

WebSphere MQ for Solaris V5.3 - Performance Evaluations

Version 1.1

5th August 2002

Ben Daly
Alexander Russell. M.Eng. IBM Certified

WebSphere MQ Performance
IBM UK Laboratories
Hursley Park
Winchester
Hampshire
SO21 2JN

Property of IBM

Take Note!

Before using this report be sure to read the general information under “Notices”.

First Edition, July 2002

Second Edition, August 2002

This edition applies to V1.1 of WebSphere MQ for Solaris V5.3 – Performance Evaluations and to all subsequent releases and modifications until otherwise indicated in new editions.

(C) Copyright International Business Machines Corporation 2002. 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

This report is intended to help the reader understand the performance characteristics of WebSphere MQ for Solaris V5.3. Information is not intended as the specification of any programming interfaces that are provided by WebSphere MQ.

References in this report to IBM products or programs do not imply that IBM intends to make these available in all countries in which it operates.

Information contained in this report has **not** been submitted to any formal IBM test and is distributed “**as-is**”. The use of this information and the implementation of any of the techniques is the responsibility of the customer. Much depends on the ability of the customer to evaluate the data and project the results to their operational environment.

The performance data contained in this report was measured in a controlled environment and the results obtained in other environments may vary significantly.

Trademarks and service marks

The following terms used in this publication are trademarks of the IBM Corporation in the United States or other countries or both:

IBM

MQSeries

WebSphere MQ

SupportPac

FFST

AIX

The following term, used in this publication is a trademark of the Sun Microsystems Corporation:

Solaris

UNIX is a registered trademark and licensed exclusively through X/Open Company Limited.

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

Preface

Target audience

This SupportPac is designed for people who:

- Will be designing and implementing environments using WebSphere MQ for Solaris
- Want to understand the performance limits of WebSphere MQ for Solaris V5.3
- Want to understand what actions may be taken to tune WebSphere MQ for Solaris

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

The Contents of this SupportPac

This SupportPac includes:

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

Feedback on this SupportPac

We welcome constructive feedback on this report. Does it provide the sort of information you want? Do you feel something important is missing? Is there too much technical detail, or not enough? Could the material be presented in a manner more useful to you? Please direct any comments of this nature to **WMQPG@uk.ibm.com**

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

Acknowledgements

The author is very grateful to Alexander Russell and Richard Eures for their help in producing this report.

Introduction

The three scenarios in this report used to generate the performance data are classified into: the local queue manager scenario, the client channel scenario, and the distributed queuing scenario. The performance improvements in WebSphere MQ V5.3 can be divided into two areas:

- queue manager enhancements, and
- channel capacity enhancements.

The enhancements to the queue manager are apparent through many of the measurements in this report where WebSphere MQ V5.3 is compared to Version 5.2. Channel capacity enhancements are covered briefly in the *release highlights* section and in more detail towards the end of the report.

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

A Solaris (model Z801) 4-way UltraSparc-II 440MHz with 4GB of RAM was used as the device under test for all the measurements in this report.

How this document is arranged

Release highlights

Pages: 1-2

Section one outlines the performance *improvements* achieved in WebSphere MQ V5.3 compared to Version 5.2. The highlights are a subset of the results shown in the performance *headlines* section.

Performance headlines

Pages: 5-16

Section two of the document contain the performance *headlines* for each of the three scenarios, with MQI applications connected to:

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

The headline tests show:

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

Large messages

Pages: 17-27

Section three of the document contains performance measurements for *large messages*. This includes *MQI response times* of 50byte to 2MB messages, and 20K and 200K messages using the same scenarios as for the performance *headlines*.

Trusted server application

Pages: 28

Section four contains performance measurements for a *trusted* server application, using the three scenarios as for the performance *headlines*.

Short sessions

Pages: 30-31

Section five contains performance measurements for *short sessions*. That is, an MQI application connecting to the queue manager, processing a few messages between connecting to and disconnecting from the queue manager.

Capacity measurements

Pages: 32-36

Section six of this document shows:

- the total number of MQI-client channels that were connected into a single queue manager, with a server application processing 1 nonpersistent round trip per MQI-client per *minute*.
- the total number of server channel pairs that were connected between two queue managers on separate server machines, with a server application processing 1 nonpersistent round trip per server channel pair per *minute*.

Tuning recommendations

Pages: 37-

In previous SupportPacs tuning recommendations have been in a separate section. In this document queue manger parameters are mentioned with the measurements they are appropriate to, with detailed discussion in '**Tuning recommendations**'—**Pages 37-**

Measurement environment

Pages: 41-42

A summary of the way in which the workload is used in each test scenario is given in the performance headlines section. For a more detailed description of the workload, hardware and software specifications, refer to the '**Measurement environment**'—**Page 41**.

Glossary

Pages: 43

A short glossary of the terms used in the tables throughout this document can be found in the '**Glossary**'—**Page 43**.

CONTENTS

| | | |
|-------|--|----|
| 1 | Release highlights..... | 1 |
| 1.1 | Improvements to nonpersistent and persistent messaging | 1 |
| 1.2 | Peak message throughput – Local queue manager | 1 |
| 1.3 | Peak message throughput – client channels | 2 |
| 1.4 | Peak message throughput – distributed queuing | 2 |
| 1.5 | Improvements to channel capacity limits | 3 |
| 1.6 | Capacity limits – client channels..... | 3 |
| 1.7 | Capacity limits – distributed queuing (server channels) | 4 |
| 2 | Performance headlines | 5 |
| 2.1 | Local queue manager test scenario..... | 5 |
| 2.1.1 | Nonpersistent messages – local queue manager | 6 |
| 2.1.2 | Persistent messages – local queue manager | 7 |
| 2.2 | Client channels test scenario..... | 8 |
| 2.2.1 | Nonpersistent messages – client channels..... | 9 |
| 2.2.2 | Persistent messages – client channels..... | 10 |
| 2.2.3 | 'runmqtsr' vs. inetd 'amqcrsta' listener – client channels | 11 |
| 2.3 | Distributed queuing test scenario | 13 |
| 2.3.1 | Nonpersistent messages – server channels | 14 |
| 2.3.2 | Persistent messages – server channels | 15 |
| 2.3.3 | 'runmqtsr' vs. inetd 'amqcrsta' listener – server channels..... | 16 |
| 3 | Large messages | 17 |
| 3.1 | MQI response times: 50bytes to 2MB – local queue manager | 17 |
| 3.2 | Large messages: 20K and 200K – local queue manager | 19 |
| 3.3 | Large messages: 20K and 200K – client channels..... | 22 |
| 3.4 | Large messages: 20K and 200K – distributed queuing | 25 |
| 4 | Trusted server application..... | 28 |
| 5 | Short Sessions | 30 |
| 6 | Performance and capacity limits | 32 |
| 6.1 | Client channels scenario – capacity measurements..... | 32 |
| 6.2 | Distributed queuing scenario – capacity measurements | 35 |
| 7 | Tuning recommendations..... | 37 |
| 7.1 | Tuning the queue manager..... | 37 |
| 7.1.1 | Queue disk, Log disk, and message persistence | 37 |
| 7.1.2 | Log buffer size, Log file size, and number of log extents..... | 38 |
| 7.1.3 | Channels: process or <i>thread</i> , standard or <i>fastpath</i> ? | 38 |
| 7.2 | Application design and configuration | 39 |
| 7.2.1 | <i>Standard</i> or <i>fastpath</i> ?..... | 39 |
| 7.2.2 | Parallelism, batching, and triggering | 39 |
| 7.3 | Tuning the Operating System (Solaris 5.8)..... | 40 |
| 8 | Measurement environment | 41 |
| 8.1 | Workload description | 41 |
| 8.1.1 | MQI response time tool | 41 |
| 8.1.2 | Test scenarios workload | 41 |
| 8.2 | Hardware | 42 |
| 8.3 | Software..... | 42 |
| 9 | Glossary..... | 43 |

TABLES

| | |
|--|----|
| Table 1 – Performance headline, nonpersistent messages, local queue manager..... | 6 |
| Table 2 – Performance headline, persistent messages, local queue manager..... | 7 |
| Table 3 – Performance headline, nonpersistent messages, client channels..... | 9 |
| Table 4 – Performance headline, persistent messages, client channels..... | 10 |
| Table 5 – 1 round trip per driving application per second, client channels..... | 12 |
| Table 6 – Nonthreaded vs. threaded listener, swap requirement, client channels.... | 12 |
| Table 7 – Performance headline, nonpersistent messages, server channels..... | 14 |
| Table 8 – Performance headline, persistent messages, server channels..... | 15 |
| Table 9 – 1 round trip per driving application per second, server channels..... | 16 |
| Table 10 – 2K, 20K and 200K messages, local queue manager..... | 19 |
| Table 11 – 2K, 20K and 200K messages, client channels..... | 22 |
| Table 12 – 2K, 20K and 200K messages, server channels..... | 25 |
| Table 13 – Trusted server application, local queue manager..... | 28 |
| Table 14 – Trusted server application, client channels and server channels..... | 29 |
| Table 15 – Short sessions, nonpersistent messages, client channels..... | 31 |
| Table 16 – Capacity measurements, client channels..... | 32 |
| Table 17 – Client capacity, swap requirement..... | 33 |
| Table 18 – Capacity measurements, server channels..... | 35 |
| Table 19 – Distributed queuing capacity, swap requirement..... | 36 |
| Table 20 – Sizing the kernel tuning parameters for System V IPC..... | 40 |

