>

**Table 6 – Performance headline, persistent messages, client channels**

*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.*

### 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.

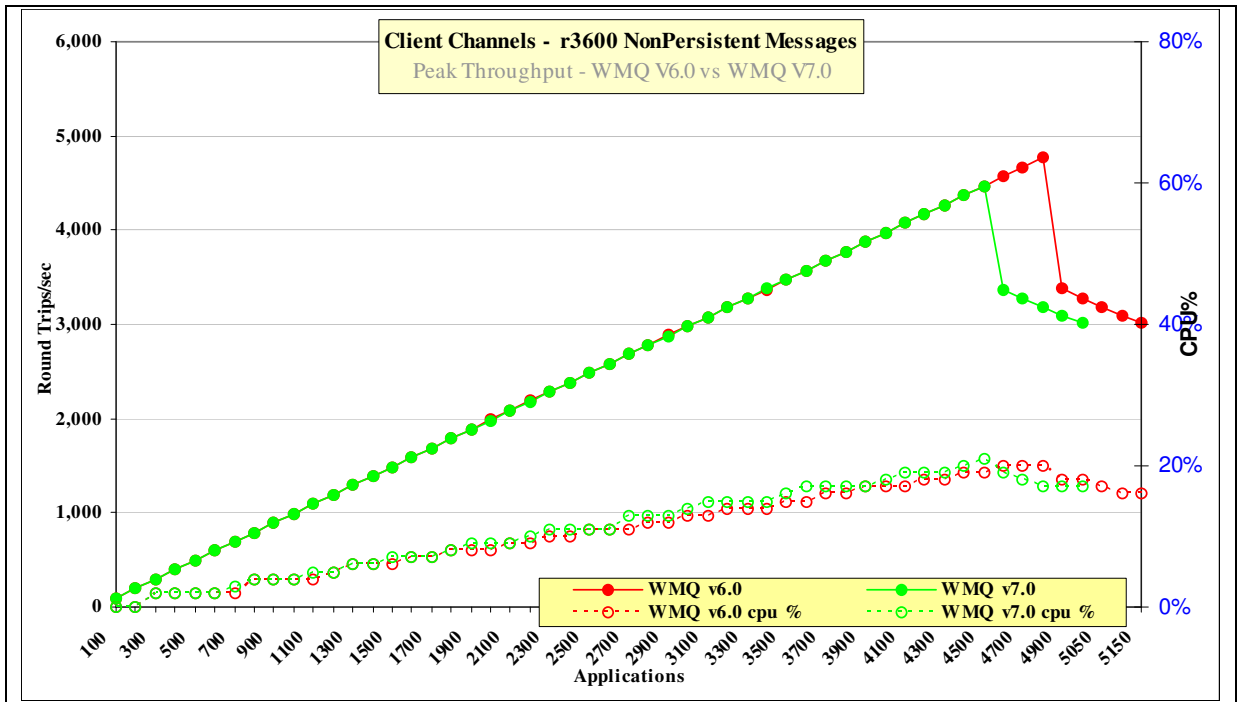


Figure 10 – 1 round trip per driving application per second, client channels, nonpersistent messages

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

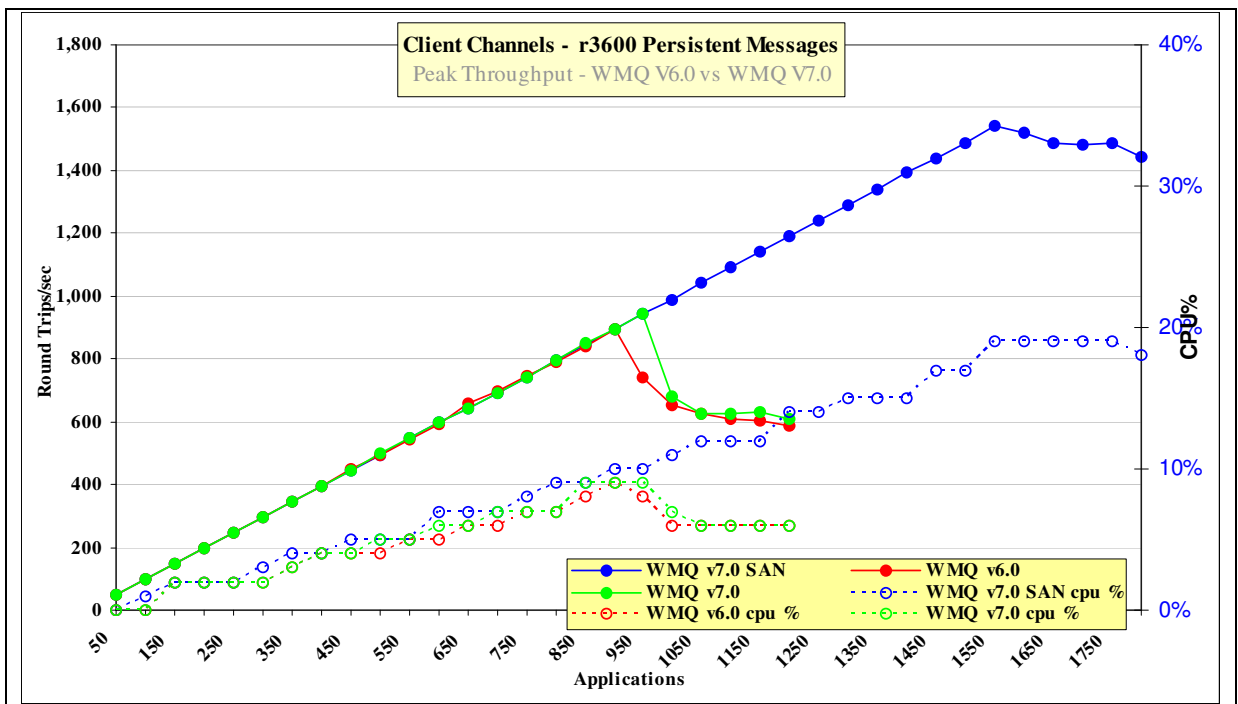


Figure 11 – 1 round trip per driving application per second, client channels, persistent messages

Figure 10, Figure 11 and Table 7 shows that WebSphere MQ V7 has similar maximum throughput with WebSphere MQ V 6.0.

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
<b>clnp_r3600</b> (WebSphereMQ v6.0)	<b>4500</b> (4800)	3,600	<b>4469</b> (4766)	<b>0.022</b> (0.027)	<b>21%</b> (20%)
<b>clpm_r3600</b> (WebSphereMQ v6.0)	<b>950</b> (900)	3,600	<b>942</b> (894)	<b>0.104</b> (0.067)	<b>5%</b> (5%)
<b>clpm_r3600 SAN</b>	<b>1500</b>	2,600	<b>1488</b>	<b>0.093</b>	<b>19%</b>

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

*Note: The large bold numbers in the table above show the WebSphere MQ V6 number of round trips per second, and the number of driving applications used to achieve the peak throughput. The numbers in brackets are included in the table to provide meaningful comparison with Version 6.*

## 2.3 Distributed Queuing Test Scenario

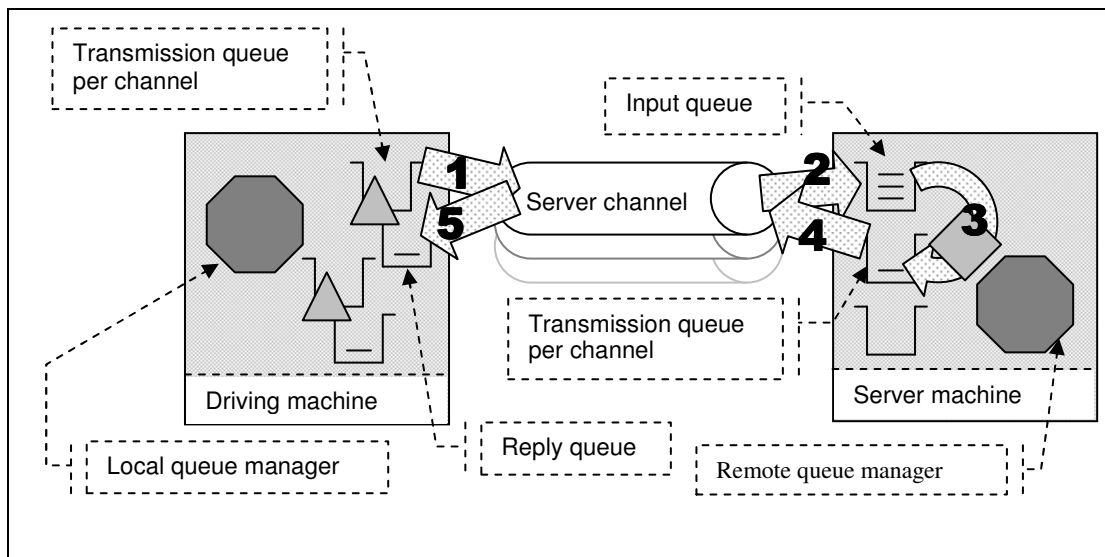


Figure 12 – Server channels between two queue managers

- 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

Nonpersistent and persistent messages were used in the distributed queuing tests, with a message size of 2K. The effect of message throughput with larger messages sizes is investigated 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 Nonpersistent Messages – Server Channels

Figure 13, Figure 14 and Figure 15 show the nonpersistent 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 V6 compared to Version 7.

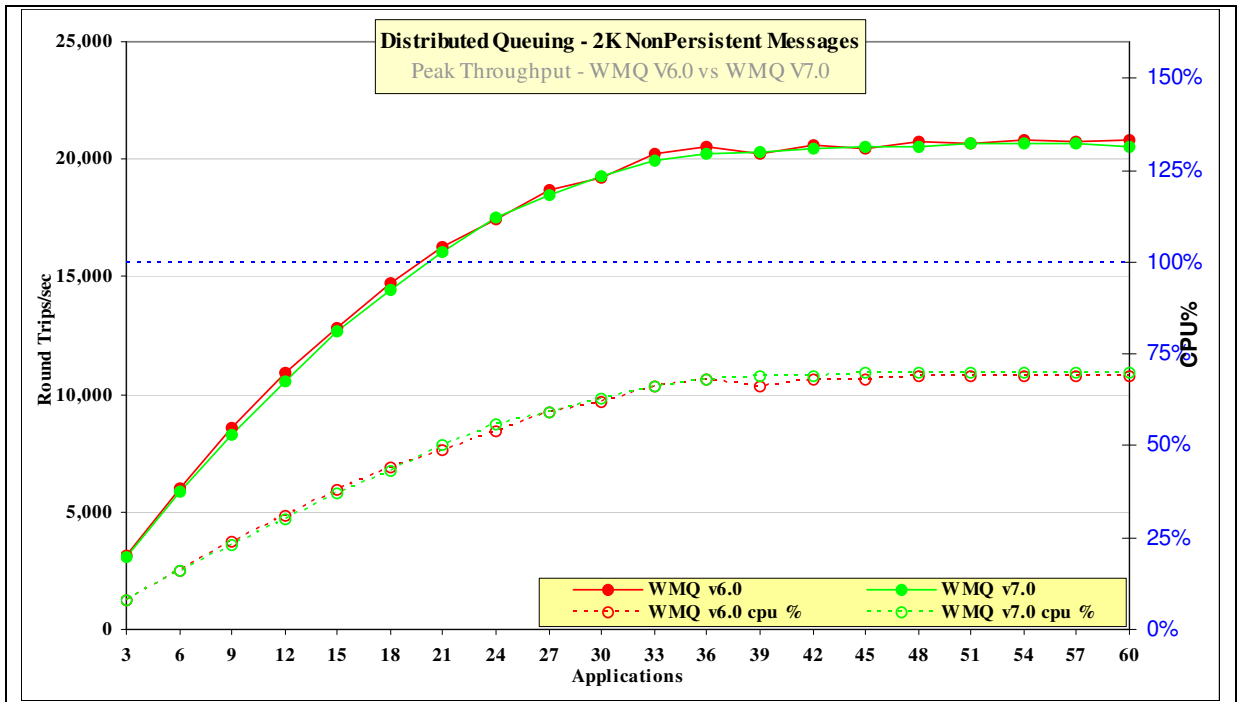


Figure 13 – Performance headline, nonpersistent messages, server channels

Figure 13 and Table 8 show that the throughput of nonpersistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>dqnp</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>54</b>	<b>20809</b>	<b>0.003</b>	<b>69%</b>
WebSphere MQ V7	<b>57</b>	<b>20704</b>	<b>0.003</b>	<b>70%</b>

