

IBM MQ Appliance HA and DR Performance Report

Version 2.0 - July 2016

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, December 2015.

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

© Copyright International Business Machines Corporation 2015, 2016. 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.

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

1	Notices	2
2	Contents	4
3	Report Highlights	5
4	HA Introduction	5
4.1	Comparison of M2001A to M2000A	8
4.2	Test Scenario HA1 – 10 Applications per QM, 1 QM, Persistent.....	9
4.3	Test Scenario HA2 – 10 applications per QM, 10 QM, Persistent	10
4.4	How does HA perform over larger distances?	11
4.5	How long does it take before the QM on the standby appliance becomes available?	14
5	DR Introduction	15
5.1	Test Scenario DR1 – 10 Applications per QM, 1 QM, Persistent.....	17
5.2	Test Scenario DR2 – 10 Applications per QM, 10 QM, Persistent.....	18
5.3	How does DR perform over larger distances?.....	19
6	HA and DR Introduction.....	20
7	Frequently Asked Questions.....	21
8	Appendix A – Client machine specification.....	23

3 Report Highlights

This report contains data points that illustrate the performance of the HA and DR functions of the IBM MQ Appliance. It is worth noting the following highlights:

- Over 42,000 round trips/second peak messaging rate achieved in HA enabled scenario (~84,000 messages produced and ~84,000 messages consumed). See section 4.2.
- Over 2x increase in peak messaging rate for very large message (200K+) in HA scenarios when compared with M2000A. See section 4.1.
- Performing a manual failover with 500,000 queued messages between a pair of HA appliances took only 11 seconds. See section 4.5.
- DR performance degrades by less than 25% across all message sizes when the latency between the paired appliances is under 50ms (compared with the directly connected DR scenario). See section 5.3.

4 HA Scenarios

This performance report at version 2.0 now contains updated performance data based on the MQ Appliance M2001A and M2001B.

M2001 introduces a hardware update to two components:

- The disk subsystem has been upgraded from a pair of 1.2TB HDD disks to a pair of 3.2TB SSD disks
- The 10Gb Ethernet subsystem has been upgraded from a 2 port 10Gb module to a 4 port 10Gb module

Section 4.1 has been added to illustrate the performance benefit provided by the SSD disks in M2001A in comparison to M2000A.

This report accompanies the previously released MQ Appliance Performance report which has also been updated for the M2001 appliances and is available at the following locations:

<http://www-01.ibm.com/support/docview.wss?uid=swg24040125>

<https://ibm-messaging.github.io/mqperf/>

High Availability (HA) can be enabled by pairing two MQ Appliances together to provide continuous availability in the event of one of the appliances suffers a failure. The Queue Manager (QM) log and queue files are synchronously replicated across the pair of appliances.

If separate networks and switches (if required) are used to connect the pair of appliances, then the pair can also continue to operate in the event of a partial network outage.

To ensure clients reconnect to the QM on either of the pair of appliances, the clients should be made aware of the IP addresses assigned to the workload interfaces of both appliances; or a Virtualised IP address in the case that a suitable load balancer component is employed.

To illustrate the performance profile of enabling the HA infrastructure, tests will be performed on two of the scenarios featured in the base MQ Appliance performance report.

- 1) Request Responder 1QM Persistent (Test C2)
- 2) Request Responder 10QM Persistent (Test C4)

Each test will be conducted with both a standalone QM and a QM incorporated into an appliance HA group, so that the cost of the synchronous replication can be evaluated.

As you will see in the results, there is a large difference in performance if you compare the 1QM with the 10QM results (both HA and non HA). This is due to the fact that for the single QM test, all of the circular logs can be held within the RAID cache (backed by battery) and thus there is no wait for acknowledgement by the IO subsystem. In most deployments where multiple queue managers (or single QM with a large set of log files) are employed, we would encourage you to use the data from the 10QM scenario to accurately perform your capacity planning.

The M2001A and M2001B appliances are supplied with 4x10Gb Ethernet Network links and 8x1Gb Ethernet network links. When the appliances are configured for redundant HA, 1x10Gb link and 2x1Gb links are reserved for use by the appliance, leaving a total of 36Gb for customer workloads. In Non-HA mode, all 48Gb connectivity can be utilised for workload traffic. There are a further two separate 1Gb links that are explicitly reserved for appliance administration. This section utilises the following connections:

Appliance A	Appliance B	Notes
eth13	eth13	Connected directly between appliances with 1Gb copper patch cable
eth17	eth17	Connected directly between appliances with 1Gb copper patch cable
eth20	eth20	Unused
eth21	eth21	Connected directly between appliances with 10Gb copper cable
eth22		Workload driven via this interface
eth23		Workload driven via this interface

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 section is M2001A 8.0.0.5+APAR IT 15823 (which we will refer to subsequently as M2001A).

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 looking to evaluate the performance of the MQ Appliance with HA enabled.