FIGURES

| | |
|--|----|
| Figure 1 – Peak message throughput, local queue manager..... | 1 |
| Figure 2 – Peak message throughput, client channels..... | 2 |
| Figure 3 – Peak message throughput, distributed queuing..... | 2 |
| Figure 4 – Maximum number of client connections..... | 3 |
| Figure 5 – Maximum number of server channels..... | 4 |
| Figure 6 – Connections into a local queue manager..... | 5 |
| Figure 7 – Performance headline, nonpersistent messages, local queue manager.... | 6 |
| Figure 8 – Performance headline, persistent messages, local queue manager..... | 7 |
| Figure 9 – MQI-client channels into a remote queue manager..... | 8 |
| Figure 10 – Performance headline, nonpersistent messages, client channels..... | 9 |
| Figure 11 – Performance headline, persistent messages, client channels..... | 10 |
| Figure 12 – ‘runmqslr’ vs. inetd ‘amqcrsta’ listener, client channels..... | 11 |
| Figure 13 – ‘runmqslr’ vs. inetd ‘amqcrsta’ listener, swap requirement, client channels | 12 |
| Figure 14 – Server channels between two queue managers..... | 13 |
| Figure 15 – Performance headline, nonpersistent messages, server channels..... | 14 |
| Figure 16 – Performance headline, persistent messages, server channels..... | 15 |
| Figure 17 – ‘runmqslr’ vs. inetd ‘amqcrsta’ listener, server channels..... | 16 |
| Figure 18 – The effect of message size on MQI response time (50byte - 32K)..... | 17 |
| Figure 19 – The effect of message size on MQI response time (32K - 2MB)..... | 18 |
| Figure 20 – 2K and 20K nonpersistent messages, local queue manager..... | 20 |
| Figure 21 – 2K and 20K persistent messages, local queue manager..... | 20 |
| Figure 22 – 200K nonpersistent and persistent messages, local queue manager.... | 21 |
| Figure 23 – 2K and 20K nonpersistent messages, client channels..... | 23 |
| Figure 24 – 2K and 20K persistent messages, client channels..... | 23 |
| Figure 25 – 200K nonpersistent and persistent messages, client channels..... | 24 |
| Figure 26 – 2K and 20K nonpersistent messages, server channels..... | 26 |
| Figure 27 – 2K and 20K persistent messages, server channels..... | 26 |
| Figure 28 – 200K nonpersistent and persistent messages, server channels..... | 27 |
| Figure 29 – Trusted server application, local queue manager..... | 28 |
| Figure 30 – Short sessions, ‘runmqslr’ vs. inetd ‘amqcrsta’ listener, client channels | 30 |
| Figure 31 – Effect of number of client channels on round trips..... | 33 |
| Figure 32 – Effect of number of client channels on round trips..... | 34 |
| Figure 33 – Effect of number of server channels on round trips..... | 36 |

1 Release highlights

1.1 Improvements to nonpersistent and persistent messaging

- Maximum nonpersistent message throughput increased by 18% in a local queue manager environment, 16% in a client channel scenario, and 14% in a distributed queuing environment.
- Maximum persistent message throughput increased by 79% in a local queue manager environment, 61% in an client channel environment, and 66% in a distributed queuing environment.

1.2 Peak message throughput – Local queue manager

Figure 1 below illustrates the peak message throughput achieved for nonpersistent and persistent messages with a local queue manager, MQSeries V5.2 vs. WebSphere MQ V5.3.




Figure 1 – Peak message throughput, local queue manager

Note: messaging in these tests is with no think-time.

1.3 Peak message throughput – client channels

Figure 2 below illustrates the peak message throughput achieved for nonpersistent and persistent messages with MQI-client channels, MQSeries V5.2 vs. WebSphere MQ V5.3.




Figure 2 – Peak message throughput, client channels

Note: messaging in these tests is with no think-time.

1.4 Peak message throughput – distributed queuing

Figure 3 below illustrates the peak message throughput achieved for nonpersistent and persistent messages with server channels, MQSeries V5.2 vs. WebSphere MQ for V5.3.

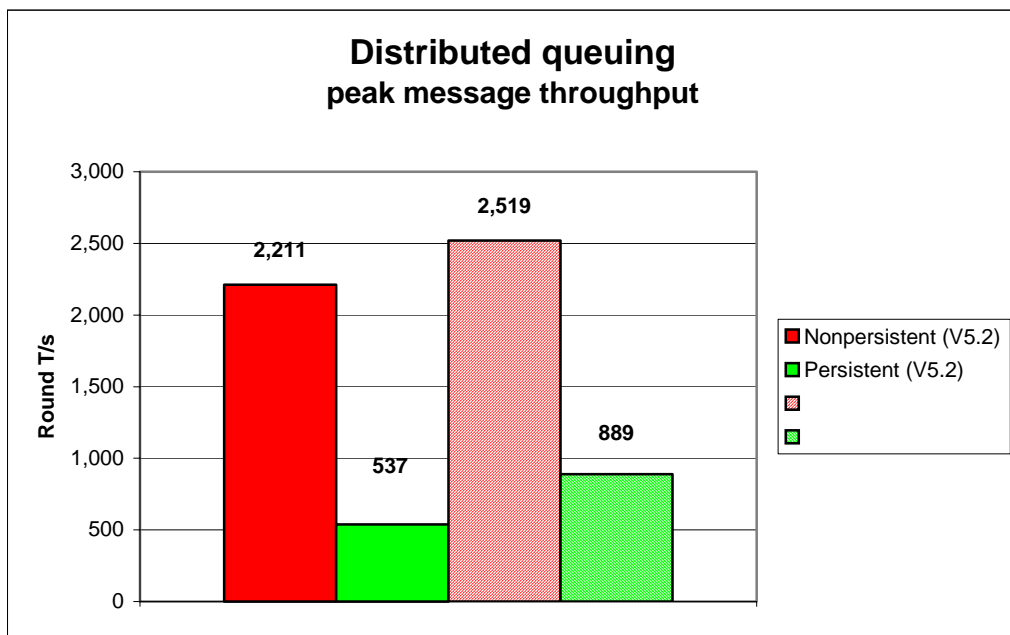


Figure 3 – Peak message throughput, distributed queuing

Note: messaging in these tests is with no think-time.

1.5 Improvements to channel capacity limits

- Client channels (trusted MQI-client connections) increased from 7,800 to 31,500 using one reply queue for all clients, and can support 21,500 using one reply queue per client.
- Distributed queuing (trusted server channels) increased from 1,100 to 6,000 pairs.

1.6 Capacity limits – client channels

Figure 4 below shows the maximum number of MQI-client connections made concurrently into a single queue manager on the server machine.




Figure 4 – Maximum number of client connections

Note: these tests use a rate of 1 nonpersistent round trip per MQI-client channel per *minute*.

1.7 Capacity limits – distributed queuing (server channels)

Figure 5 below shows the maximum number of server channel pairs achieved between two queue managers on separate server machines.




Figure 5 – Maximum number of server channels

Note: these tests use a rate of 1 nonpersistent round trip per server channel pair per *minute*.

2 Performance headlines

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

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

In the rated messaging tests, the rate used is 1 round trip per driving application per *second*. In the client channel test scenarios, each driving application uses a dedicated MQI-client channel, and in the distributed queuing test scenarios, one or more applications submit messages over a fixed number of server channels.

2.1 Local queue manager test scenario

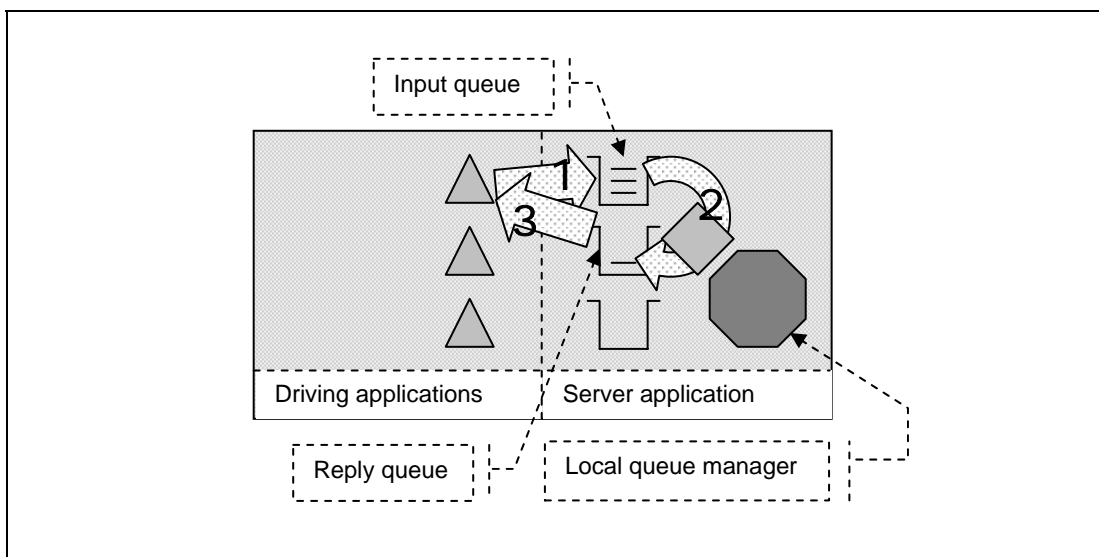


Figure 6 – Connections into a local queue manager

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

Nonpersistent and persistent messages were used in the local queue manager tests, with a message size of 2K. The effect of message throughput with larger messages sizes is investigated in '**MQI response times: 50bytes to 2MB – local queue manager**'—Page 17, and '**Large messages: 20K and 200K – local queue manager**'—Page 19.

Figure 8 and **Figure 8** below illustrate the nonpersistent and persistent message throughput achieved using an *increasing* number of driving applications in the local queue manager test scenario (see **Figure 6** above), and WebSphere MQ V5.3 compared to Version 5.2.

