



IBM MQ V8.0 for Linux on x86_64

Performance Report

Version 1.1 – December 2015

Paul Harris, Sam Massey, Rowan Lonsdale, Arundeeep Bahra

IBM MQ Performance
IBM UK Laboratories
Hursley Park
Winchester
Hampshire



**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 2015.

This edition applies to *IBM MQ V8.0* (and to all subsequent releases and modifications until otherwise indicated in new editions).

© Copyright International Business Machines Corporation 2014,2015. 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 was 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.

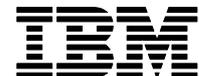
In Germany and Austria, notwithstanding the above exclusions, IBM's warranty and liability are governed only by the respective terms applicable for Germany and Austria in the corresponding IBM program license agreement(s).

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 V8.0*. The information is not intended as the specification of any programming interface that is provided by IBM MQ. 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 their respective companies in the United States, other countries or both:

- **IBM Corporation:** IBM
- **Intel Corporation:** Intel, Xeon
- **Red Hat:** Red Hat, Red Hat Enterprise Linux

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

EXPORT REGULATIONS

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

Preface

Target audience

The report is designed for people who:

- Will be designing and implementing solutions using IBM MQ v8.0 for Linux on x86_64.
- Want to understand the performance limits of IBM MQ v8.0 for Linux on x86_64.
- Want to understand what actions may be taken to tune IBM MQ v8.0 for Linux on x86_64.

The reader should have a general awareness of the Linux 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 with performance charts.
- Performance measurements with figures and tables to present the performance capabilities of IBM 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?

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

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

Additions and changes to second edition (V1.1)

- SSL performance data added. (see section 3.7)
- Improved client connection time in V8.0.0.3 (see section 1.4)



Introduction

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

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

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

CONTENTS

1	Overview	10
1.1	Locking Optimisations	11
1.1.1	Object catalog lock scope reduction	11
1.1.2	Logger optimisations	12
1.2	Improved TCP Buffer Defaults	13
1.3	Improved SHARECNV(1) Client Performance	14
1.4	Reduced Client Connection Time	15
2	Workloads	16
2.1	Workload Topologies	16
2.1.1	Requester-responder with local queue manager	17
2.1.2	Requester-responder with remote queue manager (local responders)	18
2.1.3	Requester-responder with remote queue manager (remote responders)	19
2.1.4	Requester-responder with distributed queuing	20
2.1.5	Requester-responder with distributed queuing across n queues with multiple channels	20
2.1.6	One-way messaging with local queue manager	22
2.1.7	One-way messaging with remote queue manager	23
3	Performance Test Results	25
3.1	MQI Local Bindings (Non-Persistent)	25
3.1.1	Test setup	25
3.2	MQI Local Bindings (Persistent)	26
3.2.1	Test setup	26
3.3	MQI Client Bindings (Non-Persistent)	26
3.3.1	Test setup	27
3.4	MQI Distributed Queuing (Non-Persistent)	27
3.4.1	Test setup	28
3.5	JMS (Non-Persistent)	28
3.5.1	Test setup	29
3.6	Memory Usage	29
3.7	MQI Client Bindings (Non-Persistent) - SSL	31
3.7.1	Test setup	32
	Scaling Measurements	33
3.8	MQI Rated Scaling	33
3.8.1	MQI rated scaling test setup	33
3.8.2	MQI client bindings (non-persistent)	33
3.8.3	MQI client bindings (non-persistent) with 24 request/reply queues	34
3.8.4	MQI client bindings (persistent)	35
3.8.5	MQI client bindings (persistent) with 24 request/reply queues	36
3.9	JMS Rated Scaling with Dynamic Reply Queues	36
3.9.1	JMS rated scaling test setup	36
3.9.2	JMS client bindings (Non-persistent) with 24 request queues	36
3.9.3	JMS client bindings (persistent) with 24 request queues	37
3.9.4	JMS distributed queuing (non-persistent) with 24 request queues	38
3.9.5	JMS distributed queuing (persistent) with 24 request queues	39
4	How Am I Constrained?	40
4.1	Disk Contention	40
4.2	CPU Saturation	41
4.3	Queue Lock Contention	41
4.4	Network Saturation	42
4.5	Channel Saturation	43
5	Planning for Performance	45
5.1	Persistence, High Availability, & Disaster Recovery	45
5.1.1	Latency	45
5.1.2	Bandwidth	46
5.2	Slow Networks	46
5.2.1	Many disparate clients communicating infrequently	47
5.2.2	High message rate to/from a client machine	47
5.2.3	Communicating between regions	47
5.2.4	Optimising QM to QM channel communication	47
5.3	Ultra-High Message Rates	48
5.3.1	Maximising the throughput of a single queue	48



5.3.2	Adding queues to increase throughput	49
5.4	Transparent Scalability	49
5.4.1	Scaling across machines & queue managers	49
5.4.2	Scaling across queues	50
5.5	Testing for Performance	51
5.5.1	It's a dress rehearsal, not an audition	51
5.5.2	Know your tools.....	51
5.5.3	When it doesn't work, simplify	52
5.5.4	Test your environment	52
5.6	Memory Requirements	52
6	Tuning Recommendations	53
6.1	Tuning The Queue Manager	53
6.1.1	Queue disk, log disk, and message persistence.....	53
6.1.2	Channels: process or thread, standard or fastpath?.....	56
6.2	Applications: Design and Configuration	56
6.2.1	Standard (shared or isolated) or fastpath?.....	56
6.2.2	Parallelism, batching, and triggering.....	56
6.3	Virtual Memory, Real Memory, & Paging	57
6.3.1	BufferLength	57
6.3.2	MQIBindType	57
6.3.3	Paging.....	58
Appendix A:	Workloads	59
A.1	JMSPerfHarness	59
A.2	C Performance Harness (CPH).....	59
Appendix B:	Further Resources.....	60
Appendix C:	Test Setup.....	61
C.1	Hardware/Software	61
C.1.1	Hardware.....	61
C.1.2	Software	61
C.2	Tuning Parameters Set for Measurements in This Report	61
C.2.1	Operating System	62
C.2.2	IBM MQ	62



TABLES

Table 1 - Workload types	16
Table 2 - Peak rates for MQI local bindings (non-persistent)	25
Table 3 - Peak rates for MQI local bindings (persistent).....	26
Table 4 - Peak rates for MQI client bindings (non-persistent).....	27
Table 5 - Peak rates for distributed queuing (non-persistent)	28
Table 6 - Peak rates for JMS (non-persistent).....	29
Table 7 - Memory usage	30
Table 8 - Peak rates for MQI client bindings (2KB non-persistent) - SSL.....	31



FIGURES

Figure 1 - Locking optimisations	11
Figure 2 – Performance of MQI Local Bindings, One-way (2KB Persistent)	12
Figure 3 - Large message distributed queuing improvements.....	13
Figure 4 – SHARECNV(1) performance.	14
Figure 5 – MQI Client connect times	15
Figure 6 - Requester-responder with local queue manager.....	17
Figure 7 - Requester-responder with remote queue manager (local responders).....	18
Figure 8 - Requester-responder with remote queue manager (remote responders).....	19
Figure 9 - Requester-responder with distributed queuing	20
Figure 10 - Requester-responder with distributed queuing across multiple channels	21
Figure 11 - Multi-channel distributed queuing, sample deployment.	22
Figure 12 – One-way messaging with local queue manager.....	23
Figure 13 - One-way messaging with remote queue manager	24
Figure 14 – Performance results for MQI Local Bindings (2KB Non-persistent).....	25
Figure 15 - Performance results for MQI local bindings (2KB persistent).....	26
Figure 16 - Performance results for MQI client bindings (2KB non-persistent)	27
Figure 17 - Performance results for distributed queuing (2KB non-persistent)	28
Figure 18 - Performance results for JMS local bindings (2KB non-persistent)	29
Figure 19 – Performance results for MQI client bindings (2KB non-persistent) - SSL	31
Figure 20 - Performance results for MQI client binding, rated scaling (2KB non-persistent)	33
Figure 21 - Performance results for MQI client binding, rated scaling across 24 queues (2KB non-persistent)	34
Figure 22 - Performance results for MQI client binding, rated scaling (2KB persistent)	35
Figure 23 - Performance results for MQI client binding, rated scaling across 24 queues (2KB persistent)	36
Figure 24 - Performance results for JMS client bindings, rated scaling across 24 queues with dynamic reply queues (2KB non-persistent)	37
Figure 25 - Performance results for JMS client bindings, rated scaling across 24 queues with dynamic reply queues (2KB persistent)	37
Figure 26 - Performance results for JMS local binding, rated scaling across 24 remote queues with dynamic reply queues (2KB non-persistent).....	38
Figure 27 - Performance results for JMS local binding, rated scaling across 24 remote queues with dynamic reply queues (2KB persistent)	39
Figure 28 - Performance results for MQI local bindings (2KB persistent)	40
Figure 29 - Performance results for MQI local bindings (2KB non-persistent)	42
Figure 30 - Performance results for MQI client bindings (20KB non-persistent)	43
Figure 31 - Single channel vs multi channel distributed queuing.....	44



1 Overview

IBM MQ V8.0 for Linux on x86_64 has improved performance over V7.5 in almost every area.

- For 2KB non-persistent messages V8.0 is up to 77% better than V7.5. (see section 3.3)
- For 2KB persistent messages V8.0 is up to 10% better than V7.5. (see section 3.2)
- For 2MB non-persistent messages V8.0 is up to 3% better than V7.5. (see section 3.4)
- For 2MB persistent messages V8.0 is at parity with V7.5. (see section 3.2)

Specific areas of optimisation in V8 include

- Improved queue manager scaling through lock optimisations in logging and internal catalog code.
- Improved TCP buffer defaults.
- Improved client binding performance using SHARECNV(1)
- Reduced client connection time.



1.1 Locking Optimisations

1.1.1 Object catalog lock scope reduction

Every object known to the queue manager is registered with the object catalog. Every time an object is referenced by name the object catalog is queried, therefore most MQ operations result in a number of queries on the object catalog.

Prior to MQ V8 the object catalog locking was coarse, for example an *MQOPEN* could not complete while other catalog operations were being performed. Some of the operations performing catalog queries were also involved in forced I/O and so potentially common operations that should have been running at CPU speeds could end up running at I/O speeds.

In V8 we now have implemented fine grained locking of the object catalog, catalog access does need to wait for any forced I/O.

Figure 1 below demonstrates how a client/server application using *MQPUT1* to send the reply messages has improved through these changes (an *MQPUT1* effectively implies an *MQOPEN/MQPUT/MQCLOSE* sequence). The test has been designed to highlight the improvements in the object catalog locking.

Customers who use *MQPUT1* or similar patterns which result in intensive use of the object catalog, might see significant gains in their scenarios.

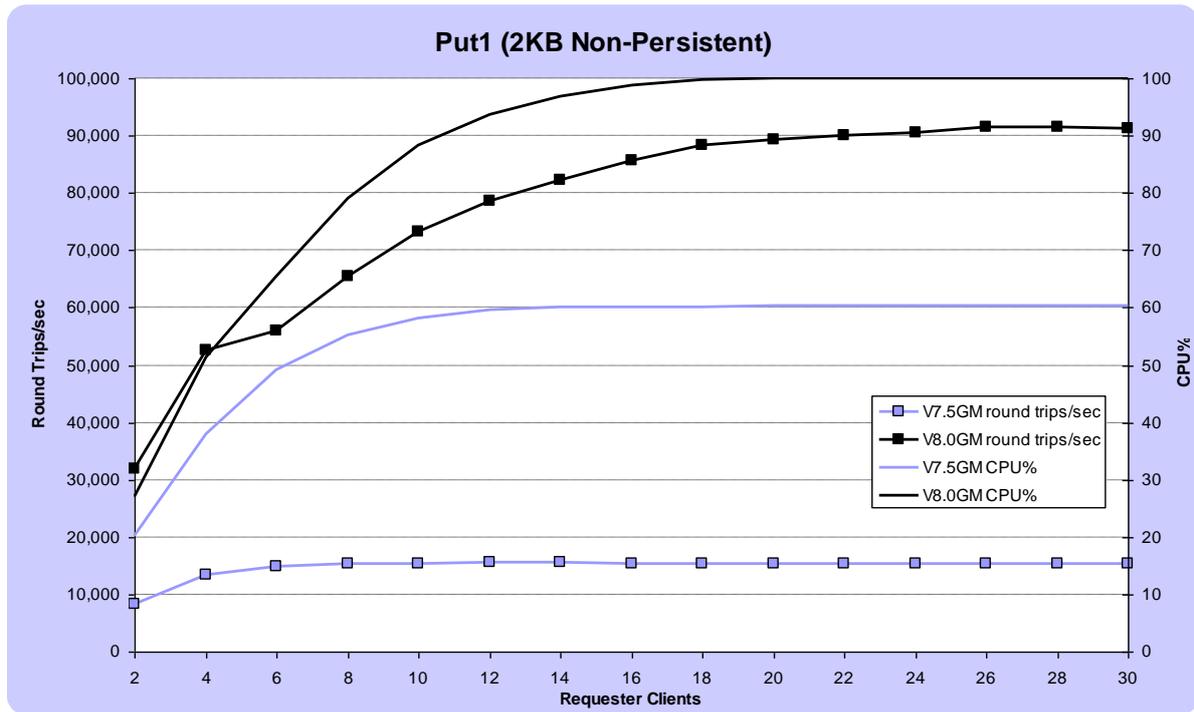


Figure 1 - Locking optimisations



1.1.2 Logger optimisations

This is an artificial environment designed to show that the logger is less of a constraint on throughput on a relatively wide SMP and uses a single request/reply application cloned across multiple queues.

The changes to the logger include:

- post offload – reduce cycles on the (single) log writing thread
- improved locking – eliminate lock contention by using lock free algorithms
- more efficient buffer usage – use writev() to eliminate extra write at buffer wrap around

In an environment with one (or another small number relative the number of cores) relatively heavily loaded queue(s), the throughput limitations caused by the coarse grained queue lock are likely to be more of a constraint than the logger, however an environment with a workload distributed across multiple queues should see significant potential benefits.

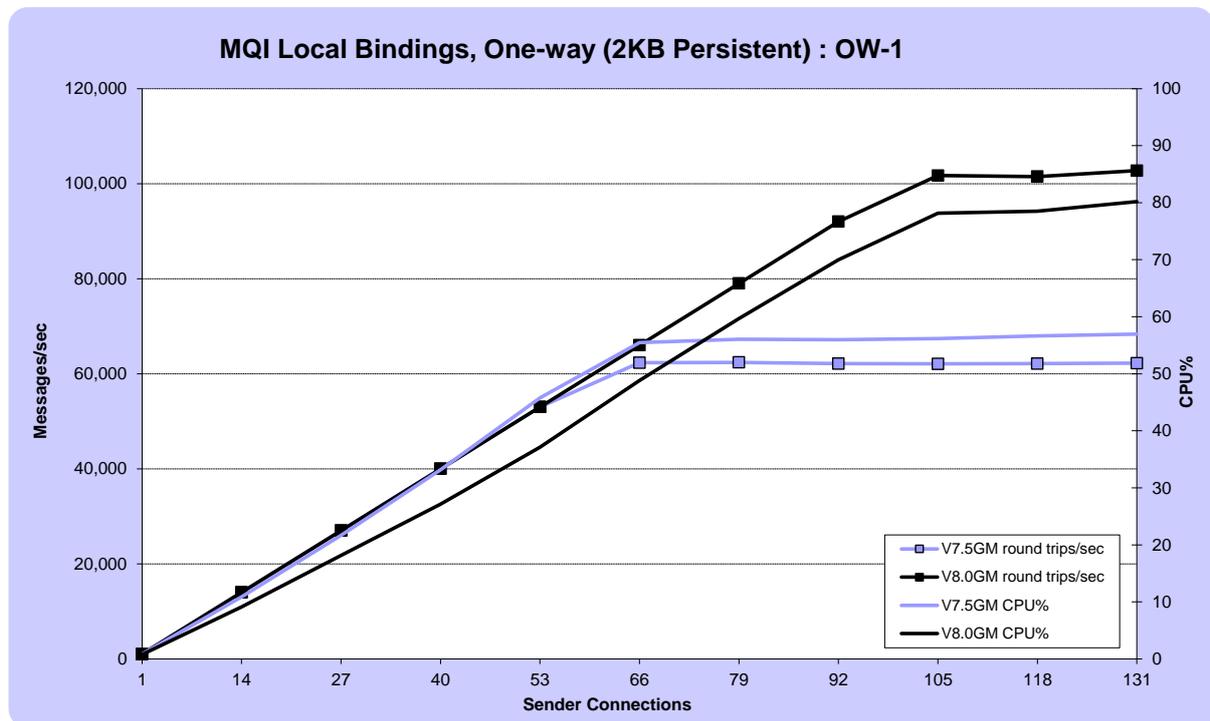


Figure 2 – Performance of MQI Local Bindings, One-way (2KB Persistent).



1.2 Improved TCP Buffer Defaults

In MQ V8.0, a newly created queue manager will set TCP socket buffers in the qm.ini file for a queue manager to 0. This indicates that the operating system will manage the buffer sizes, as opposed to the buffer sizes being fixed by the MQ configuration (the default size for TCP buffers was 32K prior to V8).

A migrated queue manager from previous versions of MQ will not have these buffers set to a value in the qm.ini by default unless you have previously modified them. This is to ensure that there is no unexpected change in performance when migrating to MQ V8 but does mean there can be differences in performance throughput for a migrated queue manager when compared with a new queue manager in V8.0, depending on the system buffer settings.

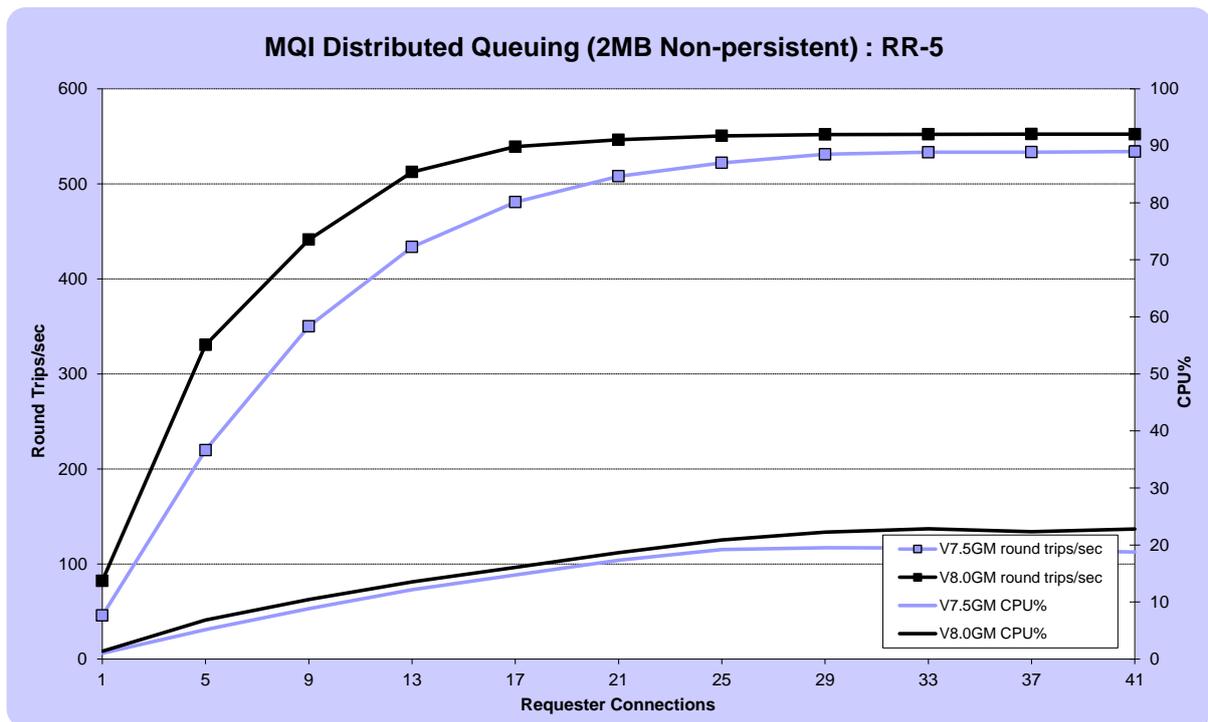


Figure 3 - Large message distributed queuing improvements.

