IBM MQ Appliance Performance Report

Version 1.0 - June 2015

Sam Massey
IBM MQ Performance
IBM UK Laboratories
Hursley Park
Winchester
Hampshire

1 Notices

Please take Note!

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

First Edition, June 2015.

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

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 Appliance*. The information is not intended as the specification of any programming interface that is provided by IBM. It is assumed that the reader is familiar with the concepts and operation of IBM MQ Appliance.

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
- **Oracle Corporation :** Java

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.

2 Contents

3 Introduction

The MQ appliance combines all of the core MQ functionality with the convenience, ease of install and simplified maintenance of an appliance.

The MQ Appliance has been primarily introduced to support three types of messaging deployment:

1) Data centre consolidation

   Provides a more powerful singular hardware platform to replace an existing diverse deployment environments

2) Edge of network installation

   Provides a simpler install and configuration for satellite MQ locations (branch/factory/business partner etc.)

3) New MQ application development

   Provides a powerful deployment platform for new MQ deployments

The scenarios featured in this report will represent use cases from the data centre consolidation and edge of network messaging deployments.

There are local disks within the appliance to enable efficient persistent messaging by the local Queue Managers. The two 1.2TB drives are configured in a RAID1 configuration so that data is protected should one of the drives suffer a failure. High Availability (HA) may be achieved by the pairing of two MQ appliances which results in the Queue Manager (QM) log and queue files being distributed across the pair of appliances. An update to this report with HA performance data will be released at a later date. The use of external storage (SAN via fibre connection) is not available in the tested version of the appliance.

The MQ appliance can be purchased in two variants:

Appliance Version	**Enabled cores**
M2000A	All
M2000B	Restricted

The majority of the tests use the all core variant of the MQ Appliance and this is the default hardware unless stated otherwise. A number of tests were also conducted using the restricted core variant and provide comparative data points to the main testing to provide appropriate capacity planning information.

The M2000A and M2000B appliances are supplied with 2x10Gb Ethernet Network links and 8x1Gb Ethernet network links. If the appliances are configured for redundant HA, 1x10Gb link and 2x1Gb links would be reserved for use by the appliance, leaving a total of 16Gb for customer workloads. In Non-HA mode, all 28Gb connectivity can be utilised for workload traffic. There are a further two separate 1Gb links that are explicitly reserved for appliance administration. This report utilises 2 of the 10Gb links for workload traffic.

All of the scenarios featured in this report utilise Request Responder messaging scenarios and the published messaging rate is measured in Round Trips/sec, which involves 2 message puts and 2 message gets. If you are only utilising one-way messaging (using a message sender, queue and message receiver to perform 1 message put and 1 message get), and you can avoid queue-lock contention, then you may achieve up to double the published rates.

The version of the MQ Appliance as tested in this report is M2000A FP3 and where a comparison is made to the restricted appliance configuration, this uses the MQ Appliance M2000B FP3.

4 Scenarios

The scenarios that will be presented in this report reflect the most common usage patterns that customers are anticipated to use with the MQ appliance and provide guidance for those customers performing capacity planning or migration activities.

Each test was initially conducted using a 2K (2048 byte) message size. Additional tests were conducted using 256byte, 20K and 200K to provide further data for capacity planning.

4.1 Data Centre Consolidation

This section looks at how customers might consolidate their existing MQ infrastructure onto the IBM MQ appliance.

As customers replace their existing MQ QM infrastructure, they may consolidate their MQ configuration from separate MQ QM servers (possibly running on different hardware and different MQ Versions) onto a single MQ appliance. They may have a mix of applications tightly bound to their existing QM and also a number of applications that connect using the MQ client API. To migrate to the MQ appliance all applications will need to connect via the MQ client API.

The following tests use MQ client connections and present the performance of MQ as deployed on the Appliance.

4.1.1 Test Scenario C1 – 10 Applications per QM, 1 QM, Non-persistent

The test scenario in Figure 1 is a request-responder scenario that simulates a number of applications that interact with a single QM. A request queue and a reply queue will be created for each application, so ten pairs of queues are created for this test. One or more requester applications will send messages to one of the application request queues and will wait for a reply on the associated reply queue. Responder applications will listen for messages on the request queues before sending them to the correct reply queue.