2.1.1 Nonpersistent messages – local queue manager




Figure 7 – Performance headline, nonpersistent messages, local queue manager

Note: messaging in these tests is with no think-time.

| Test name | Product version | Apps | Round T/s | % | Resp time (s) |
|-----------|-------------------|------------------|-------------------------|--------------|------------------|
| local_np1 | MQSeries V5.2 | 5 (20) | 2,571 (2,386) | n/a | 0.002 (0.011) |
| local_np1 | WebSphere MQ V5.3 | 6 (20) | 3,044 (2,812) | +18 (+18) | 0.002 (0.008) |

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

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

Figure 7 and **Table 1** above show that the peak throughput of nonpersistent messages has increased in Version 5.3 compared to Version 5.2, with the advantage of improved scalability when using less than 5 driving applications. Using 20 driving applications nonpersistent throughput is up by 18% (2,386 cf. 2,812 RT/s).

2.1.2 Persistent messages – local queue manager

Queue manager log configuration:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512




Figure 8 – Performance headline, persistent messages, local queue manager

Note: messaging in these tests is with no think-time.

| Test name | Product version | Apps | Round T/s | % | Resp time (s) |
|-----------|-------------------|----------------------------|------------------------------|--------------------------------|------------------------------------|
| local_pm3 | MQSeries V5.2 | 16 (92) (120) | 504 (395) (394) | n/a | 0.034 (0.244) (0.421) |
| local_pm3 | WebSphere MQ V5.3 | (16) 92 (120) | (801) 903 (864) | (+59) +129 (+119) | (0.022) 0.120 (0.166) |

Table 2 – Performance headline, persistent messages, local queue manager

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

Figure 8 and **Table 2** above show that the peak throughput of persistent messages has increased by 79% (504 cf. 903 RT/s) comparing Version 5.2 with Version 5.3. Using 92 driving applications persistent throughput is up by 129% (903 cf. 395 RT/s). Version 5.3 has the advantage of improved scalability when using less than 10 driving applications.

2.2 Client channels test scenario

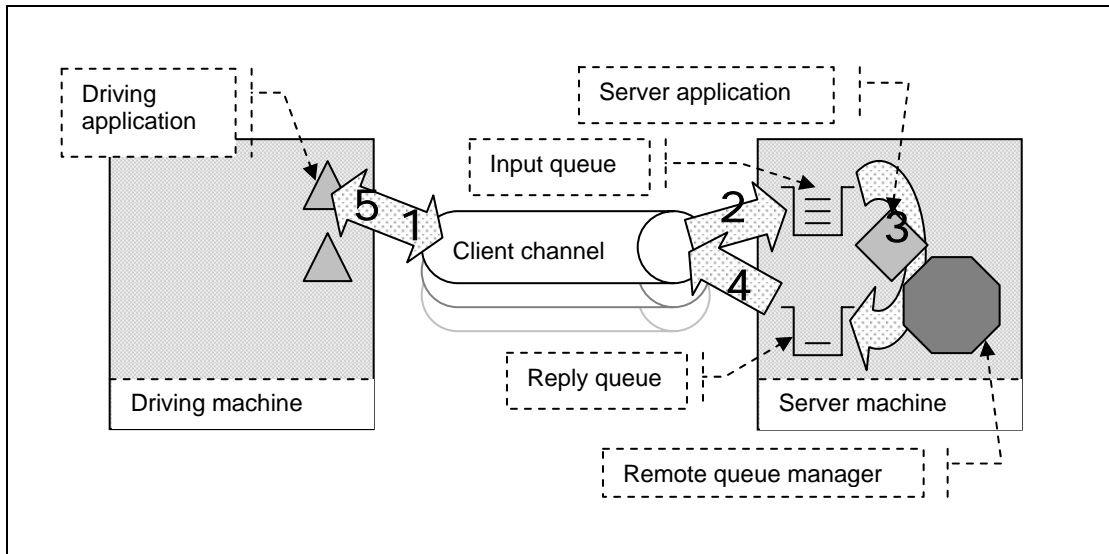


Figure 9 – MQI-client channels into a remote queue manager

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

Nonpersistent and persistent messages were used in the client channel tests, with a message size of 2K. The effect of message throughput with larger messages sizes is investigated in **'Large messages: 20K and 200K – client channels'—Page 22.**

Figure 10 and **Figure 11** below illustrate the nonpersistent and persistent message throughput achieved using an increasing number of driving applications in the client channel test scenario (see **Figure 9** above), and Version 5.2 compared with Version 5.3.

2.2.1 Nonpersistent messages – client channels




Figure 10 – Performance headline, nonpersistent messages, client channels

Note: messaging in these tests is with no think-time.

| Test name | Product version | Apps | Round T/s | % | Resp time (s) |
|---------------|-------------------|-------------------|-------------------------|--------------|------------------|
| clnp1 (inetd) | MQSeries V5.2 | 10 (20) | 1,400 (1,336) | n/a | 0.008 (0.018) |
| clnp1 (inetd) | WebSphere MQ V5.3 | 13 (20) | 1,636 (1,558) | +17 (+16) | 0.009 (0.014) |

Table 3 – Performance headline, nonpersistent messages, client channels

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

Figure 10 and **Table 3** above show that the peak throughput of nonpersistent messages is up by 17% (1,400 cf. 1,636 RT/s) comparing Version 5.2 to Version 5.3 with as few as 10 driving applications. Using 20 driving applications throughput is also improved by 16% : 1,336 cf. 1,558 RT/s).

2.2.2 Persistent messages – client channels

Queue manager log configuration:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512




Figure 11 – Performance headline, persistent messages, client channels

Note: messaging in these tests is with no think-time.

| Test name | Product version | Apps | Round T/s | % | Resp time (s) |
|---------------|-------------------|----------------------------|------------------------------|------------------------------|----------------------------------|
| clpm3 (inetd) | MQSeries V5.2 | 16 (96) (120) | 452 (370) (377) | n/a | 0.039 (0.296) (0.35) |
| clpm3 (inetd) | WebSphere MQ V5.3 | (16) 96 (120) | (665) 729 (711) | (+47) +97 (+89) | (0.028) 0.155 (0.2) |

Table 4 – Performance headline, persistent messages, client channels

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

Figure 11 and **Table 4** above show that the peak throughput of persistent messages is up by 61% (452 cf. 729 Round T/s) comparing Version 5.2 with Version 5.3, with the advantage of improved scalability giving the most performance improvement using one hundred or so driving applications (for 96 driving applications; throughput up by 97% : 370 cf. 729 RT/s).

2.2.3 ‘runmqslr’ vs. inetd ‘amqcrsta’ listener – client channels

For the following client channel measurements, the messaging rate used is 1 round trip per second per MQI-client channel, i.e. a request message outbound over the client channel and a reply message inbound over the channel per second. These tests also compare the difference between *nonthreaded* channels (the ‘amqcrsta’ process started by inetd) with *threaded* channels (the ‘runmqslr’ process started by the user).




Figure 12 – ‘runmqslr’ vs. inetd ‘amqcrsta’ listener, client channels

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

Figure 12 above and **Table 5** below illustrate how the ‘runmqslr’ and inetd ‘amqcrsta’ listeners in WebSphere MQ V5.3 give improved scalability by permitting a larger number of MQI-client connections into a single queue manager. In Version 5.2, either the ‘runmqslr’ listener physically breaks (a well know and documented problem), or the ‘amqcrsta’ listener response time exceeds one second.

In Version 5.3 it is now possible to connect more than 500 or so driving application into a single ‘runmqslr’ (31,500 fastpath MQI-client connections: refer to ‘**Capacity limits – client channels**’—Page 3). Furthermore, the ‘runmqslr’ has a reduced resource utilisation (one thread per connection compared to a process per connection for the ‘amqcrsta’ listener, a smaller memory footprint, less System V IPC), so is now the **preferred** method of running client channels and server channels. At 1,200 connections the ‘runmqslr’ does *not* break, it is the round trip response time exceeding one second that stops the test from using more driving applications.

| Test name | Apps | Rate/App/hr | Round T/s | % | Resp time (s) |
|---|------------------|-------------|------------------|------|------------------|
| clnp1_r3600 (inetd) (MQSeries V5.2) | 1,400 (1,100) | 3,600 | 1,395 (1,089) | +28 | 0.951 (0.021) |
| clnp1_r3600_runmqslr (MQSeries V5.2) | 1,350 (200) | 3,600 | 1,330 (193) | +589 | 0.764 (0.011) |
| clpm3_r3600 (inetd) (MQSeries V5.2) | 600 (350) | 3,600 | 589 (334) | +76 | 0.184 (0.173) |
| clpm3_r3600_runmqslr (MQSeries V5.2) | 550 (200) | 3,600 | 518 (182) | +185 | 1.000 (0.046) |

Table 5 – 1 round trip per driving application per second, client channels

Figure 13 below illustrates the reduced swap requirement of an MQI-client connection comparing the inetd ‘amqcrsta’ listener to the ‘runmqslr’ listener in WebSphere MQ V5.3.