In general, leaving the operating system to decide on the TCP socket buffer sizes is preferable. This option has been present in MQ from V7.5.0.3 onwards.

Figure 3 shows the relative performance of MQ V7.5 to MQ V8.0 when routing 2MB messages through a distributed queuing requester-responder scenario. The two queue managers utilise 10 server channels to send and receive the messages. For lower number of clients V8 performs much better than V7.5 out of the box (remember that from V7.5.0.3 the TCP settings can be configured the same way). Both test runs hit a network limit around 2,400 round trips/sec.

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.



1.3 Improved SHARECNV(1) Client Performance

MQ V7 introduced support for *SHARECNV*. Multiple client conversations (i.e. MQ connections within the same application process) can share the same TCP/IP socket (channel instance)

The extra capabilities reduced the top line performance even with a single conversation on the socket.

Prior to V8, *SHARECNV(0)* enabled the top performance but prevented many V7 features E.g. Heart beating and asynchronous consume.

In MQ V8 *SHARECNV(1)* has been optimized for more comparable performance to *SHARECNV(0)*

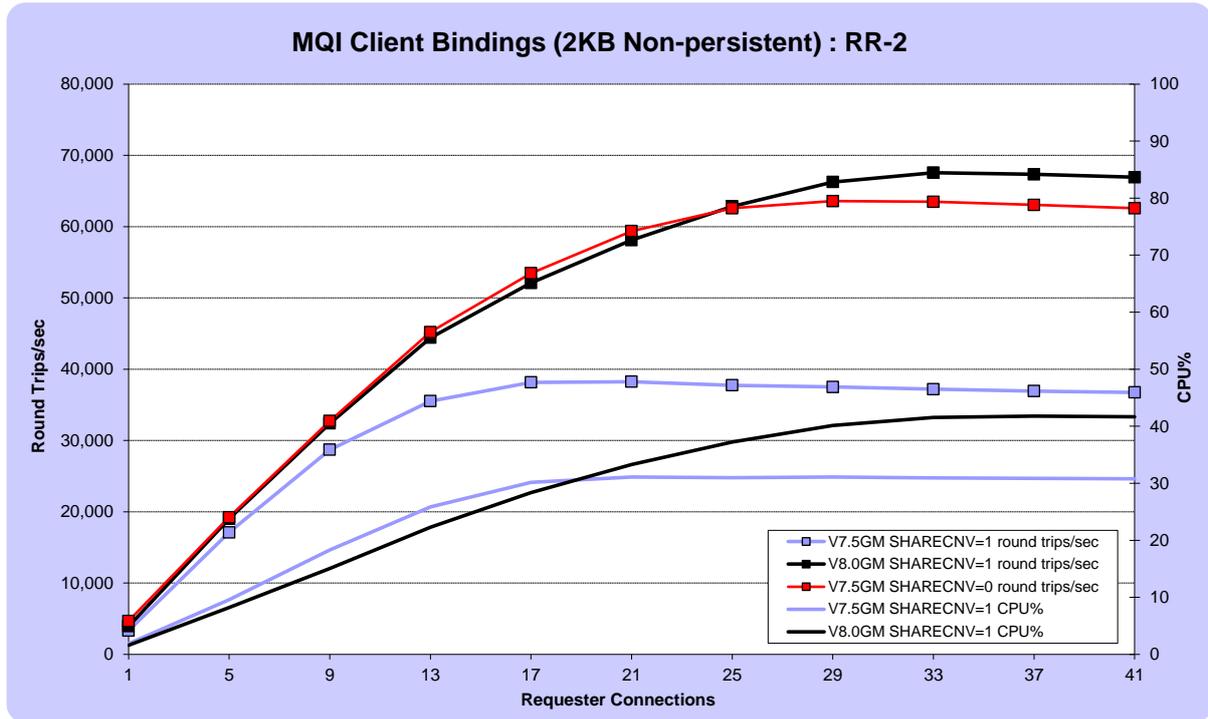


Figure 4 – SHARECNV(1) performance.

Figure 4 shows the relative performance of V7.5 and V8.0 for a MQI requester-responder scenario with client bound applications. In V7.5 *SHARECNV(1)* lagged behind *SHARECNV(0)* as the number of clients increased. Optimisations in V8.0 mean that *SHARECNV(1)* is now closer to *SHARECNV(0)* as the number of clients increases.

1.4 Reduced Client Connection Time

Note: The data for connection time improvements applies to **V8.0.0.3** onwards, where improvements were made.

Figure 5 reflects the rate at which the queue manager can accept inbound MQI (native) clients. Note that the V8.0.0.3 rates are now believed to be constrained by our test infrastructure, rather than the rate at which the queue manager can accept new clients.

In the test, 60,000 MQI clients connected to the queue manager and opened a shared queue. The time taken to do this reduced from 548 seconds in V7.5 to 60 seconds in V8.0.0.3

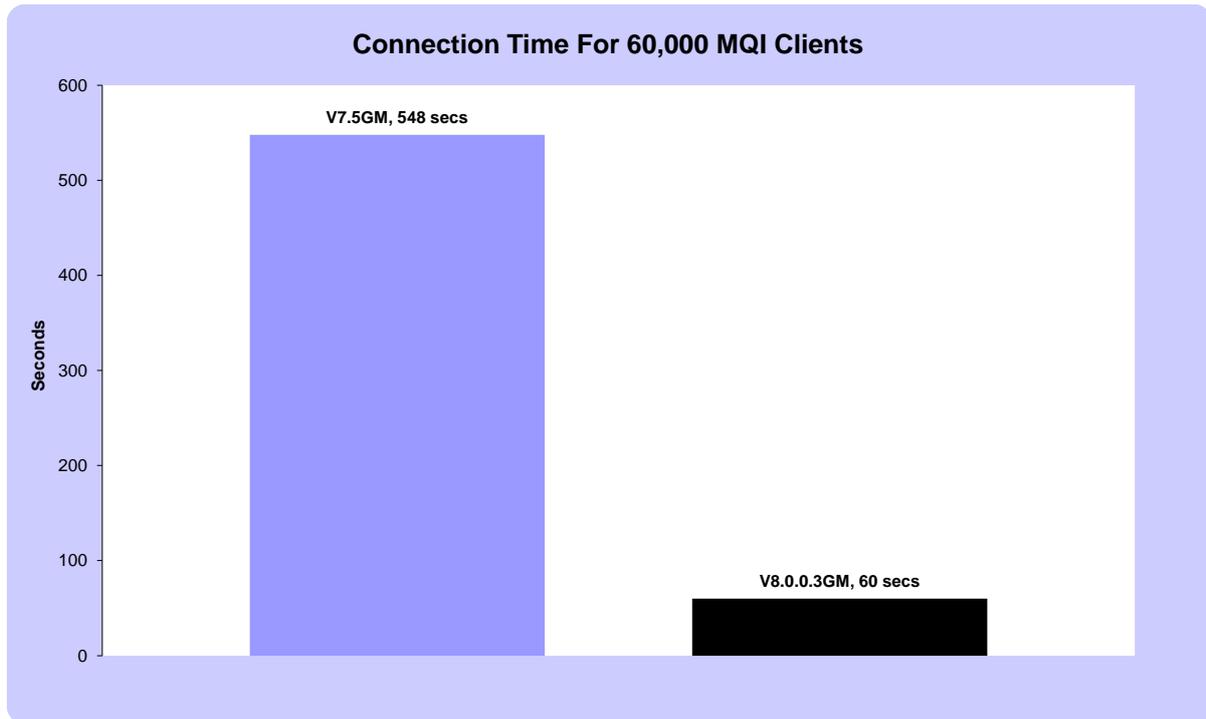


Figure 5 – MQI Client connect times

2 Workloads

Table 1 (below) lists the common workloads used in the generation of performance data for this report.

Each workload has a topology that is detailed in the following sections and is further categorised by the type of applications that run against this topology.

Where a measurement uses one of the common workloads, its reference (e.g. RR-1) is given in the chart, table or text.

Workload	Topology	Requester (or producer)	Responder (or consumer)
RR-1	Requester-responder with local queue manager. (section 2.1.1)	MQI, local bindings	MQI local bindings
RR-2	Requester-responder with remote queue manager (local responders). (section 2.1.2)	MQI, client bindings	MQI, local bindings
RR-2b	Requester-responder with remote queue manager (local responders) across n Queues. (section 2.1.2.1)	MQI, client bindings	MQI, local bindings
RR-3	Requester-responder with local queue manager. (section 2.1.1)	JMS, local bindings	JMS, local bindings
RR-4	Requester-responder with remote queue manager (remote responders). (section 2.1.3)	JMS, client bindings	JMS, client bindings
RR-4b	Requester-responder with remote queue manager (local responders) across n request queues with dynamic reply queues. (section 2.1.2.2)	JMS, client bindings	JMS, client bindings
RR-5	Requester-responder with distributed queuing across multiple channels. (section 2.1.5)	MQI, local bindings	MQI, local bindings
RR-5b	Requester-responder with distributed queuing across n queues with dynamic reply queues. (section 2.1.4.1)	JMS, local bindings	JMS, local bindings
OW-1	One-way messaging with local queue manager. (section 2.1.6)	MQI, local bindings	MQI, local bindings

Table 1 - Workload types

2.1 Workload Topologies

Three basic topologies of workload are detailed in the following sections:

- Requester-responder to (local or client bound applications).
- Requester-responder using distributed queuing.
- One-way (point to point) with local queue manager.

Most of the results in this report use the requester-responder topology. One-way tests are used in a few cases to alleviate queue locking and show the upper limit of a queue manager where the bottleneck is CPU or I/O (disk or network).

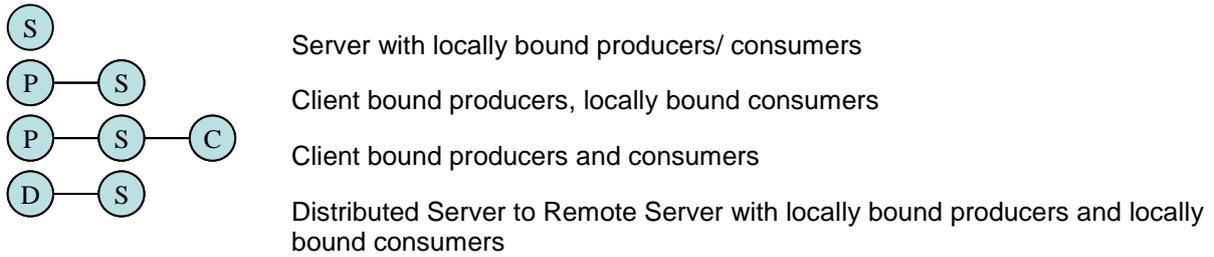
Requester-responder tests are synchronous in style because the application putting a message on a queue will wait for a response on the reply queue before putting the next message. They typically run 'unrated' (no think time between getting a reply and putting the next message on the request queue).

One-way tests are simple point-to-point tests where one application puts messages on a queue and a separate application consumes them with no reply. As such, they need to be rated to avoid filling the queue.

You will find pictorial representations that illustrate the topology in use for each test. They use the symbols:

Symbol	Description	Hardware
P	Producer Machine	Client
S	Server Machine	Server
C	Consumer Machine	Client
D	Distributed Server with local clients	Client

If the Producer or Client symbol is not illustrated, they will be co-located on the server hardware e.g.



Note that for any topology there may be more than one of each type of machine, e.g. for the client bound producers the test may use 2 physical machines to generate the load. This is made clear in the report.

2.1.1 Requester-responder with local queue manager. S

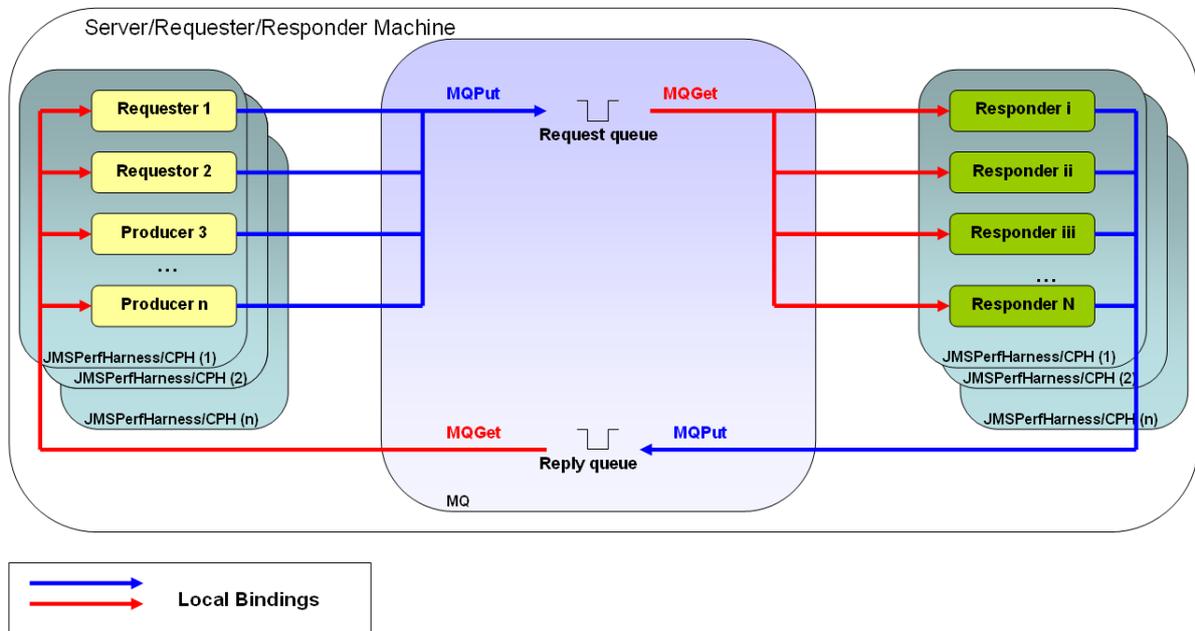


Figure 6 - Requester-responder with local queue manager

The measurements for the requester-responder with local queue manager topology are generally for processing messages with no think-time (scaling tests may have a think time to enable large numbers of clients to be connected).

Figure 6 shows the topology of this test. All applications and the queue manager run on a single machine. The applications use standard local bindings.

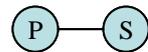
The test simulates multiple ‘requester’ applications which all put messages onto a single request queue. Each requester is a thread running in an MQI (CPH) or JMS (JMSPerfHarness) application. The number of requesters is dictated by the number of threads per application multiplied by the number of application instances started. Additional machines may be used to drive the requester applications where necessary.

Another set of ‘responder’ applications retrieve the message from the request queue and put a reply of the same length onto a shared reply queue. Each responder is a thread of CPH or JMSPerfHarness and there may be multiple instances of these MQI or JMS applications, similar to the responders. The number of responders is set such that there is always a waiting ‘getter’ for the request queue.

The flow of the test is as follows:

1. The requester application puts a message to the common request 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 request 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.

2.1.2 Requester-responder with remote queue manager (local responders).



This topology is similar to requester-responder but the requester applications connect through client channels.

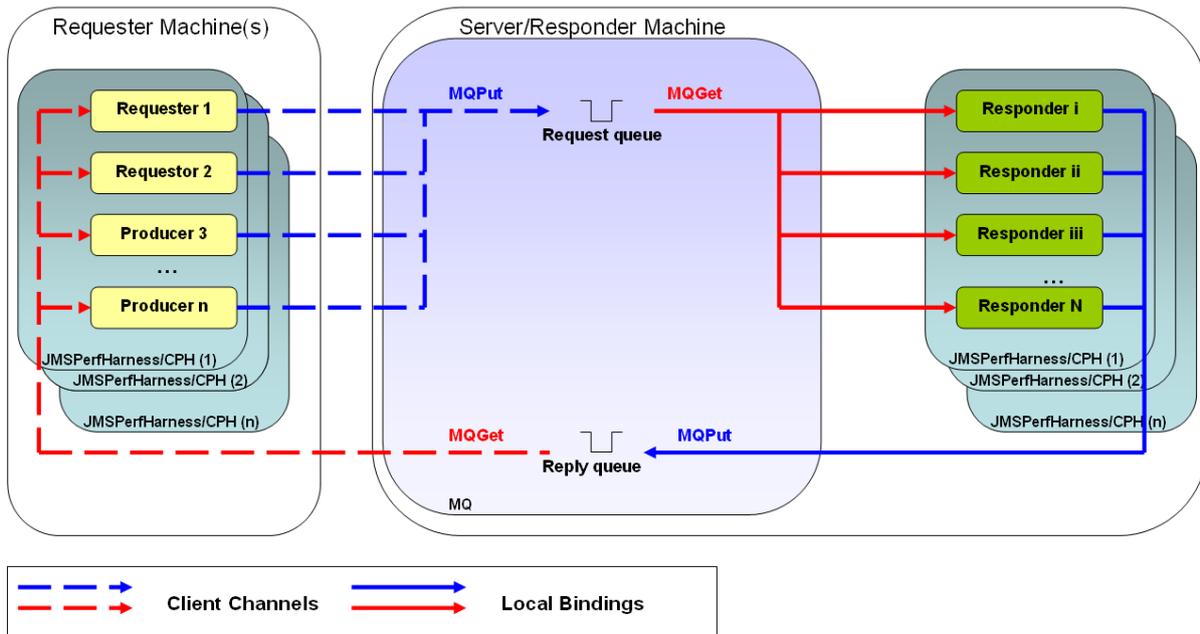
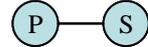


Figure 7 - Requester-responder with remote queue manager (local responders).

Figure 7 above shows the topology of the test. The requester applications are now hosted on a separate machine to the queue manager and use client connections with *FASTPATH* enabled (see

section 6.1.2). Responders are locally bound. The flow of the test is the same as that detailed in section 2.1.1

2.1.2.1 Requester-responder with remote queue manager (local responders) across n Queues.



This topology is a variant of the previous one. Instead of there being a single request queue and a single response queue the load is spread across n queues. Each producer may have its own request and reply queue or there may be n request queues with n reply queues where the number of producers $\geq n > 1$. Each request queue will still have a set of responders associated with it.

Running the test with n request and n reply queues alleviates queue locking.

2.1.2.2 Requester-responder with remote queue manager (local responders) across n request queues with dynamic reply queues.



This topology is a variant of the n queue topology described in section 2.1.2.1 but responses are sent via dynamic queues instead of using pre-defined reply queues.

2.1.3 Requester-responder with remote queue manager (remote responders).



This topology is similar to requester-responder but both the requester *and* responder applications connect through client channels.

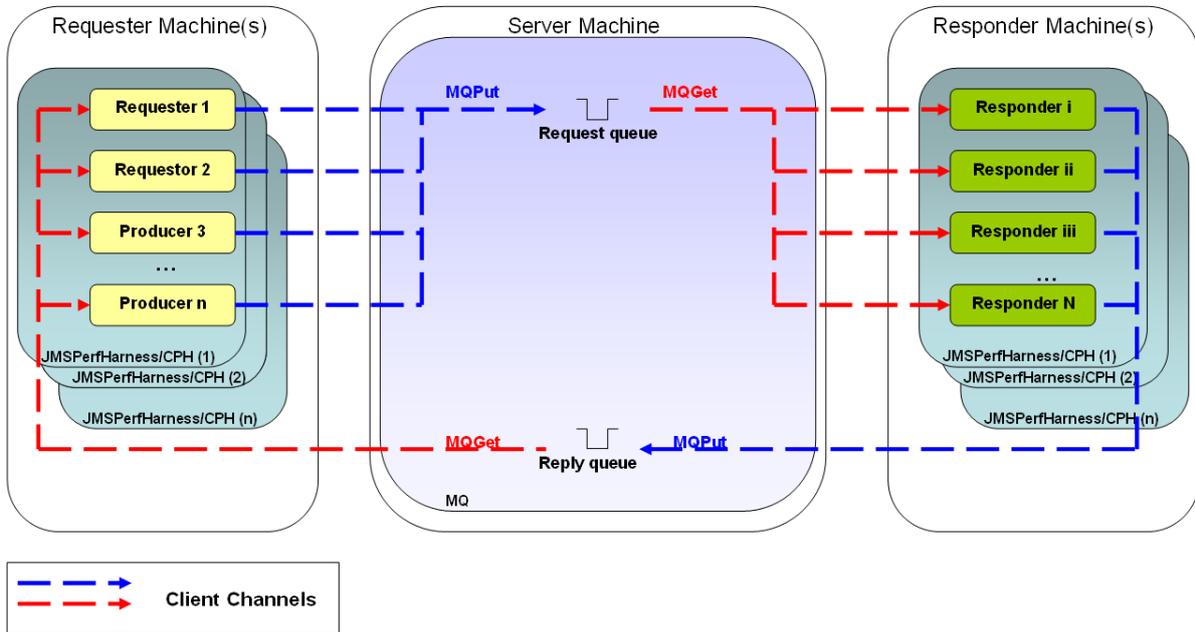


Figure 8 - Requester-responder with remote queue manager (remote responders).