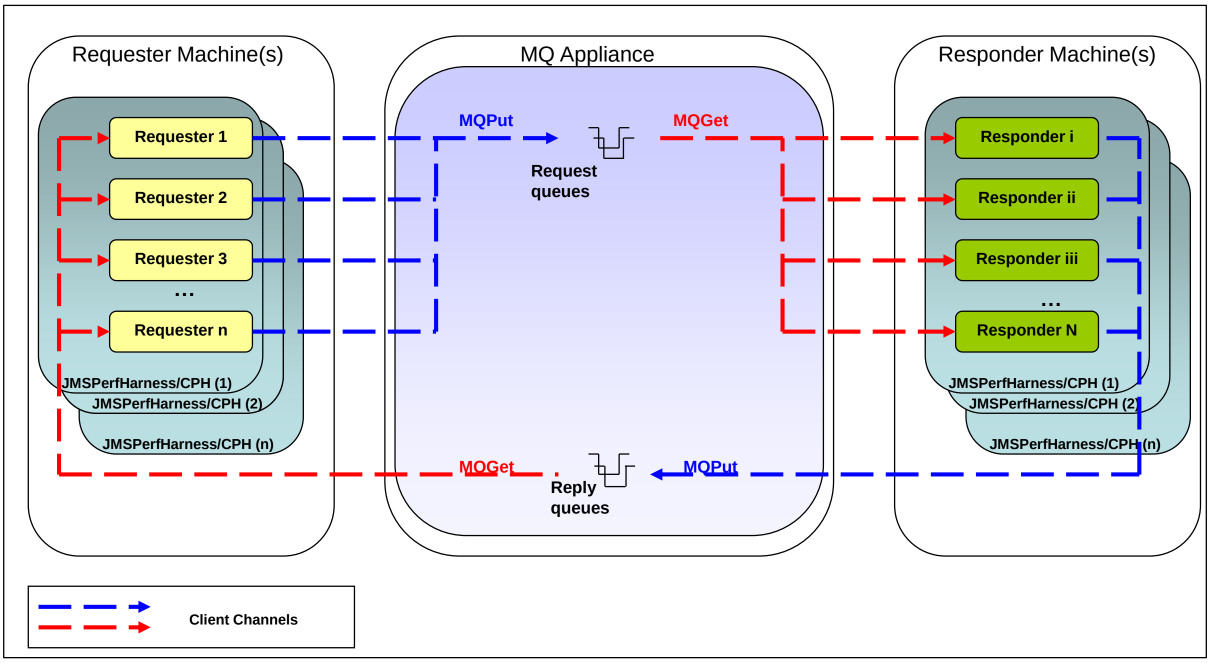

FIGURE 1 - REQUESTER-RESPONDER WITH REMOTE QUEUE MANAGER ON MQ APPLIANCE

Subsequent requester applications will send and receive messages from the set of application queues on a round-robin basis i.e. distributing the messages produced and consumed across the set of application queues.

Results are presented for various numbers of producer threads distributed across the 10 applications (using 10 pairs of queues), 200 fixed responder threads (20 responders per request queue) will send the replies to the appropriate reply queue, and the report will show the message rates achieved (in round trips/second) as the number of producers is increased.

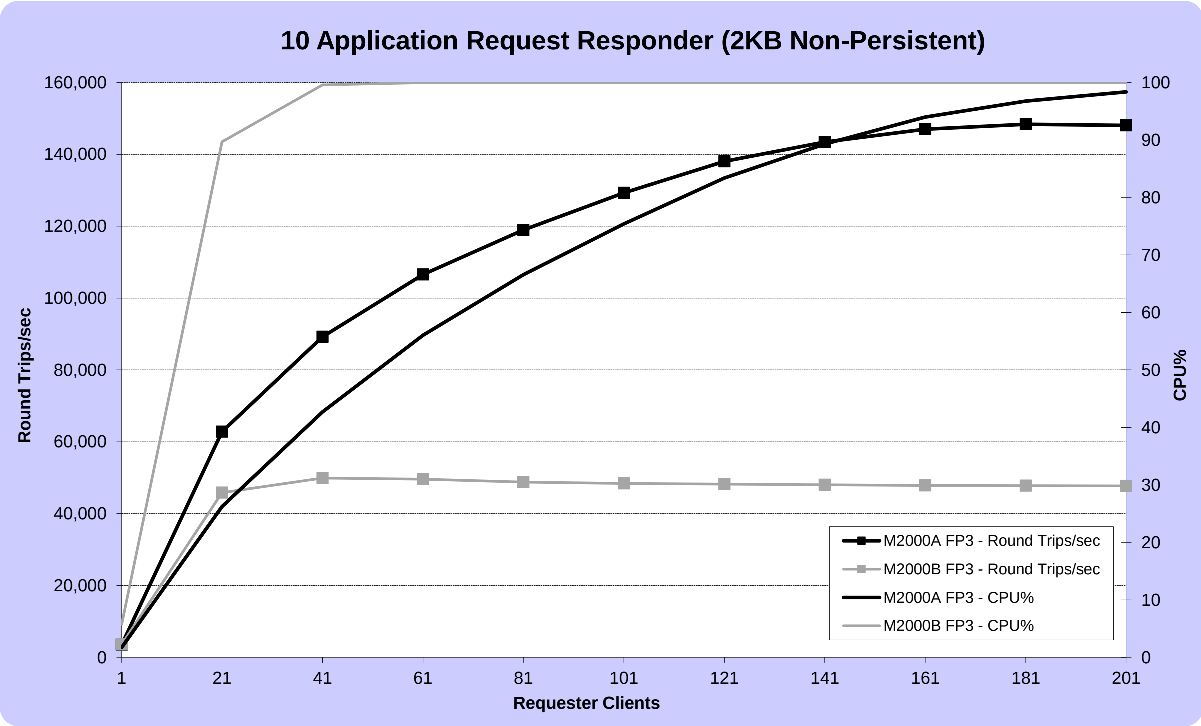

FIGURE 2 – PERFORMANCE RESULTS FOR 2KB NON-PERSISTENT MESSAGING

Figure 2 shows how by increasing the workload on the appliance (by increasing the number of concurrent requester clients), the throughput rate increases until the CPU capacity of the appliance is exhausted. If using a message size of 2KB, the M2000A appliance can achieve approximately 3 times the throughput of the M2000B appliance.