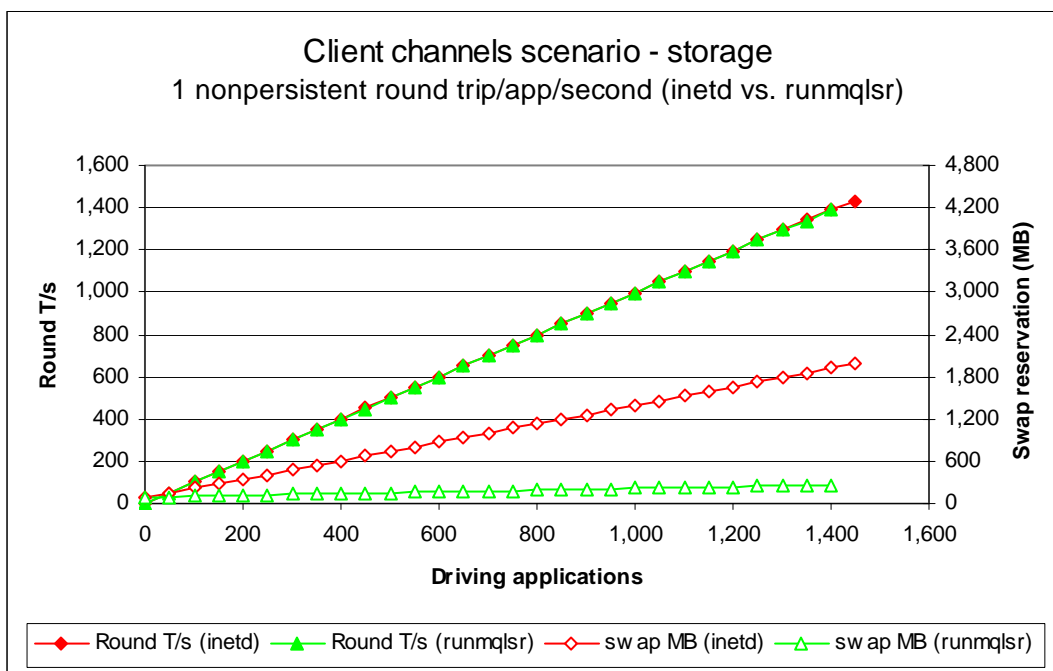


Figure 13 – ‘runmqslr’ vs. inetd ‘amqcrsta’ listener, swap requirement, client channels

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

| Test name | Apps | Swap | Free (4K pages) |
|----------------------|----------------|----------------------------|----------------------------------|
| clnp1_r3600 (inetd) | 100 (1,000) | 1.31MB/App (1.31MB/App) | 262 pages/App (279 pages/App) |
| clnp1_r3600_runmqslr | 100 (1,000) | 0.13MB/App (0.12MB/App) | 28 pages/App (27 pages/App) |

Table 6 – Nonthreaded vs. threaded listener, swap requirement, client channels

For further calculations on the swap requirement and shared memory utilisation for client channels refer to ‘Table 17 – Client capacity, swap requirement’—Page 33, and ‘Table 19 – Distributed queuing capacity, swap requirement’—Page 36.

2.3 Distributed queuing test scenario

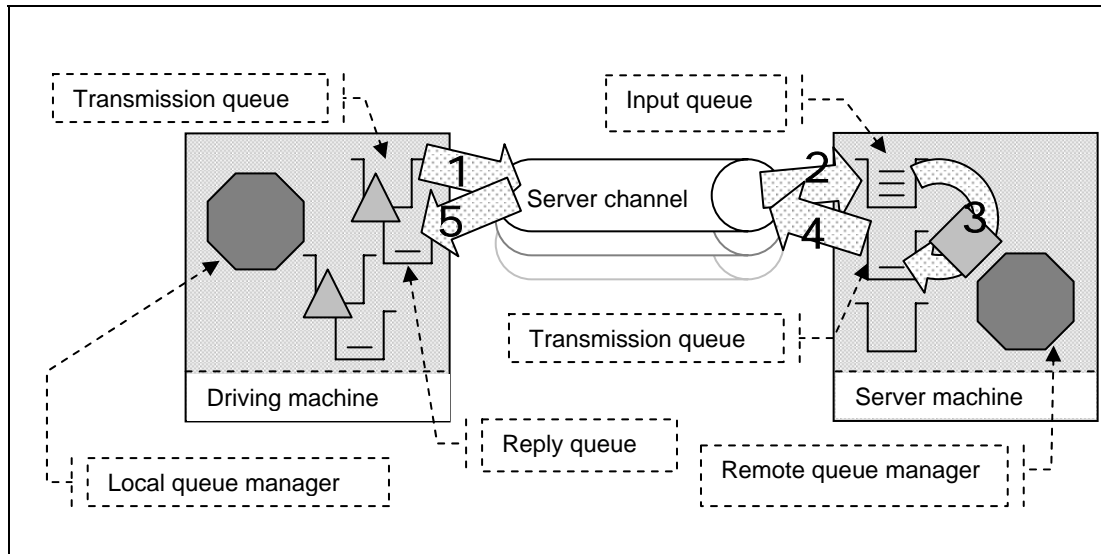


Figure 14 – Server channels between two queue managers

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

Nonpersistent and persistent messages were used in the distributed queuing tests, with a message size of 2K. The effect of message throughput with larger messages sizes is investigated in '**Large messages: 20K and 200K – distributed queuing**'—Page 25.

Figure 15 and **Figure 16** below show the peak nonpersistent and persistent message throughput achieved using an *increasing* number of driving applications in the distributed queuing scenario (see **Figure 14** above), and WebSphere MQ V5.3 compared to Version 5.2

2.3.1 Nonpersistent messages – server channels




Figure 15 – Performance headline, nonpersistent messages, server channels

Note: messaging in these tests is with no think-time.

| Test name | Product version | Apps | Round T/s | % | Resp time (s) |
|---------------|-------------------|-------------------|-------------------------|-------------|------------------|
| dqnp1 (inetd) | MQSeries V5.2 | 13 (20) | 2,211 (2,141) | n/a | 0.006 (0.011) |
| dqnp1 (inetd) | WebSphere MQ V5.3 | 10 (20) | 2,519 (2,229) | +14 (+4) | 0.005 (0.009) |

Table 7 – Performance headline, nonpersistent messages, server channels

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

Figure 15 and **Table 7** above show that the peak throughput of nonpersistent messages is up by 14% (2,211 cf. 2,519 RT/s) comparing Version 5.2 with Version 5.3.

2.3.2 Persistent messages – server channels

Queue manager log configuration:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512




Figure 16 – Performance headline, persistent messages, server channels

Note: messaging in these tests is with no think-time.

| Test name | Product version | Apps | Round T/s | % | Resp time (s) |
|---------------|-------------------|------------|------------|-----|---------------|
| dqpm1 (inetd) | MQSeries V5.2 | 116 | 537 | n/a | 0.236 |
| dqpm1 (inetd) | WebSphere MQ V5.3 | 116 | 889 | +66 | 0.159 |

Table 8 – Performance headline, persistent messages, server channels

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

Figure 16 and **Table 8** above show that the peak throughput of persistent messages is up by 66% (537 cf. 889 RT/s) comparing Version 5.2 with Version 5.3 also with the advantage of improved scalability after 20 driving applications.

The persistent message tests in **Figure 16** above do not constrain on response time. This shows that in the distributed queuing scenario used for the tests in this section (using 2 server channel pairs between two queue managers running on separate server machines), the Version 5.2 and Version 5.3 queue managers can maintain a message throughput of 537 and 889 round trips per second per server channel pair respectively.

2.3.3 'runmqslr' vs. inetd 'amqcrsta' listener – server channels

For the following distributed queuing measurements, the rate used is 1 round trip per driving application per *second*, i.e. a request message outbound over the sender channel, and a reply message inbound over the receiver channel per second. Note that there are a fixed number of 4 server channel pairs for the nonpersistent messaging tests, and 2 pairs for the persistent message tests. These tests also compare the different between nonthreaded channels (the 'amqcrsta' process started by inetd, and the 'runmqchl' process started by the queue manager) with *threaded* channels (the 'runmqslr' process started by the user, and the 'runmqchl' process started with the queue manager).




Figure 17 – 'runmqslr' vs. inetd 'amqcrsta' listener, server channels

Note: messaging in these tests is 1 round trip per driving application per *second*

| Test name | Apps | Rate/App/hr | Round T/s | % | Resp time |
|---|------------------|-------------|------------------|-----|------------------|
| dqnp1_r3600 (inetd) (MQSeries V5.2) | 1,750 (1,650) | 3,600 | 1,713 (1,629) | +5 | 0.510 (0.089) |
| dqnp1_r3600_runmqslr (MQSeries V5.2) | 1,800 (1,450) | 3,600 | 1,793 (1,447) | +24 | 0.208 (0.051) |
| dqpm1_r3600 (inetd) (MQSeries V5.2) | 650 (700) | 3,600 | 645 (635) | +2 | 0.252 (0.973) |
| dqpm1_r3600_runmqslr (MQSeries V5.2) | 650 (500) | 3,600 | 634 (482) | +32 | 0.346 (0.674) |

Table 9 – 1 round trip per driving application per second, server channels

3 Large messages

All tests are automatically stopped after the response time of 1 round trip exceeds 1 second.

3.1 MQI response times: 50bytes to 2MB – local queue manager

Queue manager log configuration:

LogPrimaryFiles=3, LogFilePages=2048

Figure 18 and **Figure 19** below show that the response for MQPUT/GET pairs is improved for all persistent message sizes between 50bytes and 2MB. For nonpersistent messages the response time for MQPUT/GET pairs is slightly quicker for all message sizes.




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

Note: messaging in these tests is with no think-time.