Table 8 – Performance headline, nonpersistent messages, server channels

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



### 2.3.2 Non Persistent non Trusted – Server Channels

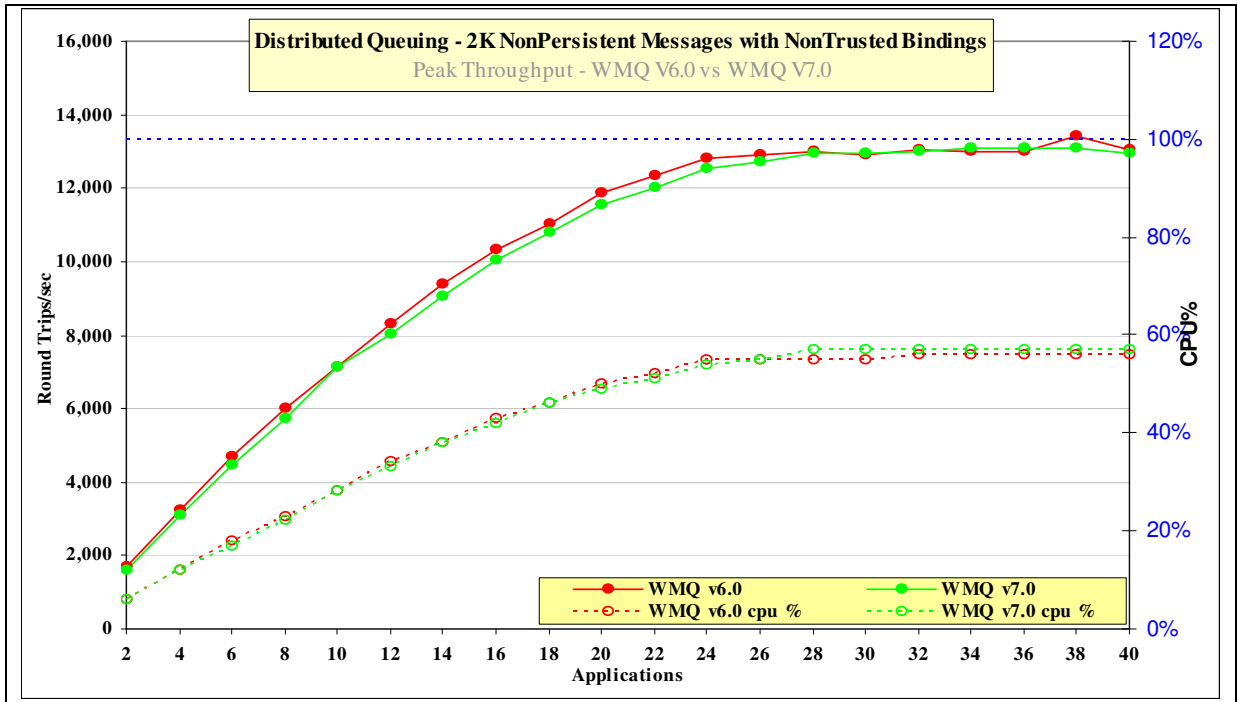


Figure 14 – Performance headline, nonpersistent messages, non-trusted, server channels

Test name: <b>Dqnp_nt</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	38	13426	0.003	56%
WebSphere MQ V7	38	13107	0.003	57%

Table 9 – Performance headline, nonpersistent messages, server channels

The throughput of nonpersistent messages has improved by 3% when comparing Version 6 to Version 7.

### 2.3.3 Persistent Messages – Server Channels

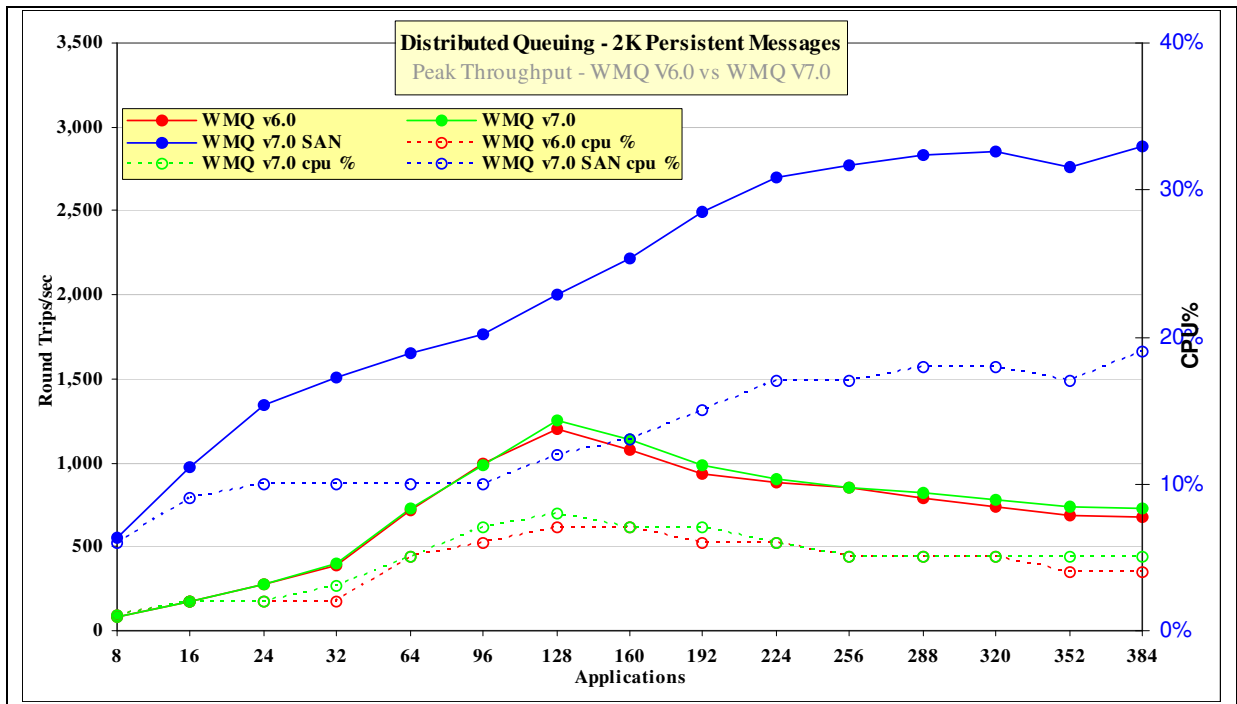


Figure 15 – Performance headline, persistent messages, server channels

Note: Messaging in these tests is with no think-time

Figure 15 and Table 10 show that the throughput of persistent messages has increased by 3% when comparing Version 6 to Version 7.

Test name: dqpm	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>128</b>	<b>1208</b>	<b>0.108</b>	<b>7%</b>
WebSphere MQ V7	<b>128</b>	<b>1249</b>	<b>0.108</b>	<b>8%</b>
WebSphere MQ V7 SAN	(128) <b>384</b>	(2001) <b>2883</b>	(0.072) <b>0.139</b>	(12%) <b>19%</b>

Table 10 – Performance headline, persistent messages, server channels

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

### 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 nonpersistent messaging tests, and 2 pairs for the persistent message tests.

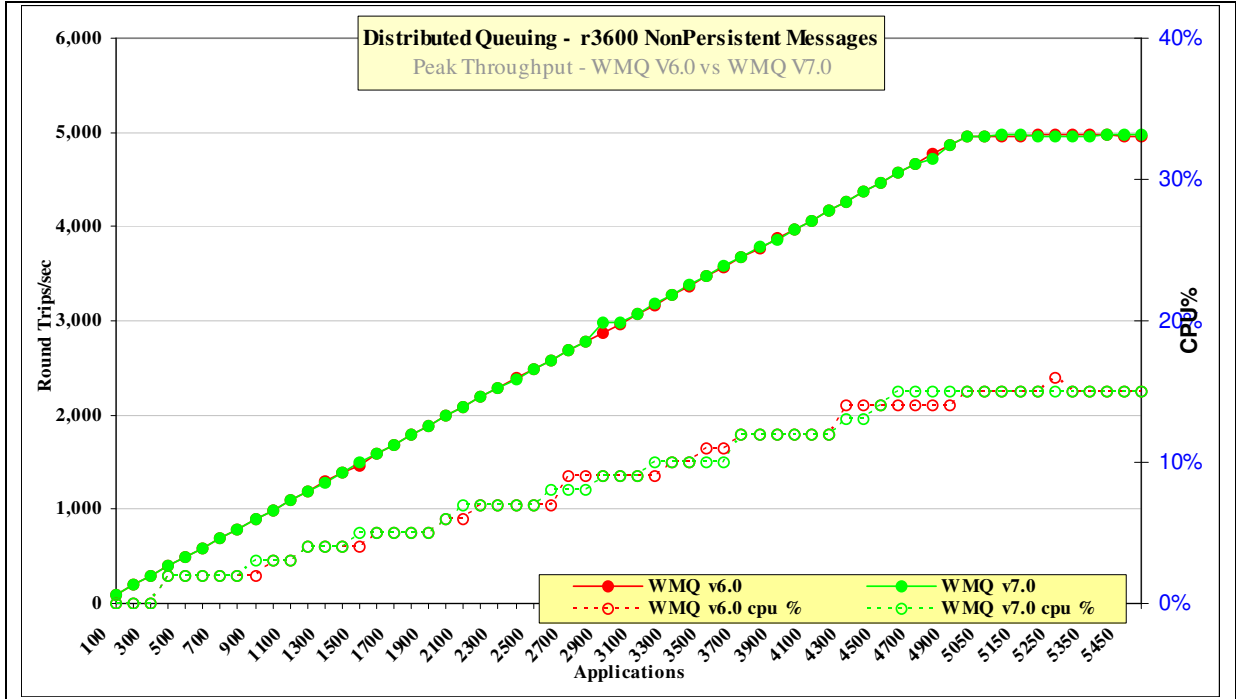


Figure 16 – 1 round trip per driving application per second, server channel, nonpersistent messages