Test	M2000A FP3 Max Rate*	CPU%	Clients	M2000B FP3 Max Rate*	CPU%	Clients
10Q Request Responder (256b Non-persistent)	166,409	96.91	161	57,452	99.66	41
10Q Request Responder (2KB Non-persistent)	148,357	96.76	181	49,891	99.57	41
10Q Request Responder (20KB Non-persistent)	54,124	30.91	51	37,479	99.75	51
10Q Request Responder (200KB Non-persistent)	5,647	14.17	41	5,599	90.38	41

*Round trips/sec

TABLE 1 - PEAK RATES FOR NON-PERSISTENT MESSAGING

4.1.2 Test Scenario C2 – 10 applications per QM, 1 QM, Persistent

This test repeats the test C1 featured in section 4.1.1, but utilises persistent messaging on the appliances local RAID-1 disk subsystem.

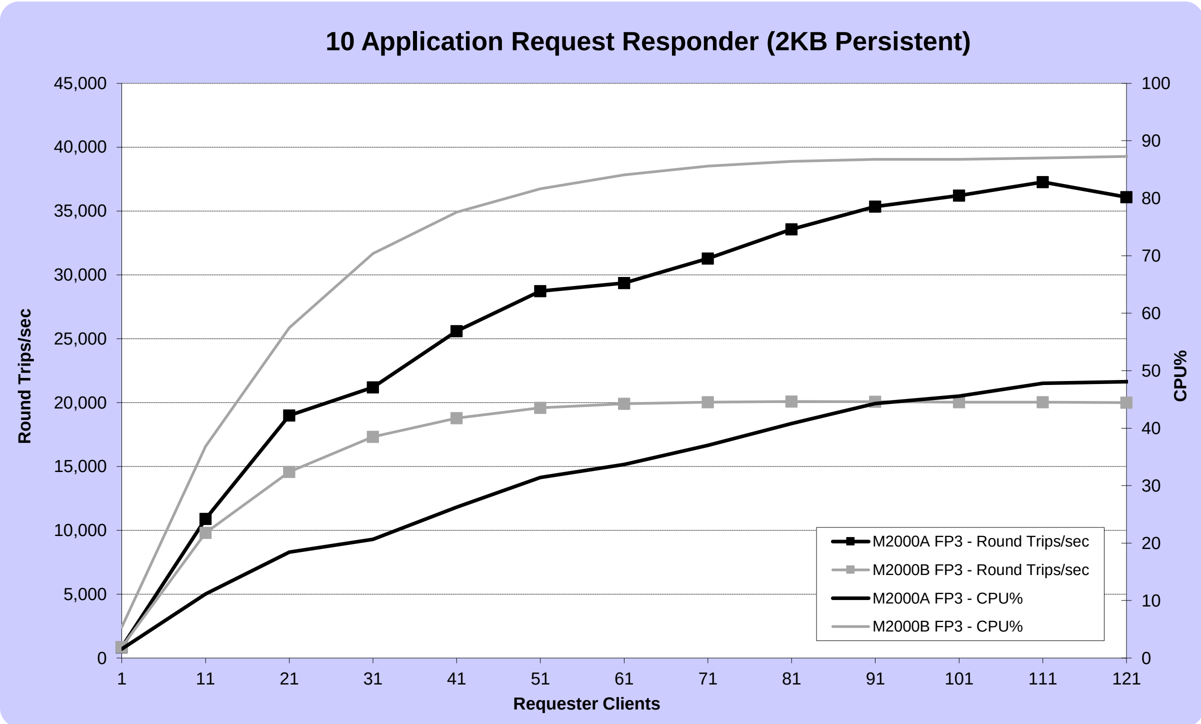

FIGURE 3 - PERFORMANCE RESULTS FOR 2KB PERSISTENT MESSAGING

Figure 3 shows that as the workload increases, a maximum throughput is achieved (~37,000 Round trips/sec for 2KB message size) and the limits of the local disk subsystem have become the limiting factor. There is large capacity to run further Non-persistent messaging workload concurrently on the M2000A appliance.

If using a message size of 2KB, the M2000A appliance can achieve almost 2 times the persistent throughput of the M2000B appliance.

