

# **IBM MQ for IBM i v8.0**

## **Performance Evaluations**

Version 1

November 2014

Craig Stirling , Michael Alexander.

WebSphere MQ Performance

IBM UK Laboratories

Hursley Park

Winchester

Hampshire

SO21 2JN

Property of IBM

**Please take Note!**

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

**First Edition, November 2014.**

This edition applies to *IBM MQ for IBM i v8.0* (and to all subsequent releases and modifications until otherwise indicated in new editions).

© Copyright International Business Machines Corporation 2014. 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 *IBM MQ for IBM i v8.0*. The information is not intended as the specification of any programming interface that is provided by IBM. It is assumed that the reader is familiar with the concepts and operation of IBM MQ v8.0.

**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
- MQ
- DB2
- IBM i

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 report is designed for people who:

- Will be designing and implementing solutions using MQ v8.0
- Want to understand the performance limits of MQ v8.0
- Want to understand what actions may be taken to tune MQ v8.0

The reader should have a general awareness of the IBM i operating system and of IBM MQ in order to make best use of this report.

## The contents of this Report

This report includes:

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

## Feedback on this Report

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 MQ system should be directed to your local IBM Representative or Support Centre.

Information in this report should be used alongside the publicly available [knowledge centre for IBM MQ V8](#).

# Introduction

Full details of the workloads and the hardware & software levels used are given in section 7 but in summary:

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

- Local queue manager.
- Client channel
- Distributed queuing

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

## **Device under test (server)**

An IBM Power 780 (7-way CPU 3.86GHz POWER7) with 25.25GB of RAM was used as the device under test.

## **Driver**

An IBM Power 570 (8-way CPU 4.2GHz POWER6) with 16GB of RAM was used as the driver for all client tests

An IBM Power 780 (8-way CPU 3.86GHz POWER7) with 32GB of RAM was used as the driver for all distributed tests.

## **Software Levels**

Operating system	: IBM i V7R1M0
MQ version	: Version 7.1, Version 8.0
Compiler	: IBM Rational Development Studio for i V7R1M0

## **How this document is arranged**

### **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.

### **Chapter 3**

Contains performance measurements for *large messages*. This includes *20K*, *200K* and *2M* byte messages using the same scenarios as for the *2KB* messages.

### **Chapter 4**

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

### **Chapter 5**

Contains information on some of the limits to IBM MQ performance and scaling.

### **Chapter 6**

Contains tuning guidance specific to v8.0 on IBM i

### **Chapter 7**

Contains a summary of the way in which the workload is used in each test scenario is given in section 2. This includes a more detailed description of the workload, hardware and software specifications.

### **Chapter 8**

Contains 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	Non-persistent Messages – Local Queue Manager	3
2.1.2	Non-persistent Messages – Non-trusted – Local Queue Manager	4
2.1.3	Persistent Messages – Local Queue Manager	5
2.2	Client Channels Test Scenario	6
2.2.1	Non-persistent Messages – Client Channels	7
2.2.2	Non-persistent Messages – Non-Trusted Client Channels	8
2.2.3	Persistent Messages – Client Channels	9
2.2.4	Client Channels	10
2.3	Distributed Queuing Test Scenario	12
2.3.1	Non-persistent Messages – Server Channels	13
2.3.2	Non-Persistent non-Trusted – Server Channels	14
2.3.3	Persistent Messages – Server Channels	15
2.3.4	Server Channels	16
<b>3</b>	<b>Large Messages</b>	<b>18</b>
3.1	20KB Messages	18
3.1.1	Local Queue Manager	18
3.1.2	Client Channel	20
3.1.3	Distributed Queuing	22
3.2	200K Messages	24
3.2.1	Local Queue Manager	24
3.2.2	Client Channel	26
3.2.3	Distributed Queuing	28
3.3	2MB Messages	30
3.3.1	Local Queue Manager	30
3.3.2	Client Channel	32
3.3.3	Distributed Queuing	34
<b>4</b>	<b>Application Bindings</b>	<b>36</b>
4.1	Local Queue Manager	36
4.1.1	Non-persistent Messages	36
4.1.2	Persistent Messages	37
4.2	Client Channels	38
4.2.1	Non-persistent Messages	38
4.2.2	Persistent Messages	39
4.3	Distributed Queuing	40
4.3.1	Non-persistent Messages	40
4.3.2	Persistent Messages	41
<b>5</b>	<b>Performance and Capacity Limits</b>	<b>41</b>
5.1	Client channels – capacity measurements	41
5.2	Distributed queuing – capacity measurements	42
<b>6</b>	<b>Tuning Recommendations</b>	<b>44</b>
6.1	Tuning the Queue Manager	44
6.1.1	Queue Disk, Log Disk, and Message Persistence	44
6.1.2	Log Buffer Size, Log File Size, and Number of Log Extents	44
6.1.3	Channels: Process or Thread, Standard or Fastpath?	46
6.2	Applications: Design and Configuration	46
6.2.1	Standard (Shared or Isolated) or Fastpath?	46
6.2.2	Parallelism, Batching, and Triggering	46
6.3	Tuning the Operating System (IBM i)	47
6.4	TCP Buffer size changes for V8.0	47
6.5	Virtual Memory, Real Memory, & Paging	47
6.5.1	BufferLength	47
6.5.2	MQIBINDTYPE	48
<b>7</b>	<b>Measurement Environment</b>	<b>49</b>
7.1	Workload description	49
7.1.1	MQI response time tool	49
7.1.2	Test scenario workload	49
7.2	Hardware	50

7.3	Device under test (Server) .....	50
7.4	Software .....	50
<b>8</b>	<b>Glossary .....</b>	<b>51</b>

# TABLES

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

# FIGURES

Figure 1 – Connections into a local queue manager .....	2
Figure 2 – Connections into a local queue manager .....	3
Figure 3 - Performance headline, non-persistent, non-trusted messages and local queue manager. ....	4
Figure 4 - Performance headline, persistent messages and local queue manager .....	5
Figure 5 - MQI-client channels into a remote queue manager .....	6
Figure 6 - Performance headline, non-persistent messages and client channels.....	7
Figure 7 - 1 round trip per driving application per second, client channels, persistent messages .....	8
Figure 8 - Performance headline, persistent messages and client channels .....	9
Figure 9 - 1 round trip per driving application per second, client channels and non-persistent messages .....	10
Figure 10 - 1 round trip per driving application per second, client channels, persistent messages .....	10
Figure 11 - Server channels between two queue managers .....	12
Figure 12 Server channels between two queue managers .....	13
Figure 13 – Performance headline, non-persistent, not trusted messages and server channels .....	14
Figure 14 – Performance headline, persistent messages and server channels .....	15
Figure 15 – 1 round trip per driving application per second, server channel, non-persistent messages .	16
Figure 16 – 1 round trip per driving application per second, server channel, persistent messages .....	16
Figure 17 – 20KB non-persistent messages, local queue manager .....	18
Figure 18 – 20KB persistent messages, local queue manager .....	19
Figure 19 – 20KB persistent messages, local queue manager .....	20
Figure 20 – 20KB persistent messages, client channels .....	21
Figure 21 – 20KB non-persistent messages, distributed queuing .....	22
Figure 22 – 20KB persistent messages, distributed queuing .....	23
Figure 23 – 200KB non-persistent messages, local queue manager .....	24
Figure 24 – 200KB persistent messages, local queue manager .....	25
Figure 25 – 200KB non-persistent messages, client channels .....	26
Figure 26 – 200KB persistent messages, client channels .....	27
Figure 27 – 200KB non-persistent messages, distributed queuing .....	28
Figure 28 – 200KB persistent messages, distributed queuing .....	29
Figure 29 – 2MB non-persistent messages, local queue manager .....	30
Figure 30 – 2MB persistent messages, local queue manager .....	31
Figure 31 – 2MB non-persistent messages, client channels .....	32
Figure 32 – 2MB persistent messages, client channels.....	33
Figure 33 – 2MB non-persistent messages, distributed queuing .....	34
Figure 34 - 2MB persistent messages, distributed queuing .....	35
Figure 35 – Application binding, non-persistent messages, local queue manager .....	36
Figure 36 – Application binding, persistent messages, local queue manage .....	37
Figure 37 – Application binding, non-persistent messages, client channels.....	38
Figure 38 – Application binding, persistent messages, client channels .....	39
Figure 39 – Application binding, non-persistent messages, distributed queuing .....	40
Figure 40 – Application binding, persistent messages, distributed queuing .....	41

# 1 Overview

IBM MQ v8.0 on IBM i has improved performance especially for smaller messages using a small number of queues. For 2KB messages, almost every test shows improvements over earlier versions of MQ.

Other improvements in IBM MQ v8.0

Faster Client Connection Time

- Elapsed time to connect a large number of clients has been reduced by up to 75%

Logger Optimisations

- Single threaded code scope has been reduced in V8

Sharecnv[1] Optimisations

- Up to 20% improvement in throughput using Sharecnv[1]

Internal MQ object processing optimised to reduce locking scopes.

- Improved performance for workloads with many object access requests (e.g. PUT1).

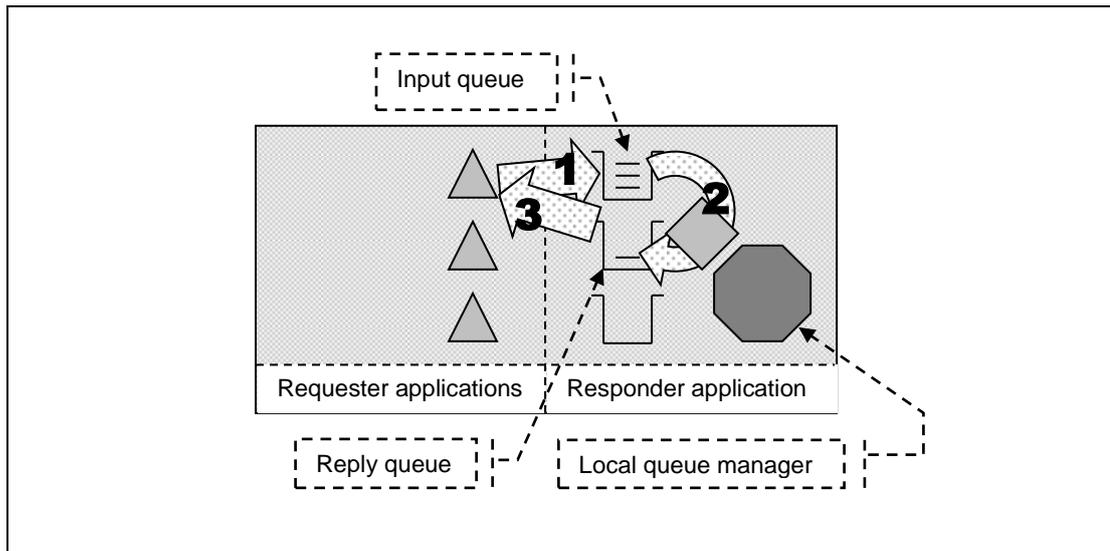
## 2 Performance Headlines

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 uses a dedicated MQI-client channel, whilst in the distributed queuing test scenarios, one or more applications submit messages over a fixed number of server channels.

### 2.1 Local Queue Manager Test Scenario



**Figure 1 – Connections into a local queue manager**

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

Non-persistent and persistent messages were used in the local queue manager tests, with a message size of 2KB. The effect of message throughput with larger messages sizes is investigated in section 3.

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

### 2.1.1 Non-persistent Messages – Local Queue Manager

Figure 2, Figure 3 and Figure 4 shows the non-persistent, non-persistent/non-trusted and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario (see Figure 1 on the previous page) for different production levels of IBM MQ (versions 7.1, 8.0).

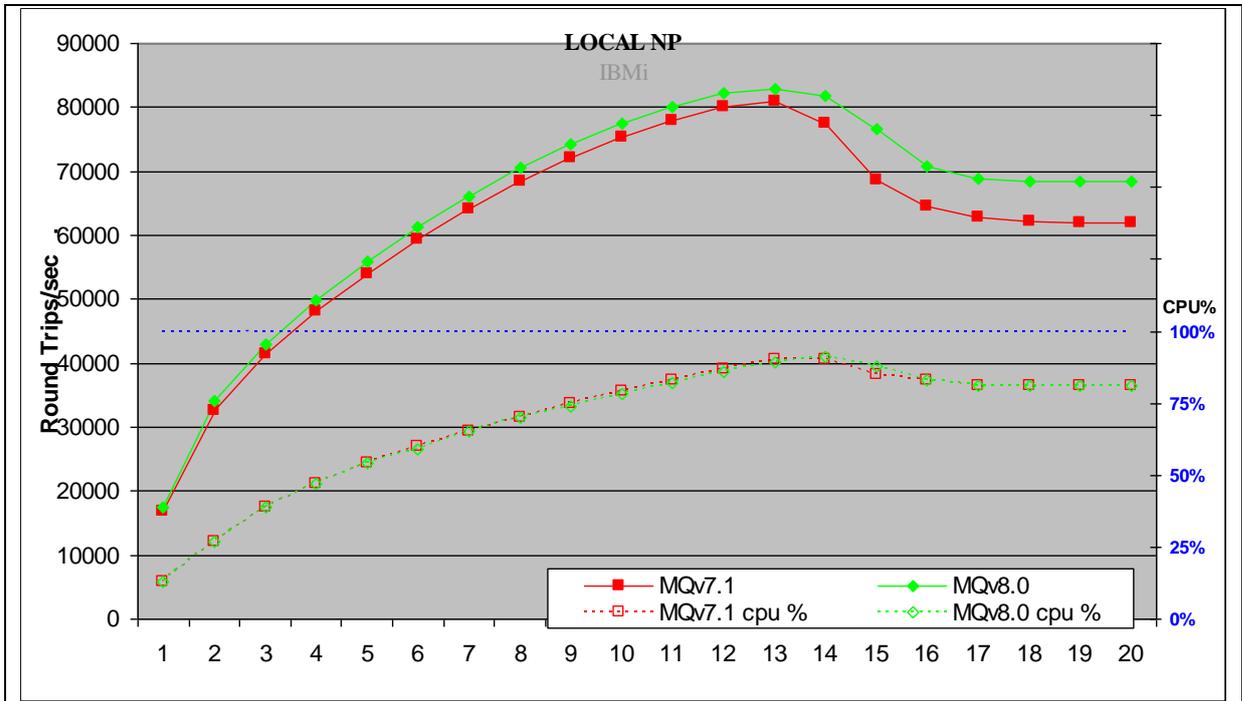


Figure 2 – Connections into a local queue manager

Figure 2 and Table 1 show that the peak throughput of non-persistent messages has increased by 2% when comparing V8.0 to V7.1

Test Name: LOCAL NP	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	13	80944	0.00018	90%
MQv8.0	13	82843	0.00017	89%

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

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

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

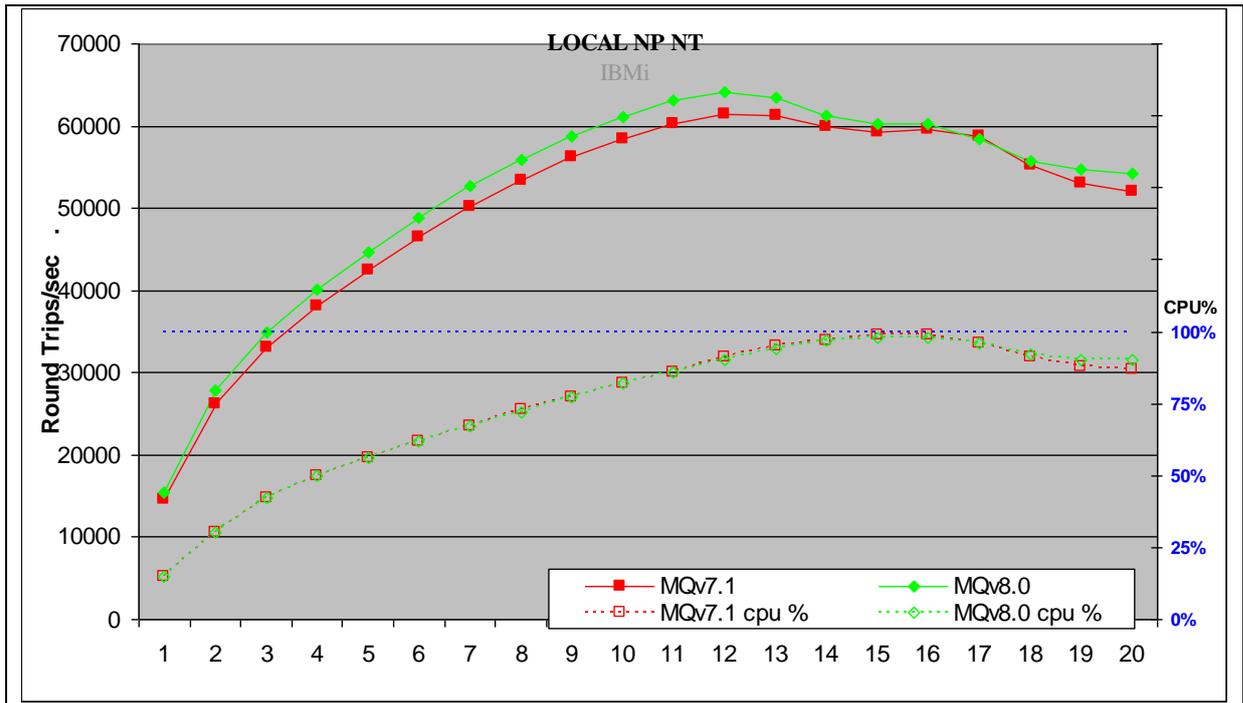


Figure 3 - Performance headline, non-persistent, non-trusted messages and local queue manager.

Figure 3 and Table 2 shows that the peak throughput of non-persistent, non-trusted messages (shared bindings - MQIBINDTYPE=STANDARD) has increased by 4% when comparing V8.0 to V7.1.