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

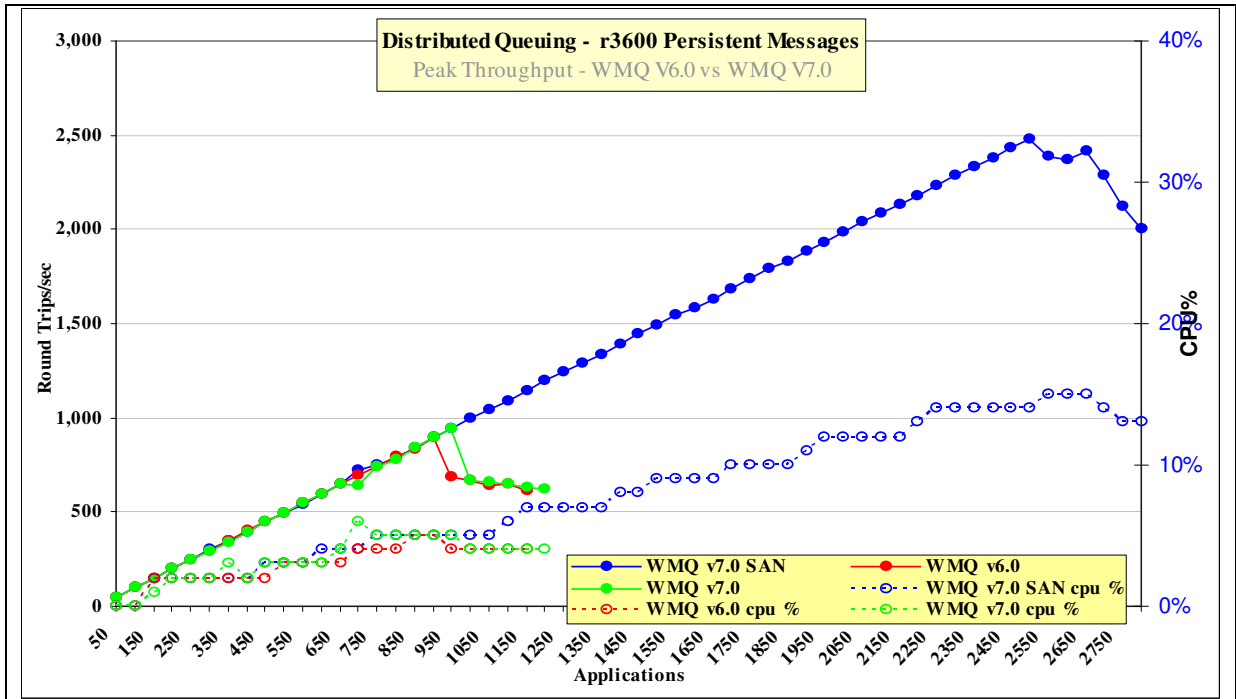


Figure 17 – 1 round trip per driving application per second, server channel, persistent messages

Figure 16, Figure 17 and Table 11 show how WebSphere MQ V7 is similar in performance as version 6.0 although at a higher CPU cost.

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
<b>dqnp_r3600</b> (WebSphereMQ v6.0)	<b>5000</b> (5000)	3,600	<b>4959</b> (4967)	<b>0.016</b> (0.011)	<b>17%</b> (15)
<b>dqpm_r3600</b> (WebSphereMQ v6.0)	<b>950</b> (900)	3,600	<b>943</b> (893)	<b>0.133</b> (0..155)	<b>5%</b> (5%)
<b>dqpm_r3600 SAN</b>	<b>2500</b>	3,600	<b>2482</b>	<b>0.154</b>	<b>14%</b>

Table 11 – 1 round trip per driving application per second, client channels

Note: The large bold numbers in the table above show the WebSphere MQ V7 peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The numbers in brackets are included in the table to provide meaningful comparison with Version 6.

### 2.3.5 SSL

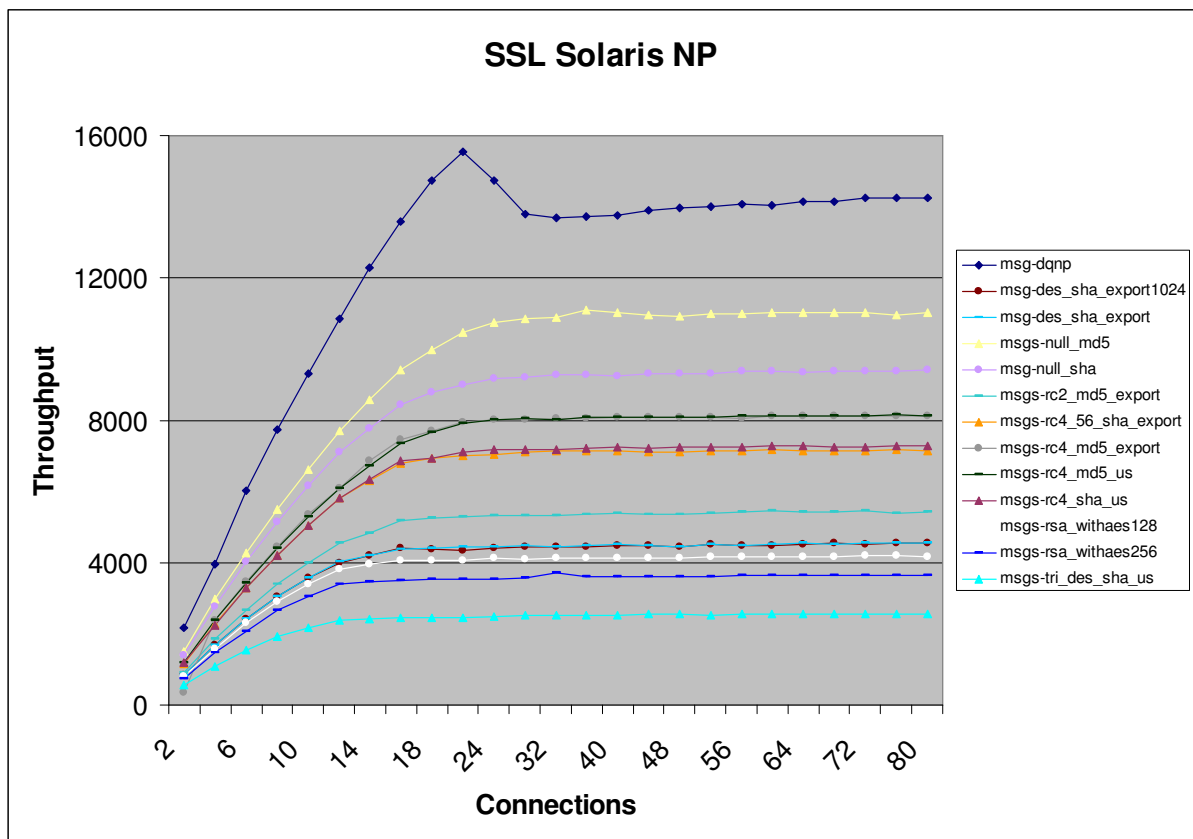


Figure 18 – effect of SSL Cipher Specifications on Throughput

There was no SSL used in the dqnp measurement. The use of cipher specifications null\_sha and null\_md5 (which do not encrypt the messages) degrade throughput by 20% and 35% respectively. The other 10 cipher specs degrade throughput by up to 80%. Table 12 shows the throughput, response time and server cpu% with 40 connections.

cipher spec	throughput	Response time	Server cpu%
DQNP (no SSL)	13775	0.003 seconds	40%
triple_des_sha_us	2532	0.016 seconds	28%
tls_rsa_with_aes_256_cbc_sha	3596	0.012 seconds	30%
rc4_sha_us	7233	0.006 seconds	37%
rc4_md5_us	8086	0.005 seconds	39%
rc4_md5_export	8094	0.005 seconds	39%
rc4_56_sha_export1024	7138	0.006 seconds	37%
rc2_md5_export	5385	0.008 seconds	33%
null_sha	9552	0.005 seconds	40%

null_md5	11023	0.004 seconds	44%
des_sha_export	4517	0.009 seconds	31%
des_sha_export1024	4475	0.010 seconds	31%
tls_rsa_with_aes_128_cbc_sha	4127	0.010 seconds	30%

**Table 12 – effect of SSL Cipher Specifications on Throughput with 40 Connections**

The GSKIT level is 7.0.4.20

### 3 Large Messages

#### 3.1 MQI Response Times: 50bytes to 100Mb – Local Queue Manager

##### 3.1.1 50bytes to 32KB

Figure 19 shows the response time for MQPut/MQGet for nonpersistent message sizes between 50bytes and 32KB.

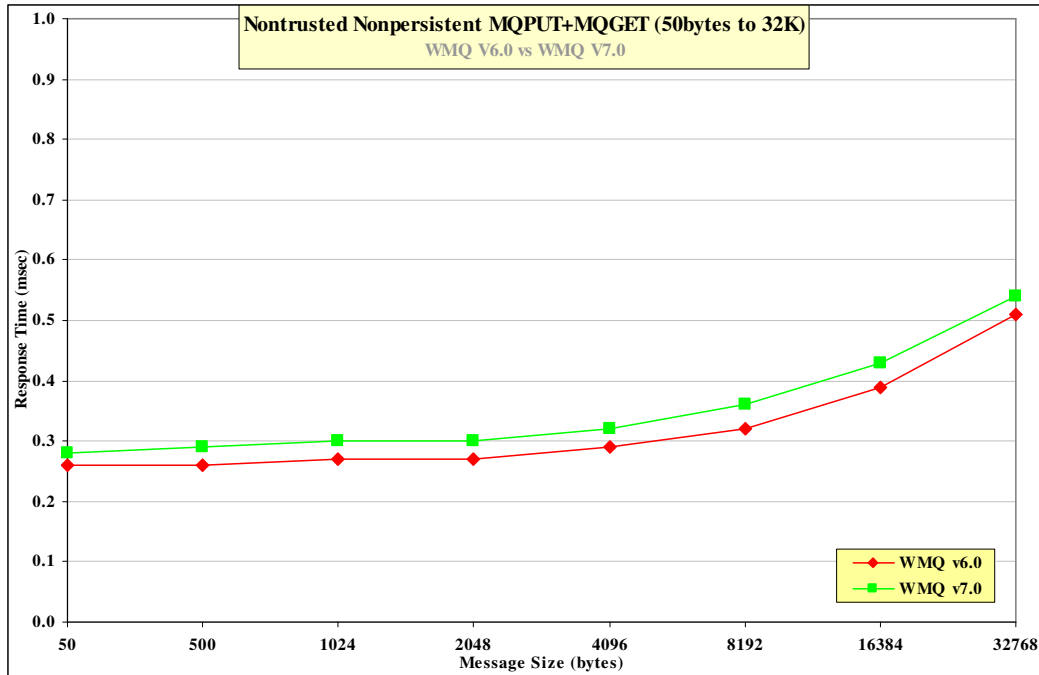



Figure 19 –The effect of nonpersistent message size on MQI response time (50byte - 32K)

Figure 20 shows the response for MQPut/MQGet pairs for persistent message sizes between 50bytes and 16KB.



**Figure 20 –The effect of persistent message size on MQI response time (50byte - 32K)**

### 3.1.2 32KB to 2Mb

Figure 21 shows that the response time for MQPut/MQGet pairs has improved for all nonpersistent message sizes between 32KB and 2Mb.




Figure 21 –The effect of nonpersistent message size on MQI response time (32K – 2Mb)

Figure 22 shows the response for MQPut/MQGet pairs for persist message sizes between 32KB and 2Mb.