Test	M2000A FP3			M2000B FP3		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
10Q Request Responder (256b Persistent)	52,078	70.78	261	22,907	91.28	101
10Q Request Responder (2KB Persistent)	37,066	47.13	111	20,082	86.4	81
10Q Request Responder (20KB Persistent)	9,855	15.52	51	8,804	47.7	61
10Q Request Responder (200KB Persistent)	1,051	5.57	31	1,009	20.03	21

*Round trips/sec

TABLE 2 – PEAK RATES FOR PERSISTENT MESSAGING

4.1.3 Test Scenario C3 – 10 applications per QM, 10 QM, Non-persistent

This test is equivalent to test C1 featured in section 4.1.1 with 10QM instead of 1QM. A total of 100 applications will be distributed across the 10 QM. This test demonstrates that there are no adverse effects from managing separate QMs within a single appliance.

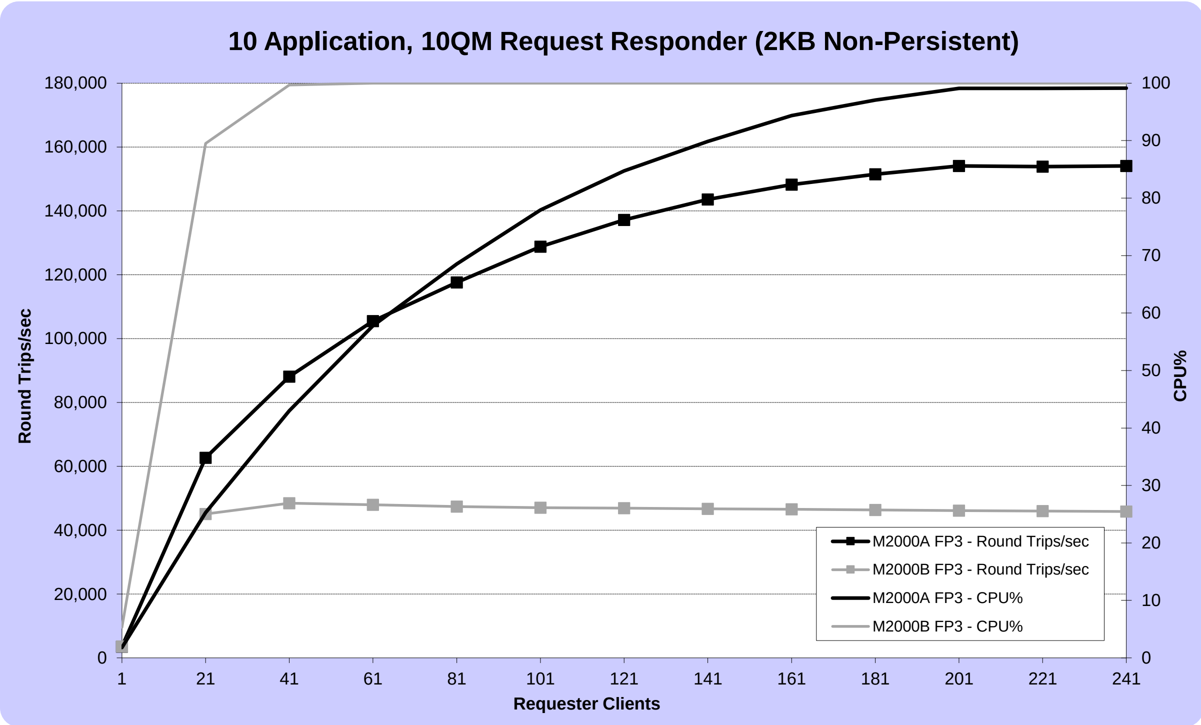

FIGURE 4 - PERFORMANCE RESULTS FOR 2KB, 10QM NON-PERSISTENT MESSAGING

Figure 4 shows that very similar performance can be obtained when running Non-persistent messaging through 10QM as compared with the single QM scenario.

If using a message size of 2KB, the M2000A appliance can achieve over 3 times the throughput of the M2000B appliance.

Test	M2000A FP3			M2000B FP3		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
10Q 10QM Request Responder (256b Non-persistent)	174,870	99.6	201	56,885	99.78	41
10Q 10QM Request Responder (2KB Non-persistent)	154,052	99.08	201	48,451	99.69	41
10Q 10QM Request Responder (20KB Non-persistent)	54,114	31.92	51	36,572	99.29	41
10Q 10QM Request Responder (200KB Non-persistent)	5,645	13.29	41	5,612	83.76	41

*Round trips/sec

TABLE 3 - PEAK RATES FOR 10QM NON-PERSISTENT MESSAGING

4.1.4 Test Scenario C4 – 10 applications per QM, 10 QM, Persistent

This test repeats the test C3 featured in section 4.1.3, but utilises persistent messaging on the appliances local RAID-1 disk subsystem.