Test Name: LOCAL NP NT	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	12	61509	0.00022	91%
MQv8.0	12	64137	0.0002	90%

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

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

### 2.1.3 Persistent Messages – Local Queue Manager

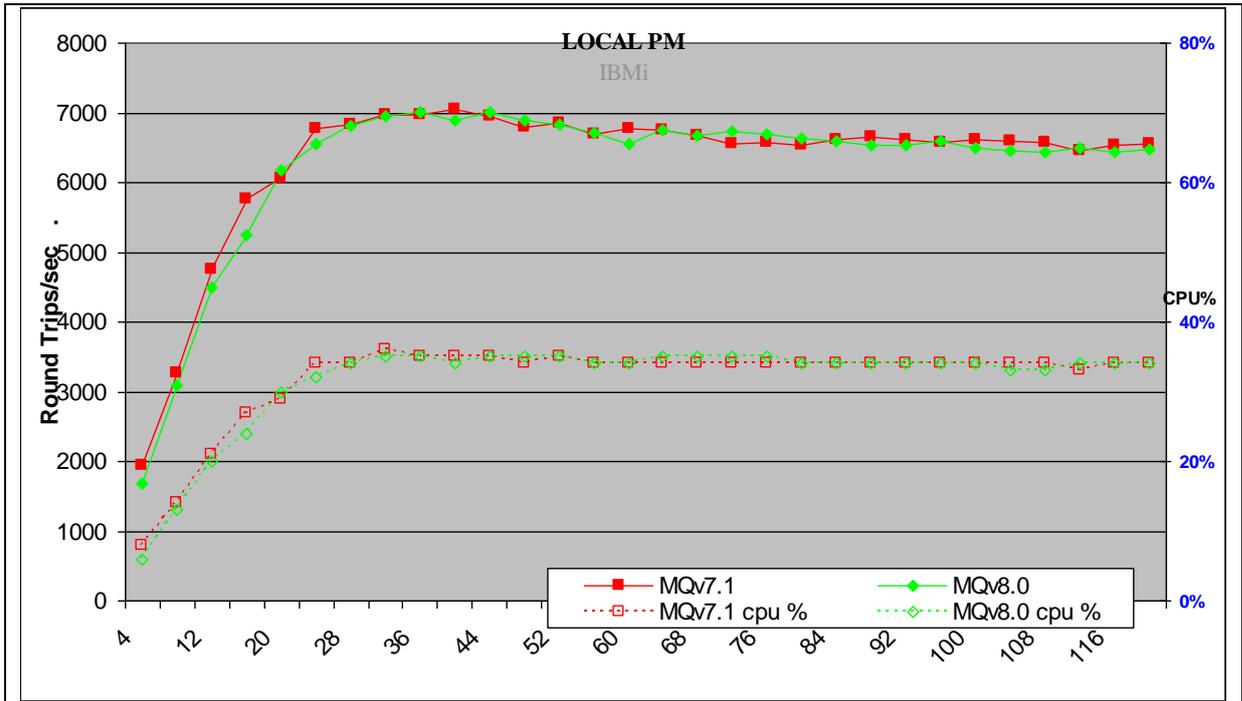


Figure 4 - Performance headline, persistent messages and local queue manager

Figure 4 and Table 3 shows that the peak throughput of persistent messages is similar when comparing V8.0 to V7.1.

Test Name: LOCAL PM	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	40	7045	0.0066	35%
MQv8.0	44	7011	0.0074	35%

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

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

## 2.2 Client Channels Test Scenario

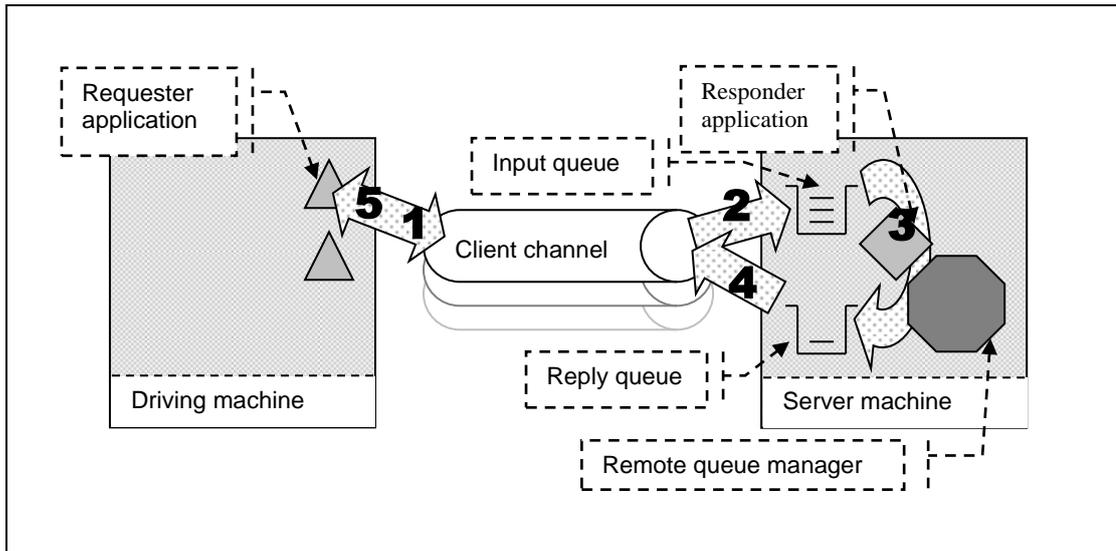


Figure 5 - MQI-client channels into a remote queue manager

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

Non-persistent and persistent messages were used in the client channel tests, with a message size of 2KB. The effect of message throughput with larger message sizes is investigated in section 3.

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

Version 7 onwards will multiplex multiple clients from the same process over one TCP socket. We have standardized all client measurements to use SHARECNV(1) since we have various tests that have between 1 and 100 clients per process and we are interested in results when all the clients come from different computers. Further information is in section 7.1

### 2.2.1 Non-persistent Messages – Client Channels

Figure 6, Figure 8 and Figure 9 shows the non-persistent, non-persistent/non-trusted and persistent message throughput achieved using an increasing number of driving applications in the client channel scenario (see Figure 5 on the previous page) for different production levels of IBM MQ (versions 7.1 and 8.0).

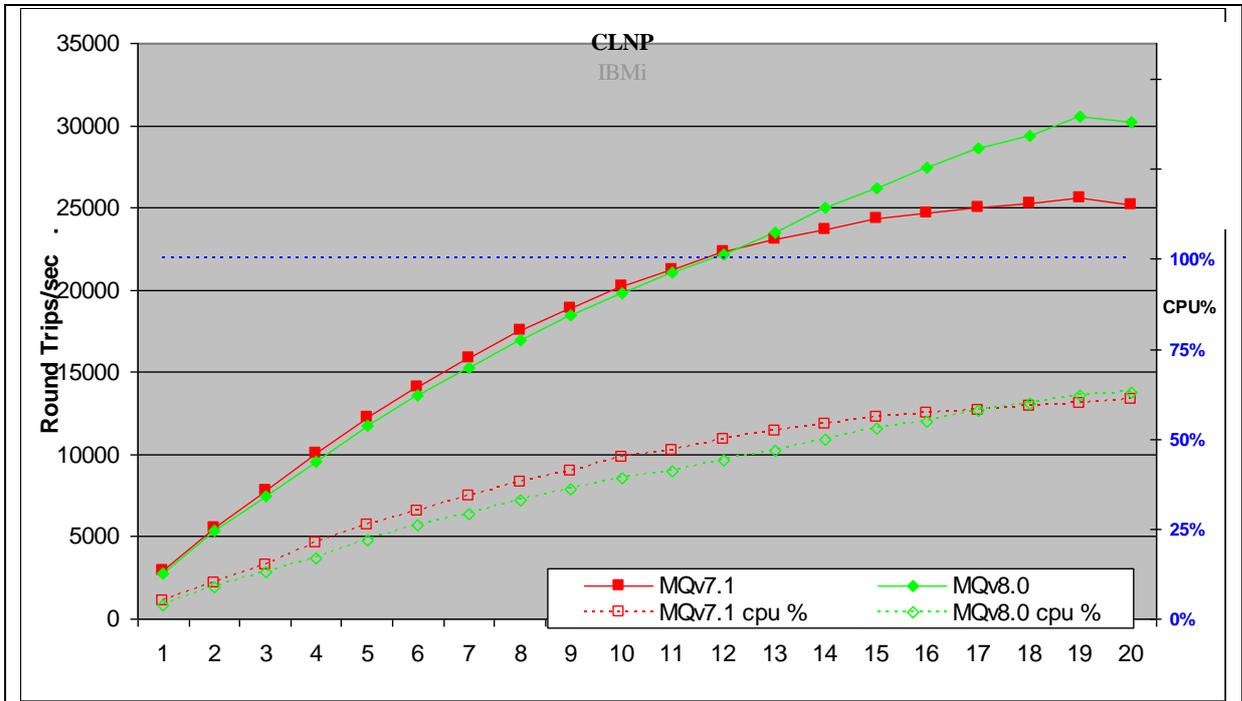


Figure 6 - Performance headline, non-persistent messages and client channels

Figure 6 and Table 4 show that the peak throughput of non-persistent messages has increased by 20% when comparing version 8.0 to 7.1.

Test Name: CLNP	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	19	25636	0.00086	60%
MQv8.0	19	30518	0.00074	62%

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

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

### 2.2.2 Non-persistent Messages – Non-Trusted Client Channels

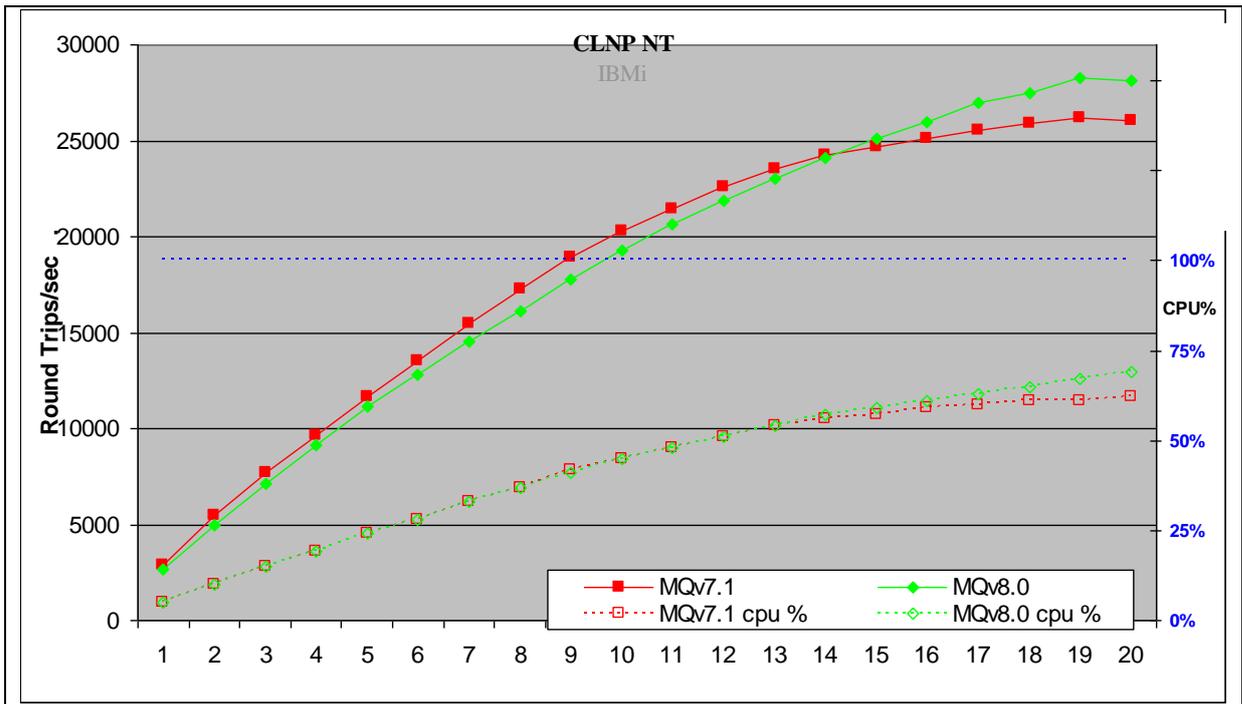


Figure 7 - 1 round trip per driving application per second, client channels, persistent messages

Figure 8 and Table 5 shows that the peak throughput of non-persistent, non-trusted messages (shared bindings - MQIBINDTYPE=STANDARD) has increased by 10% when comparing V8.0 to V7.1

Test Name: CLNP NT	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	19	26191	0.00086	61%
MQv8.0	19	28285	0.0008	67%

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

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

### 2.2.3 Persistent Messages – Client Channels

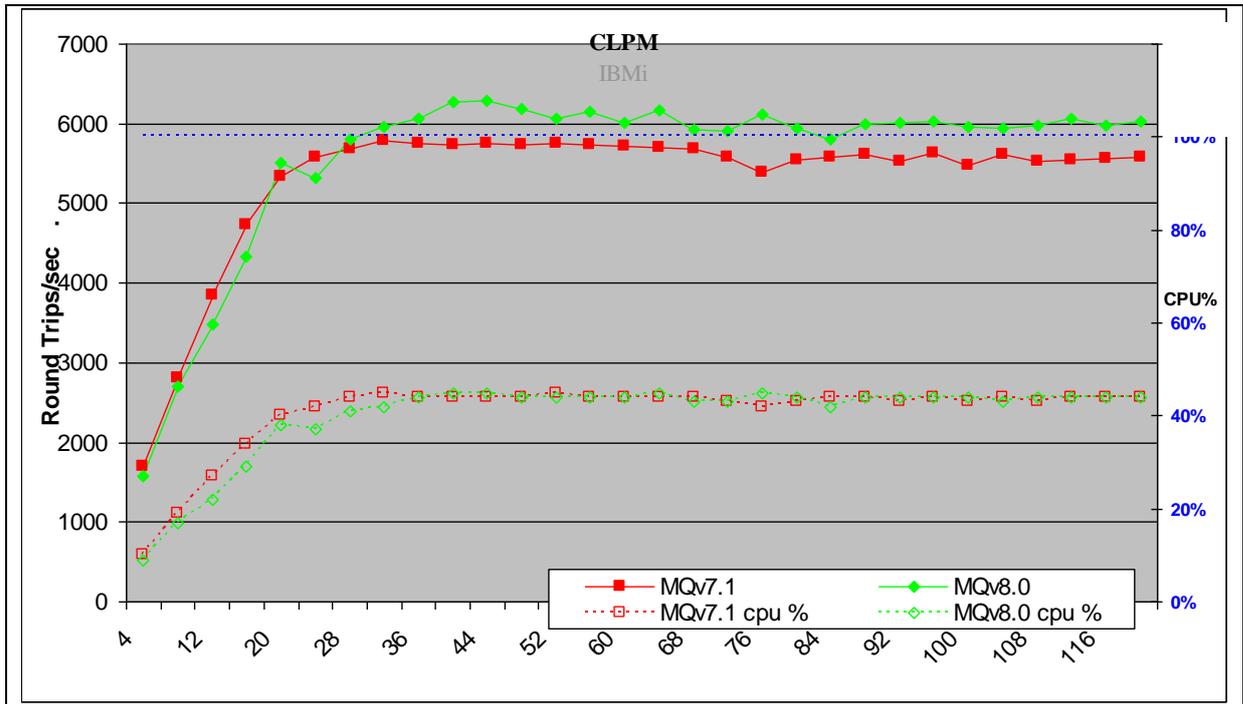


Figure 8 - Performance headline, persistent messages and client channels

Figure 8 and Table 6 shows that the peak throughput of persistent messages has improved by 8% when comparing V8.0 to V7.1

Test Name: CLPM	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	32	5780	0.0062	45%
MQv8.0	44	6291	0.012	45%

Table 6 – Performance headline, persistent messages and client channels

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

### 2.2.4 Client Channels

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

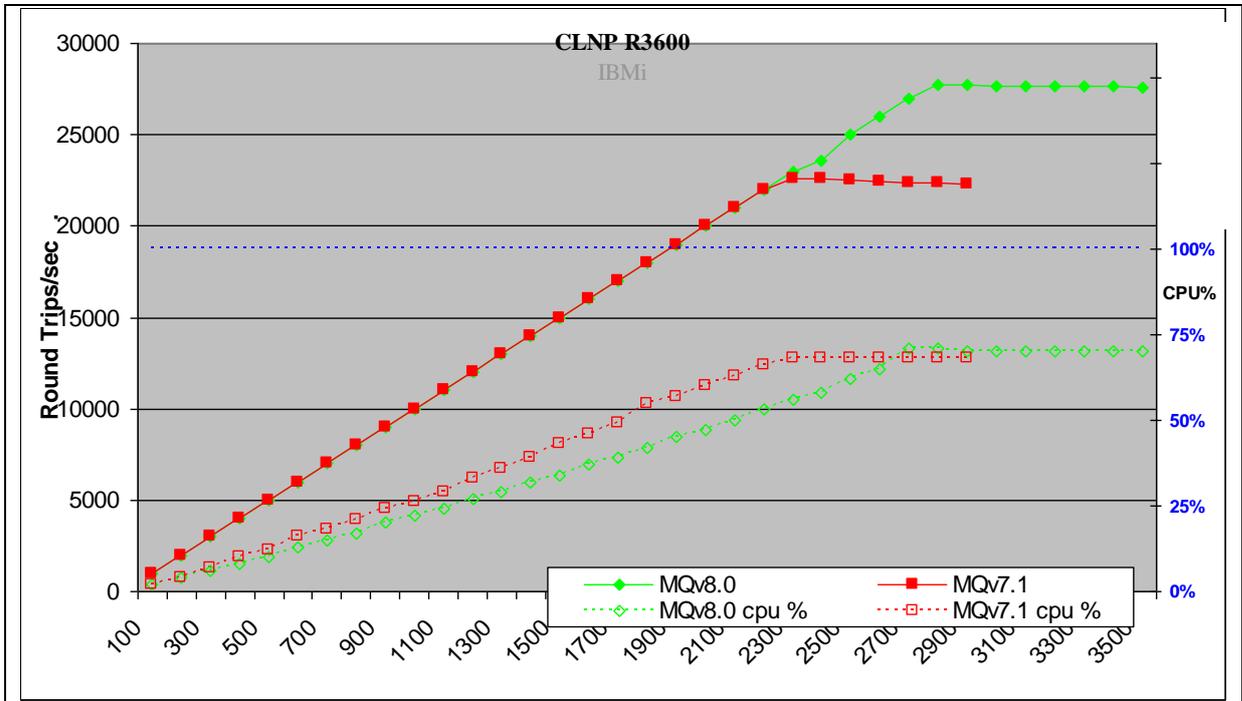


Figure 9 - 1 round trip per driving application per second, client channels and non-persistent messages