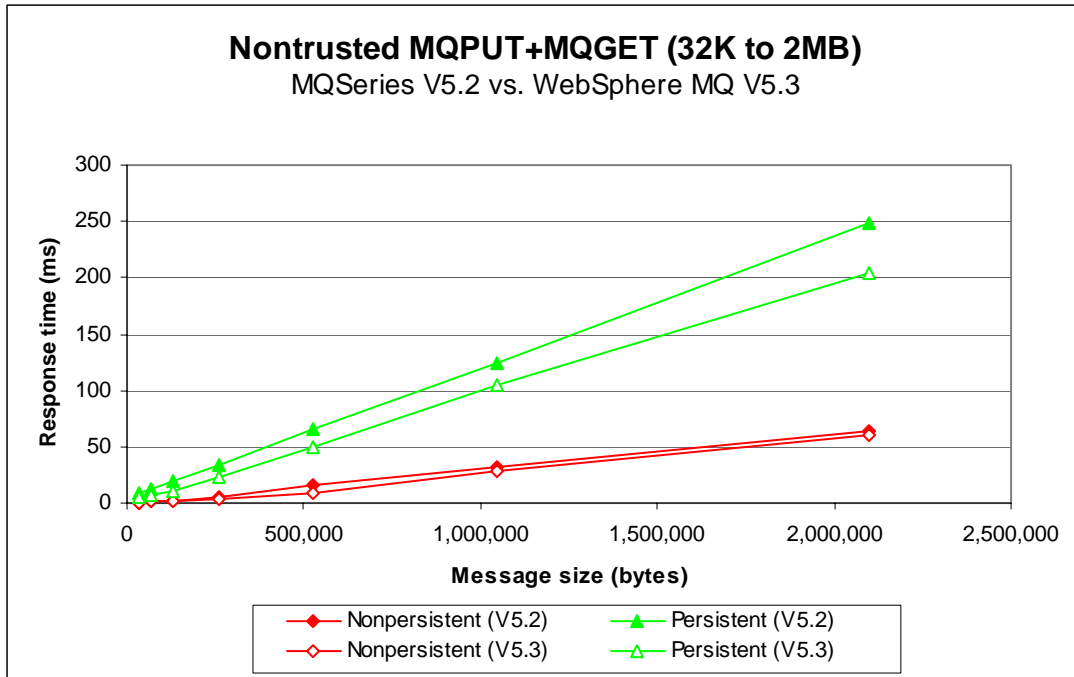


Figure 19 – The effect of message size on MQI response time (32K - 2MB)

Note: messaging in these tests is with no think-time.

3.2 Large messages: 20K and 200K – local queue manager

Queue manager log configuration for persistent tests:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512

| Test name | Apps | Msg size | Round T/s | % | Resp time (s) |
|-----------------------------------|------------|----------|------------------|-----|------------------|
| local_np1_2K (MQSeries V5.2) | 6 (5) | 2K | 3,044 (2,571) | +18 | 0.002 (0.002) |
| local_np1_20K (MQSeries V5.2) | 5 (4) | 20K | 1,899 (1,649) | +15 | 0.002 (0.002) |
| local_np1_200K (MQSeries V5.2) | 3 (3) | 200K | 263 (238) | +11 | 0.012 (0.014) |
| local_pm3_2K (MQSeries V5.2) | 92 (16) | 2K | 903 (504) | +79 | 0.120 (0.034) |
| local_pm3_20K (MQSeries V5.2) | 48 (24) | 20K | 291 (231) | +26 | 0.206 (0.129) |
| local_pm3_200K (MQSeries V5.2) | 24 (32) | 200K | 36 (34) | +6 | 0.779 (1.081) |

Table 10 – 2K, 20K and 200K messages, local queue manager

Note: messaging in these tests is with no think-time.

Note: the figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included to in the table to provide a meaningful comparison with MQSeries V5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

The measurements for 200K *persistent* messages show that there is little difference in the performance of large messages between Version 5.2 and Version 5.3—this is because most of the time taken by the queue manager is in logging the messages to disk.

Figure 20 below illustrates how the throughput of small nonpersistent messages is improved significantly in WebSphere MQ V5.3, using less than 5 driving applications.




Figure 20 – 2K and 20K nonpersistent messages, local queue manager

Figure 21 below illustrates how the throughput of 2K persistent messages is improved significantly, and the throughput of 20K persistent messages is improved slightly, in WebSphere MQ V5.3, using more than 20 driving applications.




Figure 21 – 2K and 20K persistent messages, local queue manager




Figure 22 – 200K nonpersistent and persistent messages, local queue manager

3.3 Large messages: 20K and 200K – client channels

Queue manager log configuration for persistent tests:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512

| Test name | Apps | Msg size | Round T/s | % | Resp time (s) |
|-------------------------------|------------|----------|------------------|-----|------------------|
| clnp1_2K (MQSeries V5.2) | 13 (10) | 2K | 1,636 (1,400) | +17 | 0.009 (0.008) |
| clnp1_20K (MQSeries V5.2) | 9 (9) | 20K | 1,057 (897) | +22 | 0.010 (0.012) |
| clnp1_200K (MQSeries V5.2) | 9 (8) | 200K | 130 (124) | +5 | 0.080 (0.076) |
| clpm3_2K (MQSeries V5.2) | 96 (16) | 2K | 729 (452) | +61 | 0.155 (0.039) |
| clpm3_20K (MQSeries V5.2) | 44 (32) | 20K | 227 (223) | +17 | 0.185 (0.156) |
| clpm3_200K (MQSeries V5.2) | 16 (20) | 200K | 34 (33) | +3 | 0.537 (0.699) |

Table 11 – 2K, 20K and 200K messages, client channels

Note: messaging in these tests is with no think-time, and the inetd 'amqcrsta' listener.

Note: the figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included to in the table to provide a meaningful comparison with MQSeries V5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

The measurements for 200K *persistent* messages show that there is little difference in the performance of large messages between Version 5.2 and Version 5.3—this is because most of the time taken by the queue manager is in logging the messages to disk.

Figure 23 below illustrates how the throughput of 2K and 20K nonpersistent messages has significantly improved in WebSphere MQ V5.3, using 6 or more driving applications.




Figure 23 – 2K and 20K nonpersistent messages, client channels

Figure 24 below illustrates how the throughput of 2K persistent messages has improved significantly, and the throughput of 20K persistent messages has improved slightly, in WebSphere MQ V5.3, using more than 10 driving applications. Using 20 or more driving applications Version 5.3 gives a *maintained* persistent throughput of 700 round trips or more per second (75% more than Version 5.2).




Figure 24 – 2K and 20K persistent messages, client channels




Figure 25 – 200K nonpersistent and persistent messages, client channels

3.4 Large messages: 20K and 200K – distributed queuing

Queue manager log configuration for persistent tests:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512

| Test name | Apps | Msg size | Round T/s | % | Resp time (s) |
|-------------------------------|--------------|----------------|------------------|------|------------------|
| dqnp1_2K (MQSeries V5.2) | 10 (13) | 2K (2K) | 2,519 (2,211) | +14 | 0.005 (0.006) |
| dqnp1_20K (MQSeries V5.2) | 9 (9) | 20K (20K) | 1,102 (1,011) | +9 | 0.009 (0.010) |
| dqnp1_200K (MQSeries V5.2) | 14 (11) | 200K (200K) | 128 (119) | +8 | 0.129 (0.108) |
| dqpm1_2K (MQSeries V5.2) | 116 (116) | 2K (2K) | 889 (537) | +66 | 0.159 (0.236) |
| dqpm1_20K (MQSeries V5.2) | 116 (112) | 20K (20K) | 264 (234) | +13 | 0.474 (0.670) |
| dqpm1_200K (MQSeries V5.2) | 12 (8) | 200K (200K) | 30 (11) | +172 | 0.461 (0.793) |

Table 12 – 2K, 20K and 200K messages, server channels

Note: messaging in these tests is with no think-time, and the inetd 'amqcrsta' listener.

Note: the figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included to in the table to provide a meaningful comparison with MQSeries V5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

The measurements for 200K *persistent* messages show that there is little difference in the performance of large messages between Version 5.2 and Version 5.3—this is because most of the time taken by the queue manager is in logging the messages to disk.

Figure 26 below illustrates how the throughput of small nonpersistent messages has improved slightly in WebSphere MQ V5.3, using as few as 5 driving applications.




Figure 26 – 2K and 20K nonpersistent messages, server channels

Figure 27 below illustrates how the throughput of 2K persistent messages has improved significantly, and the throughput of 20K persistent messages has improved slightly, in WebSphere MQ V5.3, using more than 10 driving applications. Using more than 40 driving applications Version 5.3 gives approximately 80% more throughput for 2K persistent messages (compared to Version 5.2).

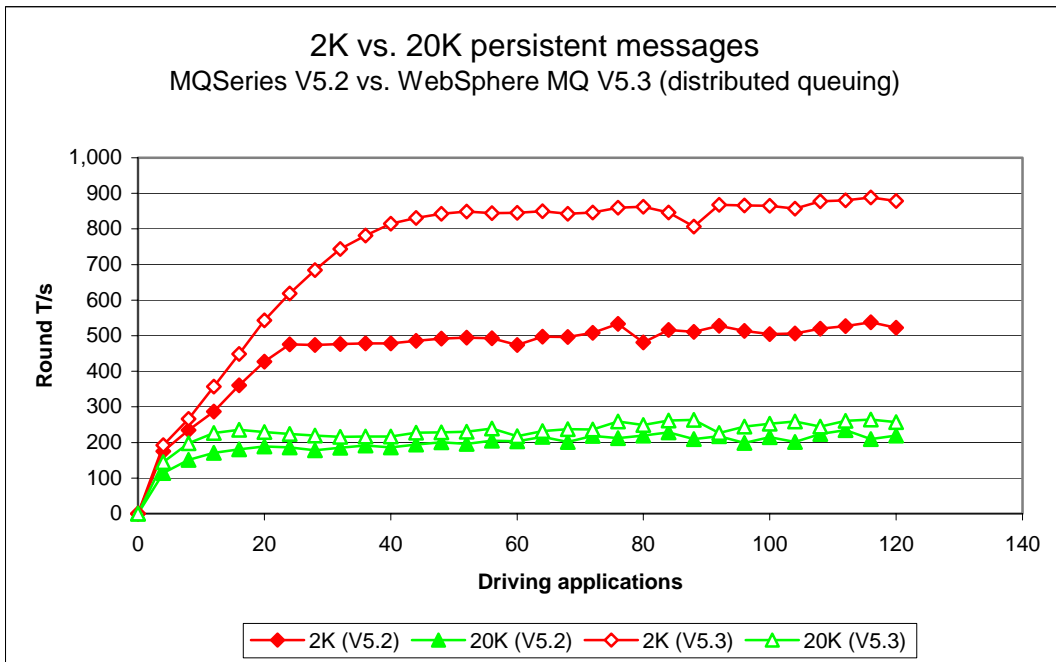


Figure 27 – 2K and 20K persistent messages, server channels




Figure 28 – 200K nonpersistent and persistent messages, server channels

With reference to **Figure 28** above, the 200K nonpersistent message tests were intentionally designed to finish at 20 driving applications. Using more than 10 driving applications, throughput neither particularly increases nor decreases. The 200K persistent message tests were designed to finish at 120 driving applications, but after only 8 applications (for Version 5.2) and 20 applications (for Version 5.3) both tests are approaching the one second response time criteria.