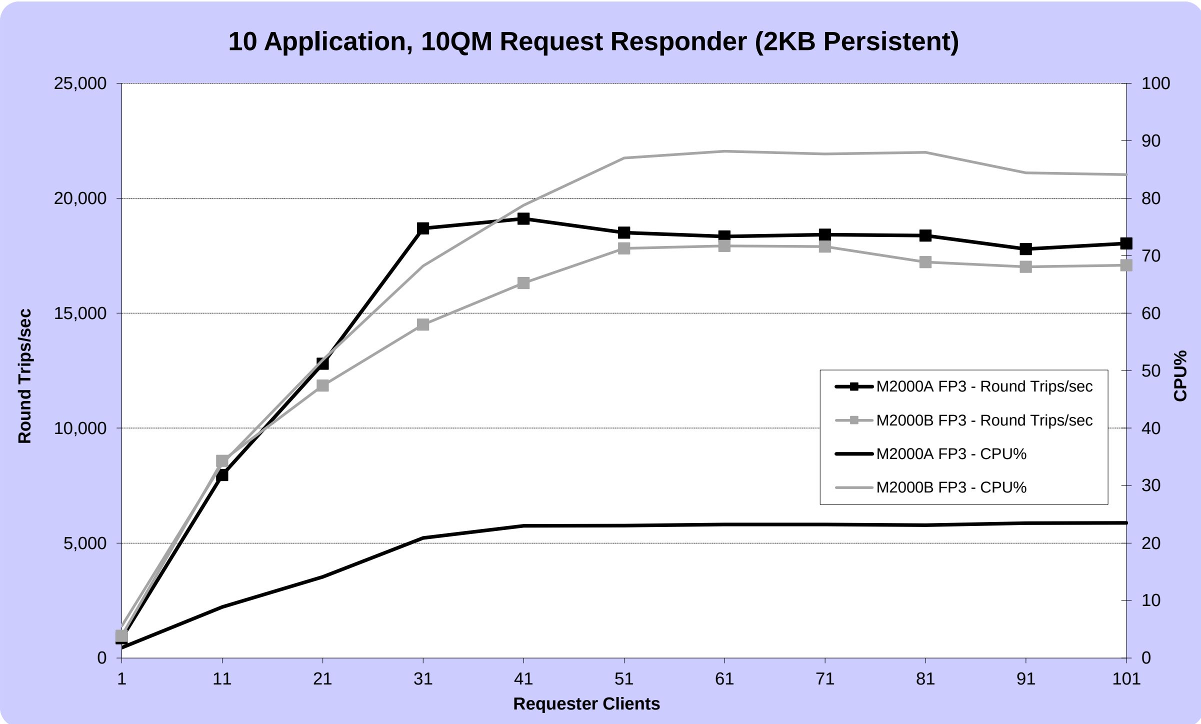

FIGURE 5 - PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING

Figure 5 shows that when we have multiple QM performing persistent messaging, the peak messaging rate obtainable on both M2000A and M2000B appliances is comparable, although there is large capacity to run further Non-persistent messaging workload concurrently on the M2000A appliance.

Test	M2000A FP3			M2000B FP3		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
10Q 10QM Request Responder (256b Persistent)	18,764	23.08	51	18,045	88.91	61
10Q 10QM Request Responder (2KB Persistent)	19,104	22.98	41	17,926	88.15	61
10Q 10QM Request Responder (20KB Persistent)	3,203	4.5	11	3,239	15.66	11
10Q 10QM Request Responder (200KB Persistent)	343	1.68	11	335	4.8	1

*Round trips/sec

TABLE 4- PEAK RATES FOR 10QM PERSISTENT MESSAGING

4.2 Edge of Network

This scenario demonstrates the use of the appliance to distribute messages from the appliance to other remote QMs using distributed queuing techniques. The test utilises client connections to the appliance and then a pair of sender and receiver channels per application to distribute the messages from within a single QM on the appliance to a remote QM on an xLinux host.

4.2.1 Test Scenario E1 – 10 applications, 10 pairs of server channels

This is a distributed queuing version of the requester-responder topology detailed in section 4.1.1. All message puts are to remote queues so that messages are now transported across server channels to the queue manager, where the queue is hosted.

Remote queue and alias definitions will be created for each application. One or more requesters will send messages to its associated request queue, and will wait for a reply on the defined reply queue. Responder applications will listen for messages on the request queues, before sending them to the correct reply queue.

The test scenario in Figure 6 shows the topology of the test. The requester applications use client bindings to send/receive message to the appliance which utilises remote queue and server channel definitions to forward the messages to the remote QM using a set of transmission queues (1 per application). The responder applications are locally bound.