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

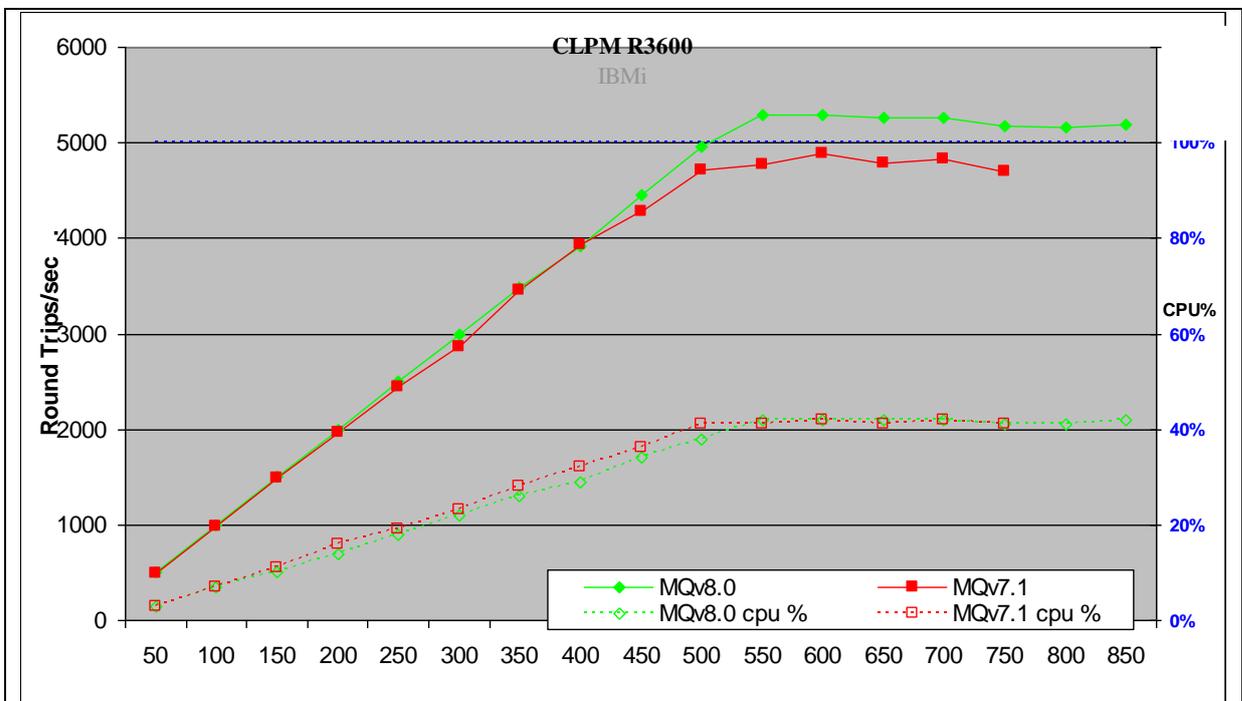


Figure 10 - 1 round trip per driving application per second, client channels, persistent messages

Figure 9, Figure 10 and Table 7 show that the peak throughput of non-persistent messages has improved by 8% when comparing V8.0 to V7.1 but Persistent messages the peak throughput is 22% better when comparing v8.0 to v7.1

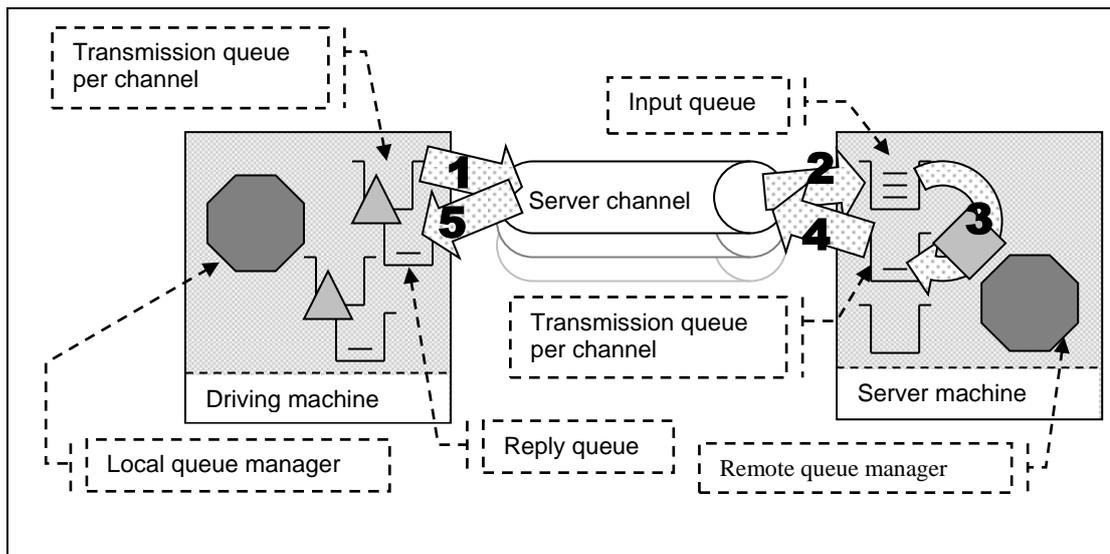
Test Name: CLNP R3600	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	2300	22623	0.069	68%
MQv8.0	2900	27731	0.077	70%

Test Name: CLPM R3600	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	600	4886	0.14	42%
MQv8.0	600	5293	0.12	42%

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

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

## 2.3 Distributed Queuing Test Scenario



**Figure 11 - Server channels between two queue managers**

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

Non-persistent and persistent messages were used in the distributed queuing tests, with a message size of 2KB. The effect of message throughput with larger messages sizes is investigated in section 3.

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

### 2.3.1 Non-persistent Messages – Server Channels

Figure 12, Figure 13 and Figure 14 show the non-persistent, non-persistent/non-trusted and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario (see Figure 11 on the previous page) and IBM MQ (versions 7.1 and 8.0).

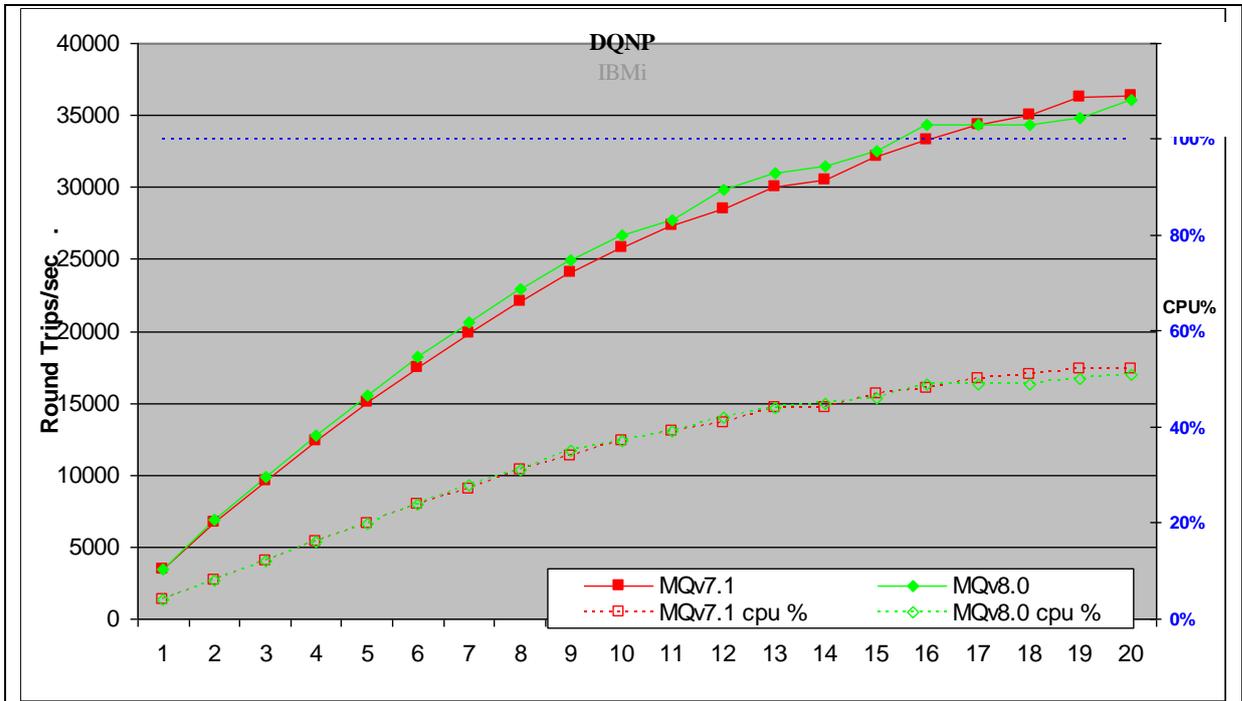


Figure 12 Server channels between two queue managers

Figure 12 and Table 8 shows that the peak throughput of non-persistent messages is similar when comparing version 8.0 to 7.1.

Test Name: DQNP	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	20	36311	0.00064	52%
MQv8.0	20	36090	0.00062	51%

Table 8 – Performance headline, non-persistent messages and server channels

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

### 2.3.2 Non-Persistent non-Trusted – Server Channels

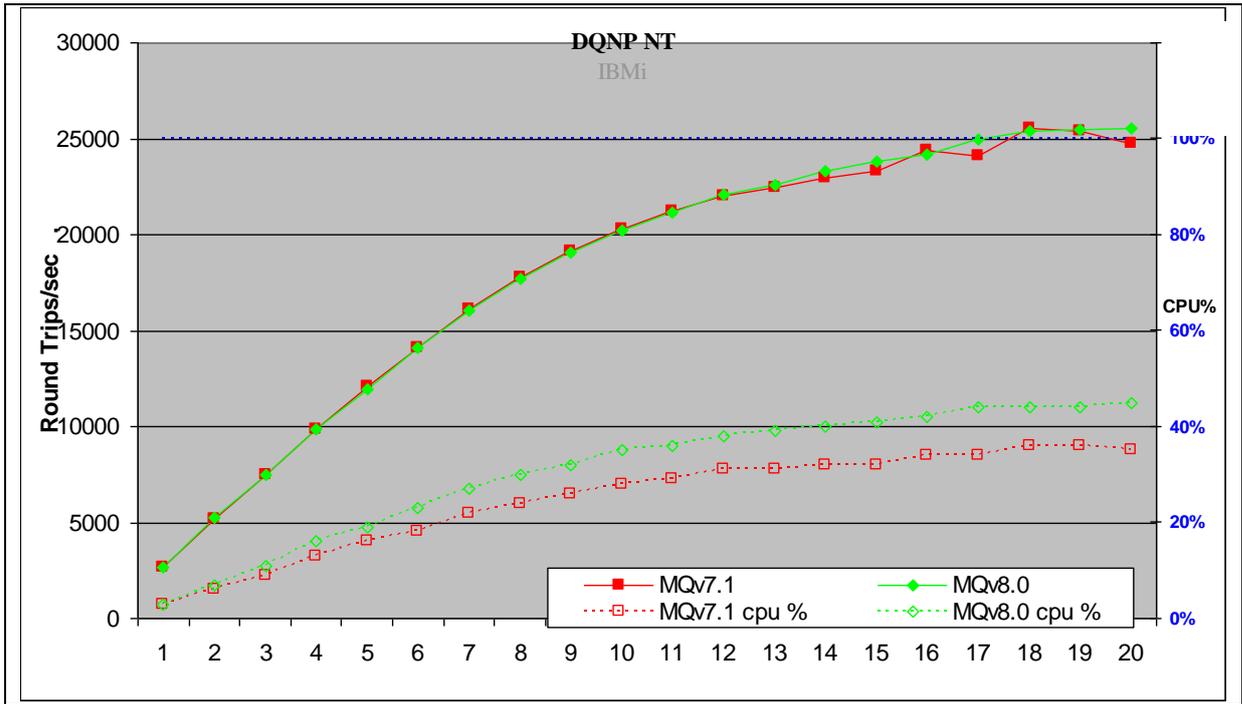


Figure 13 – Performance headline, non-persistent, not trusted messages and server channels

Figure 13 and Table 9 shows that the peak throughput of non-persistent, non-trusted messages is similar when comparing version 8.0 to 7.1

Test Name: DQNP NT	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	18	25517	0.00085	36%
MQv8.0	20	25573	0.001	45%

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

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

### 2.3.3 Persistent Messages – Server Channels

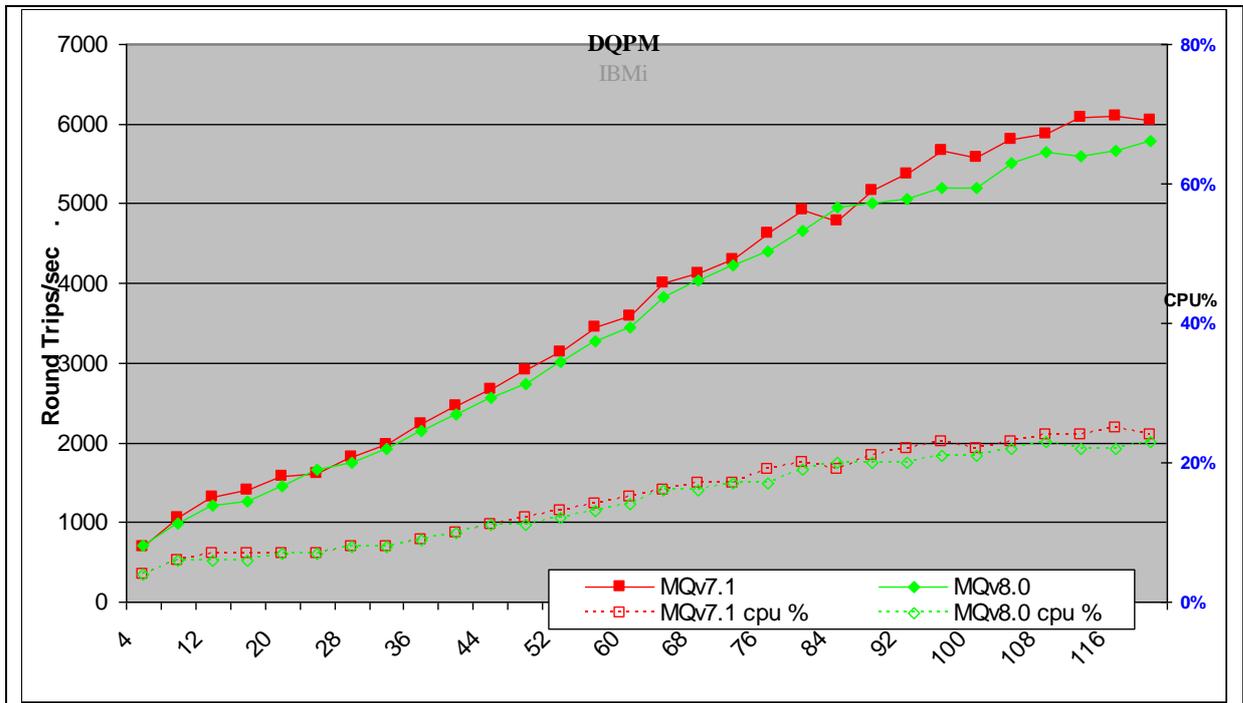


Figure 14 – Performance headline, persistent messages and server channels

Figure 14 and Table 10 shows that the peak throughput of persistent messages using 2 pairs of channels is similar when comparing V8.0 to V7.1.

Test Name: DQPM	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	116	6096	0.027	25%
MQv8.0	120	5795	0.021	23%

Table 10 – Performance headline, persistent messages and server channels

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

### 2.3.4 Server Channels

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

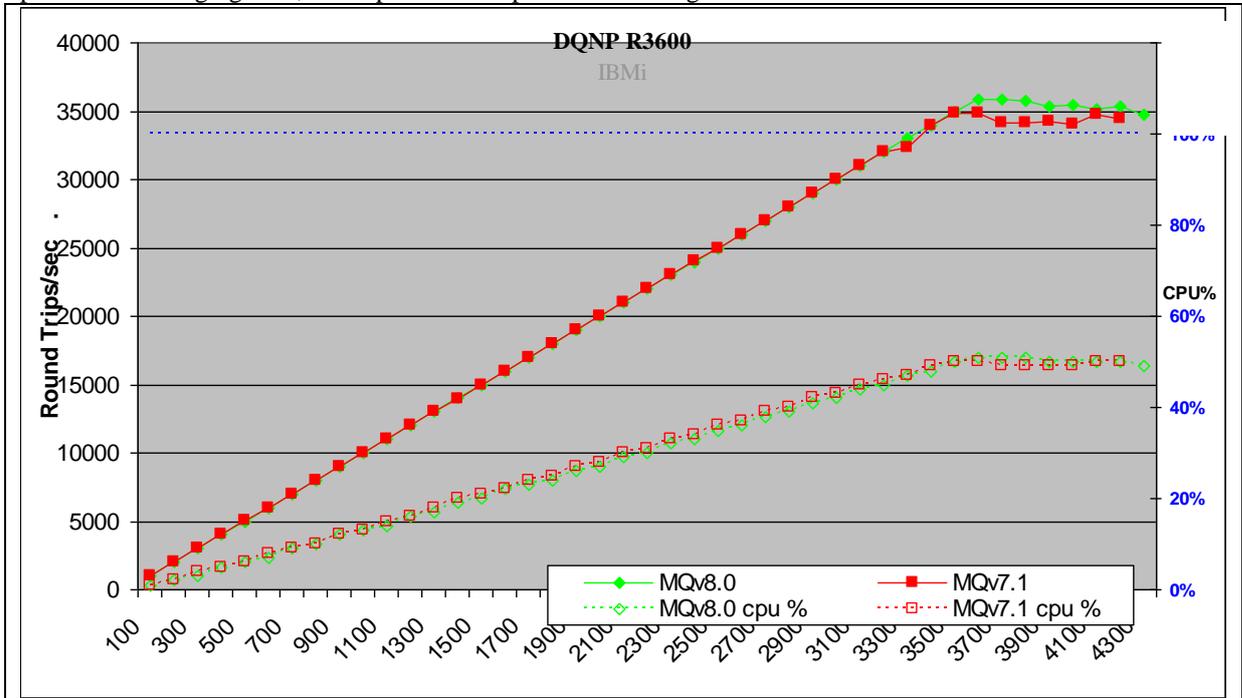


Figure 15 – 1 round trip per driving application per second, server channel, non-persistent messages

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

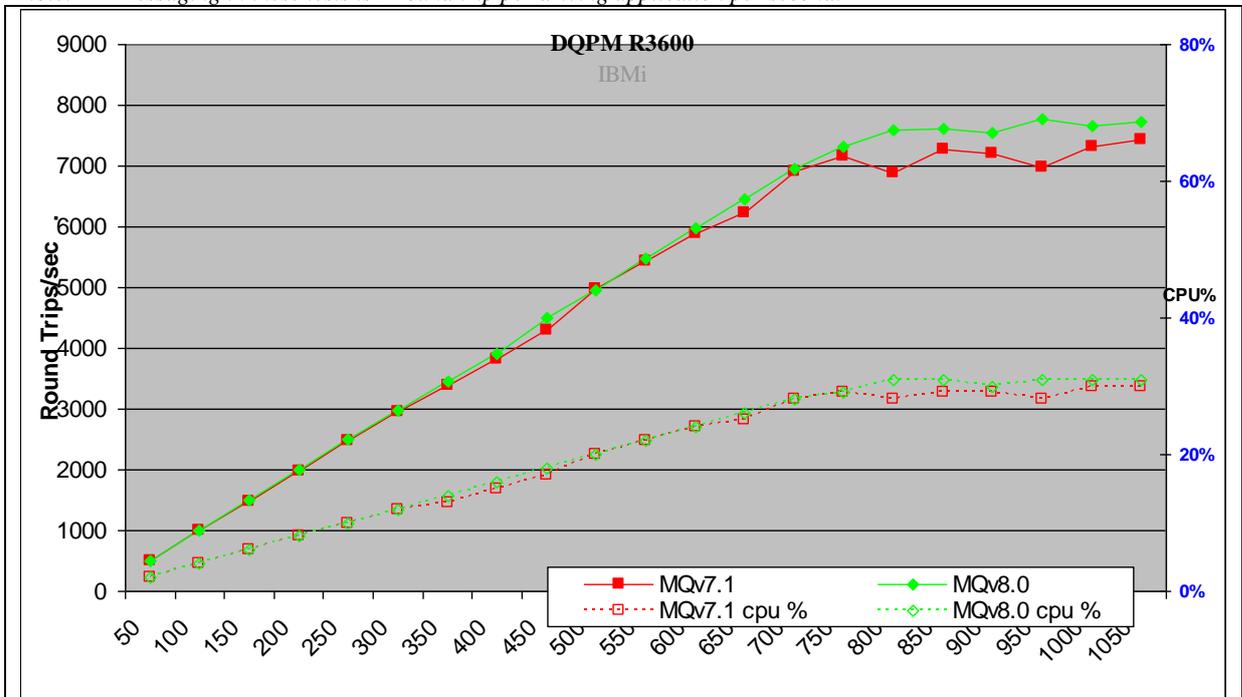


Figure 16 – 1 round trip per driving application per second, server channel, persistent messages

Figure 15 and Figure 16 shows that the throughput of non-persistent and persistent messages has improved by 3% and 5% respectively when comparing version 8.0 to 7.1.