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



### 3.1.3 2Mb to 100Mb

Figure 23. The response time for MQPut/MQGet pairs for NP message between 2Mb and 100Mb.




Figure 23 –The effect of nonpersistent message size on MQI response time (2Mb – 100Mb)

Figure 24. The response for MQPut/MQGet pairs for persistent message sizes between 2Mb and 100Mb.

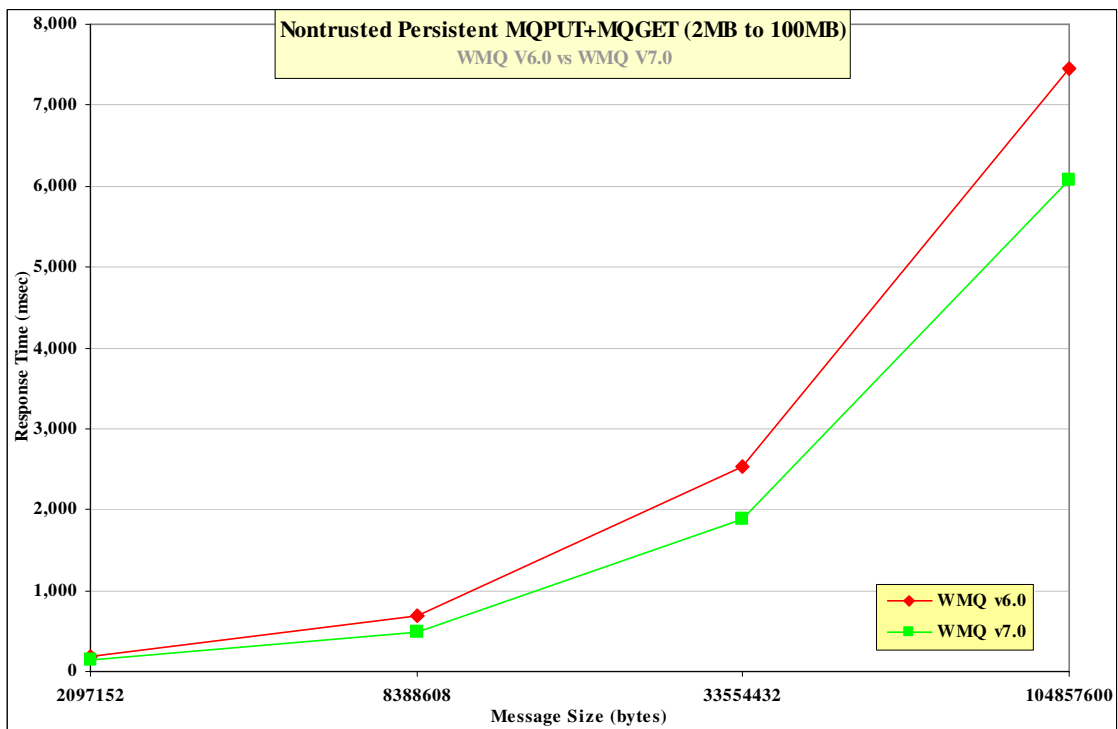


Figure 24 –The effect of persistent message size on MQI response time (2Mb – 100Mb)

### 3.2 20K Messages

#### 3.2.1 Local Queue Manager

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

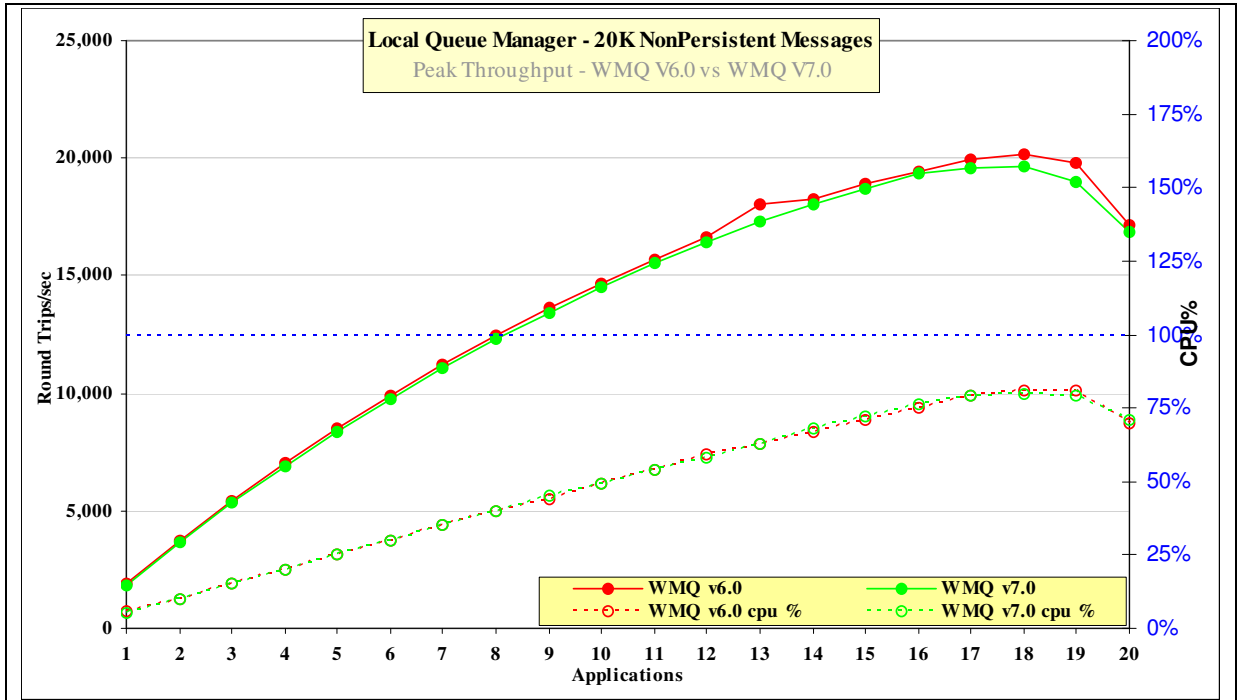


Figure 25 – 20K nonpersistent messages, local queue manager

Figure 25 and Table 13 show that the throughput of nonpersistent messages is similar comparing Version 7 to Version 6.

Test name: <b>local_np_20K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>18</b>	<b>20135</b>	<b>0.001</b>	<b>81%</b>
WebSphere MQ V7	<b>18</b>	<b>19683</b>	<b>0.001</b>	<b>80%</b>

Table 13 – 20K nonpersistent messages, local queue manager

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

### 3.2.1.1 Persistent Messages




Figure 26 – 20K persistent messages, local queue manager

Figure 26 and Table 14 show that the throughput of persistent messages has increased by 5% when comparing Version 6 to Version 7.

Test name: <b>local_pm_20K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	(60) <b>12</b>	(142) <b>219</b>	(0.485) <b>0.068</b>	(2%) <b>2%</b>
WebSphere MQ V7	(60) <b>12</b>	(148) <b>221</b>	(0.494) <b>0.067</b>	(2%) <b>2%</b>
WebSphere MQ V7 SAN	(12) <b>92</b>	(1240) <b>888</b>	(0.020) <b>0.027</b>	(12%) <b>9%</b>

Table 14 – 20K persistent messages, local queue manager

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. The numbers in brackets are included in the table to provide meaningful comparison between WebSphere MQ V6 and Version 7.

### 3.2.2 Client Channel

Figure 27 and Figure 28 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario. Clients use Optimized connections on MQ V7 in this chapter.

#### 3.2.2.1 Nonpersistent Messages

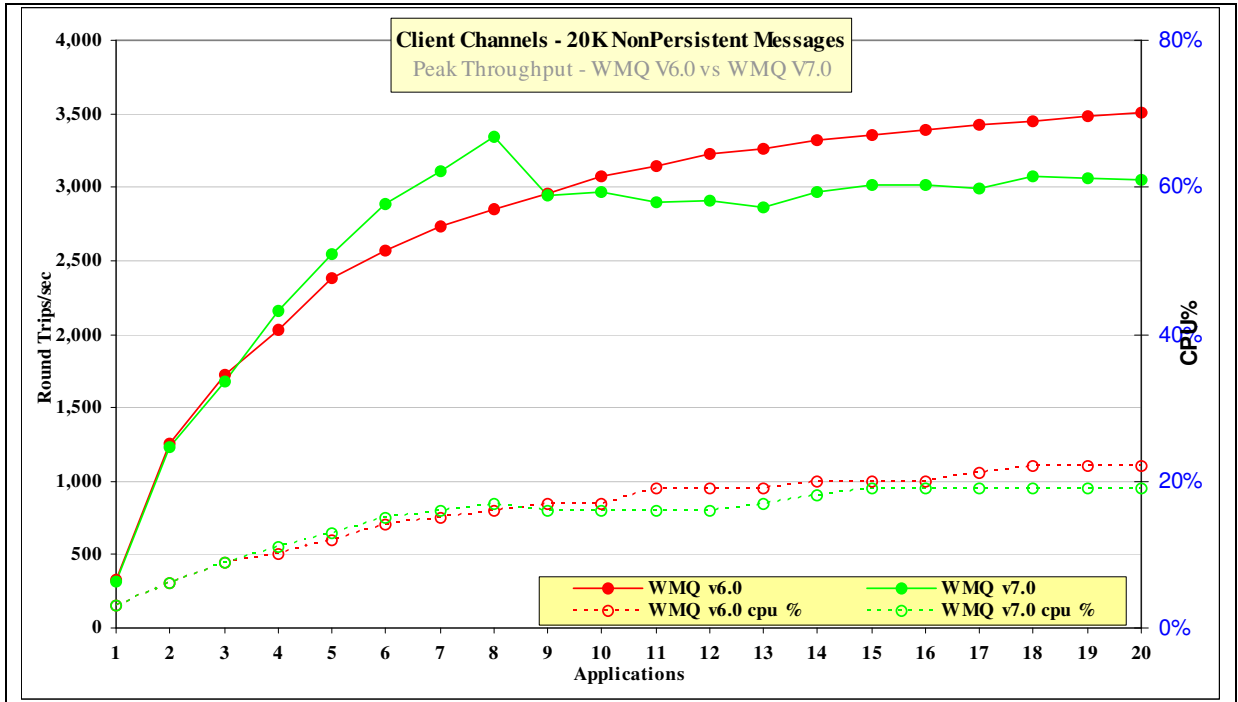


Figure 27 – 20K nonpersistent messages, client channels

Figure 27 and Table 15 show that the throughput of nonpersistent messages has increased by 3% when comparing Version 6 to Version 7.

Test name:	Apps	Round Trips/sec	Response time (s)	CPU
<b>clnp_20K</b>				
WebSphere MQ V6	<b>8</b>	<b>3021</b>	<b>0.002</b>	<b>16%</b>
WebSphere MQ V7	<b>8</b>	<b>3341</b>	<b>0.003</b>	<b>17%</b>

Table 15 – 20K nonpersistent messages, client channels

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.