Figure 8 above shows the topology of the test. The requester *and* responder applications are now hosted on separate machines to the queue manager and use client connections with *FASTPATH* enabled (see 6.1.2). The flow of the test is the same as that detailed in section 2.1.1

This topology is typically used for JMS where the multiple JVM based responders are better suited to being hosted on a separate machine to the queue manager where they will not affect the performance of the queue manager and enable the test to scale further.

2.1.4 Requester-responder with distributed queuing

D — S

This is a distributed queuing version of the requester-responder topology detailed in section 2.1.1. *MQPUTs* are to remote queues so that messages are now transported across server channels to the queue manager where the queue is hosted.

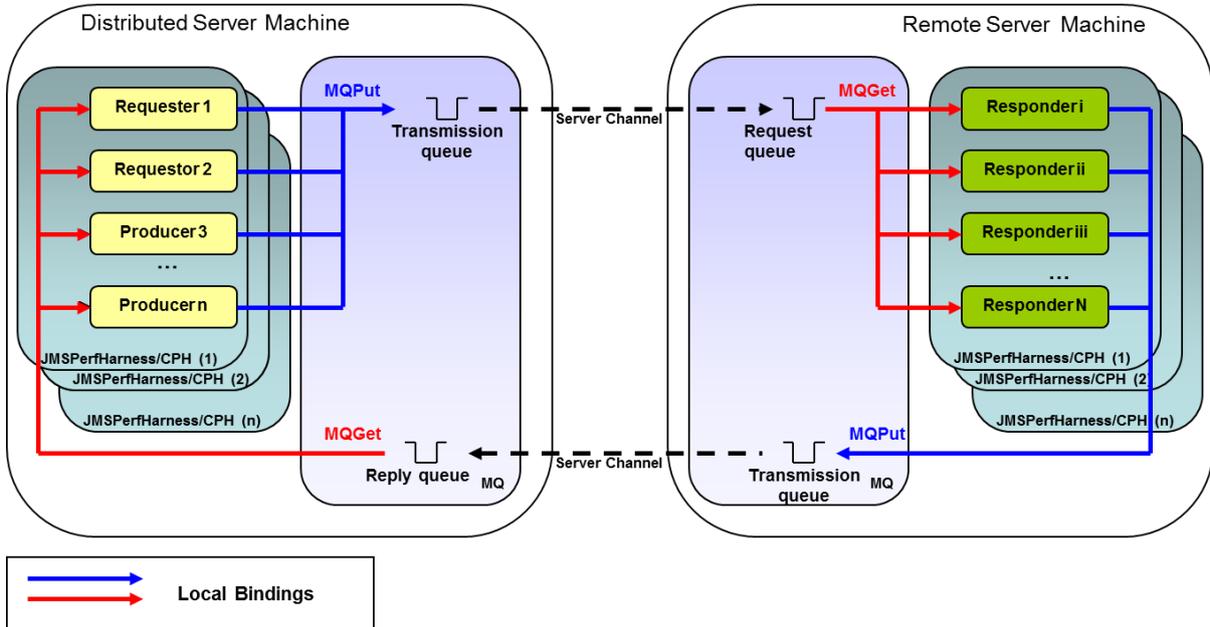


Figure 9 - Requester-responder with distributed queuing

Figure 9 above shows the topology of the test. The requester *and* responder applications are locally bound. All *MQPUTs* are to a remote queue via the local transmission queue. The flow of the test is the same as that detailed in section 2.1.1

2.1.4.1 Requester-responder with distributed queuing across n queues with dynamic reply queues.

D — S

This topology is a variant of the single server channel distributed queuing topology described in section 2.1.4. Instead of a single request queue the requesters put messages onto n request queues in a round-robin fashion and receive the reply on a dynamic response queue.

2.1.5 Requester-responder with distributed queuing across n queues with multiple channels

D — S

This topology is similar to that described in section 2.1.4 except that multiple server channels are employed to fully utilise the network bandwidth. This is particularly important for non-persistent messaging where the transmit queue may become the bottleneck. Note that the use of multiple server channels does not guarantee a performance improvement. If your deployment of MQ utilises a fast network but comprises of persistent messages you may find that the disk i/o latency is such that adding more server channels will not make any difference to throughput.

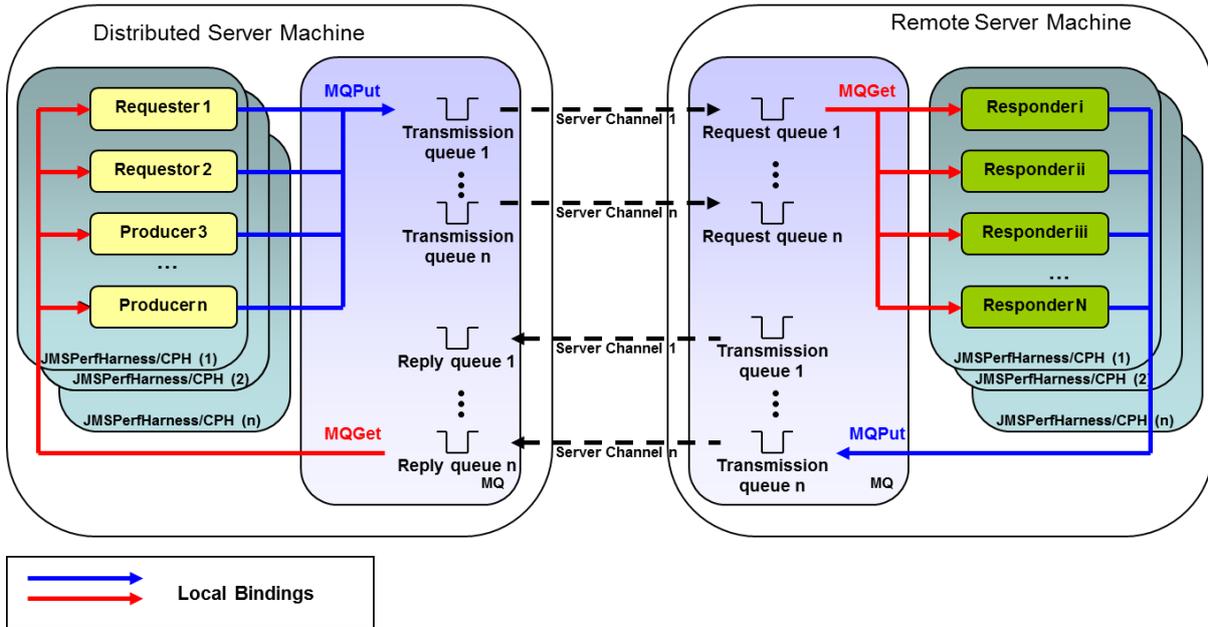


Figure 10 - Requester-responder with distributed queuing across multiple channels

Figure 10 above shows how this topology scales by adding more requesters, transmit queues, channels and responders. In practise the requesters do not all send messages to request queue one, they actually cycle round the number of request queues specified. The figure above becomes too complicated to show this effectively. Each request queue also has its own set of dedicated responders which send the response back via the reply queue specified by the requester.

Taking a simple example, we may have four requesters sending messages to two request queues. Each request queue is reached via its own dedicated channel. Requester one will send messages to request queue one and wait for a reply on reply queue one, requester 2 will send messages to request queue two (waiting for a reply on reply queue 2) and then requester three will cycle back to request queue one.

Figure 11 illustrates this simple deployment of the topology. The test can be scaled by adding more producers (e.g. increasing the producers from 4 to 8 will then result in 2 additional connections to each request queue). More channels may also be specified at the start if required (10 channels are used in most of the multi-channel distributed queuing tests in this report).

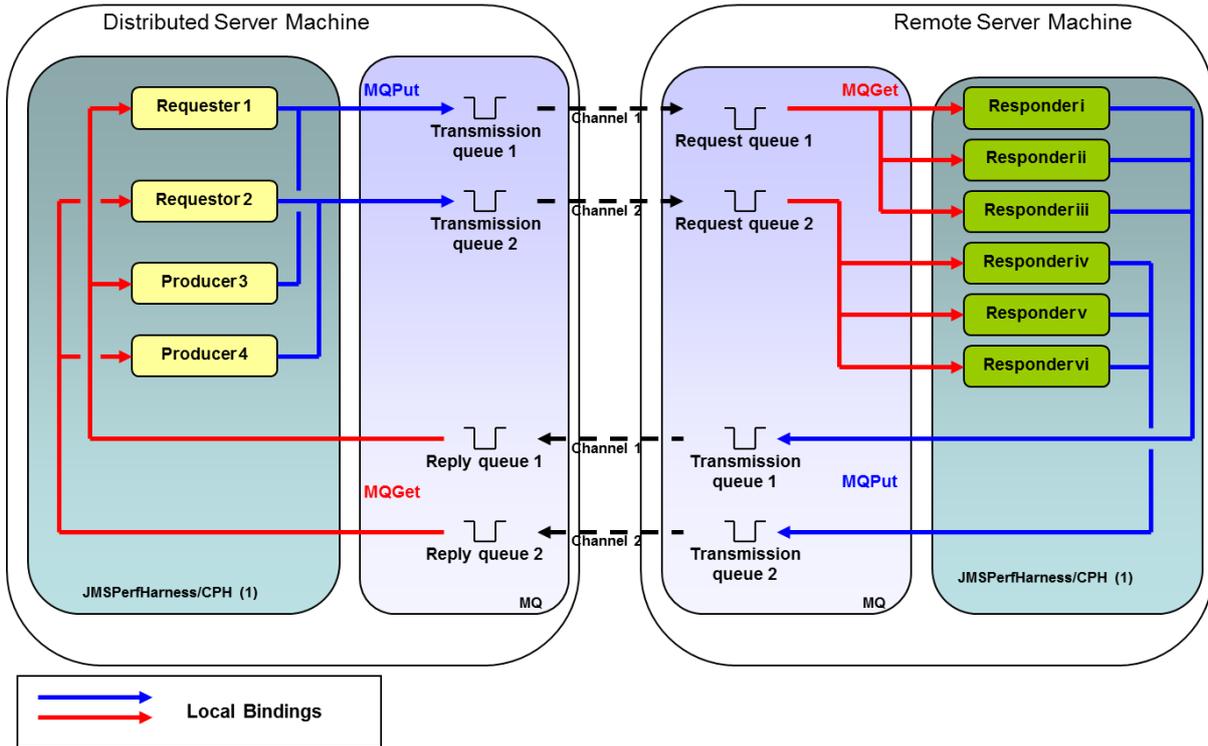


Figure 11 - Multi-channel distributed queuing, sample deployment.

2.1.6 One-way messaging with local queue manager. S

One-way tests are simple point-to-point topologies where each ‘producer’ has a unique queue with one or more dedicated ‘consumers’ per queue getting the message.

One-way tests start with one set of resources comprising of a producer, a dedicated queue and its associated consumer(s). The test is scaled up by adding additional sets (e.g. 5 consumers, 5 dedicated queues and their consumer(s)). Better performance can be achieved by having n consumers per queue (where n is typically small) so a set can be thought of as a producer, a dedicated queue and a group of associated consumers.

Each producer is a thread running in an MQI (CPH) or JMS (JMSPerfHarness) application. The number of producers is dictated by the number of threads per application * the number of application instances started. Additional machines may be used to drive the requester applications where necessary.

Similarly, each consumer is a thread of CPH or JMSPerfHarness and there may be multiple instances of these MQI or JMS applications, similar to the producers.

The flow of the test is as follows:

1. The producer application puts a message onto a queue which is for exclusive use by itself and the consumers for that queue. The producer then waits for a specified time before looping round and putting the next message on the queue. In this way the test is said to be rated.
2. One of the consumers associated with the queue gets the message and throws it away.

The rate of message delivery for each producer is determined beforehand so that as the test is scaled up by adding more sets of producer/queue/consumer(s), the queues are kept drained by the associated consumers (i.e. we don’t deliver at a rate that would cause the queue depth to increase and result in non-optimal flows in the product. An adequate number of consumers per queue are created so that there is always one waiting to get the next message. This ensures we run in message mode as much as possible (i.e. the message is copied from producer to consumer as much as possible rather than needing to put on the queue).

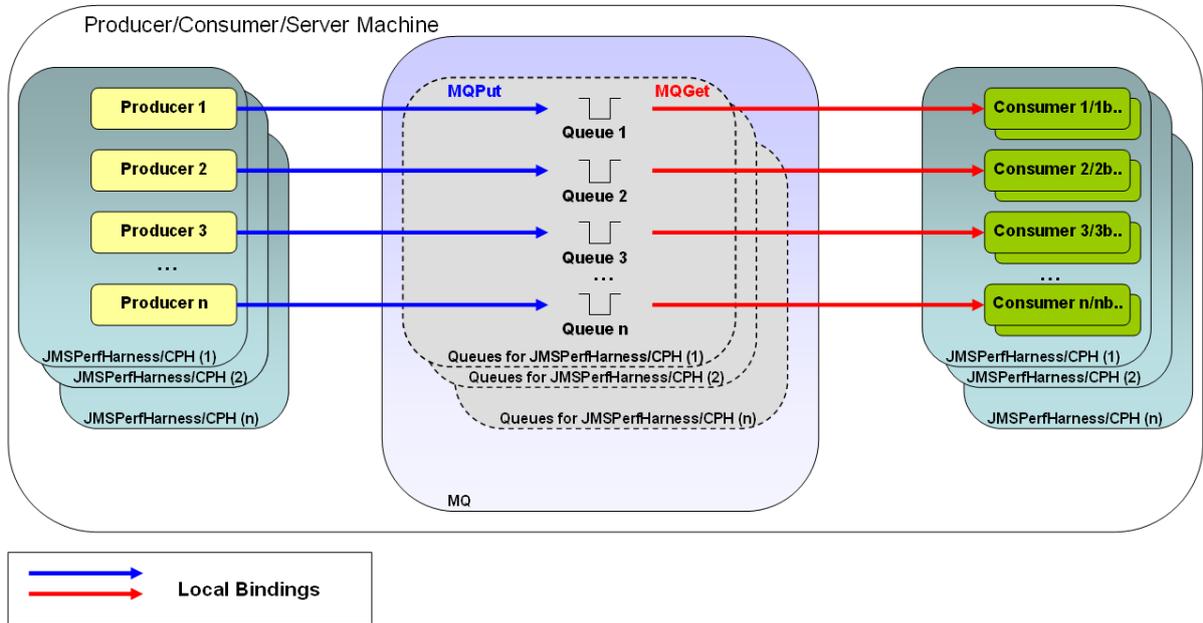


Figure 12 – One-way messaging with local queue manager

Figure 12 shows the one-way topology. The test is scaled by adding more sets of producer/queue/consumer(s).

As each queue in a one-way test has only one application putting messages on it we relieve queue locking which affects the requester-responder tests more due to the shared request/reply queues in most of those tests.

2.1.7 One-way messaging with remote queue manager.



This topology is similar to one-way messaging with a local queue manager except that in this case the producer and consumer applications are connecting to MQ from a different host across client channels.

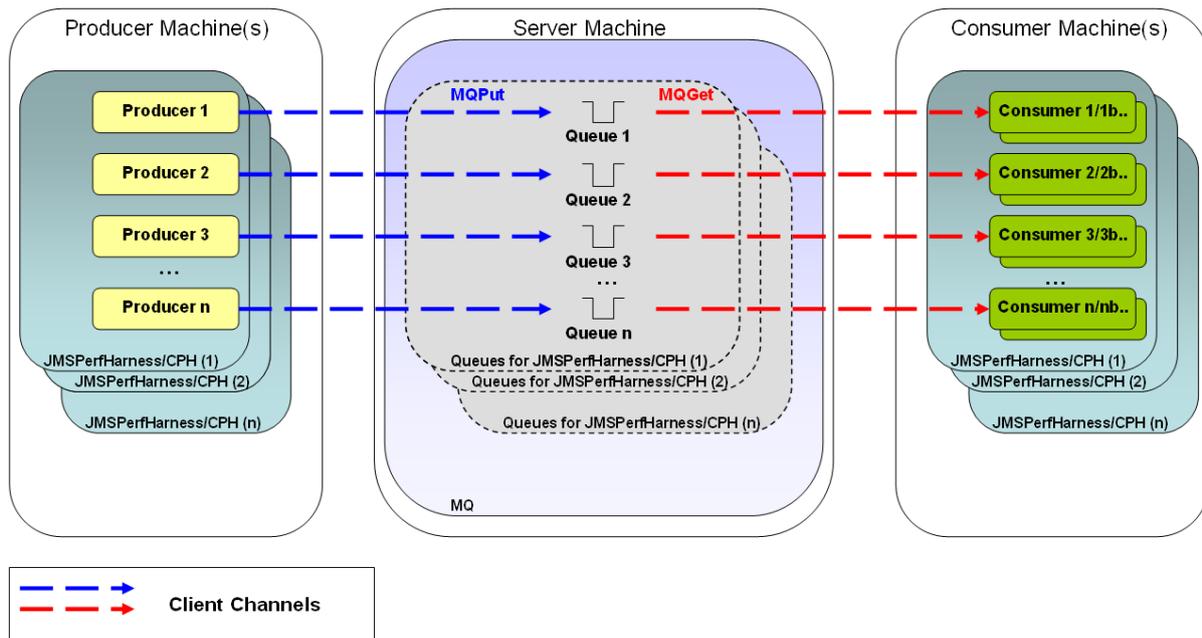


Figure 13 - One-way messaging with remote queue manager

Figure 13 above shows the topology of the test. The producer and responder applications are hosted on separate machines to the queue manager and use client connections with *FASTPATH* enabled (see section 6.1.2). The flow of the test is the same as that detailed in section 2.1.6.



3 Performance Test Results

Full performance test results are detailed below. The test results are presented by broad categories with an illustrative plot in each section followed by the peak throughput achieved for the remaining tests in that category (the remaining tests are typically for different message sizes).

For information on the hardware used in these tests please see Appendix C.: The topology used for each test is indicated alongside each chart’s caption (see section 2.1)

3.1 MQI Local Bindings (Non-Persistent)

The workload type for these measurements is RR-1 (see section 2). Requester and responder applications are on the same machine. The test is run unrated (i.e. each requester sends a new message as soon as it receives the reply to the previous one).

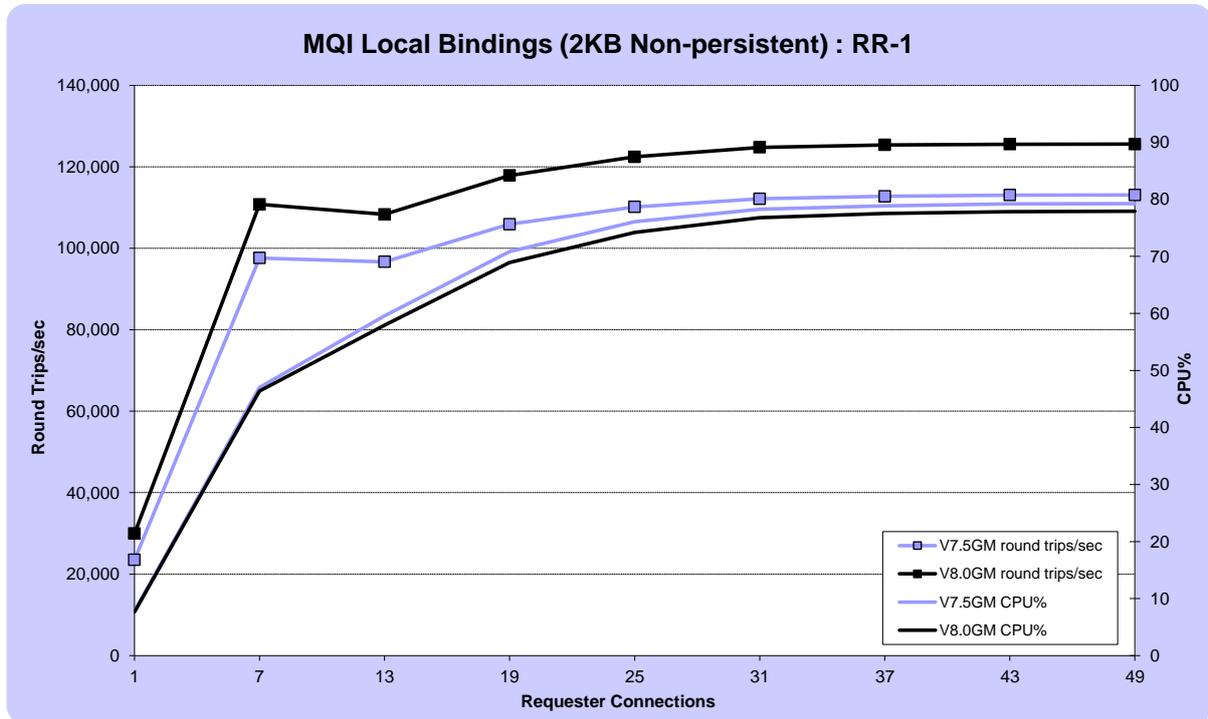


Figure 14 – Performance results for MQI Local Bindings (2KB Non-persistent)



V8.0 is up to 11% improved over V7.5 in terms of throughput.