Test Name: DQNP R3600	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	3500	34884	0.028	50%
MQv8.0	3700	35850	0.079	51%

Test Name: DQPM R3600	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	1050	7429	0.16	30%
MQv8.0	950	7777	0.14	31%

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

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

### 3 Large Messages

#### 3.1 20KB Messages

##### 3.1.1 Local Queue Manager

Figure 22 and Figure 23 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

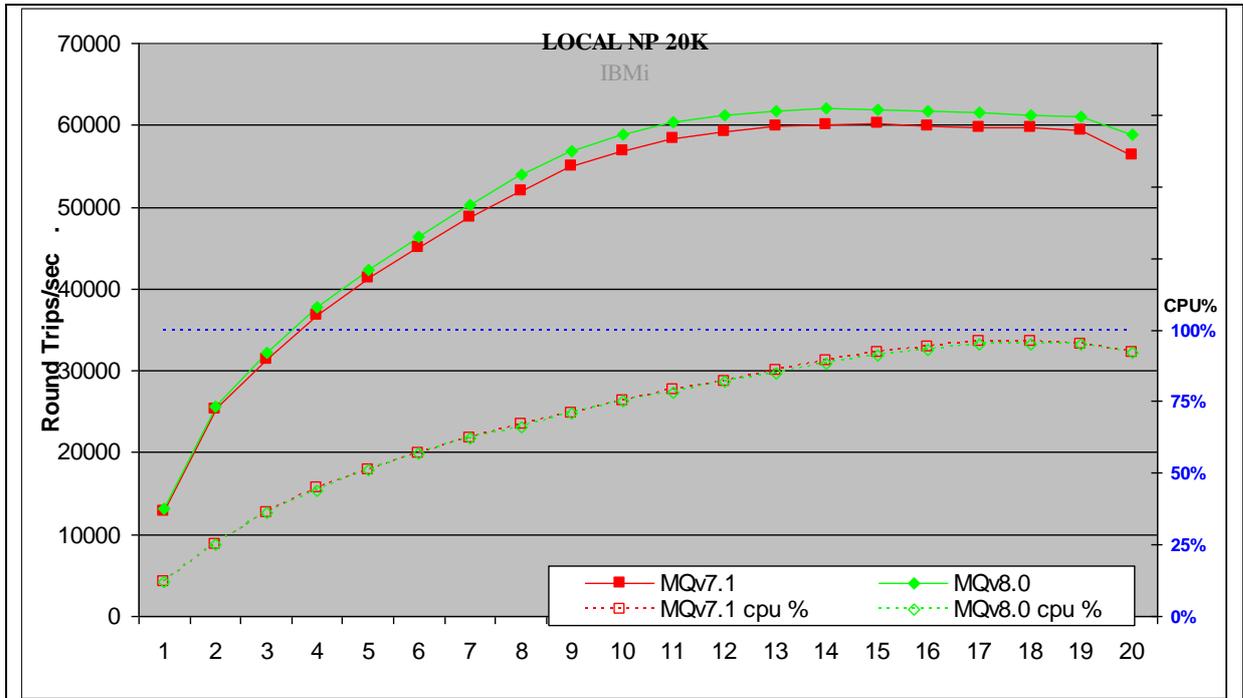


Figure 17 – 20KB non-persistent messages, local queue manager

Figure 17 and Table 12 shows that the peak throughput of non-persistent messages has increased by 3% when comparing V8.0 to V7.1.

Test Name: LOCAL NP 20K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	15	60224	0.00027	92%
MQv8.0	14	62043	0.00025	88%

Table 12 – 20KB non-persistent messages, local queue manager

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

### 3.1.1.1 Persistent Messages

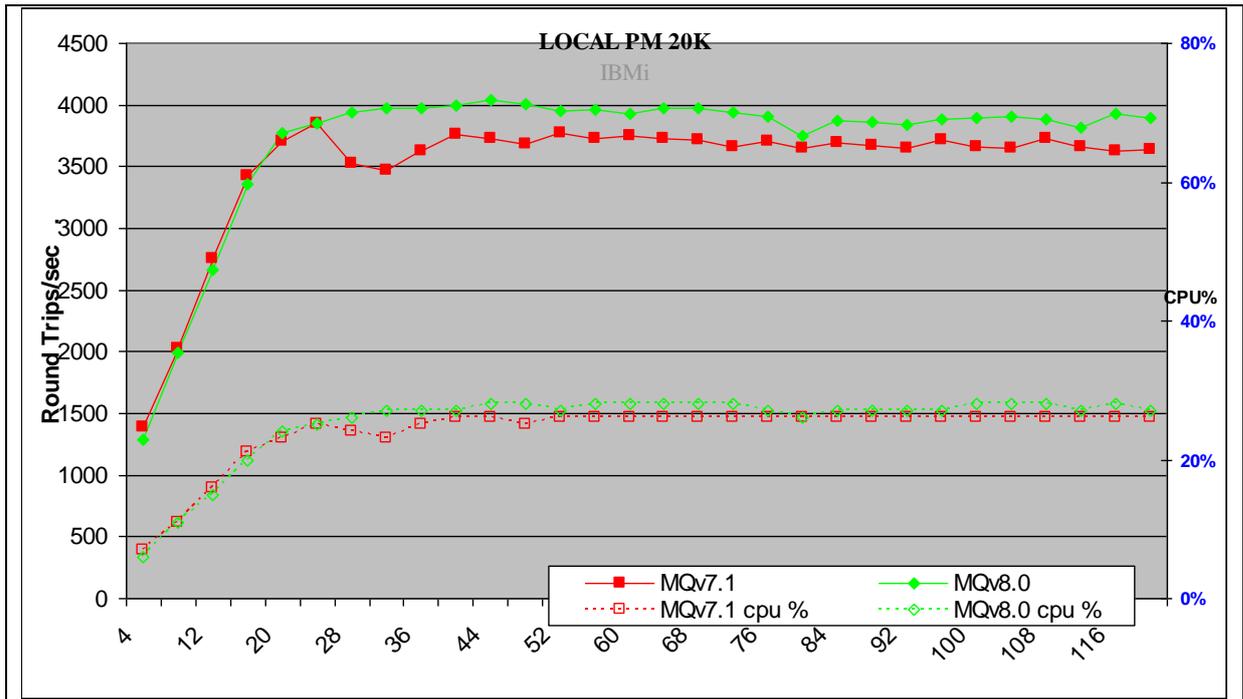


Figure 18 – 20KB persistent messages, local queue manager

Figure 18 and Table 13 shows that the peak throughput of persistent messages has increased by 5% when comparing V8.0 to V7.1.

Test Name: LOCAL PM 20K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	24	3852	0.0069	25%
MQv8.0	44	4037	0.012	28%

Table 13 – 20KB persistent messages, local queue manager

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

### 3.1.2 Client Channel

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

#### 3.1.2.1 Non-persistent Messages

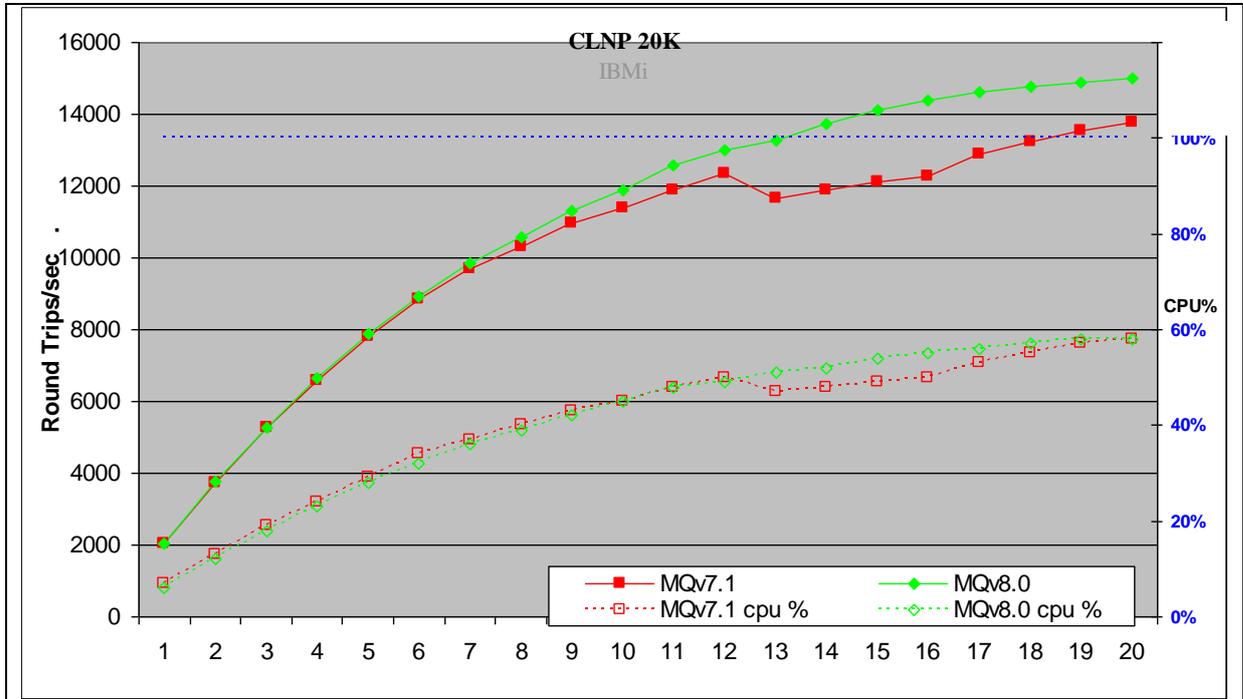


Figure 19 – 20KB persistent messages, local queue manager

Figure 19 and Table 14 shows that the peak throughput of non-persistent messages is 9% improved when comparing V8.0 to V7.1.

Test Name: CLNP 20K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	20	13760	0.0017	58%
MQv8.0	20	14996	0.0015	58%

Table 14 – 20KB non-persistent messages, client channels

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

### 3.1.2.2 Persistent Messages

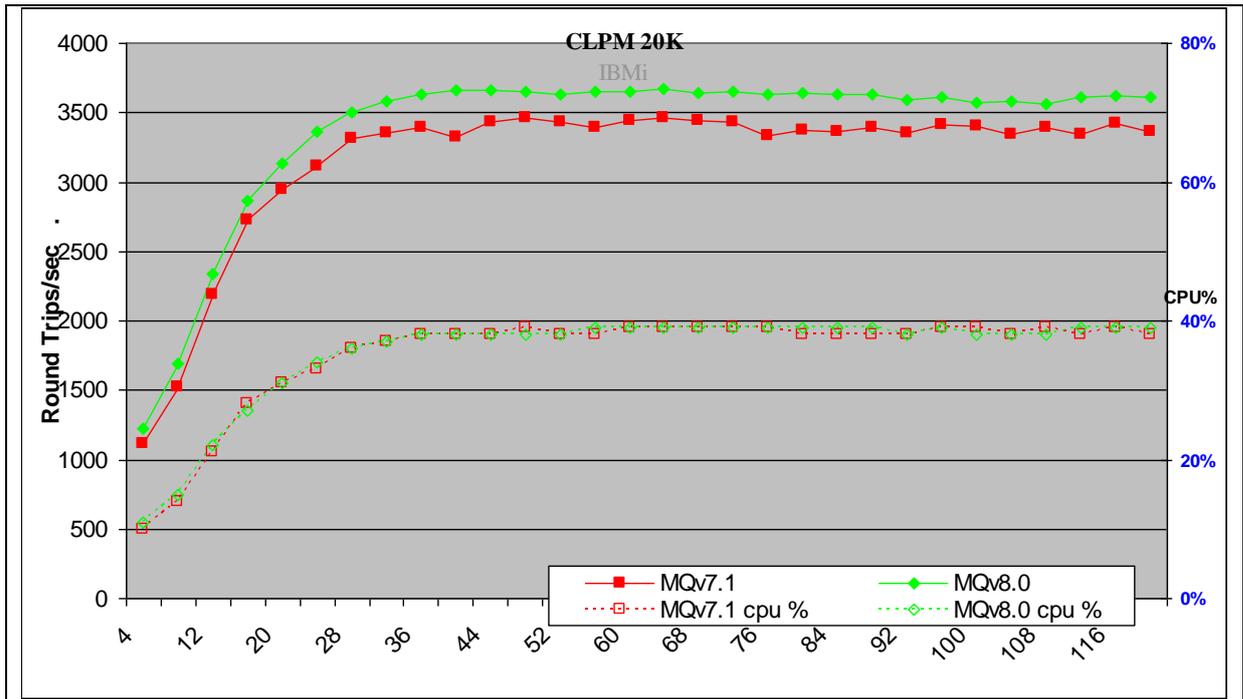


Figure 20 – 20KB persistent messages, client channels

Figure 20 and Table 15 shows that the peak throughput of persistent messages has increased by 6% when comparing version 8.0 to 7.1.

Test Name: CLPM 20K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	48	3462	0.015	39%
MQv8.0	64	3672	0.02	39%

Table 15 – 20KB persistent messages, client channels

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

### 3.1.3 Distributed Queuing

Figure 27 and Figure 28 shows the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the distributed queuing scenario.

For the non-persistent 200K distributed queuing scenario, a migrated queue manager from MQ V7.1 to MQ V8.0 was used. This was done to provide a true comparison of throughput regardless of TCP buffer settings on the system. See section 6.4 for an explanation of TCP buffer changes in MQ V8.0

#### 3.1.3.1 Non-persistent Messages

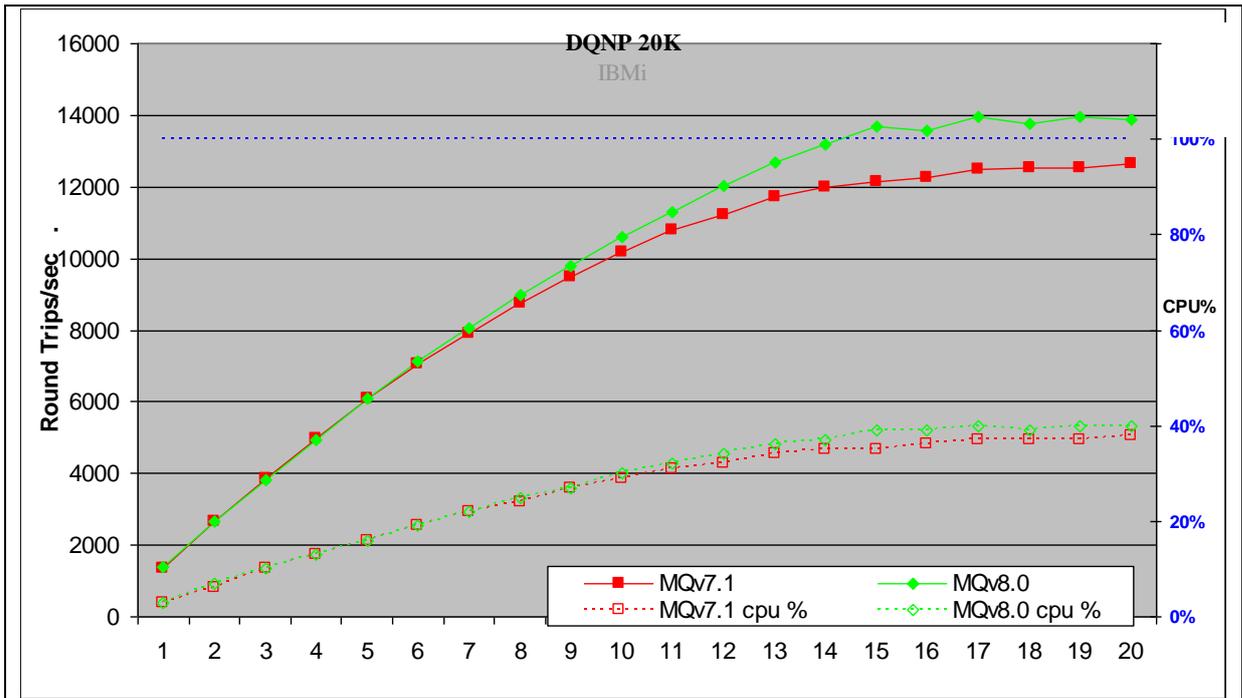


Figure 21 – 20KB non-persistent messages, distributed queuing

Figure 21 and Table 16 shows that the peak throughput of non-persistent messages has improved by 10% when comparing V8.0 to V7.1.