4 Trusted server application

Queue manager log configuration for persistent tests:

LogPrimaryFiles=4, LogFilePages=4095, LogBufferPages=512

| Test name | Apps | Msg size | Round T/s | % | Resp time (s) |
|--------------------------------------|------------|----------|-------------------------|--------------|-------------------------|
| local_np1_trusted (MQSeries V5.2) | 2 (2) | 2K | 5,404 (4,478) | +21 (n/a) | 0.002 (0.002) |
| local_pm3_trusted (MQSeries V5.2) | 44 (16) | 2K | 1,003 (567) | +77 (n/a) | 0.206 (0.129) |

Table 13 – Trusted server application, local queue manager

Note: messaging in these tests is with no think-time.

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.




Figure 29 – Trusted server application, local queue manager

| Test name | Apps | Msg size | Round T/s | % | Resp time (s) |
|---|---------------------|-----------|-------------------------|---------------------|-------------------------|
| clnp1_trusted (MQSeries V5.2) | 12 (12) | 2K | 2,597 (2,241) | +16 (n/a) | 0.005 (0.006) |
| clpm3_trusted (MQSeries V5.2) | 44 (16) | 2K | 796 (475) | +68 (n/a) | 0.064 (0.038) |
| dqnp1_trusted (MQSeries V5.2) | 12 (15) | 2K | 3,455 (2,948) | +17 (n/a) | 0.004 (0.005) |
| dqpm1_trusted (MQSeries V5.2) | 116 (120) | 2K | 978 (569) | +72 (n/a) | 0.145 (0.216) |

Table 14 – Trusted server application, client channels and server channels

Note: messaging in these tests is with no think-time, and the inetd 'amqcrsta' listener.

Note: the large bold figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The smaller figures in brackets are included in the table to provide meaningful comparison between WebSphere MQ V5.3 and Version 5.2. The percentage column shows the percentage Round T/s improvement of Version 5.3 over Version 5.2.

5 Short Sessions

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

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



“Why measure short sessions?”

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

This effect is greater than that of reducing the total messaging throughput of the queue manager by connecting thousands of MQI applications to the queue manager (refer to ‘**Figure 31 – Effect of number of client channels on round trips**’—Page 33 for an illustration).




Figure 30 – Short sessions, ‘runmqtsr’ vs. inetd ‘amqcrsta’ listener, client channels

| Test name | Apps | Short sessions/s | Resp time (s) |
|-------------------------|--------------|------------------|------------------|
| clnp1_ss_r3600 (inetd) | 140 (165) | 28 (29) | 0.941 (1.633) |
| clnp1_ss_r3600_runmqlsr | 390 | 84 | 0.230 |

Table 15 – Short sessions, nonpersistent messages, client channels

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

Note: the large figures in **Table 15** above are for WebSphere MQ V5.3 with a round trip response time of less than one second. The smaller figures in brackets are for the maximum number of round trips per second (or the maximum number of short sessions in the time indicated by 'Resp time').

The 'runmqlsr' has a much smaller overhead of connecting to and disconnecting from the queue manager because it only uses a single thread per connection rather than an entire process. Furthermore, in Version 5.3 the maximum number of connections into a single 'runmqlsr' listener has been significantly increased making it the **preferred** method of running short sessions over client channels.

6 Performance and capacity limits

6.1 Client channels scenario – capacity measurements

The measurements in this section are intended to test the maximum number of MQI-clients that can be connected into a single queue manager with a message rate of 1 round trip per client channel per *minute*. In previous SupportPacs, the rate used in capacity limit tests was 1 round trip per *hour*. For the same number of client channels, a faster message rate gives a higher total message throughput over each channel. This information is intended to be more useful to the reader and assist them in projecting the results in this section to similar scenarios.

Queue manager configuration for client channel capacity tests:

```
MaxChannels=50000
```

| Test name | Apps | Rate /App/hr | Round T/s | Resp time (s) | %cpu |
|---|--------------------|--------------|--------------|------------------|------------|
| clnp1 (inetd) | 13 | n/a* | 1,636 | 0.009 | 98 |
| clnp1_r3600 (inetd) | 1,400 | 3,600 | 1,395 | 0.951 | 71 |
| clnp1_c6000_t10 (inetd) | 6,000 | 390 | 654 | 0.935 | 63 |
| clnp1_c6000_t10_runmqlsr | 6,000 | 550 | 850 | 0.503 | 34 |
| clnp1_cmax_t10 (inetd) | 6,200 | 60 | 103 | 0.124 | 2** |
| clnp1_cmax_runmqlsr_t10 (unique reply queue per app) | 31,500 (20,700) | 60 (60) | 525 (346) | 0.011 (0.022) | 51 (64) |

Table 16 – Capacity measurements, client channels

Note: the figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The percentage column shows the CPU utilisation for the measurement point.

* Since there are **1,636 Round T/s** and 3,600 seconds in one hour, the derived throughput volume per hour is calculated to be $1,636 \times 3,600 = \mathbf{5,889,600 \text{ round trips per hour}}$. The reader should note that the number of 5,900,000 round trips in one hour was not *physically* measured over the period of one hour, however, the rates of 3,600, 390, and 550 were *actual* rates set for the measurement and obtained by the queue manager system.

With reference to **Table 16** above, it is clear to see that the inetd 'amqcrsta' listener is the more resource intensive than the 'runmqlsr' listener, both in terms of the number of Round T/s that can be achieved using a single queue manager, and the CPU utilisation required to achieve the number of Round T/s. **Table 17** below supports this observation, and also indicates the additional costs associated with using more than one reply queue per driving application.

For the 'cmax' tests, the inetd 'amqcrsta' listener appears to be efficient on CPU resource, however, other resource utilisation on the machine (not included in **Table 16 above) shows there are no free memory pages, and the Operating System swap disk is very busy (see **Figure 32 Error! Reference source not found.**)—this limits the 'clnp1_cmax_t10' test to 6,200 driving applications.

The effect of the number of client channels on maximum message throughput can be seen in **Figure 31** below.




Figure 31 – Effect of number of client channels on round trips

Note: no think-time 1st column (*tight-loop*), fixed 'slow' rate 2nd column (1 nonpersistent round trip per driving application per second), increasing rate 3rd and 4th columns (refer to **Table 16** above for rates of 390 and 550 nonpersistent round trips per driving application per second), fixed 'medium' rate 5th, 6th and 7th columns (1 nonpersistent round trip per driving application per minute).

| Test name | Apps | Swap | shm | Free (4K pages) |
|--|-------------------|-----------------------|------------------------------------|----------------------------------|
| clnp1_cmax_t10 (inetd) | 6,200 (1,000) | 8.1GB (1.32MB/App) | 148.2MB (24.5K/App) (10.2K/App) | 155 pages/App (275 pages/App) |
| clnp1_cmax_runmqsr_t10 | 31,500 (1,000) | 4.2GB (0.13MB/App) | 457.1MB (14.9K/App) (18.9K/App) | 29 pages/App 31 pages/App |
| clnp1_cmax_runmqsr_t10 Note: unique reply queue per app | 20,700 (1,000) | 3.4GB (0.14MB/App) | 1.2GB (61.1K/App) (32.0K/App) | 45 pages/App (44 pages/App) |

Table 17 – Client capacity, swap requirement

Note: the large figures in the table above show the swap requirement and shared memory measured at the given number of driving applications. The large and small figures in brackets are the proportionate swap requirement and shared memory costs per driving application—in this test scenario this relates to the cost of an MQI-client connection on the server machine.