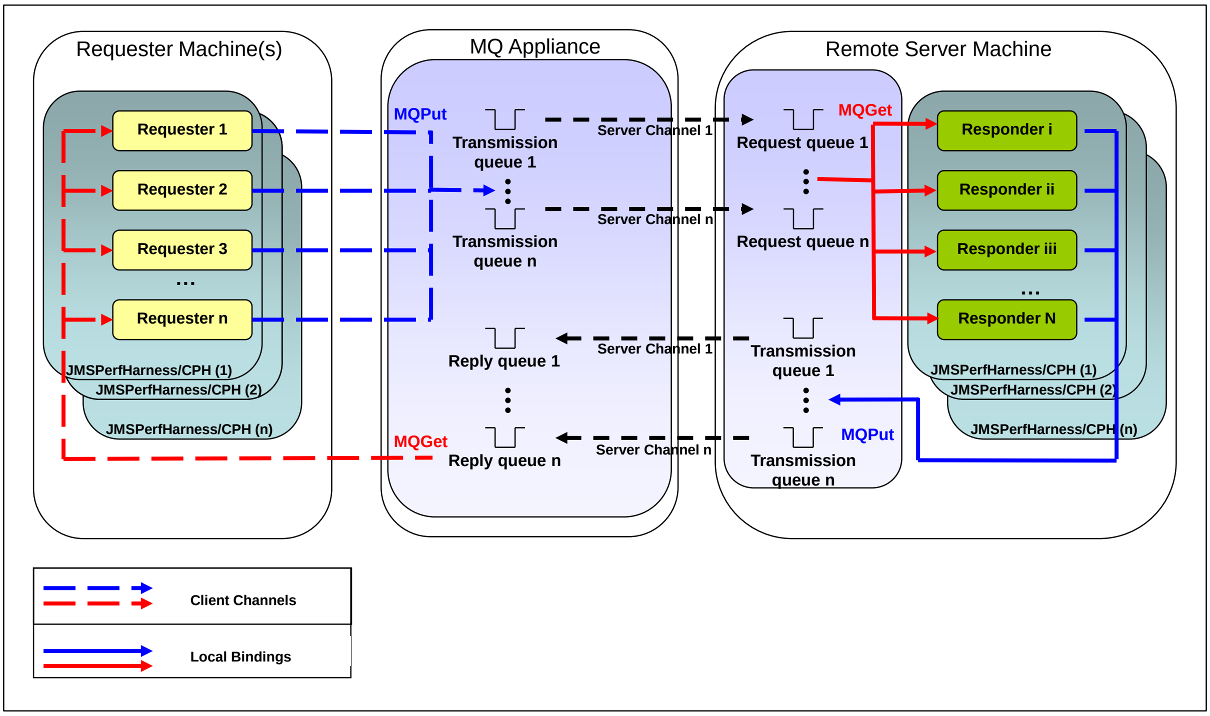

FIGURE 6 - REQUESTER-RESPONDER WITH DISTRIBUTED QUEUEING

Results are presented for a various number of producer threads distributed across the 10 applications (using 10 pairs of queues), 200 responder threads (20 responders per request queue), and will show the message rate achieved (in round trips/second) as the number of producers is increased.

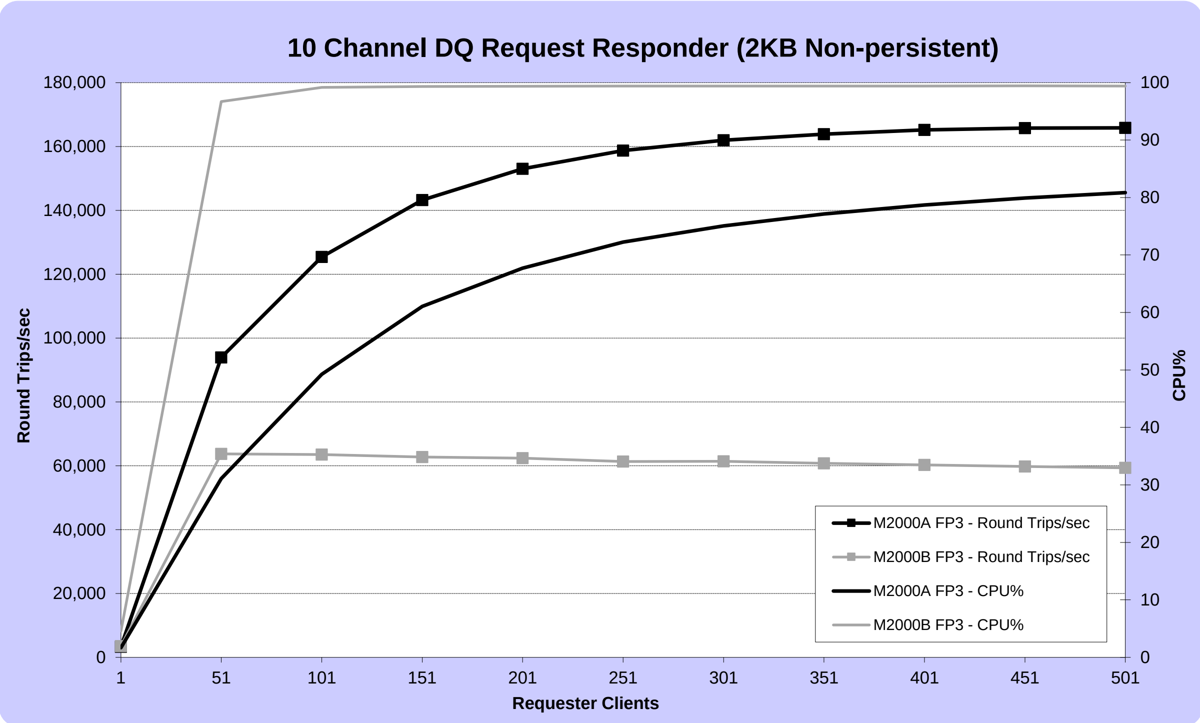

FIGURE 7 - PERFORMANCE RESULTS FOR 2KB, 10 CHANNEL DQ NON-PERSISTENT MESSAGING

Figure 7 shows that a similar peak level of performance can be obtained when using distributed queueing to that measured with the single QM scenario.