Test Name: DQNP 20K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	20	12648	0.0019	38%
MQv8.0	19	13950	0.0015	40%

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

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

### 3.1.3.2 Persistent Messages

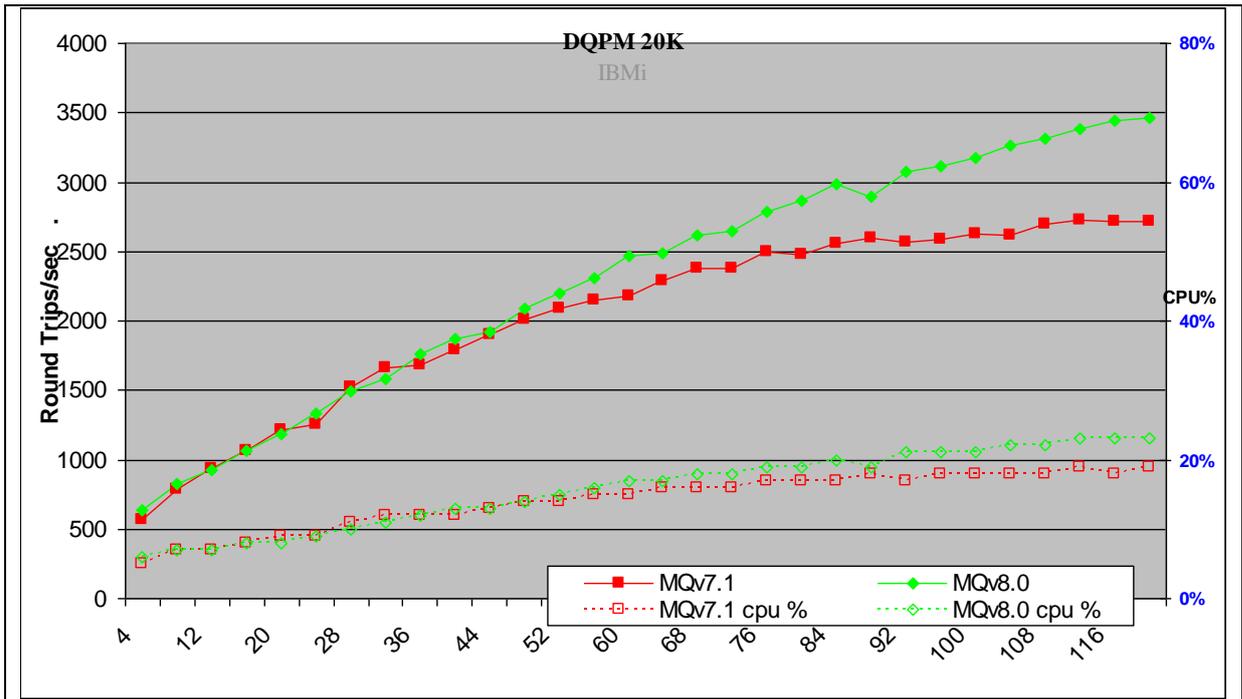


Figure 22 – 20KB persistent messages, distributed queuing

Figure 22 and Table 17 shows that the peak throughput of persistent messages has increased by 27% when comparing V8.0 to V7.1.

Test Name: DQPM 20K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	112	2724	0.051	19%
MQv8.0	120	3463	0.038	23%

Table 17 – 20KB persistent messages, client channels

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

### 3.2 200K Messages

#### 3.2.1 Local Queue Manager

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

##### 3.2.1.1 Non-persistent Messages

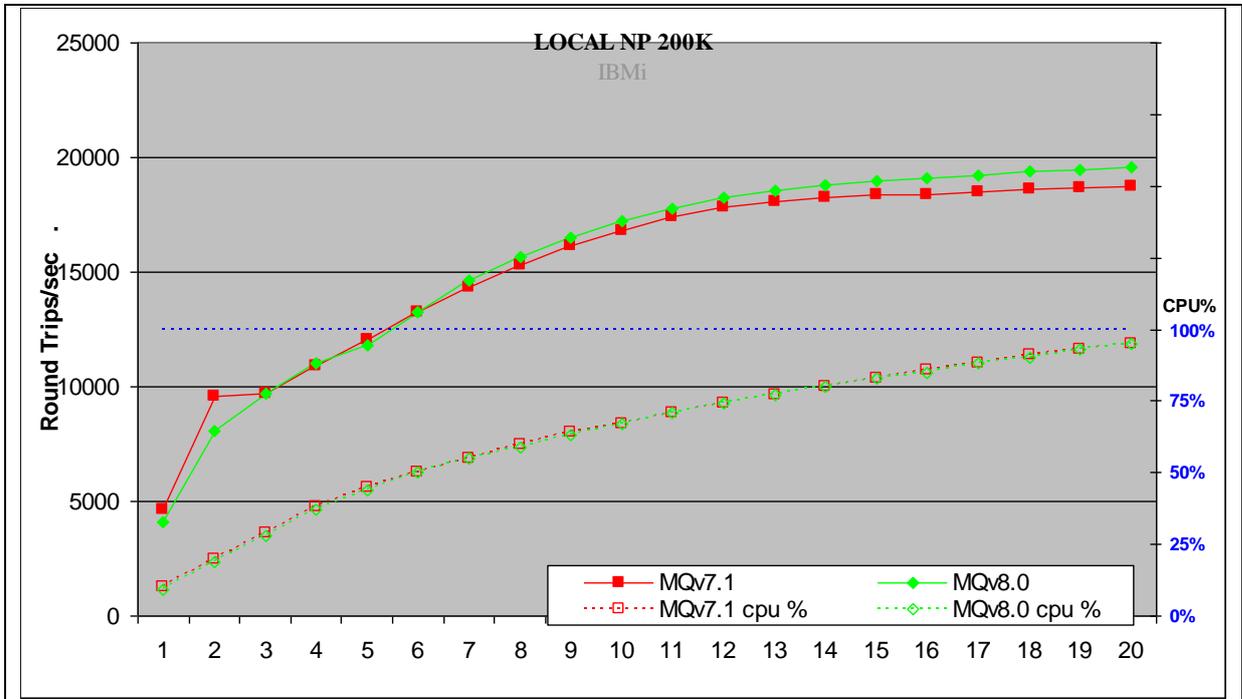


Figure 23 – 200KB non-persistent messages, local queue manager

Figure 23 and Table 18 shows that the peak throughput of non-persistent messages has improved by 4% when comparing V8.0 to V7.1.

Test Name: LOCAL NP 200K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	20	18763	0.001	95%
MQv8.0	20	19586	0.00098	95%

Table 18 – 200KB non-persistent messages, local queue manager

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

### 3.2.1.2 Persistent Messages

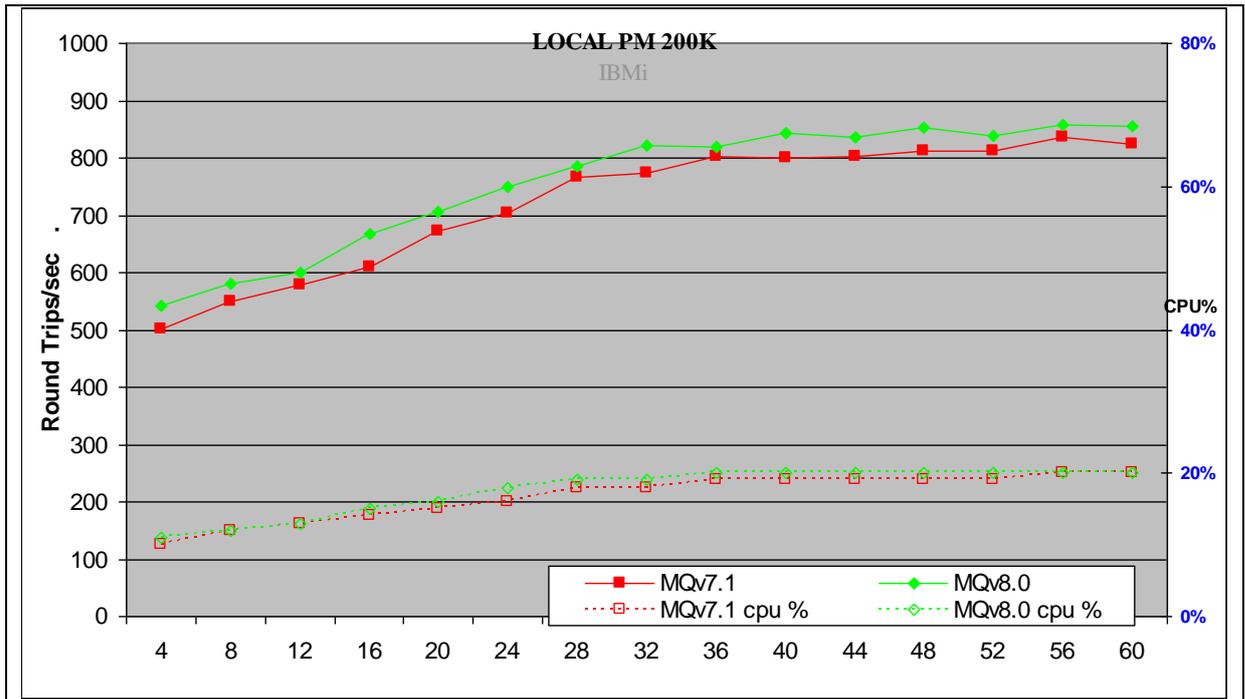


Figure 24 – 200KB persistent messages, local queue manager

Figure 24 and Table 19 shows that the peak throughput of persistent messages is similar when comparing V8.0 to V7.1.

Test Name: LOCAL PM 200K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	56	836	0.077	20%
MQv8.0	56	858	0.082	20%

Table 19 – 200KB persistent messages, local queue manager

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

### 3.2.2 Client Channel

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

#### 3.2.2.1 Non-persistent Messages

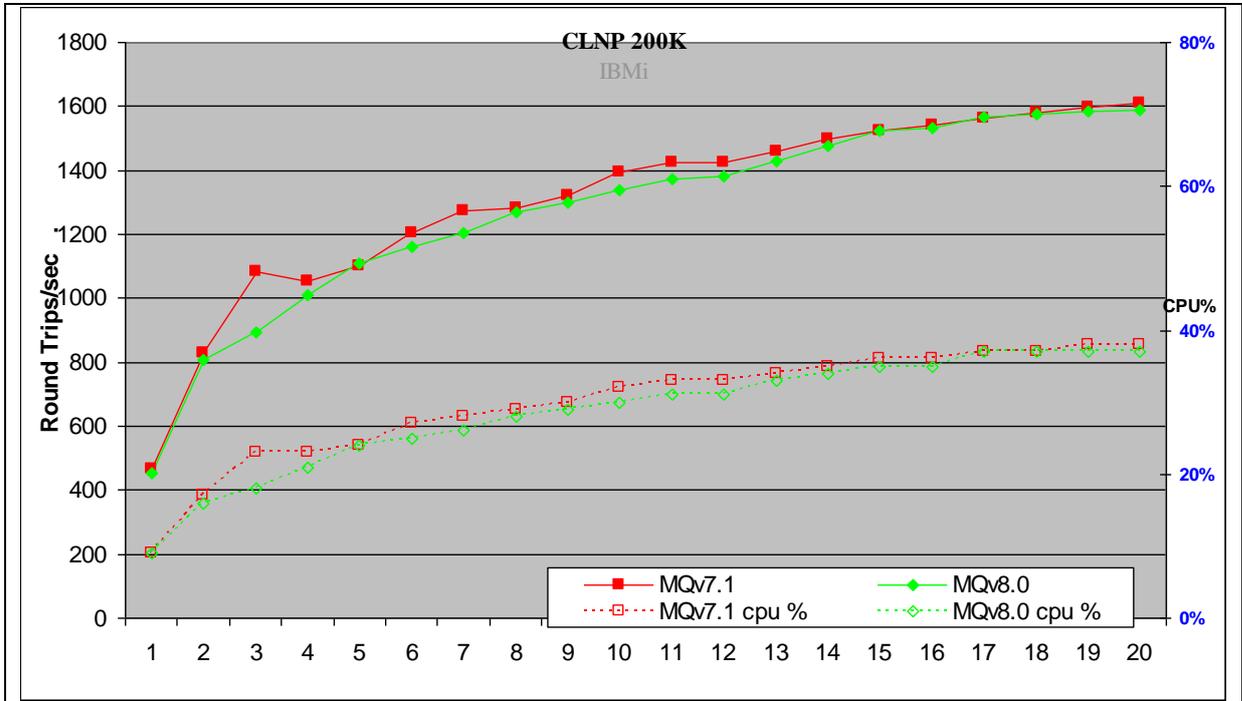


Figure 25 – 200KB non-persistent messages, client channels

Figure 25 and Table 20 shows that the peak throughput of non-persistent messages is similar when comparing V8.0 to V7.1.

Test Name: CLNP 200K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	20	1610	0.016	38%
MQv8.0	20	1588	0.01	37%

Table 20 – 200KB non-persistent messages, client channels

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

### 3.2.2.2 Persistent Messages

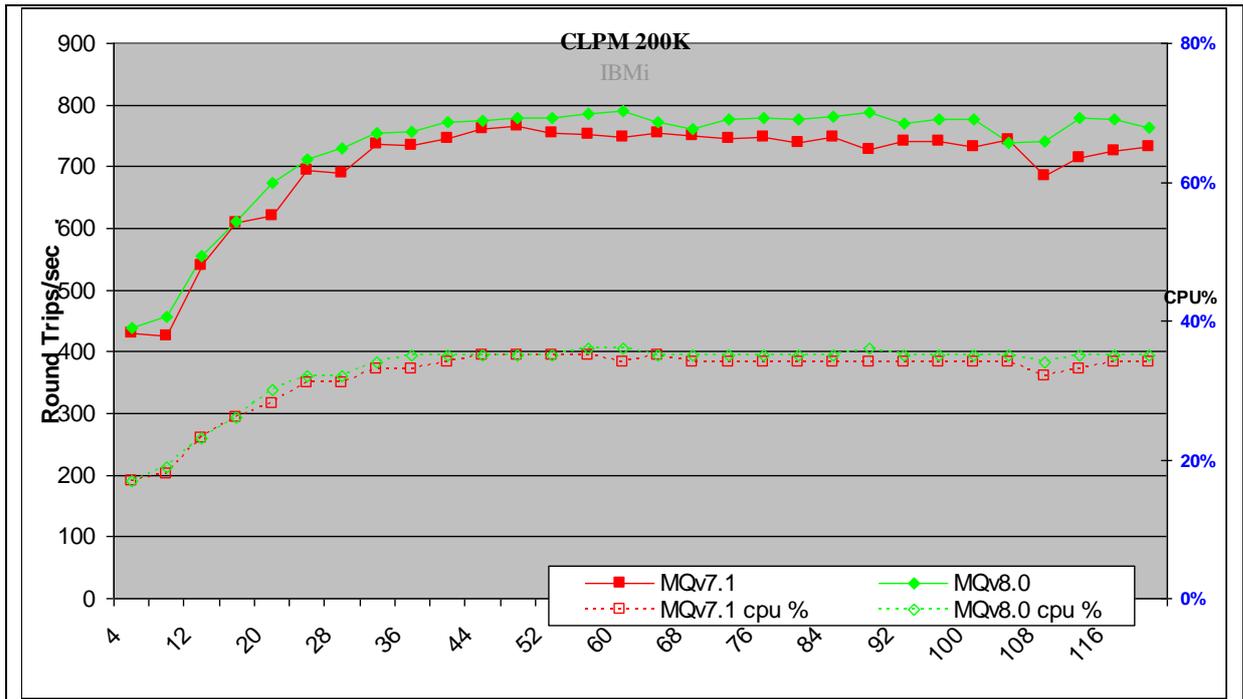


Figure 26 – 200KB persistent messages, client channels

Figure 26 and Table 21 shows that the peak throughput of persistent messages has increased by 3% when comparing V8.0 to V7.1.

Test Name: CLPM 200K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	48	765	0.077	35%
MQv8.0	60	791	0.087	36%

Table 21 – 200KB persistent messages, client channels

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

### 3.2.3 Distributed Queuing

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

For the non-persistent 2MB distributed queuing scenario, a migrated queue manager from MQ V7.1 to MQ V8.0 was used. This was done to provide a true comparison of throughput regardless of TCP buffer settings on the system. See section 6.4 for an explanation of TCP buffer changes in MQ V8.0

#### 3.2.3.1 Non-persistent Messages

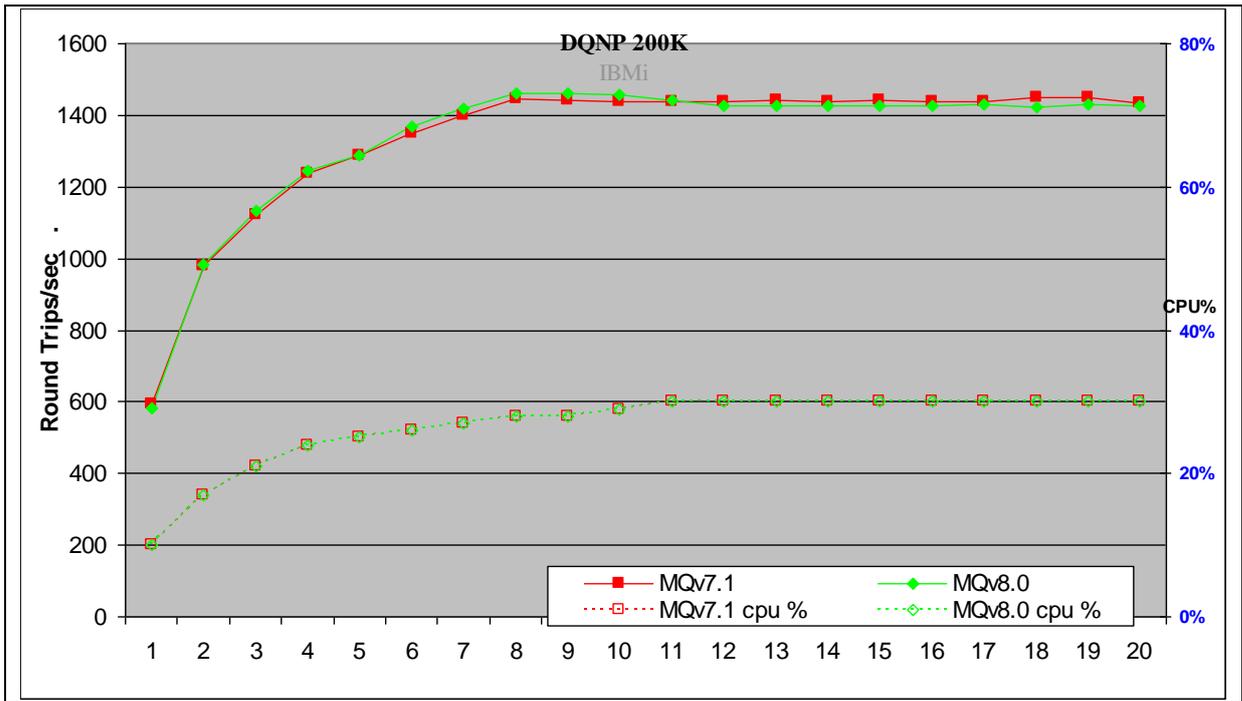


Figure 27 – 200KB non-persistent messages, distributed queuing

Figure 27 and Table 22 shows that the throughput of non-persistent messages is similar when comparing V8.0 to V7.1. At lower number of applications, V8.0 is better than V7.1.