Test	V8.0GM			V7.5GM		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
MQI Local Bindings (2KB Non-persistent) : RR-1	125,553	77.92	49	113,059	79.26	49
MQI Local Bindings (20KB Non-persistent) : RR-1	104,039	81.15	41	93,582	82.5	41
MQI Local Bindings (200KB Non-persistent) : RR-1	29,770	75.52	22	28,962	76.89	22
MQI Local Bindings (2MB Non-persistent) : RR-1	1,100	54.9	13	1,074	50.76	12

*Round trips/sec

Table 2 - Peak rates for MQI local bindings (non-persistent)

3.1.1 Test setup

Workload type: RR-1 (see section 2).
 Hardware: Server 1 (see Appendix C.1).



3.2 MQI Local Bindings (Persistent)

The workload type for these measurements is RR-1 (see section 2). Requester and responder applications are on the same machine. The test is run unrated (i.e. each requester sends a new message as soon as it receives the reply to the previous one).

This test is sensitive to the way the logger and log files have been setup and the speed of the filesystem being written to (see section 4.1).

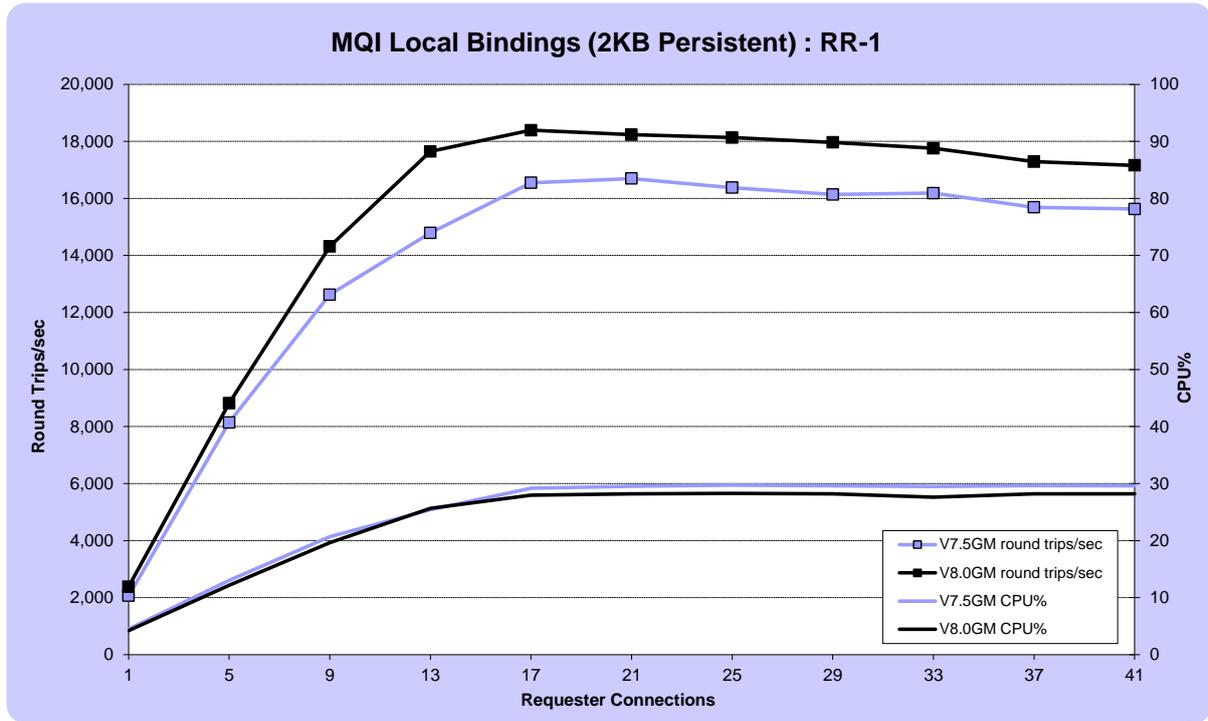


Figure 15 - Performance results for MQI local bindings (2KB persistent)



V8.0 is up to 10% improved over V7.5 in terms of throughput

Test	V8.0GM			V7.5GM		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
MQI Local Bindings (2KB Persistent) : RR-1	18,391	27.98	17	16,695	29.52	21
MQI Local Bindings (20KB Persistent) : RR-1	8,419	15.45	21	8,574	18	21
MQI Local Bindings (200KB Persistent) : RR-1	891	3.56	1	881	5.23	2
MQI Local Bindings (2MB Persistent) : RR-1	87	5.4	2	88	4.16	1

*Round trips/sec

Table 3 - Peak rates for MQI local bindings (persistent)

3.2.1 Test setup

Workload type: RR-1 (see section 2).
Hardware: Server 1 (see Appendix C.1).

3.3 MQI Client Bindings (Non-Persistent)

The workload type for these measurements is RR-2 (see section 2). The requester applications are on machines separate from the server machine. The responder applications are located on the server



machine. The test is run unrated (i.e. each requester sends a new message as soon as it receives the reply to the previous one).

This test is executed using client channels as trusted applications programs by specifying “MQIBindType=FASTPATH” in the qm.ini file. This is recommended generally, but not advised if you run channel exit programs and do not have a high degree of confidence in their robustness.

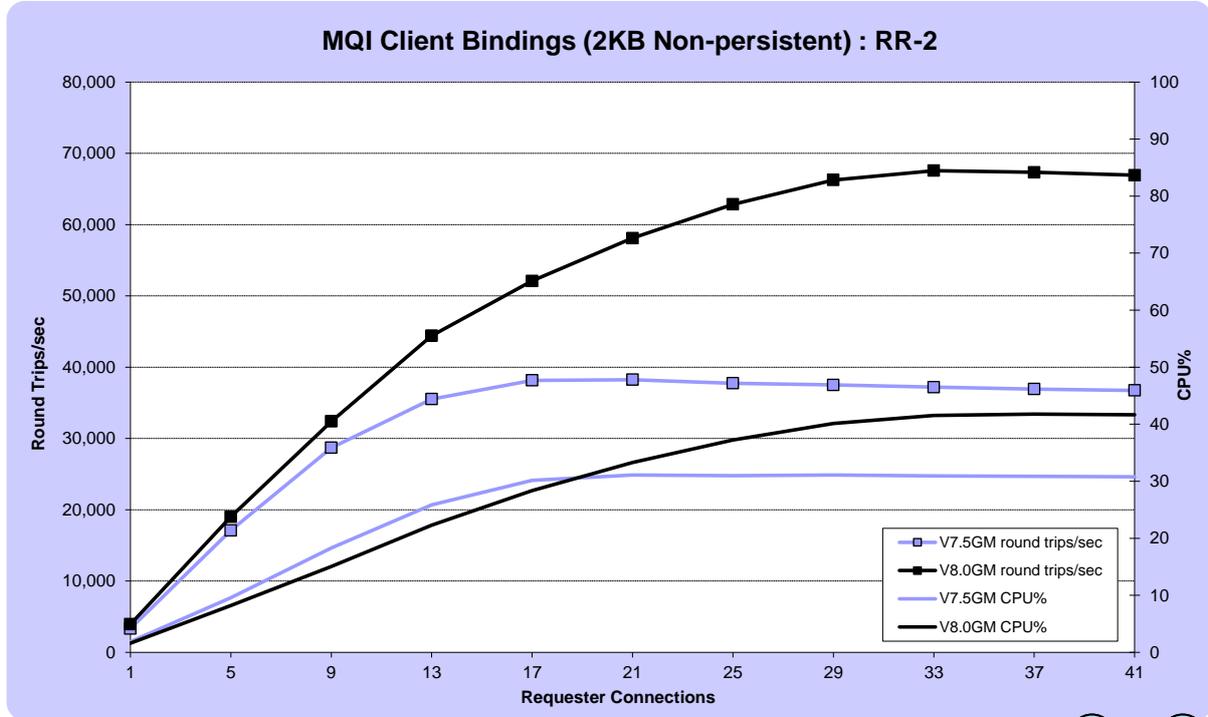
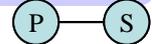


Figure 16 - Performance results for MQI client bindings (2KB non-persistent)



V8.0 is up to 77% improved over V7.5 in terms of throughput. This is mainly due to improvements in SHARECNV=1 performance, so isn't as pronounced in larger message sizes, where data transfer rates become the bottle-neck.

Test	V8.0GM			V7.5GM		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
MQI Client Bindings (2KB Non-persistent) : RR-2	67,568	41.52	33	38,239	31.1	21
MQI Client Bindings (20KB Non-persistent) : RR-2	50,956	41.5	61	29,008	29.07	25
MQI Client Bindings (200KB Non-persistent) : RR-2	5,641	14.17	45	5,592	17.82	45
MQI Client Bindings (2MB Non-persistent) : RR-2	498	14	23	497	13.97	23

*Round trips/sec

Table 4 - Peak rates for MQI client bindings (non-persistent)

3.3.1 Test setup

Workload type: RR-2 (see section 2).
 Hardware: Server 1, Client 1, Client 2 (see Appendix C.1).

3.4 MQI Distributed Queuing (Non-Persistent)

The distributed queuing scenarios use workload type RR-5 (see section 2) where locally bound requesters put messages onto a remote queue.

The throughput will be sensitive to network tuning and server channel setup amongst other things. All of the tests in this section utilise multiple send/receive channels (see section 4.5 for a single channel



comparison). This particularly helps with smaller, non-persistent messages when the network is under-utilised.

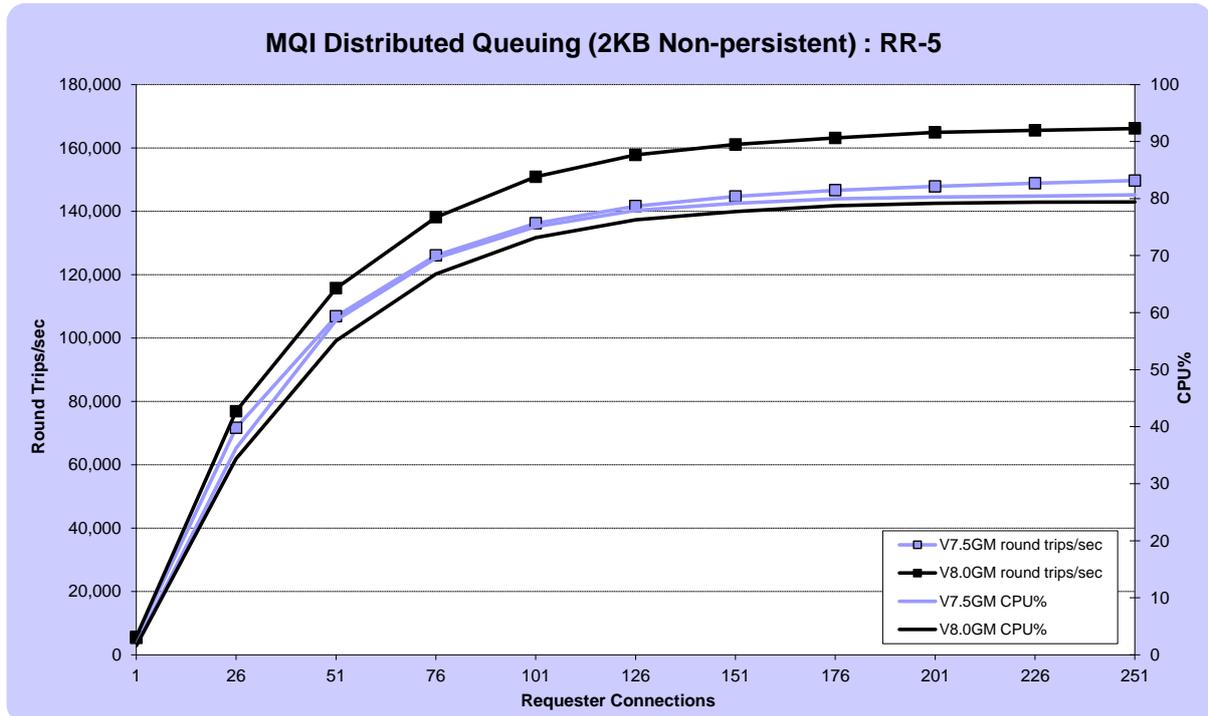
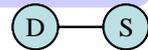


Figure 17 - Performance results for distributed queuing (2KB non-persistent)



V8.0 is up to 10% improved over V7.5 in terms of throughput. For larger message sizes there is a much greater improvement in performance at lower numbers of client connections or channels due to the improved TCP settings (though these *can* be set for V7.5.0.3 onwards), see section 1.2 for details.

Test	V8.0GM			V7.5GM		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
MQI Distributed Queuing (2KB Non-persistent) : RR-5	166,111	79.39	251	149,653	80.66	251
MQI Distributed Queuing (20KB Non-persistent) : RR-5	55,668	34.5	43	55,455	39.34	49
MQI Distributed Queuing (200KB Non-persistent) : RR-5	5,671	11.78	21	5,616	14.27	41
MQI Distributed Queuing (2MB Non-persistent) : RR-5	552	22.34	37	534	18.75	41

*Round trips/sec

Table 5 - Peak rates for distributed queuing (non-persistent)

3.4.1 Test setup

Workload type: RR-5 (see section 2).
 Hardware: Server 1, Client 1 (see Appendix C.1).

3.5 JMS (Non-Persistent)

The workload type for these measurements is RR-3 (see section 2). Requester and responder applications share the same set of machines. The test is run unrated (i.e. each requester sends a new message as soon as it receives the reply to the previous one).

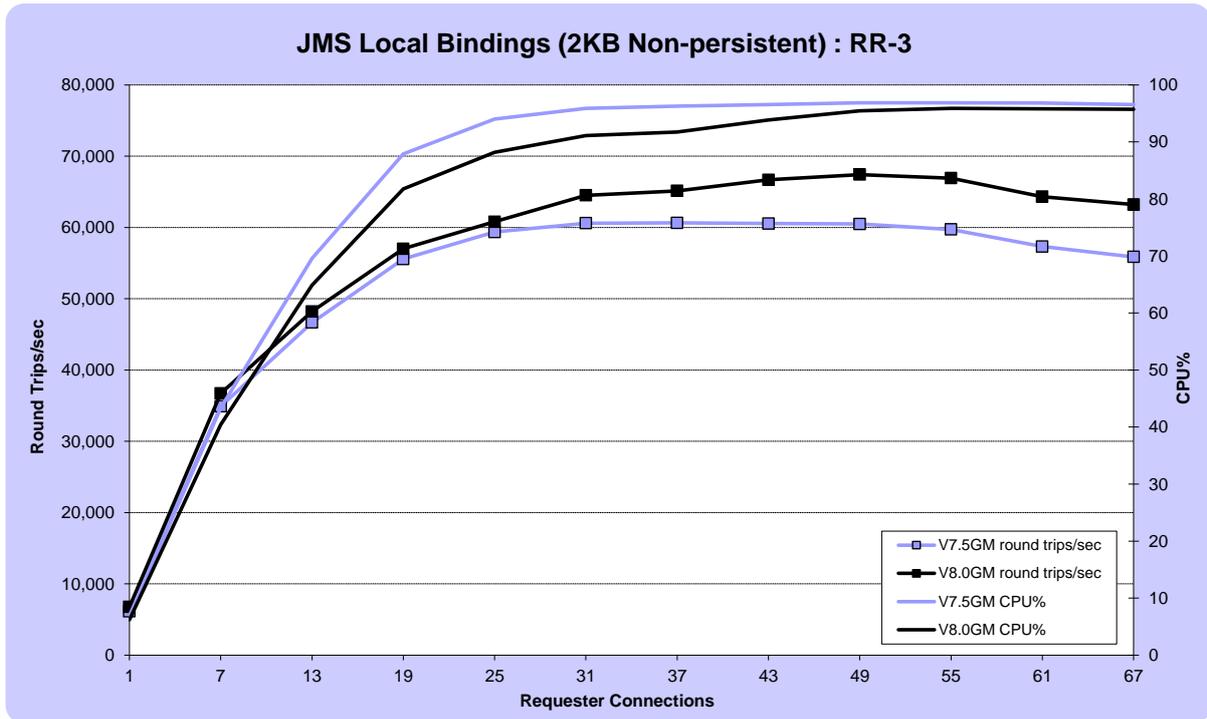


Figure 18 - Performance results for JMS local bindings (2KB non-persistent)



V8.0 is up to 29% improved over V7.5 in terms of throughput.

Test	V8.0GM			V7.5GM		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
JMS Local Bindings (2KB Non-persistent) : RR-3	67,404	95.44	49	60,640	96.26	37
JMS Client Bindings (2KB Non-persistent) : RR-4	56,159	45.13	65	43,435	46.85	49
JMS Local Bindings (200KB Non-persistent) : RR-3	14,761	77.18	17	14,416	80.49	17
JMS Client Bindings (200KB Non-persistent) : RR-4	2,752	8.98	37	2,748	11.35	37

*Round trips/sec

Table 6 - Peak rates for JMS (non-persistent)

3.5.1 Test setup

3.5.1.1 JMS local bindings

Workload type: RR-3 (see section 2).
 Hardware: Server 1, Client 1, Client 2 (see Appendix C.1).

3.5.1.2 JMS client bindings

Workload type: RR-4 (see section 2).
 Hardware: Server 1, Client 1, Client 2 (see Appendix C.1).

3.6 Memory Usage

To measure the memory usage of MQ, a simple scenario was devised in which clients connect remotely (client-bindings) to a queue manager. Each client puts a small message to a queue, then gets it back again, and repeats at a rate of 1 per second. By using a custom Linux kernel module, we ascertained which pages of real memory were mapped by MQ processes, and which of those had been accessed since all clients were connected. This experiment was repeated for varying numbers of clients spread across varying numbers of queues, and the following data were obtained:



	V8.0GM	V7.5GM
Cost of QM & 1 client accessing 1 queue	12,832.00KB	11,892.00KB
Cost of additional client	97.93KB	132.73KB
Cost of additional queue	32.85KB	38.73KB

Table 7 - Memory usage

A few points to note on these data:

1. The memory costs of additional clients and queues are the *average* increases in the memory usage of MQ seen when new clients or access to queues are added. Depending on the configuration of the queue manager, and the number of clients already connected, the additional cost of adding one more client or queue may be more or less than this. This is because certain processes and data structures within MQ only handle a limited number of objects, above which an additional instance will have to be created.
2. The cost of additional queues is not the amount of memory allocated or written to when a queue is defined. If a queue is not being accessed by an application, its associated memory is not accessed, so it can be paged out by the OS without consequence. The measured values indicate the amount of memory that is accessed regularly when one or more applications are using the queue.
3. All of these costs were measured during steady-state, light operation. They do not measure the memory used during heavy-weight operations such as QM start-up, *MQCONN*, and *MQOPEN*. During such operations, memory usage may be higher, but since they usually happen infrequently it would be reasonable to rely on a SWAP partition to fulfil these additional requirements if necessary.



3.7 MQI Client Bindings (Non-Persistent) - SSL

To illustrate the overhead of enabling SSL, results are provided comparing the performance of 4 MQ CipherSpecs with the baseline MQI Client Bindings 2KB test as seen in section 3.3.

The workload type for these measurements is RR-2 (see section 2). The Requester applications are on client machines separate from the server machine. The Responder applications are located on the server machine. The test is run unrated (i.e. each requester sends a new message as soon as it receives the reply to the previous one).