### 3.2.2.2 Persistent Messages

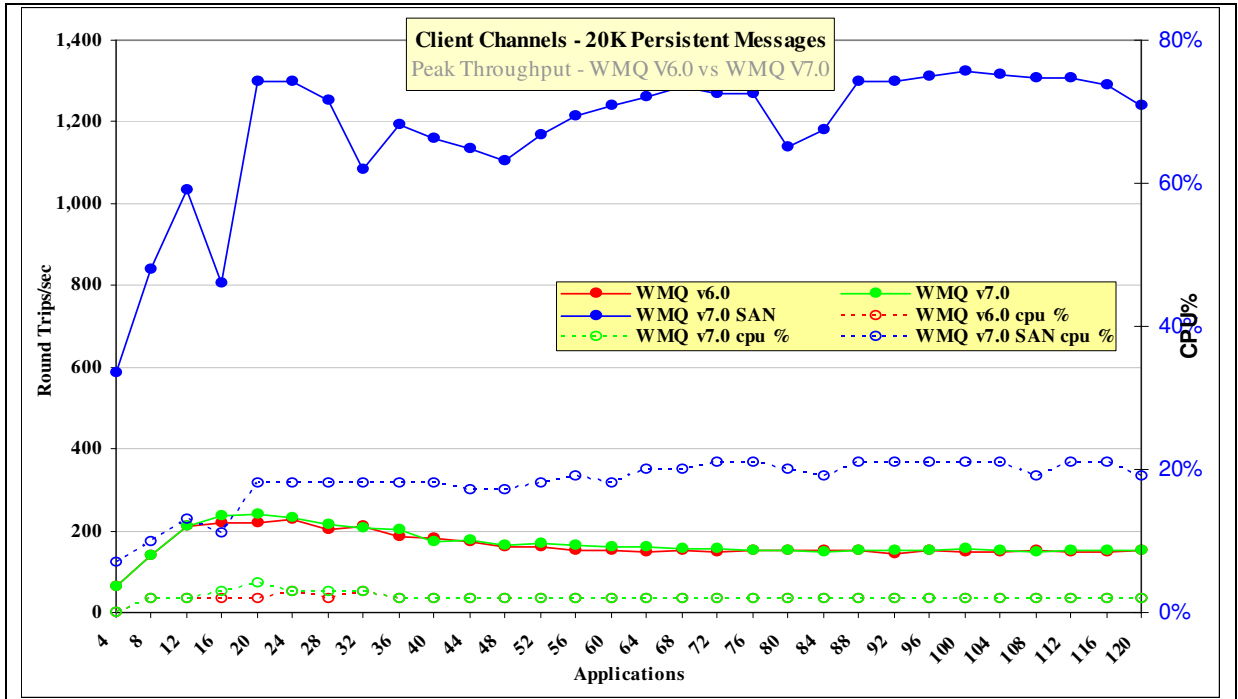


Figure 28 – 20K persistent messages, client channels

Figure 28 and Table 16 show that the throughput of persistent messages has increased by 2% when comparing Version 6 to Version 7.

Test name: clpm_20K	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	24	<b>227</b>	<b>0.117</b>	2%
WebSphere MQ V7	20	<b>242</b>	<b>0.112</b>	2%
WebSphere MQ V7 SAN	24	<b>1300</b>	<b>0.020</b>	24%

Table 16 – 20K persistent messages, client channels

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.

### 3.2.3 Distributed Queuing

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

#### 3.2.3.1 Nonpersistent Messages

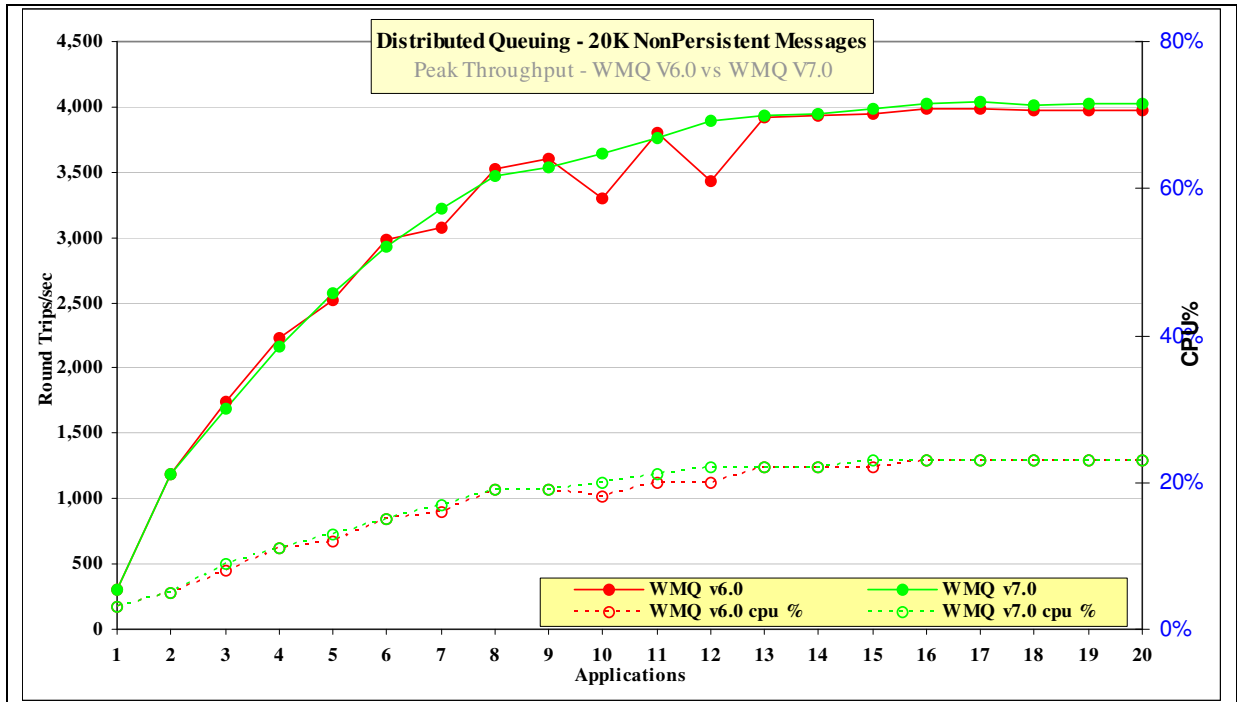


Figure 29 – 20K nonpersistent messages, distributed queuing

Figure 29 and Table 17 show that the throughput of nonpersistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>dqnp_20K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>17</b>	<b>3991</b>	<b>0.006</b>	<b>23%</b>
WebSphere MQ V7	<b>17</b>	<b>4044</b>	<b>0.005</b>	<b>23%</b>

Table 17 – 20K nonpersistent messages, client channels

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

### 3.2.3.2 Persistent Messages

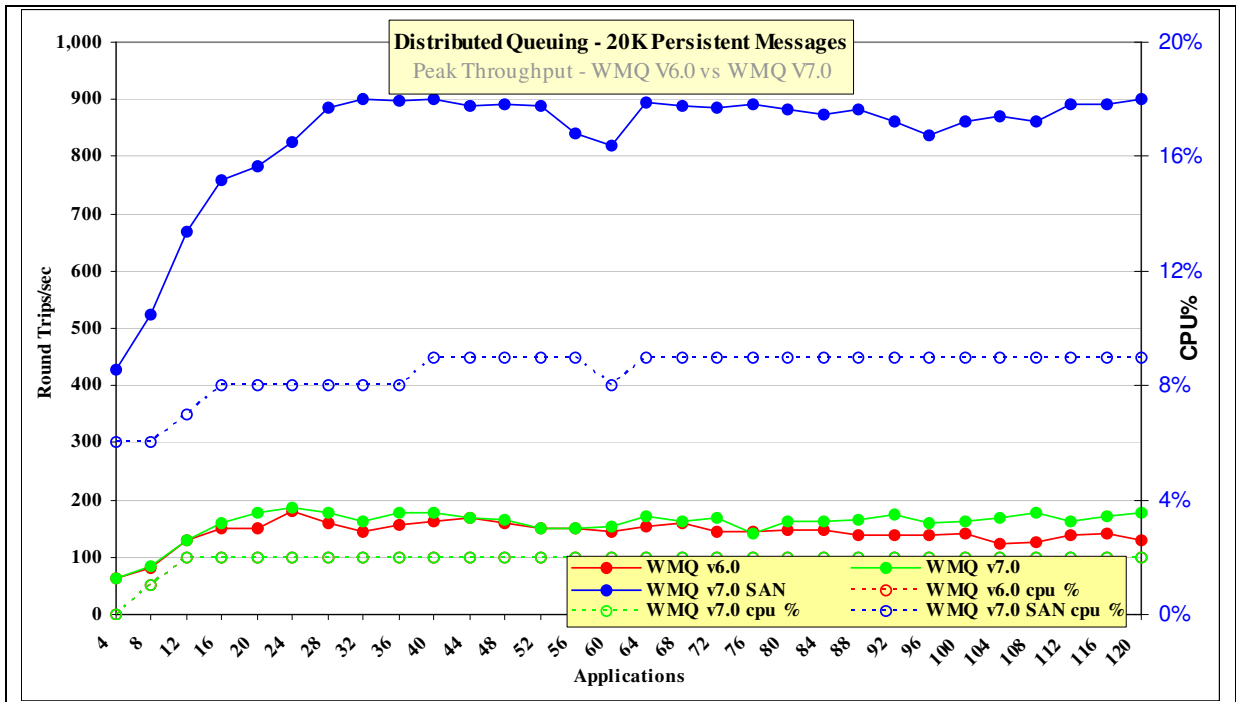


Figure 30 – 20K persistent messages, distributed queuing

Figure 30 and Table 18 show that the throughput of nonpersistent messages has increased by 11% when comparing Version 6 to Version 7.

Test name: <b>dqpm_20K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>24</b>	<b>180</b>	<b>0.146</b>	<b>2%</b>
WebSphere MQ V7	<b>24</b>	<b>188</b>	<b>0.137</b>	<b>2%</b>
WebSphere MQ V7 SAN	<b>24</b>	<b>827</b>	<b>0.032</b>	<b>9%</b>

Table 18 – 20K persistent messages, client channels

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.