If using a message size of 2KB, the M2000A appliance can achieve over 2.5 times the throughput of the M2000B appliance.

Test	M2000A FP3			M2000B FP3		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
10 Channel DQ Request Responder (256b Non-persistent)	189,083	79.94	451	75,006	98.57	51
10 Channel DQ Request Responder (2KB Non-persistent)	165,789	80.85	501	63,679	96.69	51
10 Channel DQ Request Responder (20KB Non-persistent	54,064	24.83	61	44,422	96.69	61
10 Channel DQ Request Responder (200KB Non-persister	5,656	11.41	31	5,650	76.01	41

*Round trips/sec

TABLE 5 - PEAK RATES FOR 10 CHANNEL DQ NON-PERSISTENT MESSAGING

4.2.2 Test Scenario E2 – 10 applications, 10 pairs of server channels, Persistent

This test repeats the test E1 featured in section 4.2.1, but utilises persistent messaging on the appliances local RAID-1 disk subsystem.

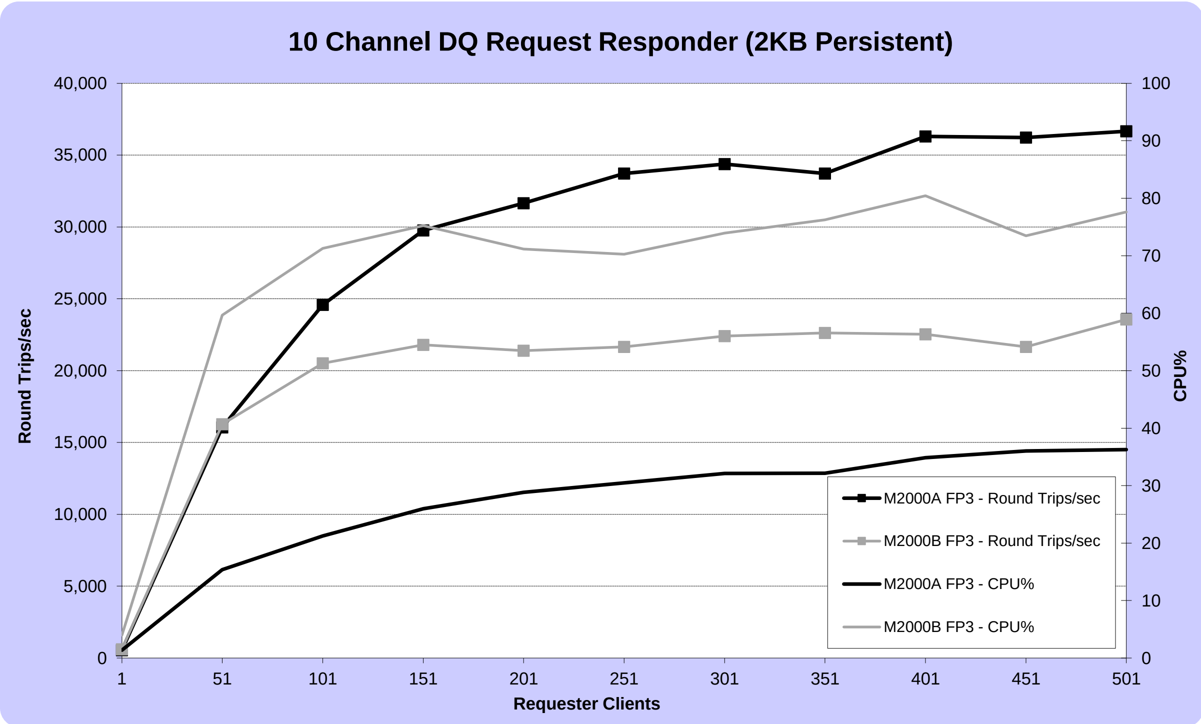

FIGURE 8 - PERFORMANCE RESULTS FOR 2KB, 10 CHANNEL DQ PERSISTENT MESSAGING

Figure 8 shows that as the workload increases, a maximum throughput is achieved (~37,000 RT/sec for 2KB message size) and as with test C2 in section 4.1.2, the limits of the local disk subsystem have become the limiting factor. There is large capacity to run further Non-persistent messaging workload concurrently on the M2000A appliance.

Test	M2000A FP3			M2000B FP3		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
10 Channel DQ Request Responder (256b Persistent)	50,373	47.56	501	29,335	91.75	401
10 Channel DQ Request Responder (2KB Persistent)	36,663	36.26	501	23,558	77.6	501
10 Channel DQ Request Responder (20KB Persistent)	9,545	11.45	81	9,538	44.91	101
10 Channel DQ Request Responder (200KB Persistent)	1,029	5.08	21	1,020	21.77	41

*Round trips/sec

TABLE 6 - PEAK RATES FOR 10 CHANNEL DQ PERSISTENT MESSAGING

5 Connection Scaling

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 throughout, these tests define an operational environment or scaling challenge to test from a performance perspective.