Queue manager authentication is used to setup the SSL conversation.

The ciphers under test are:

CipherSpec	Protocol	SuiteB	FIPS	Encryption bits
ECDHE_ECDSA_AES_256_GCM_SHA384	TLS1.2	Yes	Yes	256
TLS_RSA_WITH_AES_256_GCM_SHA384	TLS1.2	No	Yes	256
TRIPLE_DES_SHA_US	SSL3.0	No	No	168
NULL_SHA	SSL3.0	No	No	None

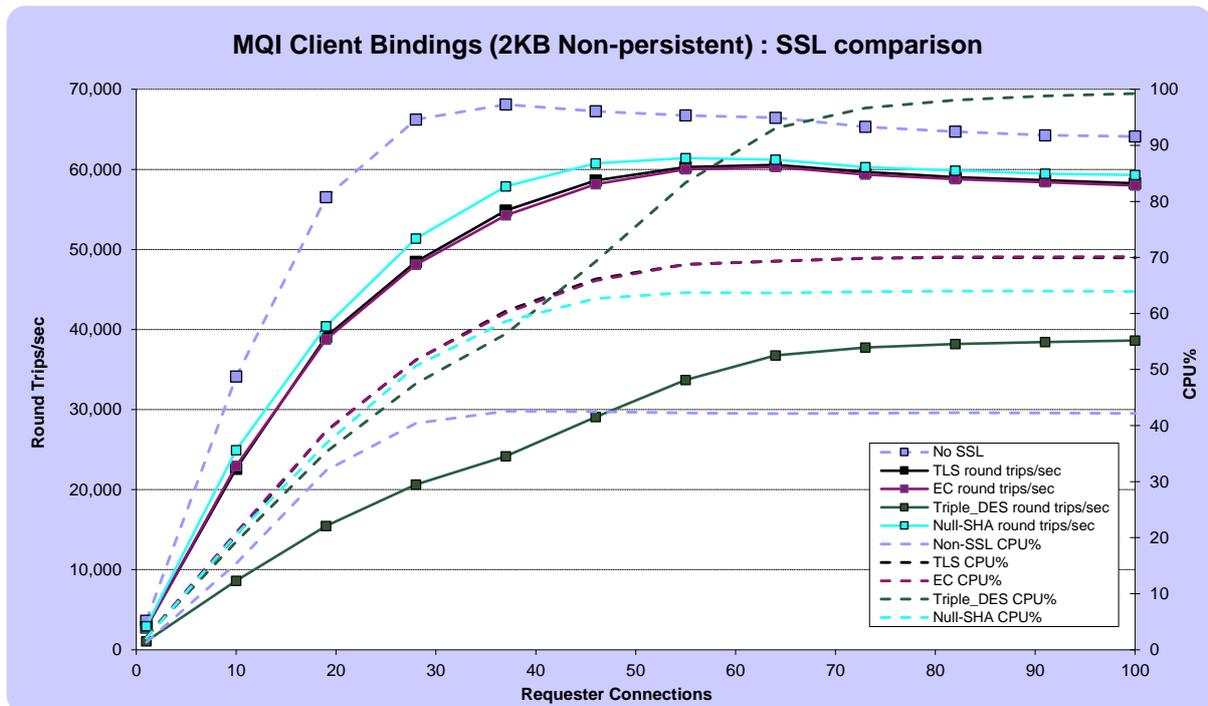


Figure 19 – Performance results for MQI client bindings (2KB non-persistent) - SSL



The newer FIPS compliant ciphers not only provide stronger encryption than some of the older CipherSpecs, but also perform better.

SSL Cipher	V8.0GM		
	Max Rate*	CPU%	Clients
ECDHE_ECDSA_AES_256_GCM_SHA384	60,317	69	64
TLS_RSA_WITH_AES_256_GCM_SHA384	60,568	69	64
TRIPLE_DES_SHA_US	38,626	99	100
NULL_SHA	61,404	64	55
No SSL	68,116	43	37

Table 8 - Peak rates for MQI client bindings (2KB non-persistent) - SSL



3.7.1 Test setup

Workload type: RR-2 (see section 2).

Hardware: Server 1, Client 1, Client 2 (see Appendix C.1).

Scaling Measurements

The scaling measurements in this section are designed to test a number of scenarios where there are a larger number of clients attached. Whereas the previous tests are optimised for throughput, these tests define an operational environment or scaling challenge to test from a performance perspective.

For information on the hardware used in these tests please see Appendix C.: The topology used for each test is indicated alongside each chart’s caption (see section 2.1)

3.8 MQI Rated Scaling

These scenarios test the point at which performance degrades when clients are continuously connected, each delivering messages at a low rate. The workload types are RR-2 & RR-2b (see section 2).

3.8.1 MQI rated scaling test setup

All of the rated scaling tests detailed in sections 3.8.n below used the same hardware and software.

Workload type: (see individual test).

Hardware: Server 1, Client 1, Client 2 (see Appendix C.1).

3.8.2 MQI client bindings (non-persistent)

For this test, clients were rated at 1.5 messages per second in order to achieve constraint before the operational limit of 65,536 clients was reached.

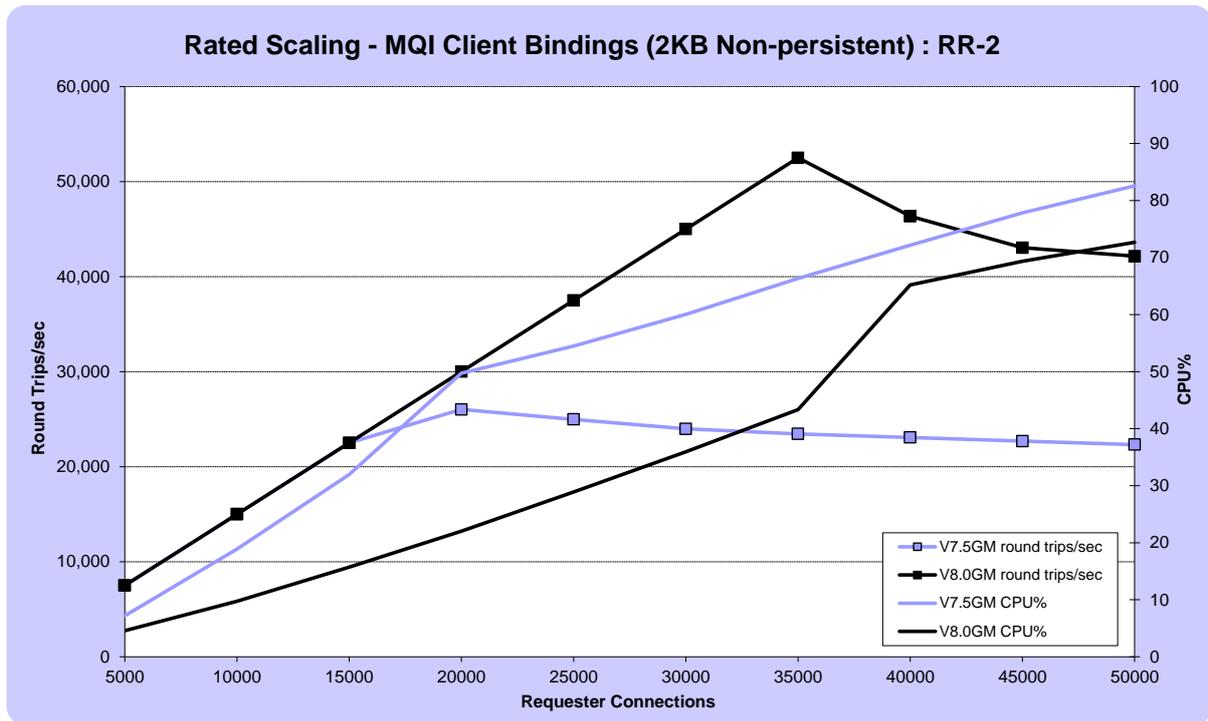
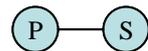


Figure 20 - Performance results for MQI client binding, rated scaling (2KB non-persistent)



For this test MQ V8.0 maintained the target messaging rate of 1.5 round trips/sec up to 35,000 clients before performance dropped off. This is in contrast to V7.5 where the significantly larger cost of processing caused the single pair of queues in use to become constrained with respect to processing power at 20,000 clients. Both versions scaled well to the point of constraint.

3.8.3 MQI client bindings (non-persistent) with 24 request/reply queues.

For this test, clients were rated at 2 messages per second in order to achieve constraint before the operational limit of 65,536 clients was reached.

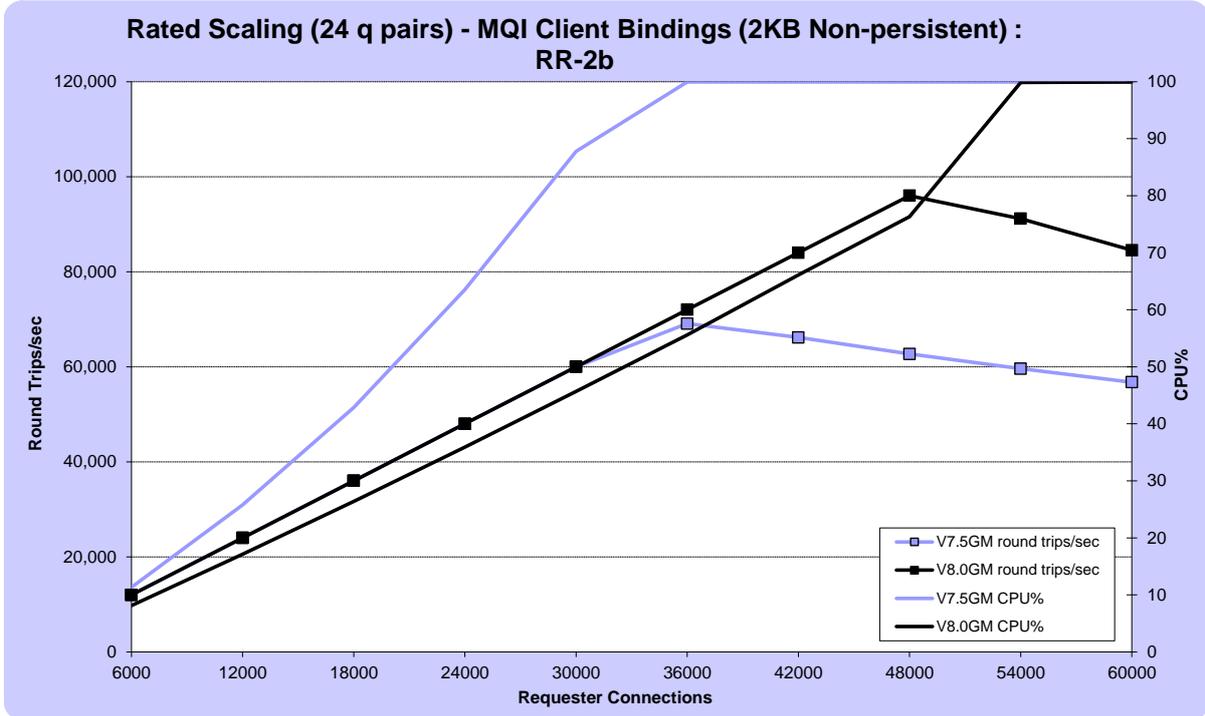
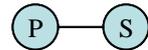


Figure 21 - Performance results for MQI client binding, rated scaling across 24 queues (2KB non-persistent)



This test is very similar to that shown in the previous section but with the workload spread out over 24 request/reply queue pairs locking is alleviated so that both V7.5 and V8 were able to saturate the processing power of the machine. V8 performance tailed off at 48,000 clients while V7.5 reached saturation at just below 36,000 clients.



3.8.4 MQI client bindings (persistent)

For this test, clients were rated at 1 message per second.

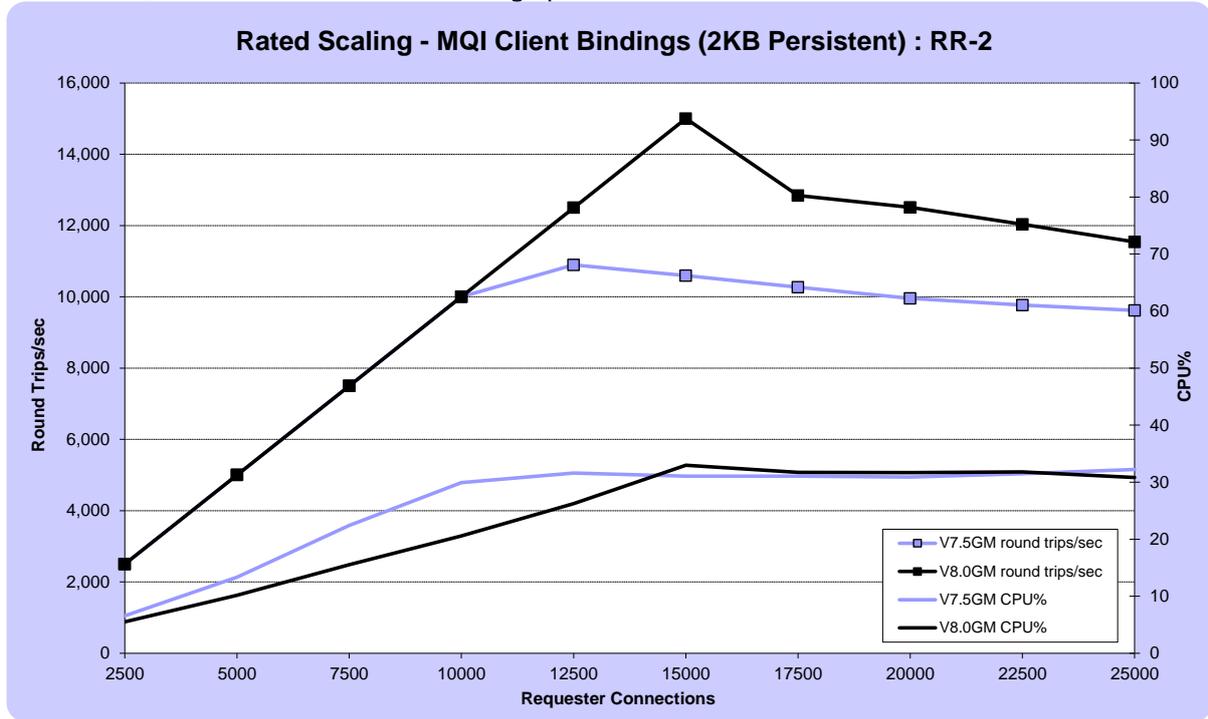
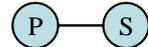


Figure 22 - Performance results for MQI client binding, rated scaling (2KB persistent)



For the persistent message variation on the single queue pair scaling test shown in section 3.8.2, V8 again out-performed V7.5 with 15,000 clients being connected before a drop-off in performance becomes apparent, compared to 10,000 clients in V7.5.

3.8.5 MQI client bindings (persistent) with 24 request/reply queues.

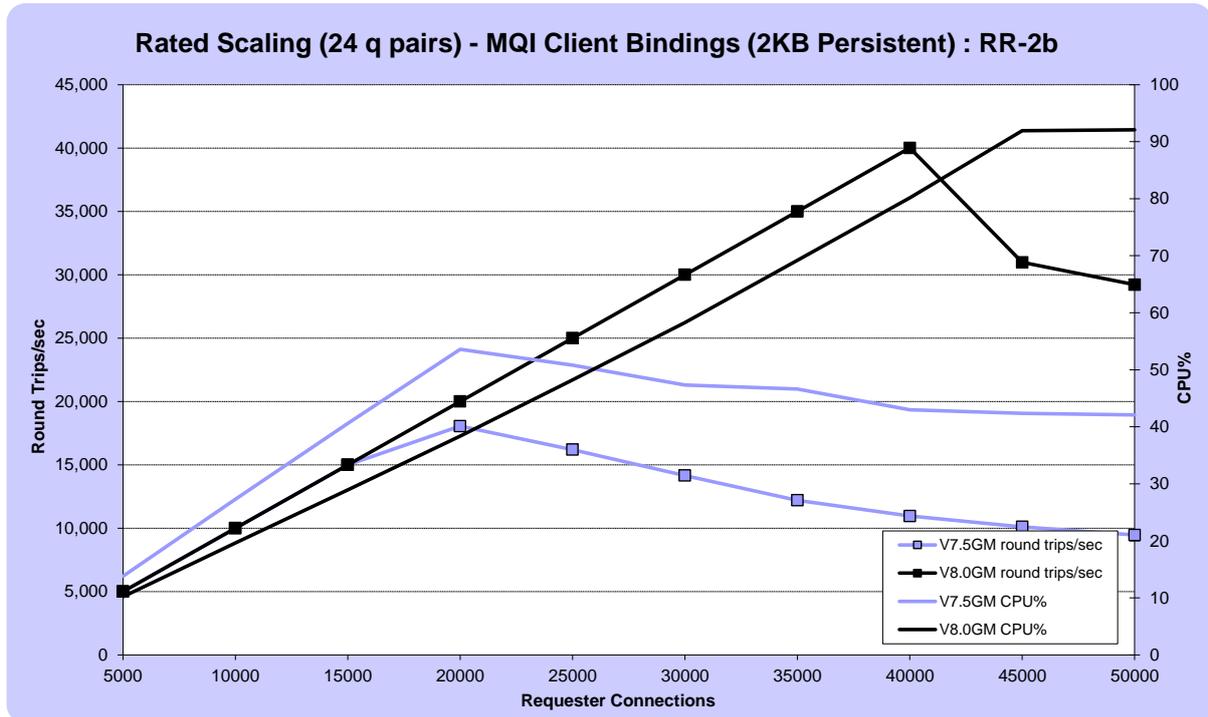
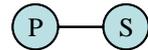


Figure 23 - Performance results for MQI client binding, rated scaling across 24 queues (2KB persistent)



With queue lock contention reduced, the improvements in the MQ logger become even more apparent, with V8 achieving over double the rate of V7.5, and nearly saturating the CPU on the machine.

3.9 JMS Rated Scaling with Dynamic Reply Queues.

3.9.1 JMS rated scaling test setup

All of the rated scaling tests detailed in section 3.9 (including distributed queuing scenarios) used the same hardware and software.

Workload type: (see individual test).

Hardware: Server 1, Client 1, Client 2 (see Appendix C.1).

3.9.2 JMS client bindings (Non-persistent) with 24 request queues.

All of the results up until this point in the report use pre-defined queues to send a response message to the requester. The tests in this section utilise dynamic reply queues, which are widely used in the field (test type RR-4b, see section 2).

Once again, each requester is rated, putting 1 message a second onto one of 24 request queues and specifying a dynamic response queue on which to receive its response.

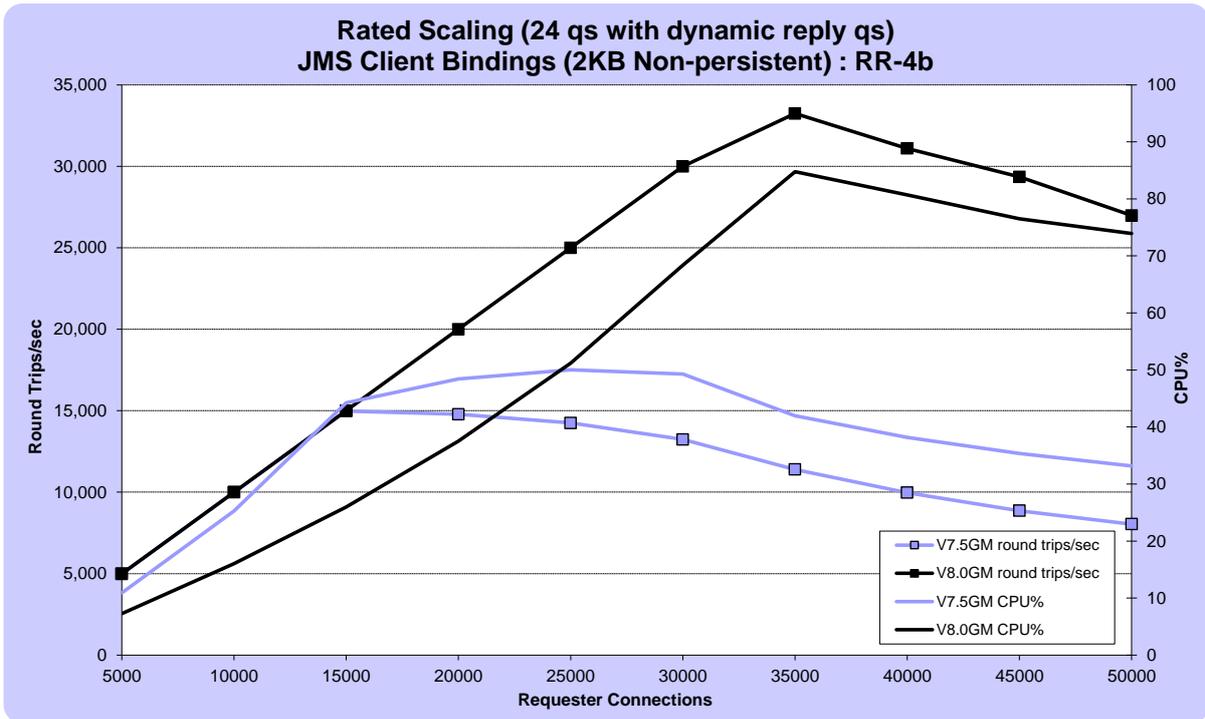


Figure 24 - Performance results for JMS client binding, rated scaling across 24 queues with dynamic reply queues (2KB non-persistent)



As with all of the scaling tests, V8 shows a significant improvement over V7.5. Both versions saturate the CPU with V8 scaling up to almost 35,000 clients before the performance drop-off compared to 15,000 for V7.5.