Test Name: DQNP 200K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	18	1449	0.015	30%
MQv8.0	8	1460	0.0068	28%

Table 22 – 200KB non-persistent messages, distributed queuing

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

### 3.2.3.2 Persistent Messages

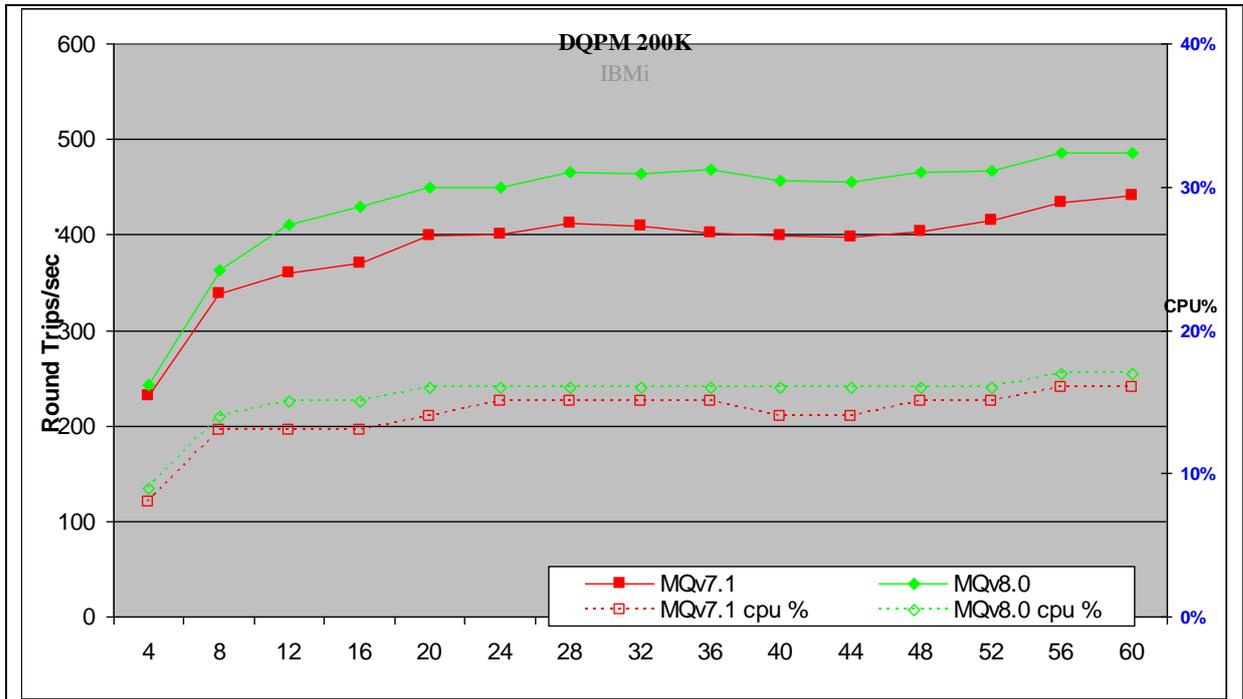


Figure 28 – 200KB persistent messages, distributed queuing

Figure 28 and Table 23 shows that the peak throughput of persistent messages has improved by 10% when comparing V8.0 to V7.1.

Test Name: DQPM 200K	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	60	441	0.15	16%
MQv8.0	56	486	0.13	17%

Table 23 – 200KB persistent messages, distributed queuing

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

### 3.3 2MB Messages

#### 3.3.1 Local Queue Manager

Figure 29 and Figure 30 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

##### 3.3.1.1 Non-persistent Messages

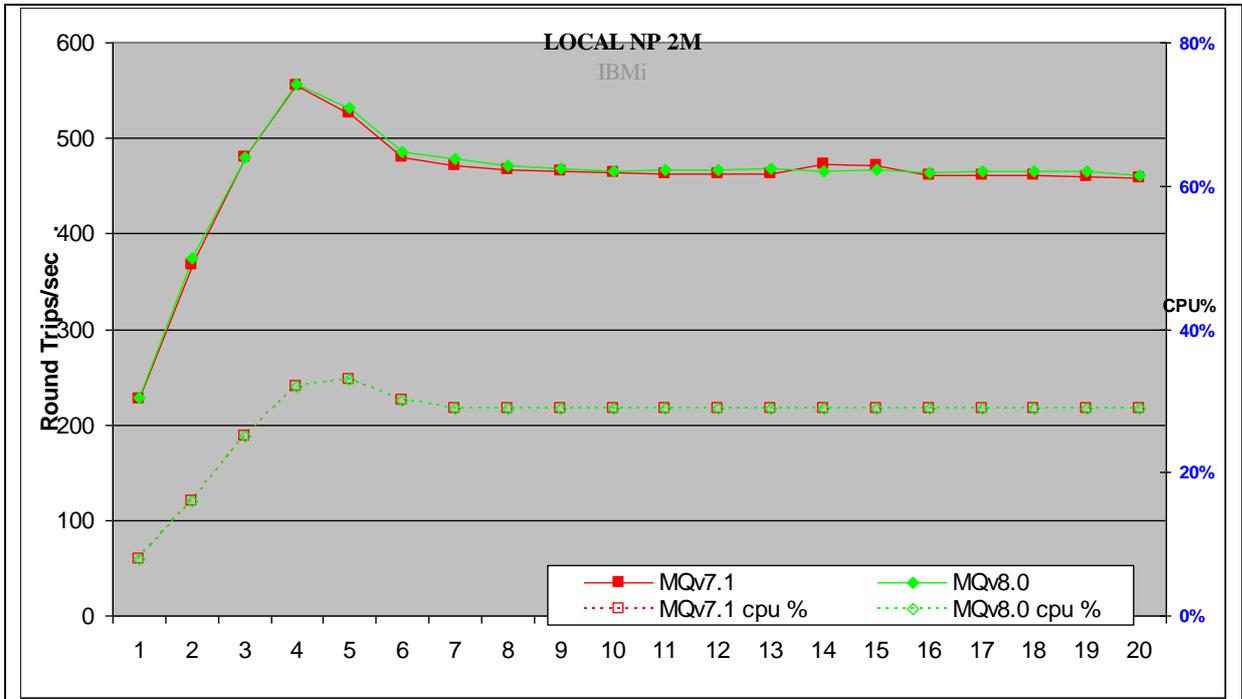


Figure 29 – 2MB non-persistent messages, local queue manager

Figure 29 and Table 24 shows that the peak throughput of non-persistent messages is similar when comparing V8.0 to V7.1.

Test Name: LOCAL NP 2M	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	4	555	0.0072	32%
MQv8.0	4	557	0.0072	32%

Table 24 – 2MB non-persistent messages, local queue manager

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

### 3.3.1.2 Persistent Messages

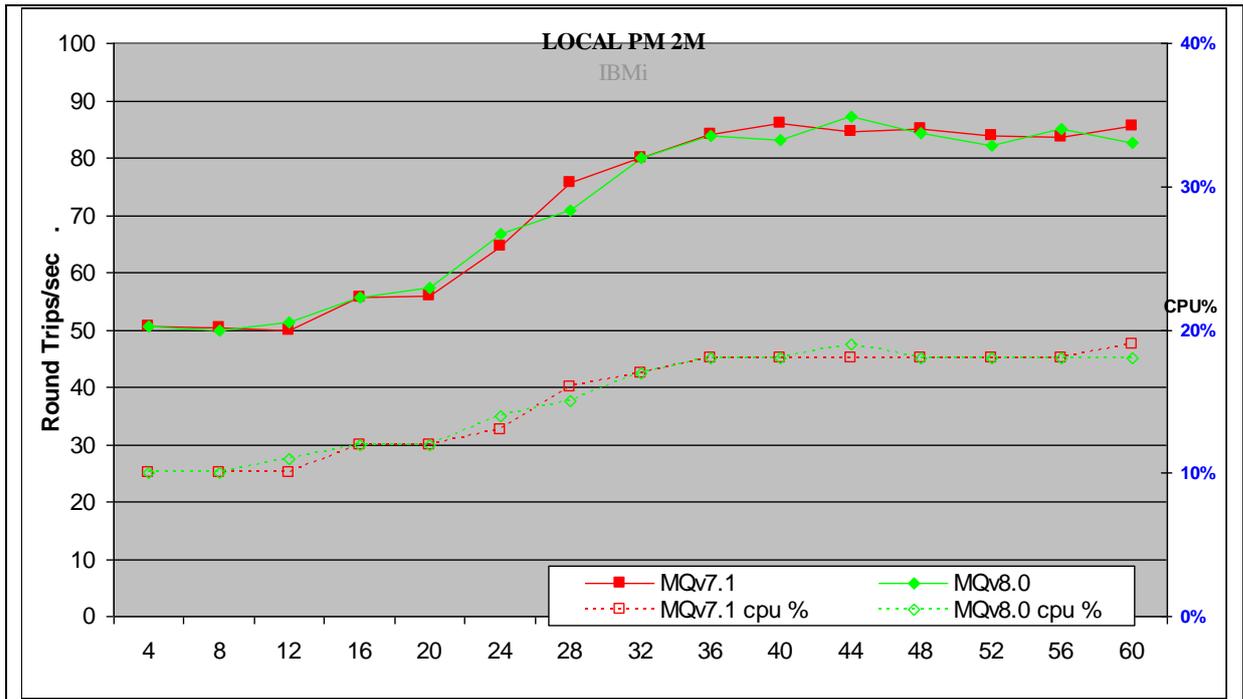


Figure 30 – 2MB persistent messages, local queue manager

Figure 30 and Table 25 shows that the peak throughput of persistent messages is similar when comparing V8.0 to V7.1.

Test Name: LOCAL PM 2M	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	40	86	0.55	18%
MQv8.0	44	87	0.61	19%

Table 25– 2MB persistent messages, local queue manager

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

### 3.3.2 Client Channel

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

#### 3.3.2.1 Non-persistent Messages

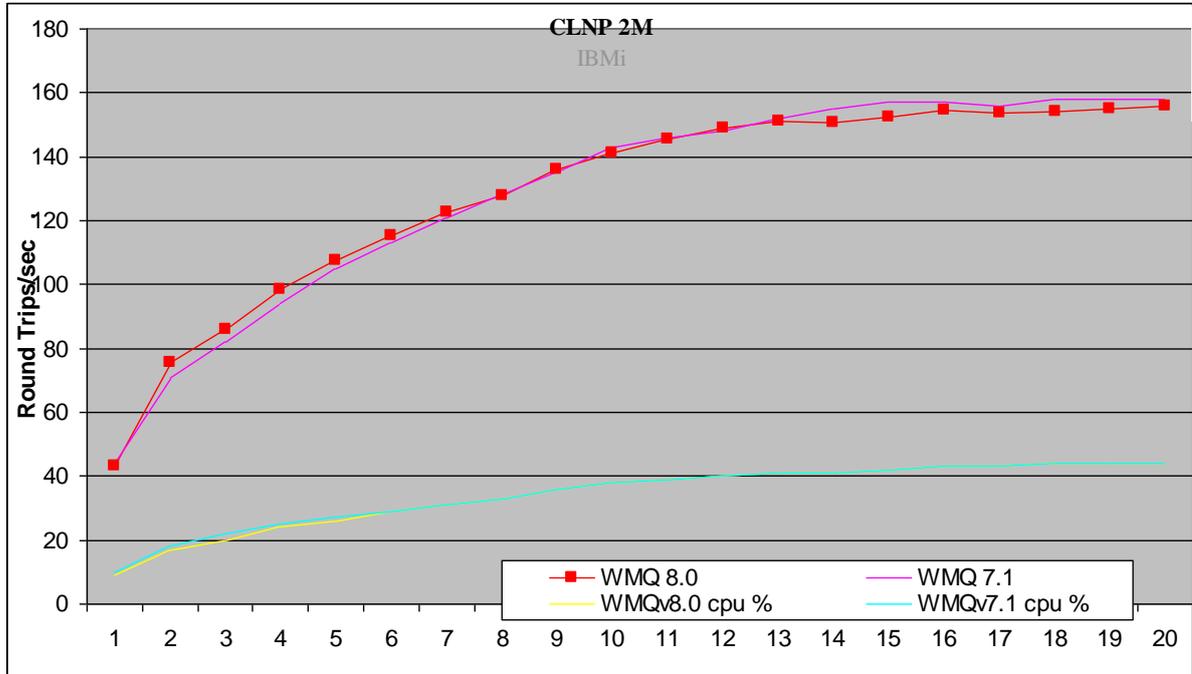


Figure 31 – 2MB non-persistent messages, client channels

Figure 31 and Table 26 shows that the peak throughput of non-persistent messages is similar when comparing V8.0 to V7.1.

Test Name: CLNP 2M	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	20	158	0.0015	44%
MQv8.0	20	156	0.15	44%

Table 26 – 2MB non-persistent messages, client channels

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

### 3.3.2.2 Persistent Messages

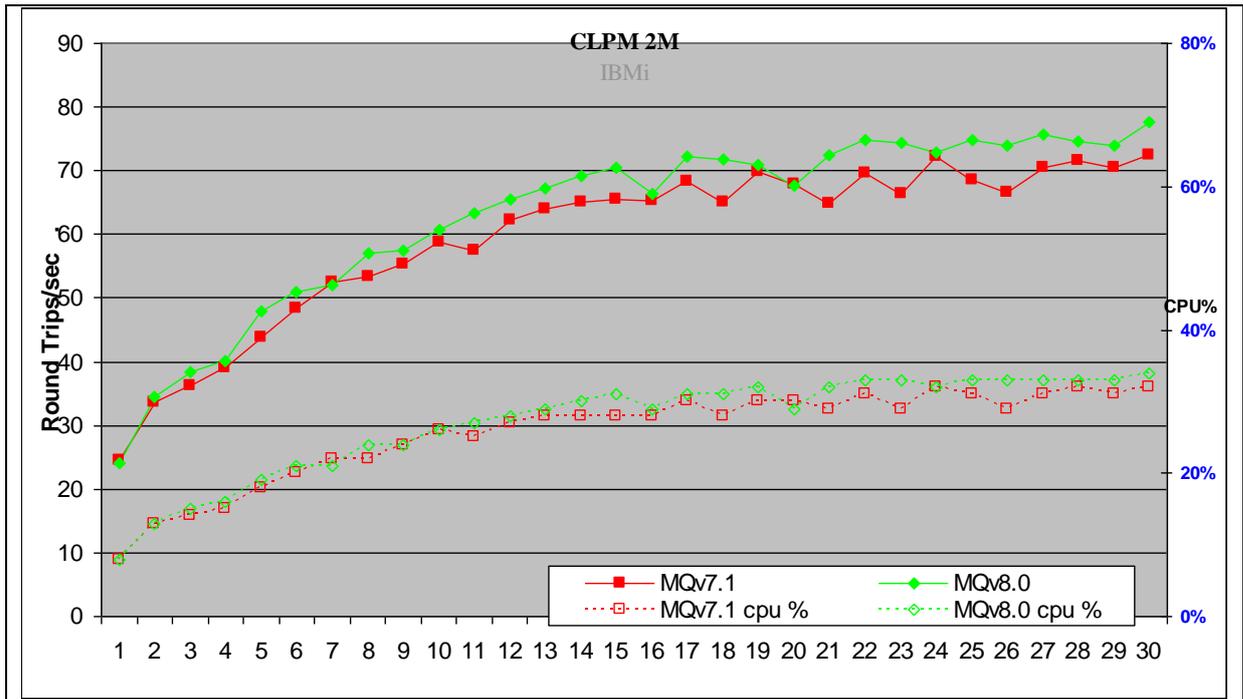


Figure 32 – 2MB persistent messages, client channels

Figure 32 and Table 27 shows that the peak throughput of persistent messages has increased by 8% when comparing V8.0 to V7.1.

Test Name: CLPM 2M	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	30	72	0.48	32%
MQv8.0	30	78	0.45	34%

Table 27 – 2MB persistent messages, client channels

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

### 3.3.3 Distributed Queuing

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

#### 3.3.3.1 Non-persistent Messages

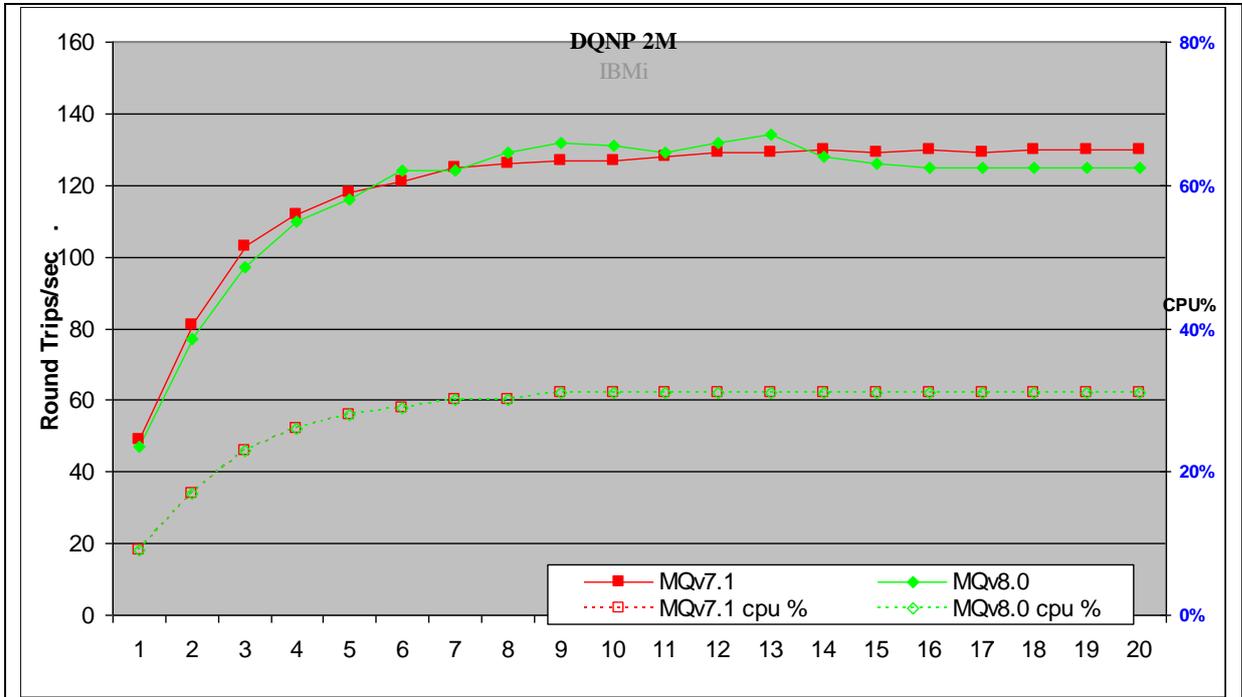


Figure 33 – 2MB non-persistent messages, distributed queuing

Figure 33 and Table 28 shows that the peak throughput of non-persistent messages is similar when comparing V8.0 to V7.1. At lower number of applications V8.0 is better than V7.1