### 3.3 200K Messages

#### 3.3.1 Local Queue Manager

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

##### 3.3.1.1 Nonpersistent Messages




Figure 31 – 200K nonpersistent messages, local queue manager

Figure 31 and Table 19 show that the throughput of nonpersistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>local_np_200K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>19</b>	<b>3615</b>	<b>0.003</b>	<b>59%</b>
WebSphere MQ V7	<b>19</b>	<b>3591</b>	<b>0.007</b>	<b>60%</b>

Table 19 – 200K nonpersistent messages, local queue manager

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



### 3.3.1.2 Persistent Messages

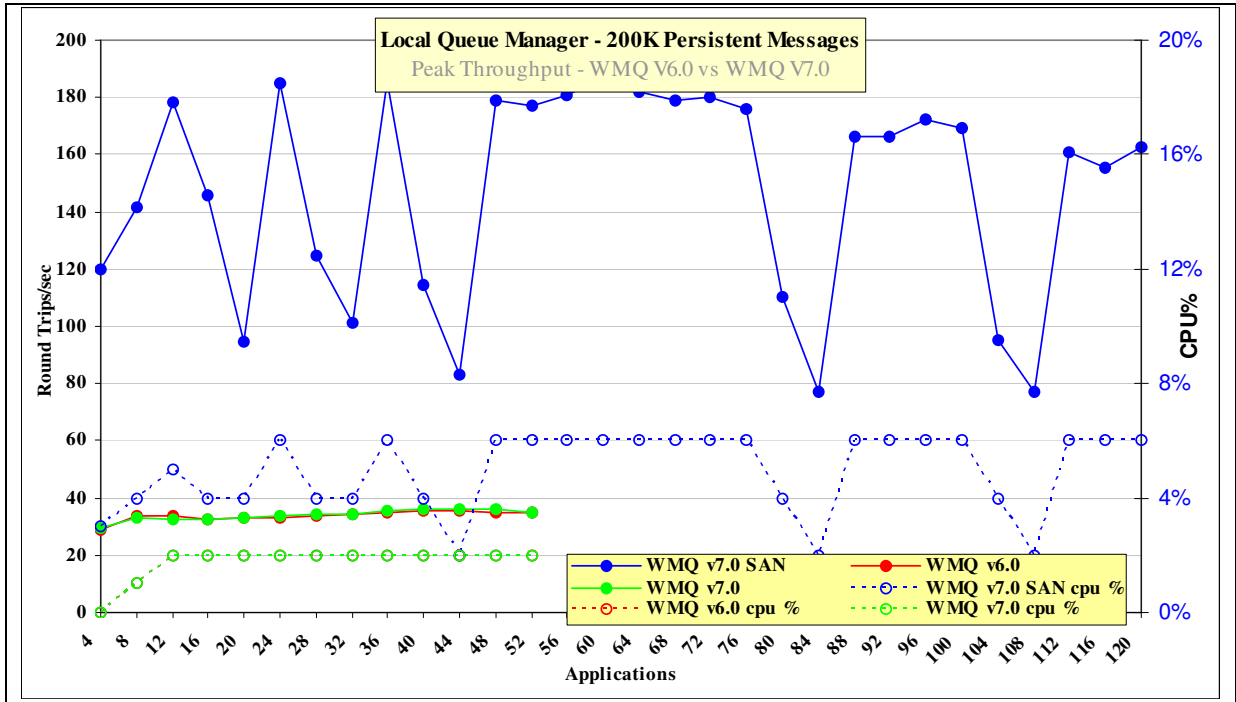


Figure 32 – 200K persistent messages, local queue manager

Figure 32 and Table 20 show that the throughput of persistent messages is similar comparing Version 6 to Version 7.

Test name: <b>local_pm_200K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>8</b>	<b>34</b>	<b>0.261</b>	<b>1%</b>
WebSphere MQ V7	<b>8</b>	<b>33</b>	<b>0.264</b>	<b>1%</b>
WebSphere MQ V7 SAN	(8) <b>64</b>	(141) <b>181</b>	(0.058) <b>0.398</b>	(4%) <b>6%</b>

Table 20 – 200K persistent messages, local queue manager

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.

### 3.3.2 Client Channel

Figure 33 and Figure 34 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario. Clients in Chapter 3 used Optimized connections on MQ V7.

#### 3.3.2.1 Nonpersistent Messages




Figure 33 – 200K nonpersistent messages, client channels

Figure 33 and Table 21 show that the throughput of nonpersistent messages has degraded by 3% when comparing Version 6 to Version 7.

Test name: <b>clnp_200K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>19</b>	<b>358</b>	<b>0.058</b>	<b>15%</b>
WebSphere MQ V7	<b>19</b>	<b>341</b>	<b>0.061</b>	<b>16%</b>

Table 21 – 200K nonpersistent messages, client channels

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.

### 3.3.2.2 Persistent Messages




Figure 34 – 200K persistent messages, client channels

Figure 34 and Table 22 only show the throughput of persistent messaging has improved by 11% when comparing Version 6 with V7 message rates

Test name: clpm_200K	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	28	34	0.890	2%
WebSphere MQ V7	28	36	0.840	3%
WebSphere MQ V7 SAN	(24) 100	(139) 186	(0.186) 0.645	(9%) 12%

Table 22 – 200K persistent messages, client channels

### 3.3.3 Distributed Queuing

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

#### 3.3.3.1 Nonpersistent Messages




Figure 35 – 200K nonpersistent messages, distributed queuing

Figure 35 and Table 23 show that the throughput of nonpersistent messages has increased by 2% when comparing Version 6 to Version 7

Test name: <b>dqnp_200K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>14</b>	<b>342</b>	<b>0.045</b>	<b>14%</b>
WebSphere MQ V7	<b>14</b>	<b>348</b>	<b>0.044</b>	<b>15%</b>

Table 23 – 200K nonpersistent 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.

### 3.3.3.2 Persistent Messages




Figure 36 – 200K persistent messages, distributed queuing

Figure 36 and Table 24 show that the throughput of nonpersistent messages has improved by 5% when comparing Version 6 to Version 7

Test name: <b>dqpm_200K</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>24</b>	<b>32</b>	<b>0.799</b>	<b>2%</b>
WebSphere MQ V7	<b>24</b>	<b>34</b>	<b>0.754</b>	<b>2%</b>
WebSphere MQ V7 SAN	<b>24</b>	<b>130</b>	<b>0.165</b>	<b>7%</b>

Table 24 – 200K 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.

### 3.4 2Mb Messages

#### 3.4.1 Local Queue Manager

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

##### 3.4.1.1 Nonpersistent Messages




Figure 37 – 2M nonpersistent messages, local queue manager

Figure 37 and Table 25 show that the throughput of nonpersistent messages has increased by 3% when comparing Version 6 to Version 7.

Test name: <b>local_np_2M</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>7</b>	<b>84</b>	<b>0.097</b>	<b>13%</b>
WebSphere MQ V7	<b>7</b>	<b>89</b>	<b>0.092</b>	<b>13%</b>

Table 25 – 2M nonpersistent messages, local queue manager

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.

### 3.4.1.2 Persistent Messages




Figure 38 – 2M persistent messages, local queue manager

Figure 38 and Table 26 show that the throughput of persistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>local_pm_2M</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>4</b>	<b>4</b>	<b>0.896</b>	<b>2%</b>
WebSphere MQ V7	<b>4</b>	<b>4</b>	<b>0.898</b>	<b>2%</b>
WebSphere MQ V7 SAN	<b>4</b>	<b>14</b>	<b>0.270</b>	<b>5%</b>

Table 26 – 2M persistent messages, local queue manager

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.

### 3.4.2 Client Channel

Figure 39 and Figure 40 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario. Clients in Chapter 3 used Optimized connections on MQ V7.

#### 3.4.2.1 Nonpersistent Messages

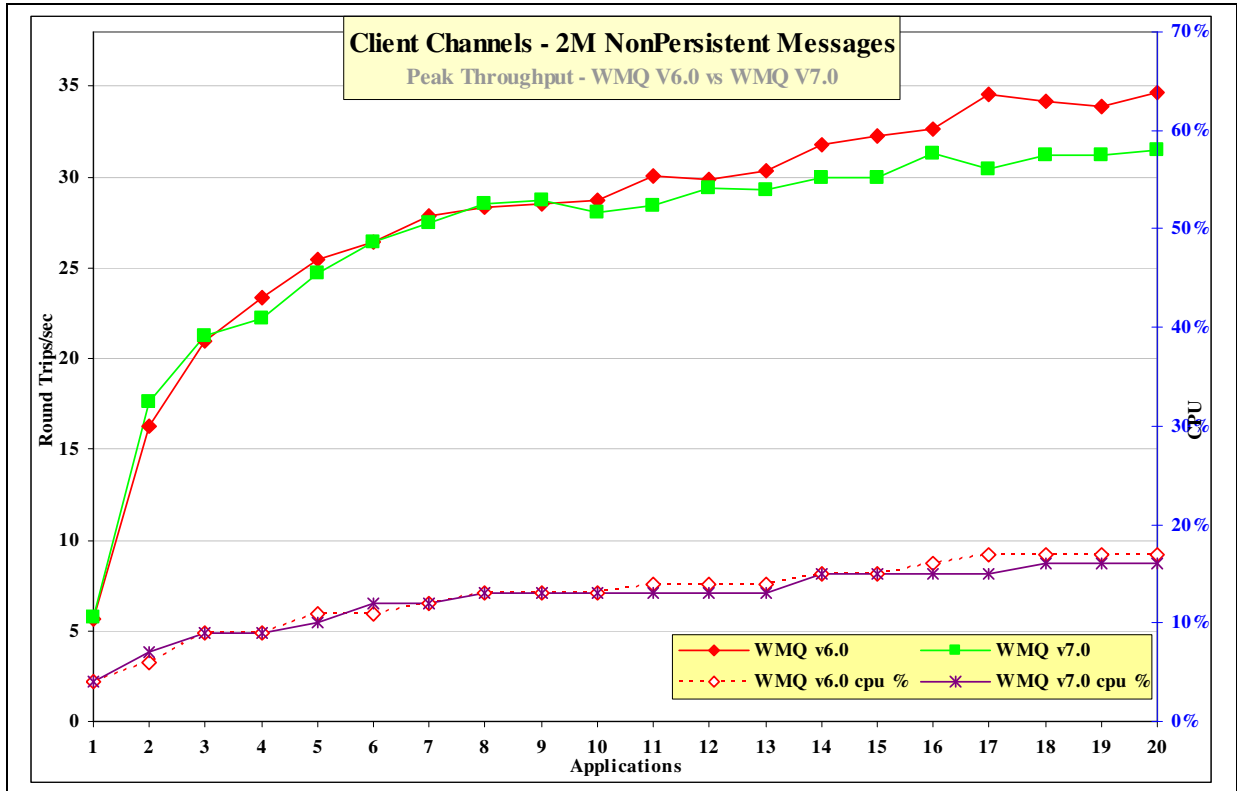


Figure 39 – 2M nonpersistent messages, client channels

Figure 39 and Table 27 show that the throughput of nonpersistent messages has degraded by 3% when comparing Version 6 to Version 7 .

Test name: <b>clnp_2M</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>18</b>	<b>34</b>	<b>0.57</b>	<b>17%</b>
WebSphere MQ V7	<b>18</b>	<b>31</b>	<b>0.63</b>	<b>16%</b>

Table 27 – 2M nonpersistent messages, client channels

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.



### 3.4.2.2 Persistent Messages




Figure 40 – 2M persistent messages, client channels

Figure 40 and Table 28 show that the throughput of persistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>clpm_2M</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>4</b>	<b>4</b>	<b>0.903</b>	<b>3%</b>
WebSphere MQ V7	<b>4</b>	<b>4</b>	<b>0.917</b>	<b>3%</b>
WebSphere MQ V7 SAN	<b>4</b>	<b>12</b>	<b>0.335</b>	<b>6%</b>

Table 28 – 2M persistent messages, client channels

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.

### 3.4.3 Distributed Queuing

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

#### 3.4.3.1 Nonpersistent Messages




Figure 41 – 2M nonpersistent messages, distributed queuing

Figure 41 and Table 29 show that the throughput of nonpersistent messages is similar when comparing Version 6 to Version 7.

Test name: <b>dqnp_2M</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>8</b>	<b>34</b>	<b>0.264</b>	<b>15%</b>
WebSphere MQ V7	<b>8</b>	<b>33</b>	<b>0.264</b>	<b>15%</b>

Table 29 – 2M nonpersistent 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.