3.9.3 JMS client bindings (persistent) with 24 request queues.

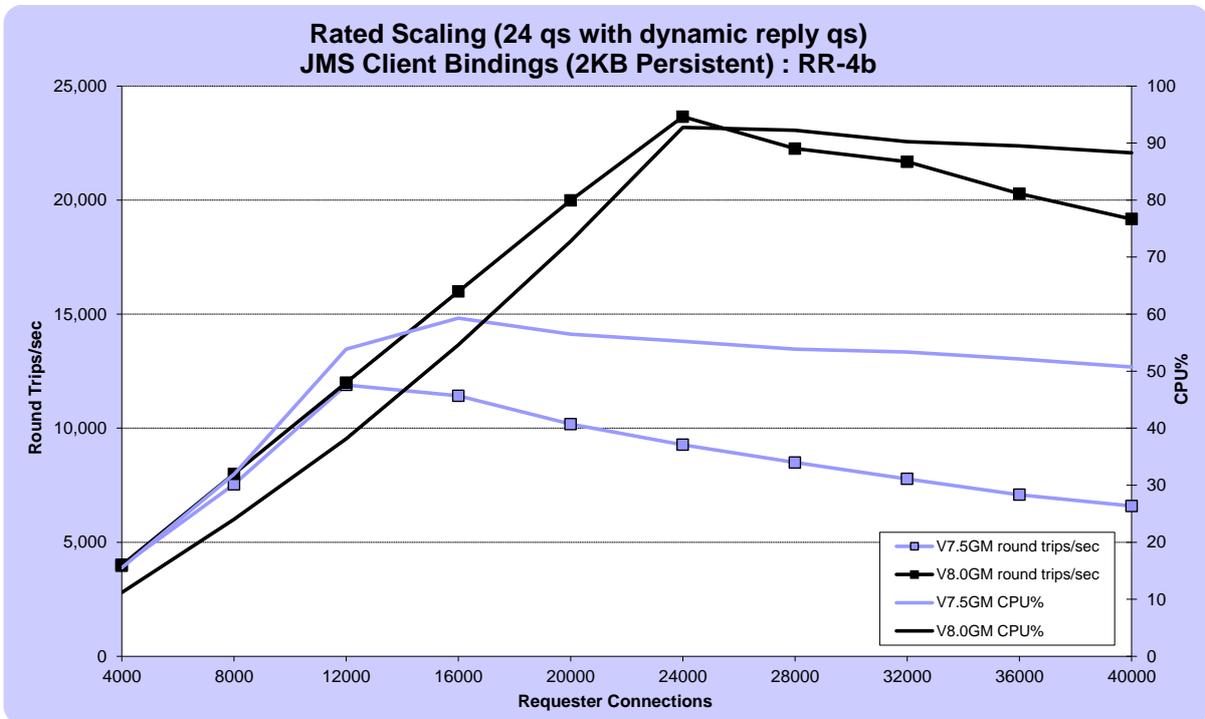
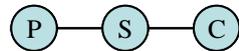


Figure 25 - Performance results for JMS client binding, rated scaling across 24 queues with dynamic reply queues (2KB persistent)



This is the persistent message version of the test shown in the previous section. V8.0 scales to 24,000 clients, double the 12,000 for V7.5.



3.9.4 JMS distributed queuing (non-persistent) with 24 request queues.

The following two tests are similar to those in sections 3.9.2 & 3.9.3 above but the reply queues are remote so distributed queuing is used. The test type is RR-5b (see section 2)

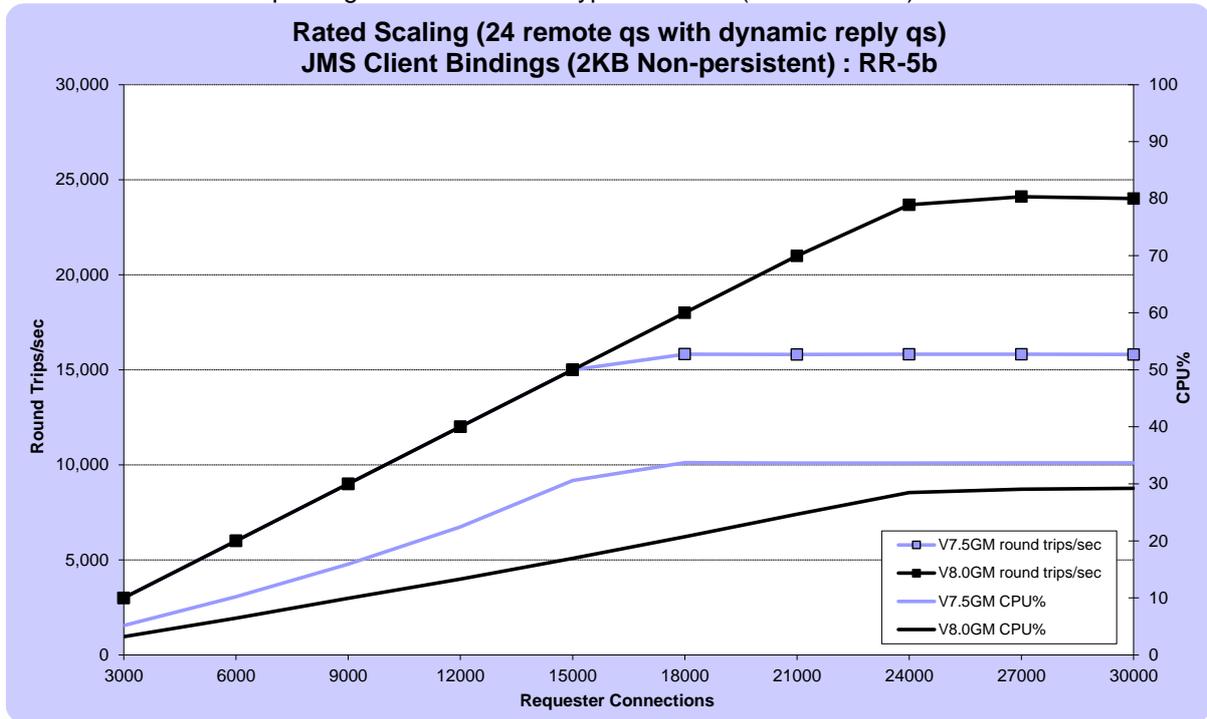


Figure 26 - Performance results for JMS local binding, rated scaling across 24 remote queues with dynamic reply queues (2KB non-persistent)



V8.0 scales significantly better than V7.5 and the CPU consumption is also much less than V7.5



3.9.5 JMS distributed queuing (persistent) with 24 request queues.

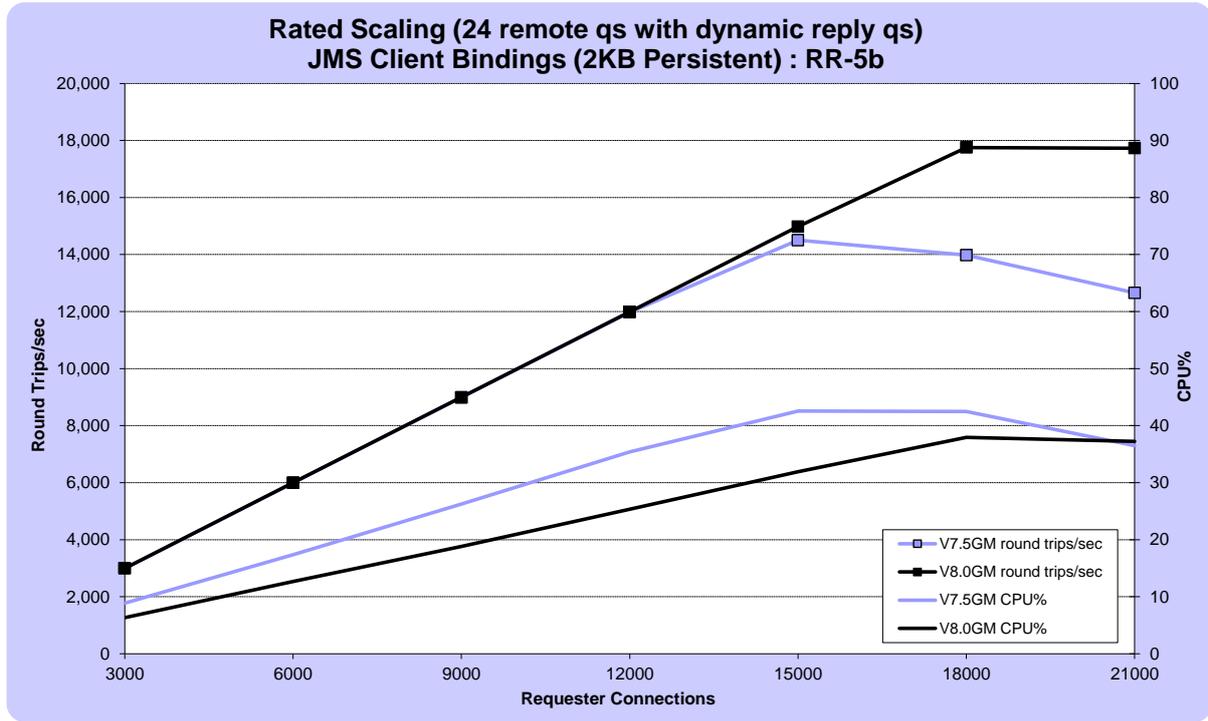


Figure 27 - Performance results for JMS local binding, rated scaling across 24 remote queues with dynamic reply queues (2KB persistent)



This is the persistent messaging variation of the test in the previous section. Once again V8 outperforms V7.5 with better throughput at lower cost.

4 How Am I Constrained?

A common question that is often asked of the MQ Performance team is:
 “What is limiting my messaging rate?”

Understanding what might be restricting your messaging rate can help to improve your messaging performance (by eliminating or reducing the constriction) or provide information to assist in planning future migration/expansion of your messaging infrastructure.

4.1 Disk Contention

For persistent messaging the likely reason that you are limited is due to the performance of the disk subsystem. This could be due to the raw speed (or lack of it) of the disk system, disk configuration, configuration of MQ logging or other system usage of the disk subsystem.

Use this checklist to assist in checking your disk contention:

- Increase the number of producers/consumers to increase concurrent messaging workload
- Use disk tools (iostat/TaskManager) to determine disk utilisation and increase speed/capacity/cache size as required to provide increased disk capacity
- Check other processes are not utilising the same disk systems
- Check log configuration
 (http://www.ibm.com/developerworks/websphere/library/techarticles/0712_dunn/0712_dunn.html)

It can often be useful to run a simple scenario (put/get across single queue) to understand the typical maximum persistent throughput through your disk system before deploying a more complex production scenario.

This test from a locally bound request/responder scenario, illustrates the maximum throughput across a pair of request/reply queues on the server hardware under test.

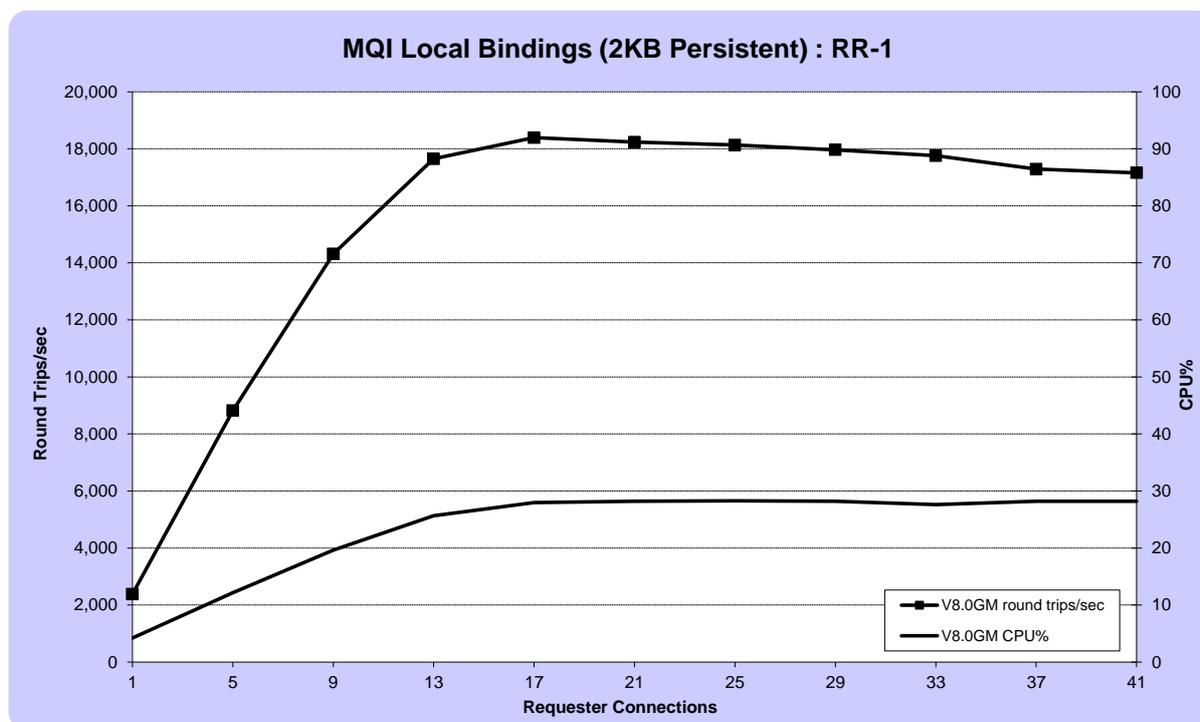


Figure 28 - Performance results for MQI local bindings (2KB persistent)



Here are some suggestions to consider when looking to improve persistent performance:

- Use fast disk sub-systems
- Locate log file on its own filesystem/disk
- Use battery-backed write back cache where possible
- Use circular logs unless linear logs are required for auditing/recovery
- Send/receive messages within syncpoint
- Use separate queues for persistent and non-persistent workloads
- Use separate channels for persistent and non-persistent workloads
- Use separate queues for out-of-band message sizes

4.2 CPU Saturation

For any type of messaging, if you have saturated the CPU, there is often little you can do to increase performance without changing hardware (or allocating more Cores in an LPAR/VM environment). Most operating systems provide a simple way of monitoring CPU usage (vmstat/top/TaskManager). If MQ processes are shown as the highest CPU consumers, it is likely that you have encountered the messaging limit on this piece of hardware.

It is possible that a badly performing application can utilise large amounts of CPU whilst interoperating with MQ (i.e. constantly polling the QM, or simply performing extensive business logic), and you would likely see non-MQ applications high on the list of CPU consumers. These processes should be investigated to determine if their behaviour is expected.

There are also a number of strategies to help reduce CPU when developing your MQ applications:

- Don't poll the QM and use sensible timeouts when receiving messages from MQ
- MQ performs optimally when the queue depth is 0 and there are consumers waiting to consume messages that are delivered to that queue
- Use *FASTPATH* channels/applications where there is no (or low) risk of QM corruption

4.3 Queue Lock Contention

There can be scenarios where you cannot use all of the CPU available and increasing the numbers of producers/consumers has no affect. If the scenario is only using a single or small number of queues, you may find that you are encountering queue lock. This can occur in scenarios where extremely high messaging throughput is distributed across a small set of queues. You can see how this scenario is encountered in the non-persistent version of the request/responder scenario (peaking at just over 120,000 messages/sec across a pair of request/reply queues):

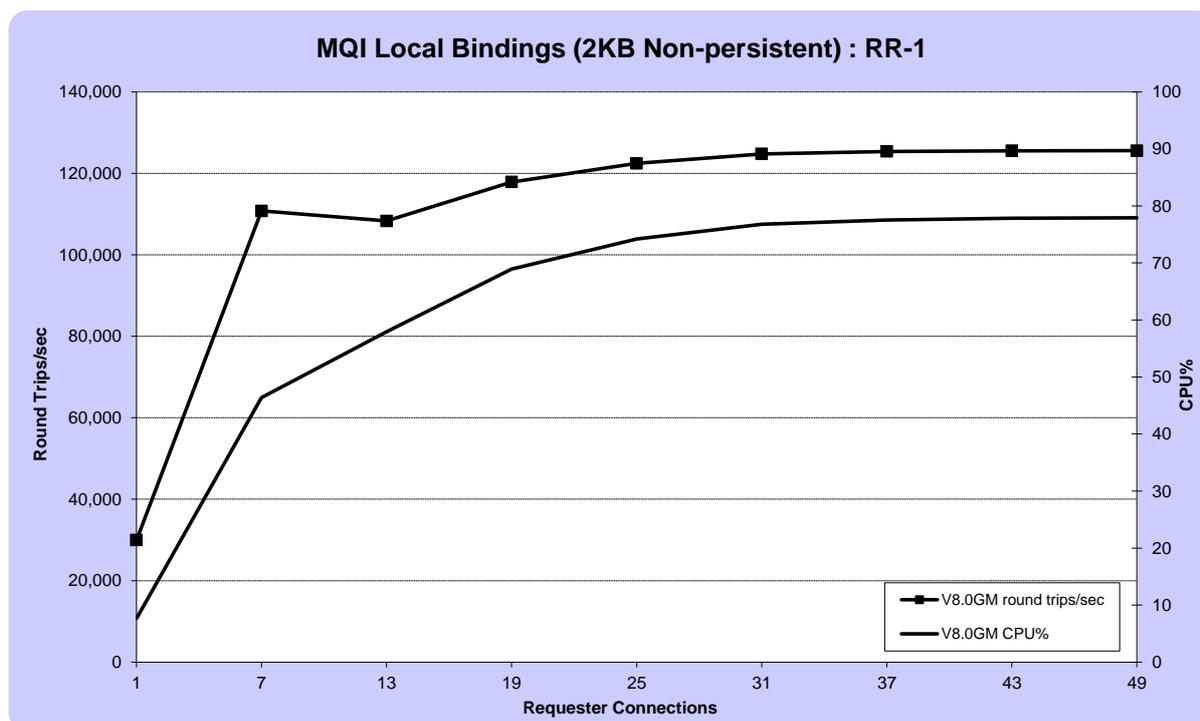


Figure 29 - Performance results for MQI local bindings (2KB non-persistent)



Queue lock contention is more likely to manifest itself on machines with more cores due to the single-threaded nature of the bottleneck.

To confirm that you are encountering queue lock and not an effect of your client applications or QM configuration, the 'MQ Local Queue Manager Debug Utility' (amqldmpa) tool can identify this behaviour:

```
amqldmpa -m <QM Name> -q <Q Name> -s <interval secs> -n <interval count> -f
<output file> -c A -d 3
i.e. amqldmpa -M QMTEST -q INPUTQ -s 10 -n 6 -f C:\ldmpa.out -c A -d 3
```

From the ldmpa output, check the following values during the period of collection:

```
hmtx.RequestCount      102954894
hmtx.WaitCount         22227930
```

The *WaitCount*, as its name suggests, is the number of times that access to the queue was initially denied and a thread had to wait before processing its work. As soon as this value reaches 10% (or higher) of the *RequestCount*, application performance will start to suffer. Distributing this work over a set of queues would alleviate this contention.

4.4 Network Saturation

Network saturation can be diagnosed by the use of monitoring tools such as dstat/TaskManager. If these tools show that you are encroaching on the bandwidth limits of the network, you may need to increase your networking capacity to 1Gb/10Gb/40Gb as appropriate or add additional network cards.