Test Name: DQNP 2M	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	14	130	0.13	31%
MQv8.0	13	134	0.15	31%

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

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

### 3.3.3.2 Persistent Messages

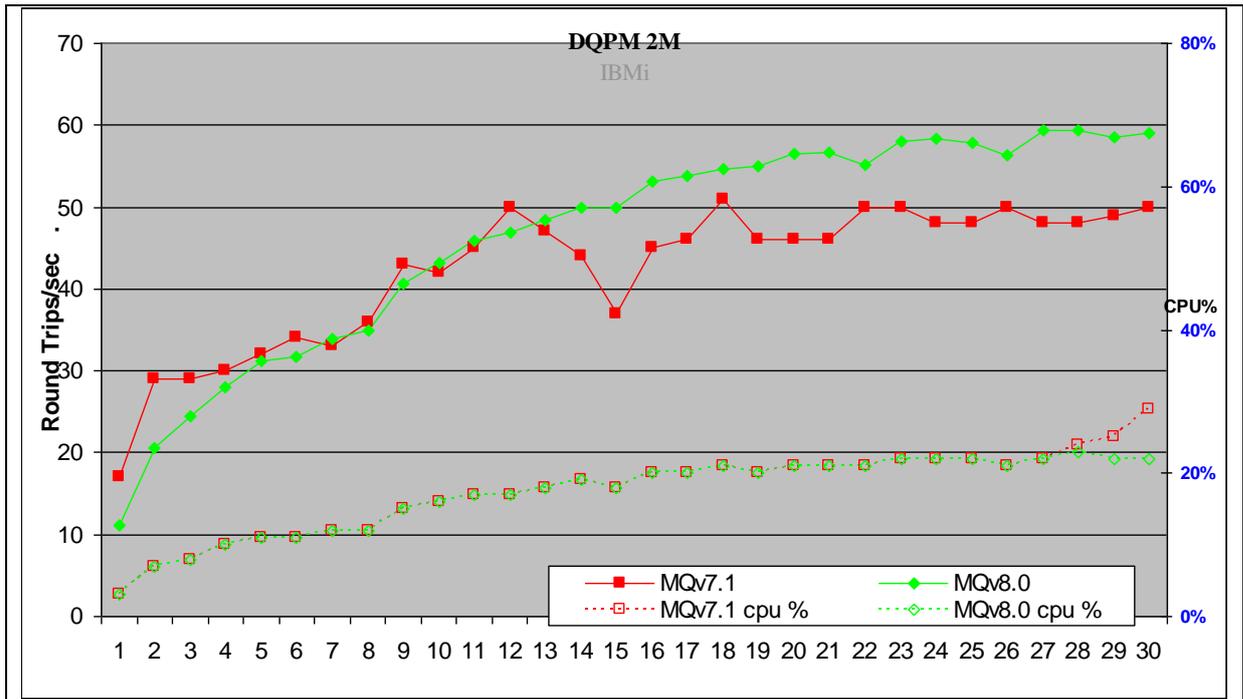


Figure 34 - 2MB persistent messages, distributed queuing

Figure 34 and Table 29 shows that the peak throughput of persistent messages has improved by 15% when comparing version 7.1 to 7.0.

Test Name: DQPM 2M	Apps	Round Trips/Sec	Response time (s)	CPU
MQv7.1	18	51	0.68	21%
MQv8.0	28	59	0.55	23%

Table 29 – 2MB persistent messages, distributed queuing

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

## 4 Application Bindings

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

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

### 4.1 Local Queue Manager

Figure 35 and Figure 36 show the non-persistent and persistent message throughput achieved using an increasing number of driving applications in the local queue manager scenario.

#### 4.1.1 Non-persistent Messages

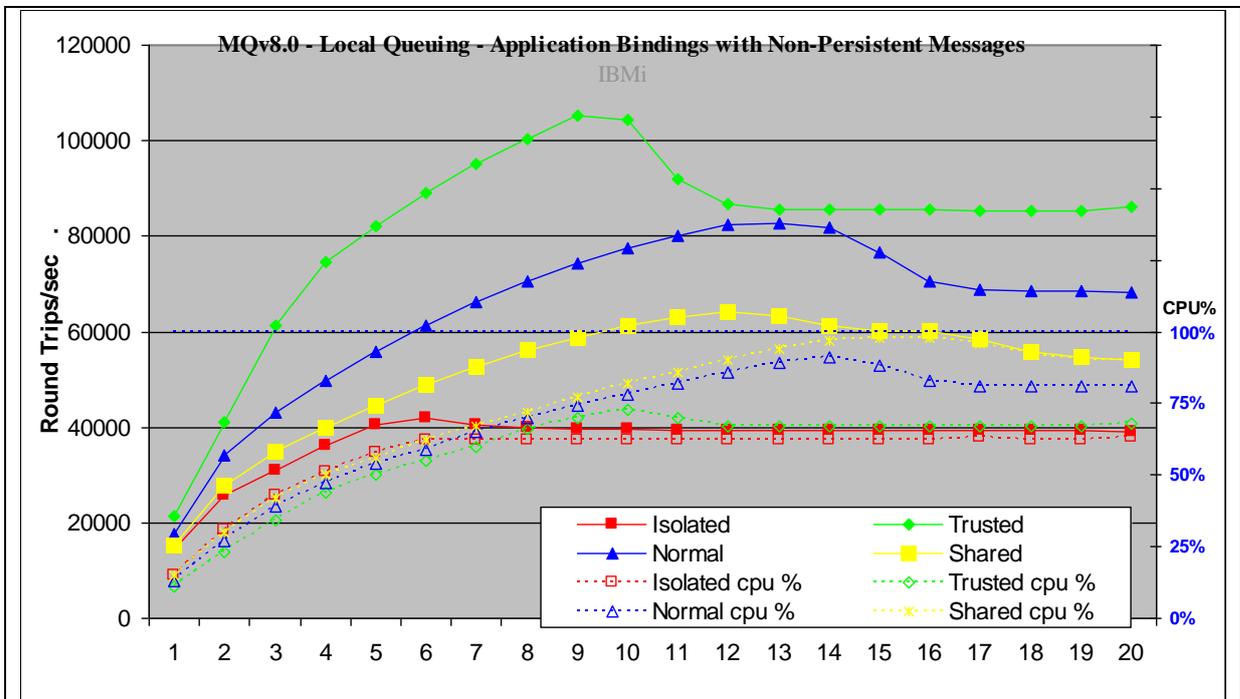


Figure 35 – Application binding, non-persistent messages, local queue manager

Figure 35 and Table 30 shows that the throughput of non-persistent messages when comparing Normal, Isolated, Trusted and Shared bindings.

Local Queuing Non-Persistent	Apps	Round Trips/Sec	Response time (s)	CPU
Isolated	6	41988	0.00016	62%
Trusted	9	105218	0.00009	70%
Normal	13	82843	0.00017	89%
Shared	12	64137	0.0002	90%

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

### 4.1.2 Persistent Messages

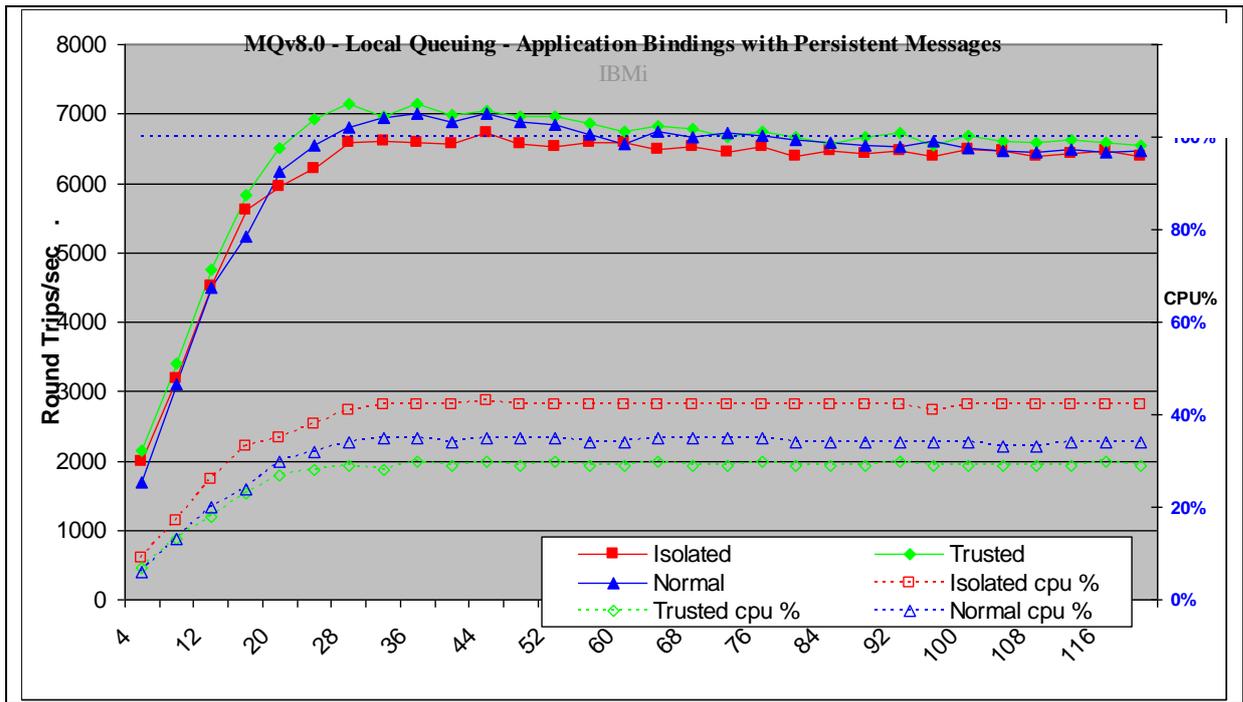


Figure 36 – Application binding, persistent messages, local queue manage

Figure 36 and Table 31 shows the throughput of persistent messages when comparing Normal, Isolated and Trusted bindings

Local Queuing Persistent	Apps	Round Trips/Sec	Response time (s)	CPU
Isolated	44	6730	0.0086	43%
Trusted	36	7145	0.0064	30%
Normal	44	7011	0.0074	35%

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

## 4.2 Client Channels

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

### 4.2.1 Non-persistent Messages

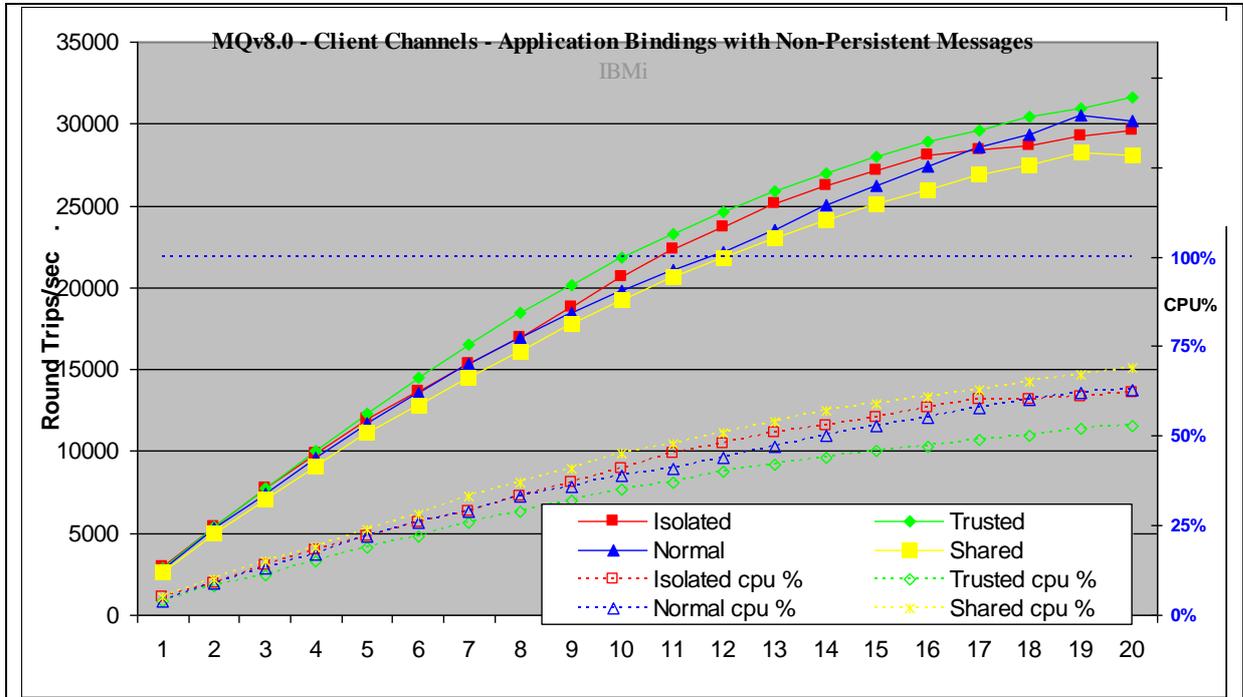


Figure 37 – Application binding, non-persistent messages, client channels

Figure 37 and Table 32 shows that the peak throughput of non-persistent messages when comparing Normal, Isolated and Trusted bindings.

Client Channels Non-Persistent	Apps	Round Trips/Sec	Response time (s)	CPU
Isolated	20	29567	0.0008	62%
Trusted	20	31665	0.00077	53%
Normal	20	31191	0.00075	63%
Shared	20	29118	0.00082	69%

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

### 4.2.2 Persistent Messages

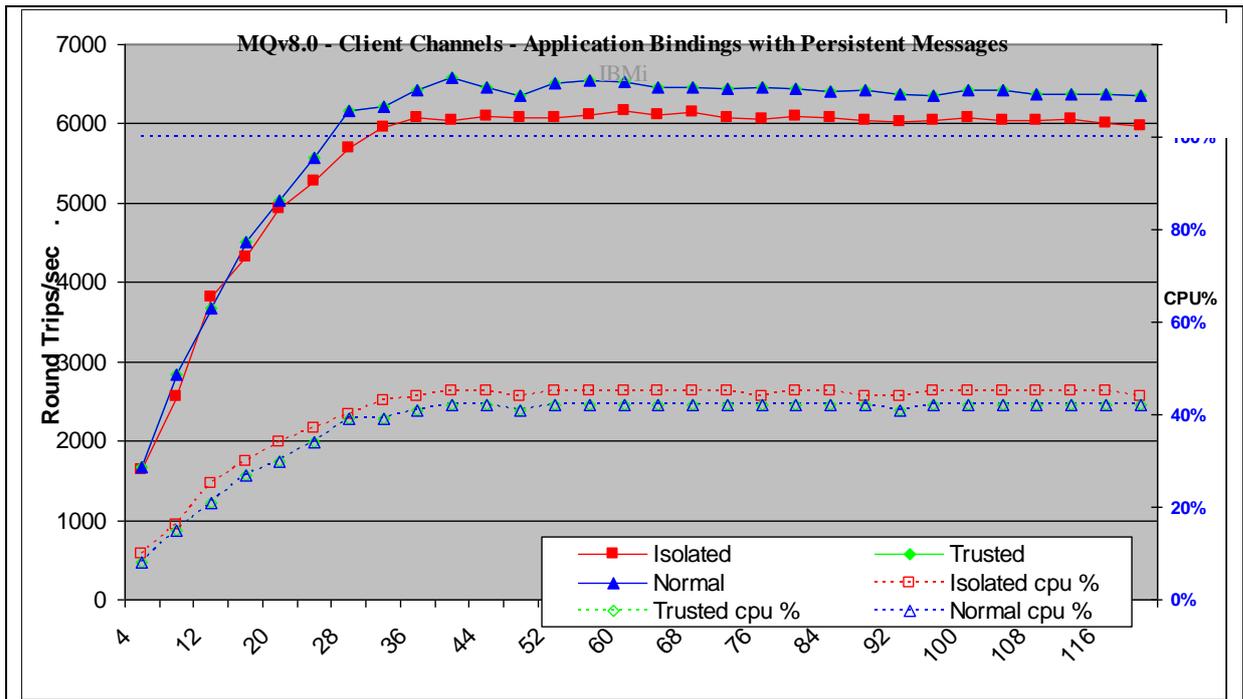


Figure 38 – Application binding, persistent messages, client channels

Figure 38 and Table 33 shows the peak throughput of non-persistent messages when comparing Isolated and Trusted bindings.

Client Channels Persistent	Apps	Round Trips/Sec	Response time (s)	CPU
Isolated	60	6160	0.011	45%
Trusted	40	6591	0.007	42%
Normal	40	6591	0.007	42%

Table 33 – Application binding, persistent messages, client channels.

### 4.3 Distributed Queuing

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

#### 4.3.1 Non-persistent Messages

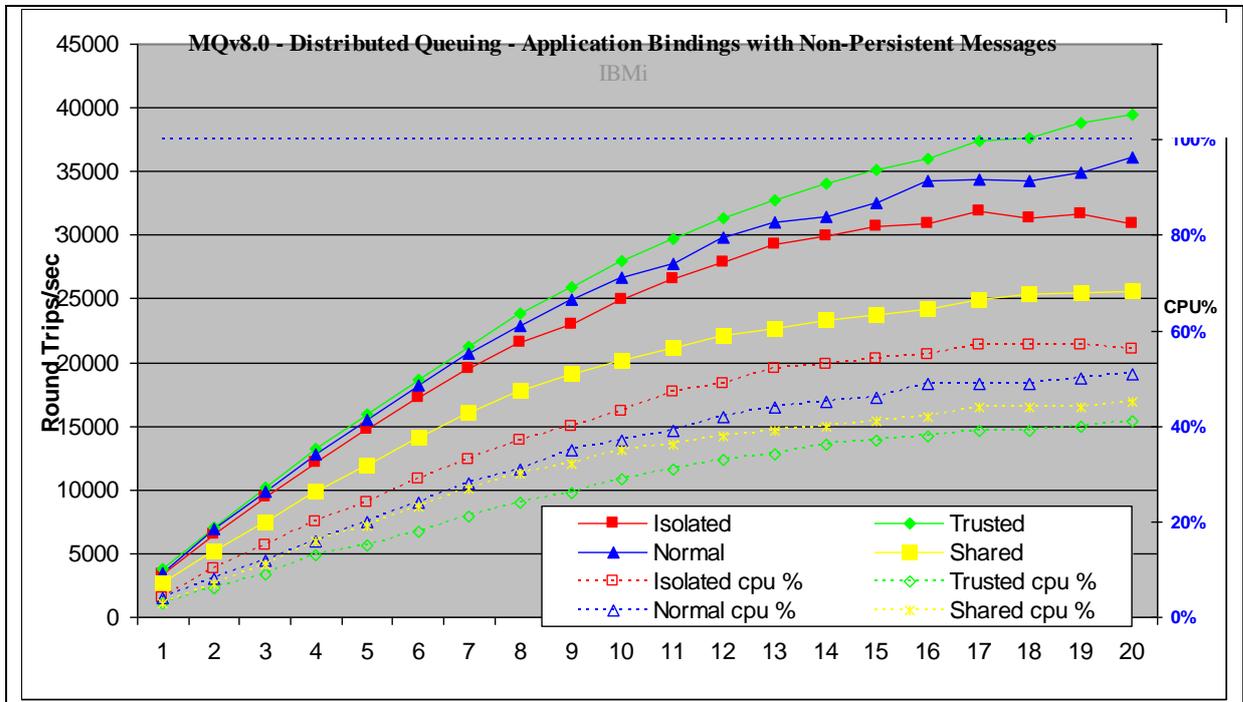


Figure 39 – Application binding, non-persistent messages, distributed queuing

Figure 39 and Table 34 show that the peak throughput of non-persistent messages when comparing Normal, Isolated and Trusted bindings.