### 3.4.3.2 Persistent Messages

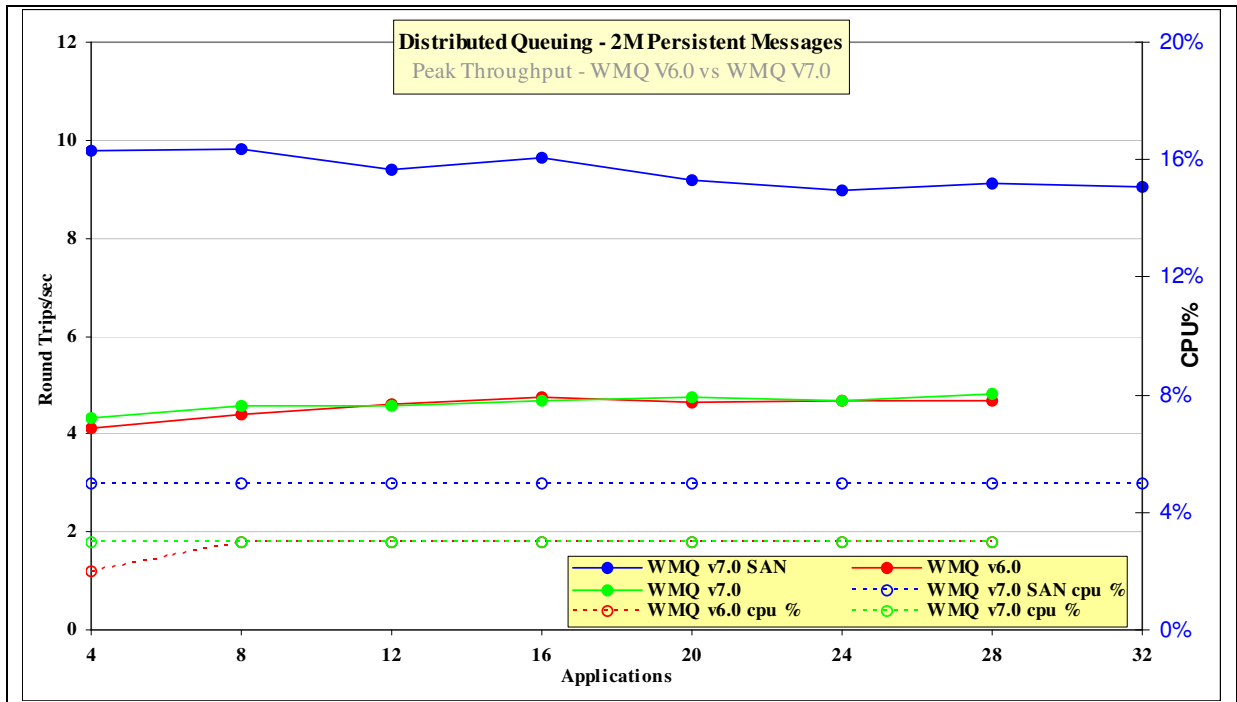


Figure 42 – 2M persistent messages, distributed queuing

Figure 42 and Table 30 show that the throughput of nonpersistent messages is similar when comparing Version 6 to Version 7. The response time with 8 or more applications exceeds one second.

Test name: <b>dqpm_2M</b>	Apps	Round Trips/sec	Response time (s)	CPU
WebSphere MQ V6	<b>4</b>	<b>4</b>	<b>0.976</b>	<b>2%</b>
WebSphere MQ V7	<b>4</b>	<b>4</b>	<b>0.938</b>	<b>3%</b>
WebSphere MQ V7 SAN	<b>4</b>	<b>9</b>	<b>0.412</b>	<b>5%</b>

Table 30 – 2M 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 Application Bindings

This report analyzes the rate that messages can be exchanged 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 43 and Figure 44 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

#### 4.1.1 Nonpersistent Messages




Figure 43 – Application binding, nonpersistent messages, local queue manager

Figure 43 and Table 31 show that the peak throughput of nonpersistent messages when comparing Normal, Isolated, Trusted and Shared bindings.

Test	Apps	Round Trips/sec	Response time (s)	CPU
Normal	13	23401	0.001	69%
Isolated	16	20530	0.001	49%
Trusted	10	29791	0.001	34%
Shared	16	21343	0.001	91%

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

### 4.1.2 Persistent Messages




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

Figure 44 and Table 32 show that the peak throughput of persistent messages when comparing Normal, Isolated and Trusted bindings.

Test	Apps	Round Trips/sec	Response time (s)	CPU
Normal	120	2023	0.072	9%
Isolated	116	2056	0.071	11%
Trusted	120	2023	0.072	9%

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

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.

The difference in throughput due to using various application bindings is obscured by the I/O necessary for Persistent messages however the smaller CPU requirement needed by Trusted bindings can still be observed. Disks are internal disks.

## 4.2 Client Channels

Figure 45 and Figure 46 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario. Clients in Chapter 4 used Optimized connections on MQ V7.

### 4.2.1 Nonpersistent Messages




Figure 45 – Application binding, nonpersistent messages, client channels

Figure 45 and Table 33 show that the peak throughput of nonpersistent messages when comparing Normal, Isolated and Trusted bindings.

Test	Apps	Round Trips/sec	Response time (s)	CPU
<b>Normal</b>	<b>30</b>	<b>14161</b>	<b>0.002</b>	<b>58%</b>
<b>Isolated</b>	<b>15</b>	<b>11596</b>	<b>0.002</b>	<b>20%</b>
<b>Trusted</b>	<b>17</b>	<b>11241</b>	<b>0.002</b>	<b>13%</b>

Table 33 – Application binding, nonpersistent messages, client channels

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.

### 4.2.2 Persistent Messages




Figure 46 – Application binding, persistent messages, client channels

Figure 46 and Table 34 show that the peak throughput of nonpersistent messages when comparing Normal, Isolated and Trusted bindings.

Test	Apps	Round Trips/sec	Response time (s)	CPU
Normal	112	1930	0.062	13%
Isolated	112	2069	0.057	11%
Trusted	112	2085	0.065	9%

Table 34 – Application binding, persistent messages, client channels

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.

The various application bindings for Client Persistent messages are limited by the same factors as the Local Persistent messages and give similar throughput.

### 4.3 Distributed Queuing

Figure 47 and Figure 48 show the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario.

#### 4.3.1 Nonpersistent Messages

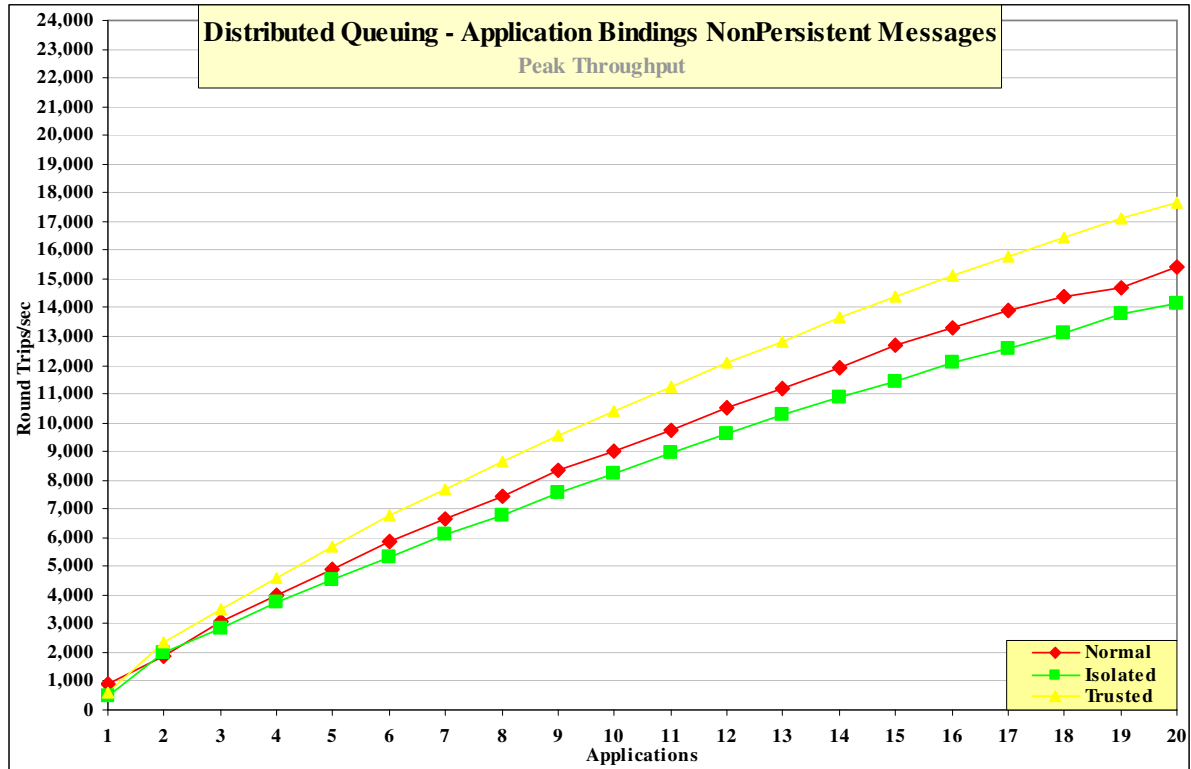


Figure 47 – Application binding, nonpersistent messages, distributed queuing

Figure 47 and Table 35 show that the peak throughput of nonpersistent messages when comparing Normal, Isolated and Trusted bindings.

Test	Apps	Round Trips/sec	Response time (s)	CPU
<b>Normal</b>	<b>57</b>	<b>27103</b>	<b>0.003</b>	<b>70%</b>
<b>Isolated</b>	<b>20</b>	<b>14150</b>	<b>0.002</b>	<b>21%</b>
<b>Trusted</b>	<b>20</b>	<b>17682</b>	<b>0.001</b>	<b>20%</b>

Table 35 – Application binding, nonpersistent 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 36 show that the peak throughput of nonpersistent messages when comparing Normal, Isolated and Trusted bindings.

Test	Apps	Round Trips/sec	Response time (s)	CPU
<b>Isolated</b>	<b>170</b>	<b>1392</b>	<b>0.135</b>	<b>5%</b>
<b>Trusted</b>	<b>170</b>	<b>1414</b>	<b>0.13</b>	<b>3%</b>

Table 36 – 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 MQCON/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 49** for an illustration).




Figure 49 – Short sessions, client channels

Test name	Apps	Round Trips/sec	Short Sessions per second	Response Time (s)	cpu%
<b>clnp_r3600</b>	<b>4,500</b>	<b>4469</b>		<b>0.022</b>	<b>21%</b>
<b>clnp_ss</b>	<b>1,070</b>	<b>1088</b>	<b>217</b>	<b>0.072</b>	<b>54%</b>
<b>clpm_r3600 SAN</b>	<b>1500</b>	<b>1488</b>		<b>0.093</b>	<b>19%</b>
<b>clpm_ss SAN</b>	<b>700</b>	<b>639</b>	<b>127</b>	<b>0.520</b>	<b>46%</b>

Table 37 – Short sessions, client channels

*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 'runmqsr' 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*. 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.