5.1 Connection Test

This test uses the requester/responder workload as described in section 4.1.1. The requester applications are rated at 1 message every 100 seconds and 60,000 client bound requester applications are connected as fast as possible to determine the overall connection time for those clients to the MQ Appliance.

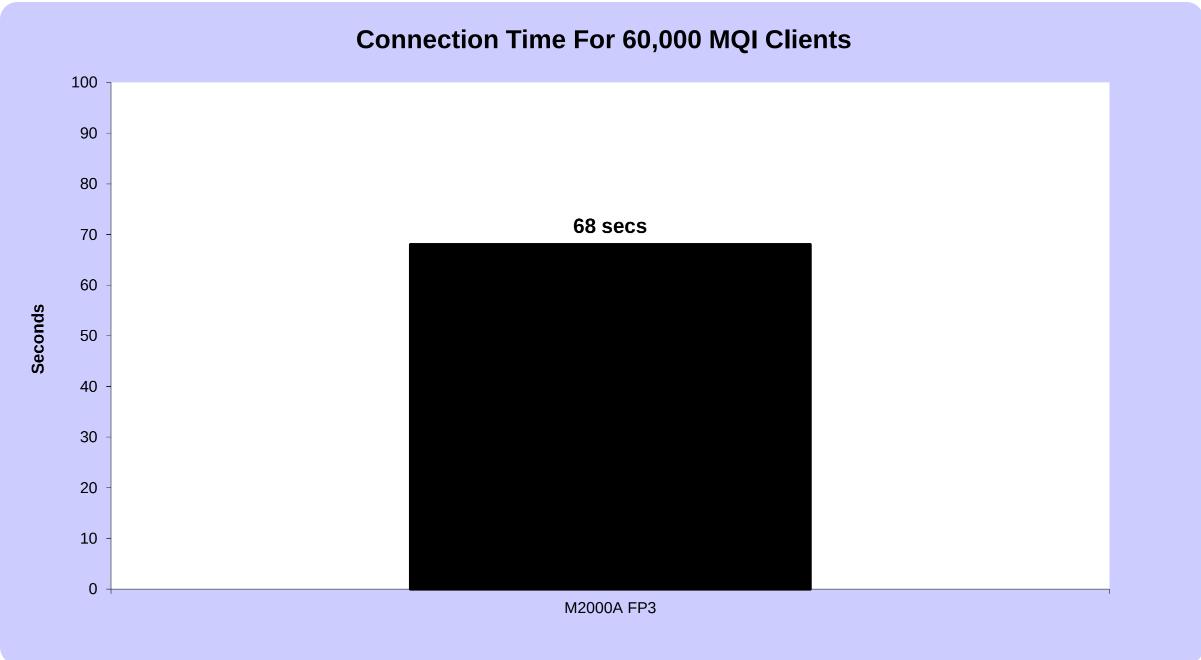

FIGURE 9 – PERFORMANCE RESULTS FOR MQI CLIENT CONNECTION TEST

There have been significant improvements in MQ with regard to the rate at which the queue manager can accept connections from client bound applications.

6 Frequently Asked Questions

Will I be able to use FASTPATH channels to send/receive messages into the MQ Appliance?

Yes - this is now the default MQIBindType as specified in the Channels stanza in the qm.ini configuration file.

How do I view and change QM settings on the MQ Appliance?

You can use the *dspmqini* command to view the QM configuration and *setmqini* to alter any configuration options. There are similar *dspmqvar* and *setmqvar* commands to view/alter environment variables.

What type of logging is used on the MQ Appliance?

Only circular logging is supported on the MQ Appliance, and thus there are no facilities to monitor/prune QM logs.

Can I run my existing user exits?

No – for appliance integrity, user exits will not be supported on the MQ Appliance. Many historic reasons for using code exits have now been resolved by product features.

What is throttling my messaging scenario?

If customers experience throttled performance when driving high throughput workloads on M2000A, they should check the following:

- Persistent workloads - Customers might encounter the disk limits as illustrated in this document
- Larger message (10K+) Non-persistent workloads - Customers might encounter network limits (especially if Ethernet ports are reserved for HA)
- Small message (2K-) Non-persistent workloads – Customers might encounter CPU saturation (Check MQ Console)

7 Appendix A – Client machine specification

The client machines used for the performance tests in this report have the following specification:

Category	Value
Machine	x3550 M5
OS	Red Hat Enterprise Linux Server 6.6
CPU	2x12 (2.6Ghz)
RAM	128GB RAM
Network	2x10Gb Ethernet
Disks	2x 300GB 15K SAS HDD
RAID	ServeRAID M5210 (4GB Flash RAID cache) MQ Logs hosted on RAID-1 partition