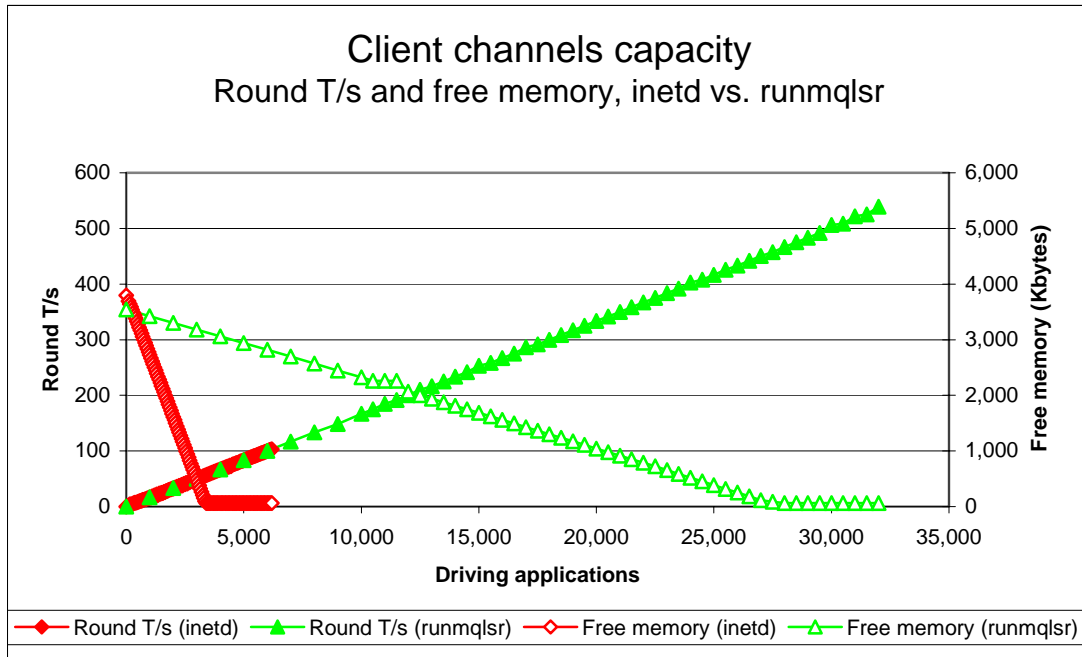


Figure 32 – Effect of number of client channels on round trips

Note: when there are a large number of MQI-clients connected to the queue manager on the server machine, the poor relative performance of the inetd 'amqcrsta' listener compared to the 'runmqslr' listener (see **Table 16** and **Figure 32** above) is due to the excessive amount of paging the Operating System must perform to page-in the process address space for each channel as messages arrive from the driving applications.

6.2 Distributed queuing scenario – capacity measurements

The measurements in this section are intended to test the maximum number of server channel pairs between two queue managers with a messaging rate of 1 round trip per server channel per *minute*. In previous SupportPacs, the rate used in capacity limit tests was 1 round trip per *hour*. For the same number of server channel pairs, a faster message rate gives a higher total message throughput over each channel pair. This information is intended to be more useful to the reader and assist them in projecting the results in this section to similar scenarios.

Queue manager and log configuration for distributed queuing capacity tests:

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

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

| Test name | Apps | Channel pairs | Rate/App/hr | Round T/s | Resp time (s) | %cpu |
|---------------------|-------|---------------|-------------|-----------|---------------|------|
| dqnp1 (inetd) | 10 | 4 | n/a* | 2,519 | 0.005 | 97 |
| dqnp1_r3600 (inetd) | 1,750 | 4 | 3,600 | 1,713 | 0.510 | 61 |
| dqnp1_q1000 (inetd) | 1,000 | 1,000 | 2,960 | 815 | 0.017 | 50 |
| dqnp1_q1000_runmqsr | 1,000 | 1,000 | 3,160 | 869 | 0.128 | 12 |
| dqnp1_qmax (inetd) | 1,560 | 6,000 | 60 | 27 | 0.005 | 1 |
| dqnp1_qmax_runmqsr | 6,000 | 6,000 | 60 | 99 | 0.142 | 5 |

Table 18 – Capacity measurements, server channels

Note: the figures in the table above show the peak number of round trips per second, and the number of driving applications used to achieve the peak throughput. The percentage column shows the CPU utilisation for the measurement point.

* Since there are **2,519 Round T/s**, and 3,600 seconds in one hour, the derived throughput volume per hour is calculated to be $2,519 \times 3,600 = \mathbf{9,068,400}$ **round trips per hour**. The reader should note that the number of 9,000,000 round trips in one hour was not *physically* measured over the period of one hour, however, the rates of 3,600, 815, and 869 were *actual* rates set for the measurement and obtained by the queue manager system.

The effect of the number of server channel pairs on maximum message throughput can be seen in **Figure 33** below.

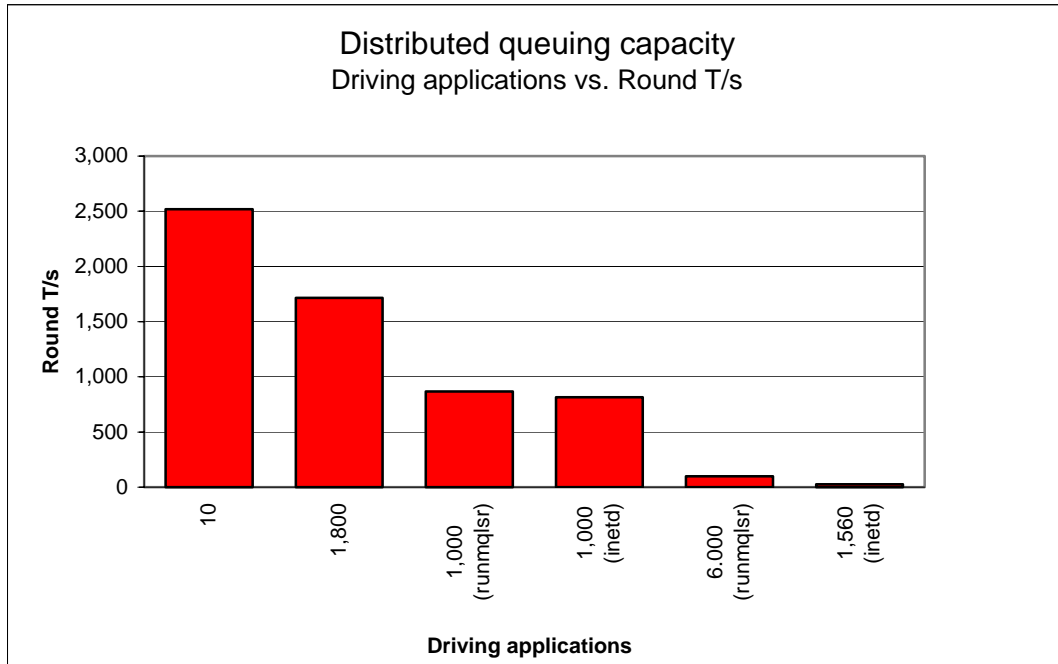


Figure 33 – Effect of number of server channels on round trips

Note: no think-time 1st column (*tight-loop*), fixed 'slow' rate 2nd column (1 nonpersistent round trip per server channel pair per *second*), increasing rate 3rd and 4th columns (refer to **Table 18** above for rates of 2,960 and ,3160 nonpersistent round trips per server channel pair per second), fixed 'medium' rate 5th and 6th columns (1 nonpersistent round trip per server channel pair *minute*).

| Test name | Apps | Swap | shm | Free (4K pages) |
|---------------------|---------------|--------------------|---------------------------------|-------------------------------|
| dqnp1_qmax (inetd) | 1,560 (1,000) | 4.3GB (2.67MB/App) | 226.7MB (39.0K/App) (29.6K/App) | 520 pages/App (589 pages/App) |
| dqnp1_qmax_runmqtsr | 6,000 (1,000) | 2.1GB (0.29MB/App) | 597.1MB (53.1K/App) (28.4K/App) | 77 pages/App (94 pages/App) |

Table 19 – Distributed queuing capacity, swap requirement

Note: the large figures in the table above show the swap requirement and shared memory measured at the given number of driving applications. The large and small figures in brackets are the proportionate swap requirement and shared memory costs per driving application—in this test scenario this relates to the cost of a server channel pair on the server machine.

7 Tuning recommendations

7.1 Tuning the queue manager

This section highlights the tuning activities that are known to give performance benefits for WebSphere MQ V5.3; all of these can be applied equally to Version 5.2. The reader should note that the following tuning recommendations **may not** necessarily **need** to be applied, especially if the message throughput and/or response time of the queue manager system already meets the required level. Some tuning recommendations that follow may degrade the performance of a previously balanced system if applied inappropriately. The reader should carefully monitor the result 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 those changes that are essential. For example, allocating several megabytes for multiple queues reduces the amount of shared and virtual memory available for other subsystems, as well as over committing real storage.

Note: the 'TuningParameters' stanza is not a documented external interface and may change or be removed in future releases.

7.1.1 Queue disk, Log disk, and message persistence

To avoid potential queue and log I/O contention due to the queue manager simultaneously updating a queue file and log extent on the same disk, it is important that queues and logs are located on *separate* and *dedicated* physical devices. With the queue and log disks configured in this manner, careful consideration must still be given to message persistence: persistent messages should only be used if the message needs to survive a queue manager restart (forced by the administrator, or as the result of a power failure, communications failure, or hardware failure). In guaranteeing the recoverability of persistent messages, the pathlength through the queue manager is three times longer than for a nonpersistent message. This overhead does not include the additional time for the message to be written to the log, although this can be minimised by using cached disks.

7.1.1.1 Nonpersistent queue buffer

The default nonpersistent queue buffer size is 64K per queue. This can be increased to 1MB using the TuningParameters stanza and the *DefaultQBufferSize* parameter. The nonpersistent queue buffer is computationally less expensive because the Operating System does not have to retrieve the message from the queue file. Increasing the queue buffer provides the capability to absorb peaks in message throughput at the expense of real storage, but it is **not** suitable as a long-term storage for nonpersistent messages as this buffer is **not** recovered after a queue manager restart. Defining queues using large nonpersistent queue buffers can degrade performance either if the system is short of real memory 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 nonpersistent queue buffer is 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 *DefaultQFileSize*. If some queues need to be defined differently to others, the values can be set in the TuningParameters stanza. When the queue manager is restarted existing queues will keep their earlier definitions and new queues will be created with the desired parameters. When a queue is opened, resources are allocated according to the definition held on disk from when the queue was created.