Queue manager configuration for client channels capacity tests:

MaxChannels=50000 (100,000 for clnp\_cmax)

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
clnp	24	n/a*	11069	0.002	53%
clnp_r3600	4500	3600	4469	0.022	21%
clnp_cmax (SC1)	33000	60	527	0.122	10%
clnp_c6000(SC1)	6000	1860	3089	0.359	18%
	6000&	2100	3484	0.380	19%
clnp_c6000_nocorrelid(SC1)	6000	1580	2608	0.395	22%
clpm_c6000(SC1)	6000	400	662	0.08	10%
		440	737	0.363	10%

**Table 38 - Client Capacity, throughput**

\* There was no delay between the response to the previous message and the insertion of the next message with 22 clients.

& DefaultQBufferSize = 25Mb was used as a Tuning Parameter

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. The clnp\_cmax\_ results show sharecnv=1 because this reflects all of the clients using a separate IP socket. The following table shows how the FREE memory pages are consumed and the size of the swap file.

Test name:	Apps	Swap	Free
clnp_cmax (SC1)	10000	1560M	62418232
	33000	25172M	44616944
		1026K per client	556K per client
clnp_cmax(SC0)	1000	1194M	63524312
	4000	2884M	62384720
		563K per client	379K per client
clnp_c6000(SC1)	1000	1696M	61374752
	6000	5391M	58690536
		739K per client	537K per client
clnp_c6000_nocorrelid(SC1)	1000	1953M	63079016
	6000	6087M	60141016
		826K per client	600K per client
clpm_c6000(SC1)	1000	1785M	61352424
	6000	5547M	58831544
		752K per client	506K per client

**Table 39 – Client capacity, memory utilisation**

*Note: The table above show the swap memory and Free pages measured at the given number of driving applications. The swap and free memory cost is the additional cost per driving application (in this test scenario this relates to the cost of an MQI-client connection on the server machine). Once the Free pages are exhausted, the swap file grows by a similar amount per client.*

The difference between this pair of measurements is the clnp\_cmax uses a Get by Correlation\_Id from a common reply queue for all the clients whereas the other case has a separate reply queue per client. Each additional Client needs a thread in the AMQRMPPA process and this accounts for most of the 256Kbytes. Using a separate queue per client needs additional shared memory per client as well as some more pages from the Free pool. The FREEMEM information is collected at low message rates. An additional 20K per client is needed at the higher rates.

## 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*. 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	57	n/a*	20704	0.003	70%
dqnp_r3600	5000	3600	4959	0.016	17%
dqnp_qmax	10000	60	166	0.002	2%
dqpm_q1000	1000	6400 7300	1739 1834	0.095 0.619	21%

**Table 40 – Capacity measurements, server channels**

\* There was no delay between the response to the previous message and the insertion of the next message with 40 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\_q1000 and dqnp\_qmax 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 and the following table shows the storage on the central hub.

Test name:	Apps	Swap	Free	Free per Applic
dqnp_qmax	200	1645	60489280	678K bytes
	10000	10694	53839712	
dqpm_q1000	100	926	62538352	949K bytes
	1000	2085	61683448	

**Table 41 – DQ capacity, memory utilisation**

*Note: The table above show the swap memory measured at the given number of driving applications. The swap and free memory cost is the additional cost per driving application (in this test scenario this relates to the cost of an MQI-Sender/Receiver pair of channels plus Transmission queue connected on the server machine).*

It can be observed that the amount of memory needed to support a DQ channel is approximately double that needed to support a client channel. Each DQ application needs both a Sender and Receiver channel as well as a Transmit queue whereas a Client connection just needs one channel.

## 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; some of these can be applied to Version 6. 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 documented external interface and may change or be removed in future releases.*

#### 7.1.1 Queue Disk, Log Disk, and Message Persistence

Nonpersistent 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 and 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 nonpersistent 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 Nonpersistent and Persistent Queue Buffer

The default nonpersistent 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 nonpersistent 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 (V6 was 128 pages) with a maximum size of 4096 pages. To improve persistent message throughput of large messages (messages size > 1M bytes) 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 (< 10K byte 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 whereas the V6 default was 128 on Windows and 1024 on Unix). 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 'rcdmqimg' 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%. The Solaris 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 nonpersistent 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 nonpersistent messages.*

### 7.1.4 Multiplexed clients

Version 6 and previous levels used a separate TCP socket for each client. Version 7 will multiplex clients from the same process over one TCP socket. Chapter 2 shows the difference in performance of these variants.

Version 6 behaviour can be obtained by using the 'sharecnv' keyword with a setting of zero. For example

```
define channel( csim_channel_TCP ) +
    chltype( svrconn ) +
    trptype( tcp ) +
    sharecnv( 0 )
```

Version 6 behaviour will also inhibit new performance features of V7 like 'ASYN Put and 'READ\_AHEAD'. This is referred to as Compatibility mode.

## 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 nonpersistent 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 'MQPutter' 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 ‘runmqslr’ 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

Systems require sufficient real memory to hold the working set otherwise paging will break the response time expectations.

- Virtual memory enables the program to address much larger amount of memory than exists as real memory.
- Real memory is the physical memory (or RAM) currently installed in the machine.
- Paging is the process of managing program access to virtual storage pages not currently resident in main memory. It locates the required page frame from auxiliary storage (disk), selects a page frame in real memory that will hold this page, copies the contents of this outgoing page frame to auxiliary storage, and retrieves the requested incoming page contents from auxiliary storage.

A simple approach is to ensure that the virtual memory of the application system does not exceed the available real memory since all memory requests will be met from the current free memory.

WebSphere MQ uses a significant amount of memory for each Queue Manager and Channel.

### 7.3.1 Queue Manager

Starting a MQ Queue manager generated using default values increases the AVM and reduces the FRE by 260M bytes.

### 7.3.2 Channels

Channels can be started by using the INETD or the RUNMQLSR listener. INETD initiated channels use between 5 and 10 times more memory than RUNMQLSR channels so the rest of this section focuses on RUNMQLSR channels.

### 7.3.3 Client Channels

Each MQ client channel uses between 380K – 620K bytes for processing 2K byte messages depending on traffic rate (Chapter 6 of the MQ V7 Performance reports provides an estimate of the storage needed when clients either share a

predefined queue with other clients or have a dynamic queue per client). 100K byte messages will use up an additional 200K bytes per client.

### 7.3.4 Server – Server Channels

Each interconnected queue manager has a pair of uni-directional channels for sending and receiving messages. The storage consumed is the same as 2 client channels plus a predefined queue (Transmission queue).

Three other aspects of storage consumption depend on type of 'Reply-Queue', MQIBINDTYPE, and BufferLength.

### 7.3.5 Reply Queue

The Queue from which the client retrieves the message can be a predefined Queue (350K bytes) probably shared among multiple clients who get messages by Correlation-id or a model (dynamic) queue (60K bytes) that is used only by one client. The model queue memory can grow by 128K bytes when more than 128K bytes of Persistent messages are held in the queue and by 192K bytes when more than 192K bytes of non persistent messages are held in the queue. This memory is not shrunk back to the underlying 60K bytes for model queues.

### 7.3.6 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.

### 7.3.7 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 80K bytes 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 V6 & V7

##### 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. Each driving application thread performed the sequence of actions as outlined in the test scenario illustrations in the “**Performance Headlines**” starting on **page 2**.

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.

Nonpersistent 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 is written as a multi-threaded program configured to use 20, 6, 6 threads for processing nonpersistent messages with Local, Client, and DQ applications, and 30, 60, 10 threads to process persistent messages with Local, Client, and DQ applications. The capacity tests in Chapters 5 and 6 use 10 server threads for processing non persistent messages. Each server thread performed the sequence of actions as outlined in the test scenario illustrations in the “**Performance Headlines**” starting on **page 2**.

Nonpersistent messaging is done outside of syncpoint control. Persistent messaging is 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.

## 8.2 Hardware

Sun T2000:	Server system (Device under test)
Model:	T201108C-64GA2G SFT2000
Processor:	1.4GHz
Architecture:	8core
Memory (RAM):	64Gb
Disk:	Internal disks for most measurements (Chapter 1-9) Fibre channel HBA PCI-X 4Gb FC Dual Port HBA (chapter 10) 2 SAN disks on DS6000 (5Gb each, 1 queue, 1 log )
Network:	1Gbit Ethernet Adapter
Sun V490:	Driver system (2 machines)
Model:	
Processor:	SPARCV9 @ 1500 HMz
Architecture:	4 CPU
Memory (RAM):	32Gb
Disk:	Fibre channel HBA PCI-X 4Gb FC Dual Port HBA) SAN with 2 partitions of 5Gb each
Network:	1Gbit Ethernet Adapter

The MQ SAN consists of a pair of 2026 model 432 (McDATA ES-4700) switches running at 4Gb/s with 32 ports each. They are connected together via two inter-switch links to form a single SAN fabric.

The MQ hosts attach via this SAN to a DS6800 disk array (1750 model 511) with one expansion drawer.

Each drawer (controller + expansion) contains 16 x 73Gb 15K fibre channel disk drives, so there are a total of 32 physical drives.

The 32 drives are configured as four RAID-5 arrays, each of which is 6+Parity+Spare (the number of spares is defined by the configuration of the DS6800).

The controller has an effective cache size of 2.6Gb plus 0.3Gb of NVS

## 8.3 Software

Solaris 64 bit:	Solaris 10
MQSeries:	Version 7, Version 6

## 9 Compiler Glossary

<b>Test name</b>	<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><b>local</b> =&gt; local queue manager test scenario</p> <p><b>cl</b> =&gt; client channel test scenario</p> <p><b>dq</b> =&gt; distributed queuing test scenario</p> <p><b>np</b> =&gt; nonpersistent messages</p> <p><b>pm</b> =&gt; persistent messages</p> <p><b>r3600</b> =&gt; 1 round trip per driving application per second</p> <p><b>runmqlsr</b> =&gt; channels using the 'runmqlsr' listener (client channel test scenario, in addition to 'runmqchi' for distributed queuing test scenarios)</p> <p><b>c6000</b> =&gt; 6,000 client driving applications (i.e. 6,000 MQI-client connections)</p> <p><b>q1000</b> =&gt; 1,000 server channel pairs</p> <p><b>max</b> =&gt; maximum number of channels (or channel pairs)</p> <p><b>no_correl_id</b> =&gt; correlation identifier not used in the response messages (as each response is placed on a unique reply-to queue per driving application)</p>
<b>Apps</b>	The number of driving applications connected to the queue manager at the point where the performance measurement is given.
<b>Rate/App/hr</b>	The target message throughput rate of each driving application.
<b>Round T/s</b>	The average achieved message throughput rate of all the driving applications together, measured by the server application.
<b>% (Round T/s)</b>	<p>The percentage change 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>
<b>CPU</b>	As reported by VMSTAT
<b>Resp time (s)</b>	The average response time each round trip, as measured and averaged by all the driving applications.
<b>CURDEPTH</b>	<p>The number of messages on the input queue as a snapshot.</p> <p><i>Note: runmqsc &lt;qmname&gt;, DISPLAY QLOCAL(&lt;qname&gt;) CURDEPTH</i></p>
<b>queue disk (kbps)</b>	The queue disk kilobytes transferred per second.
<b>Swap</b>	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).
<b>shm</b>	The amount of allocated shared memory in Mb.
<b>FREE</b>	Free memory as reported by IOSTAT
<b>SC0</b>	SHARECNV=0 specified on the def channel(x) chltype(svrconn) command. Version 6 compatibility mode
<b>SC1</b>	SHARECNV=1 specified on the def channel(x) chltype(svrconn) command. Separate socket per client