If you do not have access to these tools or direct access to the hardware, you can calculate approximate messaging limits for a requester/responder scenario with local responders given your



networking environment. We can estimate how many 20KB messages we can flow through a scenario per second in a 10Gb environment with a remote producer and local consumer¹:

$$\begin{aligned}
 \text{10 Gigabit network} &= 10 \times 1024 \times 1024 \times 1024 \\
 &= 10,737,418,240 \text{ bits} \\
 / 8 &= 1,342,177,280 \text{ bytes} \\
 * 80\% &= \sim 1,073,741,824 \text{ bytes (approx maximum switch throughput)} \\
 / 20680 &= \sim 51,922 \text{ msg} \quad (\text{each message approx } 20480 + 200 \text{ byte hdr})
 \end{aligned}$$

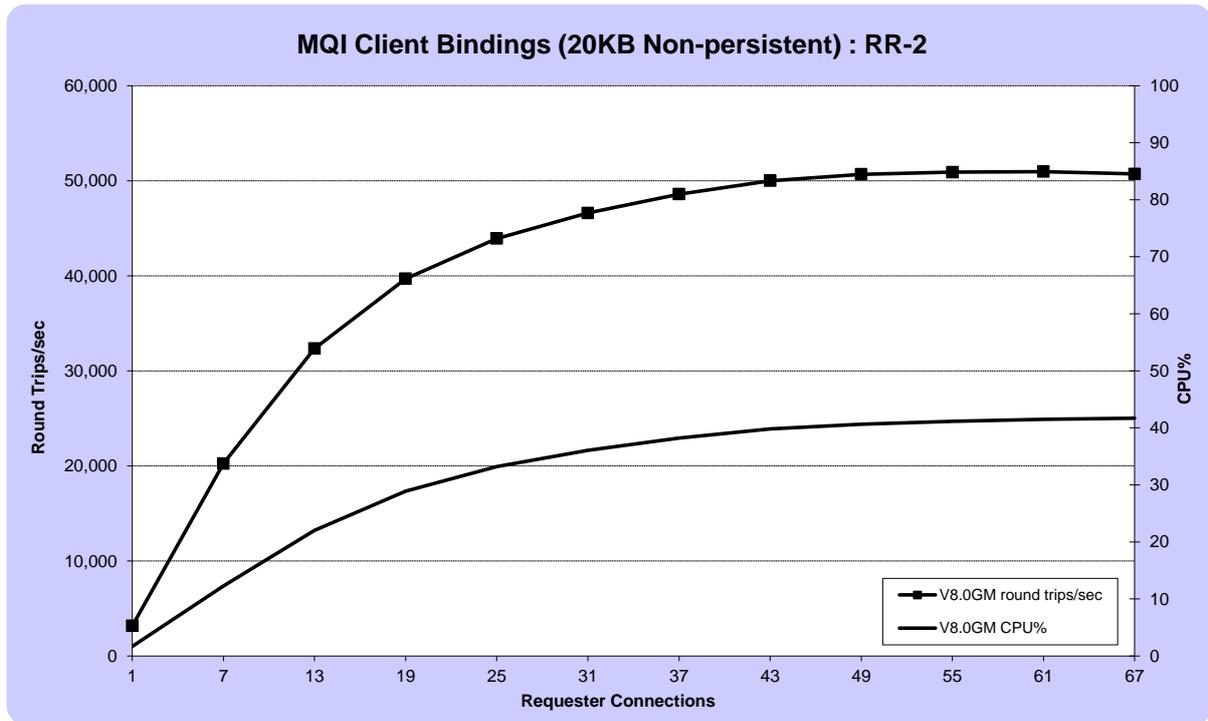


Figure 30 - Performance results for MQI client bindings (20KB non-persistent)



Note that most switches support full duplex mode so the stated switch capacity is for each direction at the network level; if the consumers were located remotely, then the maximum theoretical throughput as calculated above would need to be halved since network traffic has been doubled.

4.5 Channel Saturation

Channel saturation can occur when high throughput messaging takes place between two Queue Managers connected by a single (or small) number of channels. To determine if you are approaching channel saturation, use the technique in section 4.3 to evaluate the wait for the queue lock of the transmission queue.

Some things to try before adding additional channels:

- Increase batch size to reduce the ratio of acknowledgement flows to transmitted messages.
- Use channel attribute *NPMSPEED=FAST* for non-persistent messaging to transfer messages without the use of transactions.

If you are sending a mix of message sizes, you might consider sending very large messages on their own channel to avoid incurring any delays in the transmission of the smaller size messages.

¹ This calculation assumes an operational limit of 80% of the nominal bandwidth of the network adapter. The achievable throughput is limited by additional transport layer overheads and will vary by platform and network adapter.

If you find that you are encountering channel saturation, adding more channels can provide increased performance. Figure 31 below shows the 2KB distributed queuing non-persistent test reported on in section 3.4 when using a single and 10 channel configuration. It can be seen that in this case the use of multiple server channels enabled MQ to process significantly more messages and more fully utilise the server CPU.

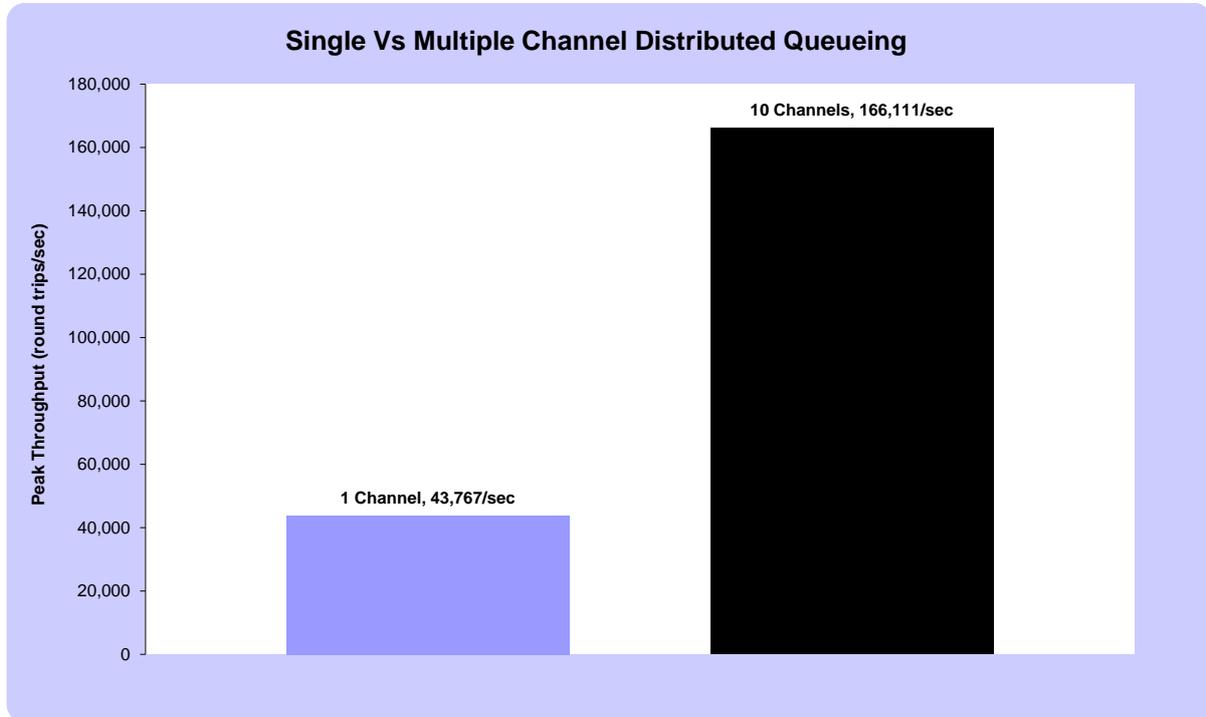
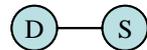


Figure 31 - Single channel vs multi channel distributed queuing



Multiple server channels are not a panacea in distributed queuing however. In our test environment they did not provide a significant benefit for persistent messaging as the i/o subsystem was already the bottleneck in that case. Larger message sizes did not gain as significantly either (another reason to consider the separation of messages across different channels based on message size)

5 Planning for Performance

There are many things that need to be taken into account when planning an MQ infrastructure. Often, performance considerations end up being given a lower priority than those of security, integrity, and even application transparency. While this is often correct (particularly in the cases of security and integrity), it's important to bear in mind that many techniques used to mitigate these concerns, if not implemented carefully, can have a negative impact on the performance of the system. In some cases, we have even seen situations where the performance of a system becomes severely degraded, rendering it unusable for the task it was designed to perform. However, with careful planning in advance, it should be possible to design a system that performs sufficiently, while being secure, scalable, transparent, and highly available.

This section illustrates some common pitfalls and best practices that should be adhered to when designing and testing an MQ infrastructure for performance.

5.1 Persistence, High Availability, & Disaster Recovery

Using persistent messaging will slow a system down. Because of this, implementations using persistent messages and high availability are some of the most common scenarios where performance issues arise. Avoiding these issues involves careful planning in advance, but can be achieved.

The first step that should always be taken is to establish whether and where persistence is actually needed. Some things to consider:

- MQ logging protects the integrity of data in transit, so that it can be recovered in the event of a failure. It is not designed to act as a historical database for audit or any other purposes.
- Since recovery often takes some time to complete in the event of a failure, it's probably not worth logging information that expires (loses its relevance) quickly.
- In a well-designed system, failures should be rare. If the in-flight information could be reconstructed and re-sent with relative ease, then this may be a preferable to logging all messages, the vast majority of which will arrive at their destination without incident. In particular, if client applications already have disaster-recovery logic built-in, then it may be possible to switch off persistence with no loss of integrity, and no additional work required.

If persistence is definitely required, then in addition to the basic recommendations given in section 4.1, there are a number of things to consider when designing an MQ architecture.

Persistence can impact performance in two ways:

1. High latency - how long it takes for a filesystem write operation to return
2. Low bandwidth - the maximum data-rate that can be achieved on the file-system

5.1.1 Latency

The latency of a filesystem will depend on the level of integrity it has, with the most robust set-ups generally suffering the most.

- For local persistent messaging, latency occurs due to the time taken to write logs to the disk, or its battery-backed cache. SAN systems may be able to ensure integrity before the data is synced to disk, but will instead have latency caused by the time taken to communicate with the SAN host.
- For highly-available persistent infrastructures (multi-instance), the requirement for networked file-systems introduces an additional latency in communicating with the filesystem host.
- Finally, in full disaster-recovery situations, file-systems may be synchronously replicated to a second data-centre, often geographically distant from the first. This additional network communication and disk write must be waited for before the write operation completes.

In all cases, there are 2 techniques that can be used to overcome latency:

1. Concurrency – running multiple application threads in parallel.
2. Batching – reducing the number of write operations each application thread performs.

5.1.1.1 Concurrency

All persistent messaging operations should be done inside syncpoint, this delays forcing messages to disk until *MQCMIT* is called by the application. If multiple applications are accessing a queue manager concurrently, this may allow the disk writes of several applications to be grouped together into a single operation, thus reducing the number of writes required.

From a client perspective, having multiple applications also means that one application can be making use of other resources (such as CPU) while another is waiting on a disk write.

5.1.1.2 Batching

As well as doing all persistent messaging inside syncpoint, it may be possible for applications to perform several messaging operations before each call to *MQCMIT*, thus reducing the number of write operations each thread requires.

- Responder type applications may be able to get their request and put their response inside the same syncpoint. This is how the MQI responders used to gather the results in this report are designed to work.
- Batch processing applications may be able to process multiple records before each *MQCMIT*. The time taken to reprocess a batch in the event of a failure should be weighed up against the likelihood of failure.

In more complex infrastructures, involving queue manager to queue manager communication (distributed queueing and clustering), one can also increase the batch size on sender/receiver channels to the same effect, using the *BATCHSZ* configuration parameter.

It's worth noting that if a disk system (particularly a networked filesystem) is unreliable, then batching may have the opposite effect. Increasing the size of write packets may make the write more likely to fail, requiring the operation to be repeated, and thus *increasing* the total number of writes.

5.1.2 Bandwidth

A bottleneck in I/O bandwidth can only be mitigated by increasing hardware capacity. It's important to first establish which aspect of the I/O subsystem might be constrained. It could be the disks themselves, in which case adding more disks in a RAID5 array may help. Where possible, upgrading the constrained component to a higher-capacity alternative should help.

When it's not feasible to upgrade a component further, it's likely that the logs will need to be split, which will require using multiple queue managers (see section 5.4.1). How this split is implemented will depend on which aspect of the I/O system is constrained. Some possibilities:

- The local file-system on the file-system host – put the logs for each QM in a directory on a separate disk or RAID device.
- For networked filesystems, other physical constraints (CPU, RAM, network) on the file-system host – put the logs for each QM in a different remote file-system, hosted on different machines.
- For networked filesystems, network bandwidth on the queue manager host – put each QM on a separate machine.
- Synchronous data replication – any of the above, synchronising each QMs log directory with a different destination.

5.2 Slow Networks

When using MQ over slow or unreliable networks, such as satellite links and the Internet, some additional design considerations may be beneficial in achieving optimal performance. If network bandwidth is contended, and there are spare CPU cycles available on the source and destination

machines, then it may be possible to obtain higher throughput by utilising header compression (*COMPHDR*)² and data compression (*COMPMSG*)³.

Aside from this, it may be possible to adjust the architecture of an MQ infrastructure to make most efficient use of the network. We'll consider a number of scenarios, and some possible techniques to increase throughput.

5.2.1 Many disparate clients communicating infrequently

In this scenario, a large number of clients, all from different places on the network, connect to a queue manager, and send or receive messages at a relatively low rate. As such, there's no constraint on the network path between the queue manager and any one of these clients, but together these clients put a strain on the network interface of the QM host.

If messages are small, it could be that, although the total data flowing to and from the host is within the supposed bandwidth of the network interface, the sheer number of networking operations cannot be processed fast enough. If this is the case, one possible solution is to introduce a number of "satellite" queue managers in different locations. Clients connect to the satellite that is "closest" to it (in terms of network topology), which then forwards communication to the main QM. Following the advice in section 5.2.4 then allows us to optimise the connections between the satellites and the main host, reducing the number of networking operations performed.

If raw data bandwidth is constrained, there's often very little that can be done here from a design perspective; ensuring that host's network interface has enough capacity for the desired message rate is probably the simplest solution (see section 4.4). You should also ensure that the IP layer is tuned optimally. If this is not feasible, then the workload will need to be split across multiple machines. See section 5.4.1 for tips on how this can be achieved.

5.2.2 High message rate to/from a client machine

It may be that certain machines in the infrastructure will have large volumes of communication with one or more QMs elsewhere on the network. This could either be one application messaging at a high rate, or a number of applications all on the same machine. Examples could include application or web servers, databases, or mainframes.

In this situation, it may be advantageous to host an additional QM on such machines, and connect it to target QMs by making it part of a cluster, or adding distributed queuing channels. Applications then connect to this QM using local bindings, and communication can be optimised by tuning the sender/receiver channels to the target QMs (see section 5.2.4 - Optimising QM to QM channel communication). This solution has the added advantage that the MQ messaging operations performed by the applications will complete as soon as the message is queued on the local QM; they are then free to continue working without having to wait for a slow network communication to complete.

5.2.3 Communicating between regions

Often it may be the case that an MQ infrastructure spans multiple regions or geographic locations. The network within each of these regions may be fast and reliable, while communication between regions is less so, possibly travelling over the Internet. It may help for MQ inter-region messaging to be amalgamated by one or more "gateway" QMs in each region. Applications connect to their local gateway, which then handles sending messages to other regions using distributed queuing or clustering. These channels of communication are then open to optimisation by the techniques mentioned in section 5.2.4.

5.2.4 Optimising QM to QM channel communication

For communication between queue managers, there are a number of things that can be done administratively to improve performance:

² http://www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.ref.con.doc/q081880.htm

³ http://www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.ref.con.doc/q081840.htm

- Adjusting the *BATCHSZ*, *BATCHINT*, and *NPMSPEED* parameters of the channel definitions – see www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.ref.con.doc/q081660.htm
- Setting *PipeLineLength=2* in the *CHANNELS* stanza of the QM's *qm.ini* file – see www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.ref.con.doc/q080700.htm
- Using multiple channels – see section 4.5

Please note, however, that the *final* option may result in messages arriving on their destination queues in a different order than they were originally put. If message ordering must be maintained, then while multiple channels may still be used, each subset of messages that must keep their order (e.g. all those from the same client application) must be configured to use the same channel for communication.

5.3 Ultra-High Message Rates

If other hardware factors (such as network and I/O) are not the constraint, but message rates are very high, then eventually the maximum rate at which messaging operations can be performed on a single queue will be reached. Section 3.1 gives an indication of the sort of message rates obtainable through a single queue, but the exact rate achievable in a set-up will depend on several factors, including the pattern in which messages are put/got, the type of message selection, if any, and the size of messages. This limitation occurs because certain parts of most messaging operations require exclusive access to certain data about the queue. We call this “queue lock” (see section 4.3).

Queue lock can be mitigated in certain ways by upgrading hardware capacity and altering the configuration of the queue manager:

- Faster processor cores (or more exclusive access to them by MQ) may often help, but more cores often won't, unless there is other work being done by the machine that can be offloaded onto these additional engines.
- If queue buffers are already occupying all the available physical memory on the machine (see section 6.1.1.1), then adding more RAM may help. However, it would be advisable to reconsider the design of the infrastructure, as MQ performs best with short queues: it is not intended to be used as a message store, except for brief periods in exceptional circumstances.

5.3.1 Maximising the throughput of a single queue

There are several measures that can be taken to ensure that a queue operates as efficiently as possible.

5.3.1.1 Ensure there are “waiting getters”

If an application requesting to get a message off of a queue makes this request before a suitable message arrives, then when one does arrive, it is passed straight to the waiting application without being written to the queue. As well as reducing the total path length required in handling the message, this has the added advantage that the getting application doesn't have to take the queue lock; all required data are updated by the application that puts the message on the queue, so less expensive lock negotiation is required. This only applies to messages outside of syncpoint control.

Unless it is necessary to process all messages from a queue in a single thread (for order integrity reasons), it is therefore advantageous to have a number of application threads receiving messages from a queue. It is far better to have “too many” (applications waiting in *MQGET* and not ever doing anything) than too few (messages waiting on a queue for getters to become available).

5.3.1.2 Be selective about message selection

Sometimes, receiving applications may only be interested in a subset of the messages that might appear on a queue. MQ caters for this with message selection. However, message selection can dramatically increase the amount of time a receiving application spends browsing a queue, holding the queue lock, and thus preventing other applications from accessing it.

If message selection can instead be replaced by putting different categories of message to different queues, such that receiving applications can get messages from a queue in an indifferent manner, then this should be preferred. It is often easier to make this design decision in the first place, rather than trying to retro-fit it to an infrastructure that's become overloaded, especially since the performance impact of having additional queues is minimal.

If message selection must be used, try to rely mainly on MQ's built in *Correlation ID* field. Selection by Correlation ID has been optimised within the queue manager, so that getters spend far less time than usual finding a suitable message. It also allows for the advantages afforded to "waiting getters", as described in section 5.3.1.1.

5.3.2 Adding queues to increase throughput

The current trend for hardware, or cloud-based systems, to get wider (more cores) rather than taller (faster cores) means you cannot rely on hardware upgrades to satisfy future increases in demand on your MQ infrastructure. Once the throughput limit of a queue has been reached, there is little that can be done to increase capacity further. It is therefore important to design systems and applications where high demand is anticipated to be able to use multiple queues if necessary. Splitting messages by their purpose, as described in 5.3.1.2, is probably the first option that should be considered.

If the purpose (or "topic") of messages that an application will produce or consume is not known in advance, or if the distinction between message purposes is complex, then using publish-subscribe messaging, rather than simple queueing, may be advantageous. Then, messages of different topics can be optimised administratively in MQ, and consuming applications can still share the workload of a particular topic by attaching them to the subscriber queue of a single administrative subscription object. Furthermore, it may be possible to reorganise the topic space to handle increased demand without requiring a change to the applications.

Beyond these options, scaling a system across multiple queues transparently (without the need to manually assign queues to applications) is difficult, but can be achieved in certain ways (see section 5.4.2).