7.1.2 Log buffer size, Log file size, and number of log extents

To improve persistent message throughput the *LogBufferPages* should be increased to its maximum configurable size of 512 x 4K pages = 2MB, the *LogFilePages* (i.e. `crtmqm -lf <LogFilePages>`) should be configured to a large size, for example: 16384 x 4K pages = 64MB, and the number of *LogPrimaryFiles* (i.e. `crtmqm -lp <LogPrimaryFiles>`) should be configured to a large number. The cumulative effect of this tuning will:

- improve the throughput of persistent messages (permitting a possible maximum 2MB of log records to be written from the log buffer to the log disk in a single write),
- 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).

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

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

7.1.3 Channels: process or *thread*, standard or *fastpath*?

It is no longer necessary to consider the system design when deciding whether it is preferable to configure `inetd` to use process channels ('`amqcrsta`', and for server channels an `MCATYPE` of '`PROCESS`'), or use threaded channels ('`runmqslr`', and for server channels an `MCATYPE` of '`THREAD`') where prior to Version 5.3 it was necessary to use more than one '`runmqslr`' listener using more than one port. '`runmqslr`' **can now be used in all** scenarios with client and server channels. Additional resource saves are available using '`runmqslr`', including a reduced requirement on: virtual memory, number of processes, file handles, and System V IPC.

Fastpath channels, and/or fastpath applications—see later paragraph for further discussion, can increase throughput for both nonpersistent and persistent messaging. For persistent messages the improvement is only for the path through the queue manager, and does not affect performance writing to the log disk. The reader should note that since the greater proportion of time for persistent messages is in the queue manager writing to the log disk, the performance improvement for fastpath channels is less apparent with persistent messages than with nonpersistent messages.

7.2 Application design and configuration

7.2.1 Standard or fastpath?

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

7.2.2 Parallelism, batching, and triggering

An application should be designed wherever possible to have the capability to run *multiple instances* or *multiple threads* of execution. Although the capacity of a multi-processor (SMP) system can be fully utilised with a small number of applications using nonpersistent messages, more applications are 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 batch of persistent messages to the log disk. The performance profile of a workload will also be subject to variability through cycles of low and high message volumes, therefore a degree of experimentation will be required to determine an optimum configuration.

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

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

Processing messages inside syncpoint (i.e. in batches) is 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—outside of syncpoint control, the queue manager must wait for each message to be written to the log before returning control to the application.

A typical triggered application follows the performance profile of a short session (refer to ‘**Short Sessions**’—Page 30). The ‘runmqsr’ has a much smaller overhead of connecting to and disconnecting from the queue manager because it does not have to create a new process. Furthermore, in Version 5.3 the maximum number of connections into a single ‘runmqsr’ listener has been significantly increased making it the preferred method of running short sessions over client channels. Nonetheless, the implementation of triggering is still worth consideration with regard to programming a disconnect interval as an input parameter to the application. This can provide the flexibility to make tuning adjustments in a production environment, if for instance, it is more efficient to remain connected to manager between periods of message processing, or disconnect to free queue manager and Operating System resources.

7.3 Tuning the Operating System (Solaris 5.8)

The System V IPC kernel tuning parameters underlying the measurements in this document were the same ones as for the 'MQSeries for Solaris V5.2 – Capacity Planning Guidance'–**SupportPac 6F**. A summary of the kernel parameters are included in **Table 20** below for reference.

The reader **must** note that these parameters were may not be sufficient for 31,500 *standard* client connections into a single queue manager. Extra resource required by standard connections will be: CPU utilisation, shared memory, and semaphores for the additional agent processes.

| Tuneable Name | Default | Max |
|---------------|---------|---------------|
| shmmax | 131,072 | 2,147,483,647 |
| shmin | 1 | 2,147,482,647 |
| shmmni | 100 | 65,535 |
| shmseg | 6 | 65,535 |
| semmap | 10 | 2,147,482,647 |
| semnmi | 10 | 65,535 |
| semmns | 60 | 65,535 |
| semmnu | 30 | 2,147,482,647 |
| semmsl | 25 | 2,147,482,647 |
| semopm | 10 | 2,147,482,647 |
| semume | 10 | 2,147,482,647 |
| semvmx | 32,767 | 65,535 |
| semaem | 16,384 | 32,767 |
| msgmap | 100 | 2,147,482,647 |
| msgmax | 2,048 | |
| msgmnb | 4,096 | |
| msgmni | 50 | |
| msgssz | 8 | |
| msgtql | 40 | |
| msgseg | 1,024 | 32,767 |

Table 20 – Sizing the kernel tuning parameters for System V IPC

8 Measurement environment

8.1 Workload description

8.1.1 MQI response time tool

The MQI tool exercises the local queue manager by measuring elapsed times of the 8 main MQSeries verbs: MQCONN(X), MQDISC, MQOPEN, MQCLOSE, MQPUT, MQGET, MQCMIT, and MQBACK. The following MQI calls are paired together inside a test application:

- MQCONN(X) and MQDISC,
- MQOPEN and MQCLOSE,
- MQPUT and MQGET,
- MQCMIT and MQBACK with MQPUT and MQGET.

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

Note: performance of MQCMIT and MQBACK is measured in conjunction with MQPUT and MQGET, putting and getting messages inside a unit of work (i.e. inside syncpoint control). The unit of work is committed at the end of each batch. The number of messages per batch is a parameter of the test.

Note: this tool is not used to measure the performance of verbs: MQSET, MQINQ, or MQBEGIN.

8.1.2 Test scenarios workload

8.1.2.1 The driving application programs

The test scenario workload simulates many driving applications running on a single driving machine. This is not typical of a customer environment and is only used to facilitate test coordination. Driving applications were multi-threaded with each thread performing a sequence of MQI calls. The number of threads in each application was adjusted according to whether the test was measuring a local queue manager, a client channel, or distributed queuing scenario. This was done to reduce storage overheads on the driving system. Each driving application thread performed the sequence of actions as outlined in the test scenario illustrations in the '**Performance headlines**' starting on **page 5**.

Message size: For the *release highlights* and *performance headlines* (including rated messaging tests) a 2K message size was used. For the large message measurements a 20K and 200K message size was used.

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

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

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

8.1.2.2 The server application program

The server application is written as a multi-threaded program configured to use 5 threads for processing nonpersistent messages, and 20 or more threads to process persistent messages. Each server thread performed the sequence of actions as outlined in the test scenario illustrations in the '**Performance headlines**' starting on **page 5**.

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

8.2 Hardware

| | |
|--------------------|--|
| Sun Ultra-80 420R: | Server system (device under test) |
| Model: | Z801 0000022003 |
| Processor: | 450MHz UltraSparc-II |
| Architecture: | 4-way SMP |
| Memory (RAM): | 4GB |
| Disk: | 2 Internal Ultra2 SCSI (18.2GB ea. 1 O/S, 1 swap) 2 External Ultra2 SCSI (18.2GB ea. 1 queue, 1 log) |
| Network: | 1GBit Ethernet |
| | |
| Sun Sun-Fire 3800: | Driving applications machine |
| Model: | Z801 0000024977 |
| Processor: | 750MHz UltraSparc-III |
| Architecture: | 6-way SMP |
| Memory (RAM): | 6GB |
| Disk: | 2 Internal Ultra2 SCSI (36GB ea. 1 O/S, 1 swap) 2 External Ultra2 SCSI (36GB ea. 1 queue, 1 log) |
| Network: | 1GBit Ethernet |
| | |
| IBM S80: | Driving applications machine (for maximum server channel pairs test) |
| Model: | 7017-S80 |
| Processor: | 375MHz PowerPC RS64-III |
| Architecture: | 24-way SMP IBM SSA 160 SerialRAID Adapter |
| Memory (RAM): | 32GB |
| Disk: | 2 Internal 16Bit LVD SCSI (9.1GB ea. 1 O/S, 1 O/S + swap) 3 SSA Logical disks (Note: /usr/samples/kernel/vmtune -c 0) (1 Physical SSA160, 9.1GB, 1 swap, 1 queue, 1 log) |
| Network: | 1GBit Ethernet |

8.3 Software

| | |
|--------------|--|
| Solaris O/S: | SunOS Version 5.8 |
| MQSeries: | Version 5.3 (B.11.530.00), and Version 5.2 (B.11.520.00) |
| Compiler: | Sun Workshop Compiler C V5.0 |

9 Glossary

| | |
|---------------|---|
| Test name | <p>The name of the test</p> <p>Note: the test names in some cases are rather long. This is done to provide a descriptive qualification of the test measurement to relate to the performance discussion in the sections throughout the document:</p> <p>local => local queue manager test scenario</p> <p>cl => client channel test scenario</p> <p>dq => distributed queuing test scenario</p> <p>np => nonpersistent messages</p> <p>pm => persistent messages</p> <p>r3600 => 1 round trip per driving application per second</p> <p>runmqtsr => channels using the 'runmqtsr' listener (client channel test scenario, in addition to 'runmqchi' for distributed queuing test scenarios)</p> <p>c6000 => 6,000 client driving applications (i.e. 6,000 MQI-client connections)</p> <p>q1000 => 1,000 server channel pairs</p> <p>max => maximum number of channels (or channel pairs)</p> <p>no_correl_id => correlation identifier not used in the response messages (as each response is placed on a unique reply-to queue per driving application)</p> |
| Apps | The number of driving applications connected to the queue manager at the point where the performance measurement is given |
| Rate/App/hr | The target message throughput rate of each driving application |
| Round T/s | The average achieved message throughput rate of all the driving applications together, measured by the server application |
| % (Round T/s) | The percentage increase in throughput from Version 5.2 to Version 5.3 Note: the nature of the comparison is noted under each table giving percentage improvements. |
| Resp time (s) | The average response time each round trip, as measured and averaged by all the driving applications |
| Swap | The total amount of swap space that is either allocated or reserved |
| shm | The amount of allocated System V IPC shared memory in MB |
| segs | The number of System V IPC shared memory segments |
| sems | The number of System V IPC semaphores |

***** end of document *****