Distributed Queuing Non-Persistent	Apps	Round Trips/Sec	Response time (s)	CPU
Isolated	17	31832	0.0006	57%
Trusted	20	39456	0.0006	41%
Normal	20	36090	0.00062	51%
Shared	20	25573	0.001	45%

Table 34 – Application binding, non-persistent messages, distributed queuing

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

### 4.3.2 Persistent Messages

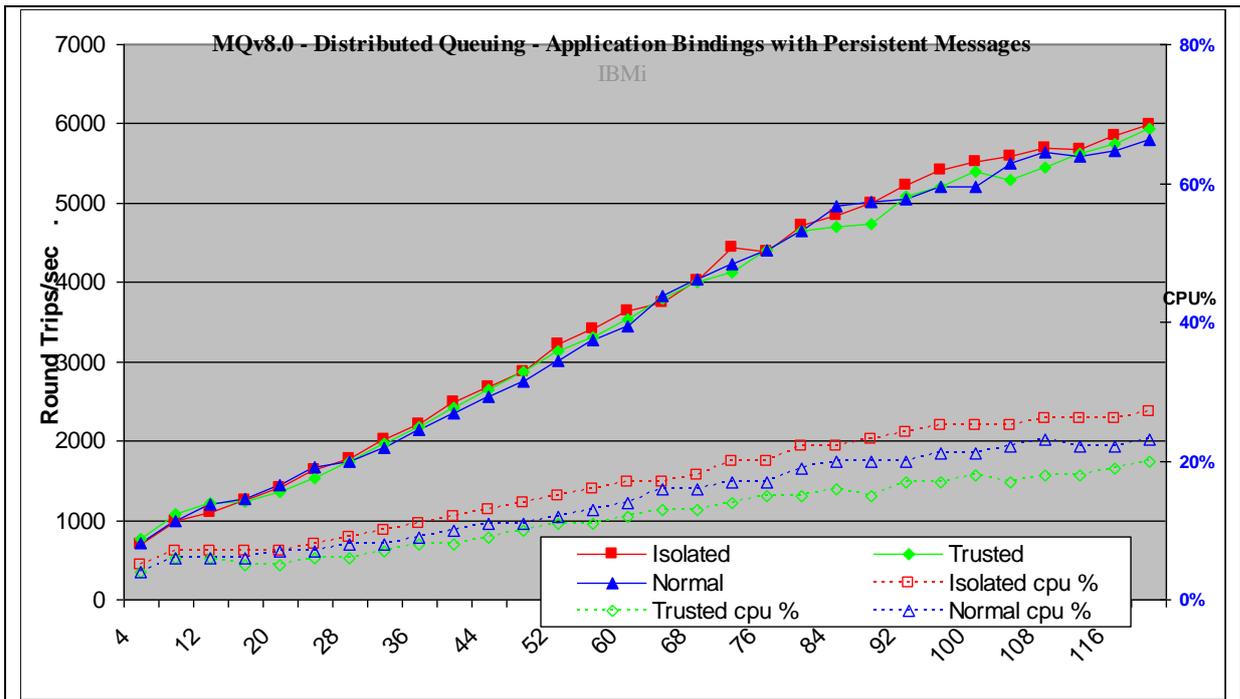


Figure 40 – Application binding, persistent messages, distributed queuing

Figure 40 and Table 35 show that the peak throughput of non-persistent messages when comparing Isolated and Trusted bindings.

Distributed Queuing Persistent	Apps	Round Trips/Sec	Response time (s)	CPU
Isolated	120	5988	0.023	27%
Trusted	120	5939	0.024	20%
Normal	120	5795	0.021	23%

Table 35 – Application binding, persistent messages, distributed queuing

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

## 5 Performance and Capacity Limits

### 5.1 Client channels – capacity measurements

The measurements in this section are intended to test the maximum number of client channels into a server queue managers with a messaging rate of 1 round trip per client channel per *minute* while additional connections are made. The maximum number of connected applications is likely to be determined by other criteria such as recovery time or manageability. Measurements are also made with smaller number of Client channels where the message insertion rate is increased until the system gets congested. This information is intended to be useful to the reader sizing a system with similar scenarios. These client measurements of V8.0 allocate a separate socket for each client (sharecnv=1 on svrcon channel).

Queue manager configuration for client channels capacity tests:

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

The tests run are as follows:

- **clnp**  
For test description see [2.2 Client Channel Test Description](#)
- **clnp\_r3600**  
For test description see [2.2.4 Client Channels](#).
- **clnp\_c6000**  
6000 remote clients connect, each delivering messages at a fixed rate of 1msg/sec and then the rate is increased until a constraint is hit. Usually response time > 1 sec. In the table below Apps \* Rate/app/hr = Round Trips/sec
- **clnp max**  
30000 remote clients connect, each delivering messages at a fixed rate of 1msg/sec and then the rate is increased until a constraint is hit. Usually response time > 1 sec. In the table below Apps \* Rate/app/hr = Round Trips/sec

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
clnp	20	n/a (unrated)	26140	0.00089	61%
clnp_r3600	2300	3600	22623	0.069	68%
clnp_c6000	6000	10000	16666	0.0007	41%
clnp max	30000	90	750	0.00204	3%

**Table 36 – Capacity measurements, client channels**

\* *There was no delay between the response to the previous message and the insertion of the next message with 20 clients.*

The maximum message throughput is achieved when there are a small number of requester applications. The clnp\_3600 measurement peaks when the queue of input messages waiting to be processed by the Server application builds up because the server application threads can no longer keep up with the demand. Although this ensures the server threads are always busy, the messages are being spilt from the Queue buffer to the file system and possibly to the disk. Each client uses a thread in the AMQRRMPA processes and the management of lots of threads and lots memory objects results in a larger CPU cost to handle each message.

Measurements normally use a Get by Correlation\_Id from a common reply queue for all clients whereas the tests labelled ‘no\_correlid’ have a separate reply queue per client. Each additional Client needs a thread in the AMQRMPPA process. Using a separate queue per client needs additional shared memory per client.

## 5.2 Distributed queuing – capacity measurements

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

Queue manager and log configuration for distributed queuing capacity tests:

MaxChannels=30000, LogPrimaryFiles=16, 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).*

- **dqnp**  
For test description see [2.3 Distributed Queuing Test Description](#)
- **dqnp\_r3600**  
For test description see [2.3.4 Server Channels](#).
- **dqnp\_qmax**  
17500 remote clients connect, each delivering messages at a fixed rate of 1msg/sec and then the rate is increased until a constraint is hit. Usually response time > 1 sec. In the table below Apps \* Rate/app/hr = Round Trips/sec.
- **dqnp\_q1000**  
This test shows the throughput experienced when 1000 queues are connected into a central hub for non persistent messages. In the table below Apps \* Rate/app/hr = Round Trips/sec.
- **dq-persist\_q1000**  
This test shows the throughput experienced when 1000 queues are connected into a central hub for persistent messages. In the table below Apps \* Rate/app/hr = Round Trips/sec.

Test name:	Apps	Rate/app/hr	Round Trips/sec	Response time (s)	CPU
dqnp	20	n/a (unrated)	36311	0.0006	52%
dqnp_r3600	3500	3600	34884	0.03	50%
dqnp_q1000	1000	25000	5908	0.02139	12%
dqnp_qmax	17500	60	291	0.00178	2%
dq-persist_q1000	1000	1500	416	0.75	5%

**Table 37 – Capacity measurements, server channels**

\* *There was no delay between the response to the previous message and the insertion of the next message with 20 driving applications..*

The dqnp and dqnp\_r3600 both used a total of 4 pairs of Sender/Receiver pairs of channels between queue managers while the dqnp\_qmax used a pair of channels per application.

## 6 Tuning Recommendations

### 6.1 Tuning the Queue Manager

This section highlights the tuning activities that are known to give performance benefits for IBM MQ V8.0; 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.

Generally only one parameter should be changed and tested at a time before considering further changes.

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

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

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

Non-persistent messages are held in main memory, spilt to the file system as the queues become deep and lazily written to the Queue file (i.e. operating system I/O buffers are used). Persistent messages are synchronously written to the log by an MQCmit (writing tho I/O the operating system buffers to maintain integrity) and are also periodically flushed (lazily) 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 significantly longer than for a non-persistent message. This overhead does not include the additional time for the message to be written to the log, although this can be minimised by using cached disks or SAN.

##### 6.1.1.1 Non-persistent and Persistent Queue Buffer

The default non-persistent queue buffer size is 64KB per queue and the default persistent is 128KB per queue for 32 bit Queue Managers and 128KB /256KB for 64 bit Queue Managers.. They can all be increased to 100MB using the `qm.ini` TuningParameters stanza and the `DefaultQBufferSize` and `DefaultPQBufferSize` parameters. (For more details see SupportPac MP01: MQSeries – Tuning Queue Limits). Increasing the queue buffer provides the capability to absorb peaks in message throughput at the expense of real storage. Once these queue buffers are full, the additional message data is given to the file system that will eventually find its way to the disk. Defining queues using large non-persistent or persistent queue buffers can degrade performance if the system is short of real memory either because a large number of queues have already been defined with large buffers, or for other reasons (e.g. a large number of channels are defined).

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.

#### 6.1.2 Log Buffer Size, Log File Size, and Number of Log Extents

The Log component is often the bottleneck when processing persistent messages. Sufficient information is stored on the log to restart the queue manager after failure. Circular logging is sufficient to recover from application, software, or power failure while linear logging will also recover from media (or disk) failure. Log records are written at each MQPut, MQGet, and MQCmit into the log buffer. This information is moved onto

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

The log buffer is a circular piece of main memory where the log records are concatenated so that multiple log records can be written to the log file in a single I/O operation. The default values used for `LogBufferPages` and `LogFilePages` have been increased in V7 and are probably suitable for most installations. The default size of the log buffer is 512 pages with a maximum size of 4096 pages. To improve persistent message throughput of large messages (messages size > 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). The larger the disk extent, the longer the elapsed times between changing disk extents. It is better to have a smaller number of large extents but long running UOW can prevent Checkpointing efficiently freeing the disk extent for reuse. The largest size (maximum 65536 pages) will reduce the frequency of switching extents. The number of `LogPrimaryFiles` (i.e. `crtmqm -lp <LogPrimaryFiles>`) can be configured to a large number and the maximum number of Primary plus Secondary extents is 255(Windows) and 511(UNIX) but it is for functional reasons rather than performance that need more than 20 primary extents for Circular logging. Circular logging should be satisfied by Primary logs because Secondary logs are formatted each time they are reused. The Active log set is the number of extents that are identified by the Checkpoint process as being necessary to be kept online. As additional messages are processed, more space is taken by the active log. As UOWs complete, they enable the next Checkpoint process to free up extents that now become available for archiving with Linear logging. Some installation will use Linear logging and not archive the redundant logs because archiving impacts the run time performance of logging. They will periodically (daily or twice daily) use ‘`rcdmqmg`’ on the main queues thus moving the ‘point of recovery’ forward, compacting the queues, and freeing up log disk extents. The cumulative effect of this tuning will:

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

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

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

### 6.1.2.1 LogWriteIntegrity: SingleWrite or TripleWrite

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

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

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

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

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

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

### 6.1.3 Channels: Process or Thread, Standard or Fastpath?

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

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

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

## 6.2 Applications: Design and Configuration

### 6.2.1 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 MQ V8 knowledge centre..

### 6.2.2 Parallelism, Batching, and Triggering

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

Queue avoidance is a feature of the queue manager that allows messages to be passed directly from an MQ 'Putter' to an MQ 'Getter' 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 MQ 'Getter' ), 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 MQ 'Getters' need to retrieve messages from the buffer rather than being received directly from an MQ 'Putter'. A secondary effect is that as messages are spilled from the buffer to the queue disk, the MQ 'Getters' 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 MQ 'Getters' (i.e processing threads in the server application), or using a larger queue buffer, it may not be possible to avoid a performance degradation.

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

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

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

## 6.3 Tuning the Operating System (IBM i)

Please refer to IBM i specific tuning literature for IBM i tuning techniques

### Journal Considerations on IBM i

Persistent messaging on IBM i uses native journaling support to ensure that messages are recoverable. To ensure that maximum rates of throughput are achieved when using persistent messages, you should consider the following:

When the queue manager is created on IBM i, a journal receiver is automatically created and located on the system ASP disk arms. To avoid contention with other IO on these arms, it is recommended that the user manually create and attach a new journal receiver, ensuring that it is located on a user ASP with dedicated disk arms. This will help improve response times for the synchronous disk writes to the journal that are needed for each persistent message.

It will also be helpful to ensure that the disk arms and IOPs used in the user ASP have good overall performance characteristics for write activity, including good write cache performance. Slower disk arms and IOPs will result in less favourable response times and less overall capacity in terms of message throughput.

## 6.4 TCP Buffer size changes for V8.0

In MQ V8.0, a newly created queue manager will set TCP buffers in the qm.ini to a value of 0. This indicates that the operating system will manage the buffer sizes, as opposed to the buffer sizes being fixed by MQ. A migrated queue manager from V7.1 will, by default, not have these buffers set to a value in the qm.ini, unless they had been previously modified.

Therefore, there can be differences in performance throughput for a migrated queue manager vs a new queue manager in V8.0, depending on the system buffer settings. The system TCP buffers are changed using the CHGTCPA command. Great care should be taken when changing either the qm.ini TCP stanza or the system settings. For further information about how this affects performance, refer to the TCP/IP documentation for your environment. For more information on qm.ini TCP properties, refer to the MQ V8.0 Knowledge Center.

## 6.5 Virtual Memory, Real Memory, & Paging

### 6.5.1 BufferLength

The AMQRMPPA process contains a thread per connected client. The BufferLength parameter of the MQGet is also used to allocate a long term piece of storage of this size in which the message is held before being retrieved by the client. If the size of the arriving messages cannot be predicted then the application should provide a buffer that can deal with 90% of the messages and redrive the MQGet after return code **2080 (X'0820') MQRC\_TRUNCATED\_MSG\_FAILED** by providing a

larger BUFFER for retrieving this particular message. There is a mechanism to gradually reduce the size of the storage in AMQRMPPA if the recent BufferLength size is significantly smaller than previous BufferLength.

### 6.5.2 MQIBINDTYPE

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

## 7 Measurement Environment

### 7.1 Workload description

#### 7.1.1 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.*

#### 7.1.2 Test scenario workload

The MQI applications use 64 bit libraries for MQ

##### 7.1.2.1 The driving application programs

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

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

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

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

##### 7.1.2.2 The server application program

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

Non-persistent messaging 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.

##### 7.1.2.3 Analysis techniques

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

## 7.2 Hardware

### 7.3 Device under test (Server)

Server system:	IBM POWER7 with IBM i
Model:	Power 780
Processor:	3.86GHz POWER7
Architecture:	7-core partition
Memory (RAM):	25.25GB
Disk:	20x 70GB disk, Nearline (NL) SAS disk, direct attach from DS887
Network:	10GBit Ethernet Adapter

#### Driver system used for client scenarios

Driver system:	IBM POWER6 with IBM p
Model:	Power 570
Processor:	4.2GHz
Architecture:	8-way CPU
Memory (RAM):	16GB
Disk:	1x locally attached 146GB SAS disk plus two SAN disks, 20GB and 5GB, from DS8870 via SVC
Network:	10GBit Ethernet Adapter

#### Driver system used for distributed scenarios

Driver system:	IBM POWER7 with IBM i
Model:	Power 780
Processor:	3.86GHz POWER7
Architecture:	8-core partition
Memory(RAM):	32GB
Disk:	40x 35GB disk, from DS8870 through VIOS and the IBM System Storage SAN Volume Controller (SVC)
Network:	10GBit Ethernet Adapter

#### Driver system used in capacity and short sessions

IBM x3850:	Driver system
Model:	x3850 M2 8864 4RG
Processor:	2.93GHz Intel Xeon x7350
Architecture:	2 x quad core CPU
Memory (RAM):	32GB
Disk:	2 SAN disks on DS8700 (5GB each, 1 queue, 1 log)
Network:	10Gbit Ethernet Adapter

## 7.4 Software

Operating system	: IBM i V7R1M0
MQ version	: Version 7.1, Version 8.0
Compiler	: IBM Rational Development Studio for i V7R1M0

## 8 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><i>local</i> =&gt; local queue manager test scenario</p> <p><i>cl</i> =&gt; client channel test scenario</p> <p><i>dq</i> =&gt; distributed queuing test scenario</p> <p><i>np</i> =&gt; non-persistent messages</p> <p><i>pm</i> =&gt; persistent messages</p> <p><i>r3600</i> =&gt; 1 round trip per driving application per second</p> <p><i>runmqslr</i> =&gt; channels using the 'runmqslr' listener (client channel test scenario, in addition to 'runmqchi' for distributed queuing test scenarios)</p> <p><i>c6000</i> =&gt; 6,000 client driving applications (i.e. 6,000 MQI-client connections)</p> <p><i>q1000</i> =&gt; 1,000 server channel pairs</p> <p><i>max</i> =&gt; maximum number of channels (or channel pairs)</p> <p><i>no_correl_id</i> =&gt; correlation identifier not used in the response messages (as each response is placed on a unique reply-to queue per driving application)</p> <p><i>Messages /sec</i> =&gt; Round Trips/sec</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 increase in the total message throughput rate.</p> <p><i>Note: The nature of the comparison is noted under each table where percentage improvements have been given.</i></p>
<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>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>FREE</b>	Free memory as reported by IOSTAT