5.4 Transparent Scalability

Designing an MQ infrastructure that scales to accommodate future (or present) demand is something that needs careful planning from the beginning. The hardware factors that normally constrain MQ infrastructures (disks and network) are often infeasible to upgrade further (make faster), and adding more of them to split workload presents administrative challenges. Traditional methods of load-balancing networked communication by using intermediaries is often not feasible when the assured delivery guarantees that MQ makes are required: the contention just shifts from the original servers to the load-balancer.

When increasing demand requires additional instances of the objects involved in an infrastructure (machines, queue managers, queues), as discussed in the previous sections of this chapter, it is often desirable to do this in the most transparent way possible. Briefly, we'll discuss some of the available techniques to do this, requiring the smallest possible change to the peripheral parts of a network, and instead focusing alterations on the core, so that changes can be administered in as few places as possible.

5.4.1 Scaling across machines & queue managers

From the perspective of an MQ client, there is little distinction between a queue manager and the machine it's hosted on. Hence, this section will cover options for scaling across either as if they were the same thing.

5.4.1.1 Clustering

One technique that's often attempted to achieve transparent scalability is to put the relevant queues and queue managers into an MQ cluster. The hope is that then the demands on networking or I/O will be spread across all the machines and log file-systems involved in the cluster, and allow additional capacity to be added simply by adding more queue managers to the cluster.

This can certainly help with relieving queue lock (see 5.4.2.1 below), or CPU contention caused by heavy processing of messages by locally-bound receiving applications. It can also relieve the load placed on systems by MQ itself, but only if application connections are spread across the queue managers in the cluster. Having all input applications connect to a single “gateway” queue manager, as is often desired to simplify the configuration of applications across an organisation, simply shifts contention to this gateway. The gateway queue manager will have the same requirements on its network and I/O interfaces as would a single server handling all messages itself.

It’s also worth noting that using clusters increases the latency of messages, and often results in a higher overall use of hardware resources than pushing the same number of messages through a single QM (if possible). Network bandwidth is consumed for each QM that a message passes through, and persistent messages will need to be written to the logs of each QM they touch as well.

5.4.1.2 CCDTs and application logic

An alternative (or complement) to using a cluster is to have applications connect using a *Client Channel Definition Table*⁴. This allows applications to select a QM to connect to from a defined group at connection time. Thus, application connections will be spread across all queue managers in that group, which could be a cluster, or could just be a set of queue managers, each having the required queues defined individually.

CCDTs are not suitable when XA distributed transactions will be involved, and may not always result in the best balancing of workloads. For this, a small amount of application logic, which can be redistributed across an organisation as part of a library, is probably preferable. The 3-part IBM developerWorks article found at the following URLs discusses these options in detail, and provides some example code:

1. www.ibm.com/developerworks/websphere/library/techarticles/1303_broadhurst/1303_broadhurst.html
2. www.ibm.com/developerworks/websphere/library/techarticles/1308_broadhurst/1308_broadhurst.html
3. www.ibm.com/developerworks/websphere/library/techarticles/1403_broadhurst/1403_broadhurst.html

5.4.2 Scaling across queues

If splitting a workload across multiple queues is sufficient to relieve contention (normally because queue lock is the contending factor), then in addition to the techniques mentioned in section 5.4.1 (which will work for queues by defining an instance of the queue on multiple queue managers), a few additional options are available.

5.4.2.1 Clustering

Clustering can more easily be used to relieve queue contention than other constraints (see section 5.4.1.1). Although the use of multiple queue managers is still required (possibly all on the same machine), a single gateway queue manager can be used to simplify application connection, provided this gateway is at MQ v7.5 or later: using multiple cluster transmission queues⁵ on the gateway (i.e. one per destination QM) allows workload to be split across multiple queues as soon as it enters the cluster.

5.4.2.2 API exits

An alternative to using a cluster to spread workload across queues is to write a small custom MQ API Exit program. This is a piece of code attached to the queue manager that can intercept MQI calls. For example, to could listen for calls to *MQOPEN* a particular queue. When such a call is intercepted, the queue being requested is silently swapped for an alternative from a pre-defined pool. As such, applications can all make identical requests to open the queue, while being transparently balanced across a number of them.

Although this requires more work to implement than setting up a cluster, it has the advantage that everything is done with a single queue manager, making maintenance simpler. It also requires less hardware resource to run, will have slightly lower latency, and since the logic of how the load

⁴ www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.dev.doc/q027490_.htm

⁵ www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.pro.doc/q118645_.htm

balancing occurs is up to the code in the API exit, it can be made as simple or as complex as necessary.

The details of writing API exits are beyond the scope of this document, but as a place to start, one can study the API exit sample program, *amqsaxe*, and read the associated documentation at [www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.dev.doc/q027920 .htm](http://www-01.ibm.com/support/knowledgecenter/SSFKSJ_8.0.0/com.ibm.mq.dev.doc/q027920.htm).

5.5 Testing for Performance

It is important to thoroughly test any new MQ infrastructure prototype for performance before going into production. This section mentions a few best practices and common pitfalls to avoid to make your testing more effective.

5.5.1 It's a dress rehearsal, not an audition

Performance testing should be done in an environment as close to that of production as possible. This means the same – or very similar – hardware, operating systems, storage systems, network topology, bandwidth, congestion, message size, message rate, protocol, MQ configuration, application messaging profile, etc.

It is easy to underestimate the impact small changes to configuration can have, or think that one can “compensate” for the lack of a feature in the test environment by requiring better-than-actual performance without it. This is usually not the case, and if an infrastructure can perform as well as required in a test environment with a constraining feature missing, then any further improvements may have no impact on the complete environment. Obviously, if a cut-down environment *doesn't* perform as required, then adding a constraining factor is unlikely to improve performance, so there may be some merit in investigating and attempting to widen the bottlenecks that already exist.

However, in some cases, apparently simpler configurations actually degrade performance, leading to unnecessary investigations of an issue that would never have occurred in production. Two of the most common examples of this are:

1. Persistent messaging outside of syncpoint (when it would be inside in production) (see section 5.1.1.1).
2. Putting a batch of messages all at once, then getting them, or relying on message expiry to remove them from the queue (in production, putting and getting would happen concurrently). This prevents optimisations for “waiting getters” from being used (see section 5.3.1.1) and could also lead to slower operation due to having to access disks instead of memory (see section 6.1.1.1).

5.5.2 Know your tools

When testing a messaging infrastructure, production applications are often swapped out for performance-oriented simulators or “harnesses”. This is fine, but it's important to know how the application you've decided to use works. Where possible, use one of the following:

- An IBM-produced tool, such as the freely available JMSPerfHarness⁶
- An application developed in-house, designed to behave similarly to the applications it's simulating.
- An open-source project

With all of these, the key factor is the ability to know exactly how the application is behaving internally. This is important, as we have discovered that many such applications use MQ in a sub-optimal way. For instance, some tools cannot make proper use of syncpoint transactions, which seriously impacts the performance of persistent messaging scenarios.

6

www.ibm.com/developerworks/community/groups/service/html/communityview?communityUId=1c020fe8-4efb-4d70-afb7-0f561120c2aa

5.5.3 When it doesn't work, simplify

Despite what was said in section 5.5.1, when an infrastructure doesn't perform as required, your first question should not be:

“What can I change to try to make things better?”

Instead, you should ask:

“How can I determine what's constraining my system?”

Often, the easiest way to do this is to start *removing* features of your system to see which is acting as a bottleneck to performance. Even if those features are necessary for functional reasons, knowing where the problem lies will help to tackle the first question more efficiently.

For instance:

- If using persistent messaging, try non-persistent. If things don't improve, then the issue is not with the logging or disks, so there's no need to investigate those.
- If using networked file-systems, switch to local disks. If an improvement is seen, then the problem likely lies in the latency or bandwidth of the networked file-system.
- If using a cluster, drop back to a single QM. If performance divided by the number of QMs improves dramatically, then something about the cluster configuration is holding the system back.

Keeping detailed, comparable records of how your system performs in various configurations can often help lead to a swifter resolution in the event that IBM assistance is needed.

5.5.4 Test your environment

When trying to establish the cause of a problem, it often helps to know how the performance of MQ compares to the theoretical limits of the system on which it's running. Stress and latency testing tools exist for CPU, I/O, and networking on most operating systems. These should be used to test all aspects of an infrastructure before you waste time trying to “tune” your way out of a hardware limitation.

5.6 Memory Requirements.

The memory requirements of MQ are heavily dependent on the nature of the workload and the objects defined to the queue manager. Modern server class machines should not encounter memory constraints but injudicious setting of very large queue buffers across multiple heavily used queues might cause a problem, for instance.

Monitor your system to ensure that you are not paging. Paging *out* of memory is not necessarily a problem, but does indicate that the system is nearing memory exhaustion. Once the ‘working set’ of the system (i.e. the pages of allocated memory that are frequently accessed) exceeds the physical memory available then performance will be significantly impacted.

Factors affecting memory use of MQ:

- Number of objects in use on the queue manager
- Number of client channels attached to the queue manager
- Queue depth, up to the buffer size for each queue

6 Tuning Recommendations

6.1 Tuning The Queue Manager

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

If several changes are to be made, it would be prudent to test the impact of each change individually. This allows one to establish exactly which changes are providing a benefit (some may even degrade performance), which may also in turn provide useful information about how the system is constrained, and so how it may be improved further.

Note: The 'TuningParameters' stanza in the queue manager's *qm.ini*, or the MQ installation's *mq.ini*, is not a documented external interface and may be 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. Persistent messages are synchronously written to the log and are also periodically flushed to the Queue file.

To avoid potential queue and log I/O contention due to the queue manager simultaneously updating a queue file and log extent on the same disk, it can help if queue files and logs are located on separate and dedicated physical devices. Storage requirements may be fulfilled by the use of a Storage Area Network (SAN), but multiple high volume queue managers can require different logical volumes, and potentially individual interface pathways 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 path length 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 which will depend on the write speed of the i/o device (including any caching at the hardware level)..

6.1.1.1 Non-persistent and persistent queue buffer

The default non-persistent queue buffer size is 128K per queue and the default persistent is 256K per queue for MQ on a 64-bit operating system. These can be increased up to 100MB using the *DefaultQBufferSize* and *DefaultPQBufferSize* parameters in the *TuningParameters* stanza of the *qm.ini* file. (For more details see SupportPac MP01: MQSeries – Tuning Queue Limits). Increasing the queue buffer provides the capability to absorb spikes in message throughput at the expense of real memory. Once these queue buffers are full, the additional message data is written to the file

system that will eventually find its way to the disk (operating system buffers will be utilised for queue files).

Defining large queue buffers is *not* a solution to avoiding disk I/O when the input message rate to MQ exceeds the output rate, for whatever reason, for sustained periods. The buffer will fill and spilling will occur.

Defining queues using large non-persistent or persistent queue buffers can degrade performance if the system is short of real memory either because a large number of queues have already been defined with large buffers, or for other reasons - e.g. large number of channels defined.

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

Queues can be defined with different values of *DefaultQBufferSize* and *DefaultPQBufferSize*. The value is taken from the *TuningParameters* stanza in use when the queue manager was last started.

6.1.1.2 Logging

6.1.1.2.1 Log type

The log component is often the bottleneck when processing persistent messages. Sufficient information is stored in the log to restart the queue manager after failure.

Both circular and linear logging are sufficient to recover from application, software, or power failure whilst linear logging can also recover from media failure resulting in the loss of a queue file (assuming some historical copy of the queue file exists along with the linear logs to perform a forward recovery from that point).

Circular logging is preferable for a performance perspective as no time is required to allocate and format new log extents or to delete or archive them. The choice of logging type will not be based on performance alone however. A useful article comparing logging types is available on developerWorks:

http://www.ibm.com/developerworks/websphere/techjournal/0904_mismes/0904_mismes.html

6.1.1.2.2 Log buffer

Log records are written to the log buffer at each *MQPUT* and *MQGET* of messages outside of syncpoint, and each *MQCMIT*. This information is synced onto the log disk. Periodically the checkpoint process will decide how many of these log-file 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 logging. 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* are probably suitable for most installations. The default size of the log buffer is 512 pages with a maximum size of 4096 pages.

To optimise the throughput of large persistent messages (> 1MB) *LogBufferPages* could be increased to improve the likelihood of messages only needing one I/O to get to the disk. Environments that

process under 100 small (< 10KB) persistent messages per second can reduce the memory footprint by using smaller values like 32 pages without impacting throughput.

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.

6.1.1.2.3 **Log files**

LogFilePages (or *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 a long running transaction can prevent checkpointing from efficiently freeing the disk extent for reuse.

Larger extents reduce the frequency of log switching (permitting a greater amount of log data to be written into one extent) and allow more time to prepare new linear logs or recycle old circular logs (especially important for long running units of work). As an initial target, one log extent should hold at least 10 seconds of log data streaming.

The number of *LogPrimaryFiles* (or. *crtmqm -lp <LogPrimaryFiles>*) can be configured to a large number (the maximum number of Primary plus Secondary extents is 255 for Windows and 511 for UNIX). For circular logging you should configure enough primary logs to cope with expected peak load as secondary logs are formatted each time they are used, so incur a performance penalty,

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 transactions complete, they enable the next checkpoint process to free up extents that now become available for archiving with linear logging or re-use with circular logging.

Some installation will use linear logging and *not* archive the redundant logs because archiving impacts the run time performance of logging. Instead, 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. This approach prevents the continuous build-up of log extents (assuming the old ones are deleted).

6.1.1.2.4 **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% but in practice, we see very few customer deployments that have gained significantly from changing this parameter.

Measurements in this report used *LogWriteIntegrity=TripleWrite*

6.1.2 Channels: process or thread, standard or fastpath?

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

Fastpath channels 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.

It is not recommended to use fastpath channels when channel exits are being used as any problem with the exit code has the potential to bring the queue manager down.

Note: 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?

There are issues associated with writing and using fastpath applications—described in the 'MQSeries Application Programming Guide'. Although it is generally 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).

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.

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 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 whilst 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 performs in the same way as a 'short session' program. The 'runmqslr' program 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 so is more suitable to triggered applications.

When programming a triggered application it may be worth exposing a disconnect interval as an input parameter to the application program. This can provide the flexibility to make tuning adjustments in a production environment to establish the best balance between reducing connection costs or freeing up queue manager and operating system resources.

6.3 Virtual Memory, Real Memory, & Paging

6.3.1 BufferLength

The *amqrmppa* process contains a thread per connected client. The *BufferLength* parameter of the *MQGET* on the client application is also used to allocate a long-term piece of storage of this size in the *amqrmppa* process, in which the message is held before being retrieved by the client. If the size of the arriving messages cannot be predicted then the application should provide a buffer than can deal with 90% of the messages and re-drive 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*.

For messages encrypted with AMS it is advisable to set a buffer *larger* than the plaintext version of the largest expected message as the encrypted form will be longer.

6.3.2 MQIBindType

MQIBindType=FASTPATH will cause the channel to run 'trusted' mode. Trusted applications do not use a thread in the agent (*amqzlaa*) 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 IPC memory (size = *BufferSize* with a maximum size of 1MB) that is dynamically obtained and only held for the lifetime of the *MQGET*. Standard channels each require an additional 80KB of memory. As the message rate increases, there will be more IPC memory used in parallel.



6.3.3 Paging

The memory available on a machine needs to handle the peaks in throughput. It is important to prevent the queue depths increasing if possible, especially if large queue buffers have been set as these will occupy memory and can cause paging in the worst circumstances.

Queue buffers will grow up to the maximum specified but do not shrink back until the queue manager is restarted.

Appendix A: Workloads

All of the workloads in this report use the JMS Performance Harness (JMS) or C Performance Harness (MQI) to put and/or get messages from queues. These drivers are described briefly in this appendix.

Both of these performance tools are run inside an in-house performance management framework (Perf) which sets up the machines required in a test including any necessary tuning of the queue managers etc.

A.1 JMSPerfHarness

The JMS Performance Harness (JMSPerfHarness) is an IBM internal version of the tool available publicly via Developer Works (see Appendix B:).

“Performance Harness for Java Message Service is a flexible and modular Java package for performance testing of JMS scenarios and providers. It provides a complete set of JMS functionality, as well as many other features such as throttled operation (a fixed rate and/or number of messages), multiple destinations, live performance reporting, JNDI, and multiple vendor plug-ins. It is one of the many tools used by performance teams for IBM MQ and IBM Integration Bus in order to conduct tests ranging from a single client to more than 10,000 clients.”

Multiple clients are simulated in a single process and the number of clients can be further scaled up by adding further processes and deploying across multiple driver machines within a suitable performance management framework.

A.2 C Performance Harness (CPH)

The C Performance Harness (CPH) is an IBM internal only, multi-threaded MQI messaging client supporting most of the options available to MQ. Functionally, it is similar to JMSPerfHarness on which it has been based so the configuration files are almost identical to the JMSPerfHarness ones (see the public site for JMSPerfHarness for details).

CPH can be scaled in exactly the same way as JMSPerfHarness using a performance management framework to deploy multiple clients and utilise multiple machines as necessary.



Appendix B: Further Resources

Performance Harness for Java Message Service:

<https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuiid=1c020fe8-4efb-4d70-afb7-0f561120c2aa>

MQ Documentation Library:

<http://www-01.ibm.com/software/integration/wmq/library/>

Configuring and tuning MQ for performance on Windows and UNIX (Developer Works article):

http://www.ibm.com/developerworks/websphere/library/techarticles/0712_dunn/0712_dunn.html

Appendix C: Test Setup

C.1 Hardware/Software

All of the testing in this document (apart from when testing results are shown from a different platform and are clearly identified as such) was performed on the following hardware and software configuration:

C.1.1 Hardware

Server 1:	IBM System x3550 M4 24 x Intel® Xeon® E5-2697 v2 CPU cores @ 2.70GHz 112GB RAM
Client 1:	IBM System x3550 M4 24 x Intel® Xeon® E5-2697 v2 CPU cores @ 2.70GHz 96GB RAM
Client 2:	IBM System x3850 40 x Intel® Xeon® E7-8850 CPU cores @ 2.40GHz 64GB RAM Dedicated links to SAN Storage DS8800 for log and queue files.

10Gb Ethernet link to isolated performance LAN

C.1.2 Software

Red Hat Enterprise Linux, version 6 (x86_64)
JMSPerfHarness test driver (see section A.1)
CPH MQI test driver (see section Appendix A)
IBM MQ V7.5 & IBM MQ V8.0

C.2 Tuning Parameters Set for Measurements in This Report

The tuning detailed below was set specifically for the tests being run for this performance report but in general follow the best practises outlines in section 6. Some scaling tests in particular require larger buffers/limits to be set.

C.2.1 Operating System

The following Linux parameters were set for measurements in this report

```
/etc/sysctl.conf  
fs.file-max = 1048576  
net.ipv4.ip_local_port_range = 1024 65535  
vm.max_map_count = 1966080  
kernel.pid_max = 655360  
kernel.sem = 1000 1024000 500 8192  
kernel.msgmnb = 131072  
kernel.msgmax = 131072  
kernel.msgmni = 2048  
kernel.shmni = 8192  
kernel.shmall = 536870912  
kernel.shmmax = 137438953472  
kernel.sched_latency_ns = 2000000  
kernel.sched_min_granularity_ns = 1000000  
kernel.sched_wakeup_granularity_ns = 400000  
net.core.somaxconn = 16384
```

```
/etc/security/limits.d/mqm.conf  
@mqm soft nofile 1048576  
@mqm hard nofile 1048576  
@mqm soft nproc 1048576  
@mqm hard nproc 1048576
```

C.2.2 IBM MQ

The following parameters are added or modified in the qm.ini files for the tests run in section 3 of this report:

Channels:

```
MQIBindType=FASTPATH  
MaxActiveChannels=5000  
MaxChannels=5000
```

Log:

```
LogBufferPages=512  
LogFilePages=16384  
LogPrimaryFiles=16  
LogWriteIntegrity=TripleWrite
```

TuningParameters:

```
DefaultPQBufferSize=104857600  
DefaultQBufferSize=104857600
```



For the tests in section 0, the following parameters were used:

Channels:

MQIBindType=FASTPATH
MaxActiveChannels=65536
MaxChannels=65536

Log:

LogBufferPages=512
LogFilePages=16384
LogPrimaryFiles=16
LogWriteIntegrity=TripleWrite

TuningParameters:

DefaultPQBufferSize=52428800
DefaultQBufferSize=52428800

QMErrorLog:

ErrorLogSize=262144

TCP:

ListenerBacklog=16384