Each test was conducted using a 2K (2048 byte) message size and this data is shown in the graphs included below. Additional tests were conducted using 256byte, 20K and 200K to provide further data points.

4.1 Comparison of M2001A to M2000A

The upgrade of the disks from 1.2TB HDD to 3.2TB SSD has provided significant performance and capacity benefits. The largest increases in performance are found in situations which heavily utilise the IO subsystem of the MQ Appliance i.e. scenarios with a requirement to support a high persistent message rate.

The scenario from section 4.3 is used to compare the performance of M2001A with M2000A and the graph below show the performance of the 200K message size on each version of the appliance.

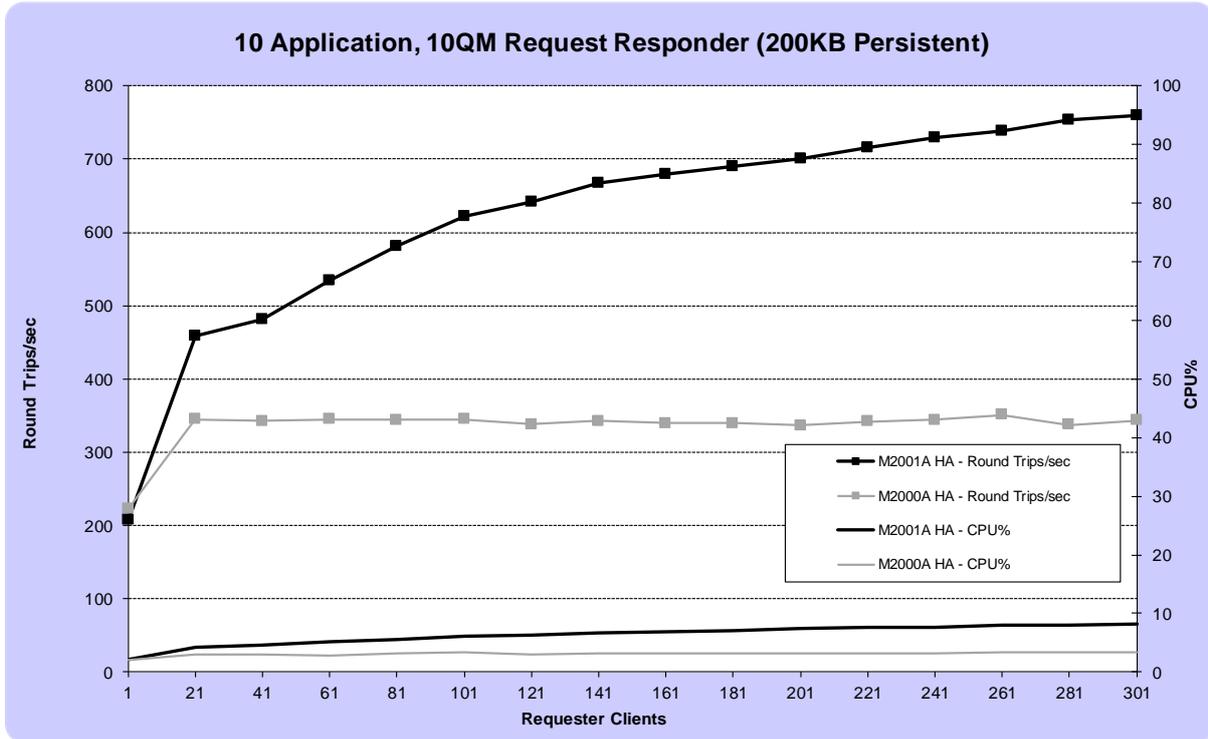


FIGURE 1 - PERFORMANCE RESULTS FOR 200KB, 10QM PERSISTENT MESSAGING, M2001A VS M2000A

Figure 1 shows that there is more than 2x performance improvement when using the 200K message size. At smaller message sizes, the demonstrated improvement is less because the IO subsystem limits have not been reached. Further additional clients or message batching would demonstrate a larger differential between the two appliances.

Test	M2001A HA			M2000A HA		
	Max Rate*	CPU%	Clients	Max Rate*	CPU%	Clients
10Q 10QM Request Responder (256b Persistent)	14,657	15.97	301	12,448	15.33	301
10Q 10QM Request Responder (2KB Persistent)	12,908	15.37	301	11,876	13.91	301
10Q 10QM Request Responder (20KB Persistent)	4,918	9.81	301	3,342	6.04	301
10Q 10QM Request Responder (200KB Persistent)	760	8.11	301	351	3.3	281

*Round trips/sec

TABLE 1 - PEAK RATES FOR 10QM PERSISTENT MESSAGING, M2001A VS M2000A

4.2 Test Scenario HA1 – 10 Applications per QM, 1 QM, Persistent

This test is identical to test C2 in the MQ Appliance performance report and is presented here with results from running tests against a standalone QM and also against a QM that is included in an HA group.

Results are presented for various numbers of requester 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 requesters is increased.

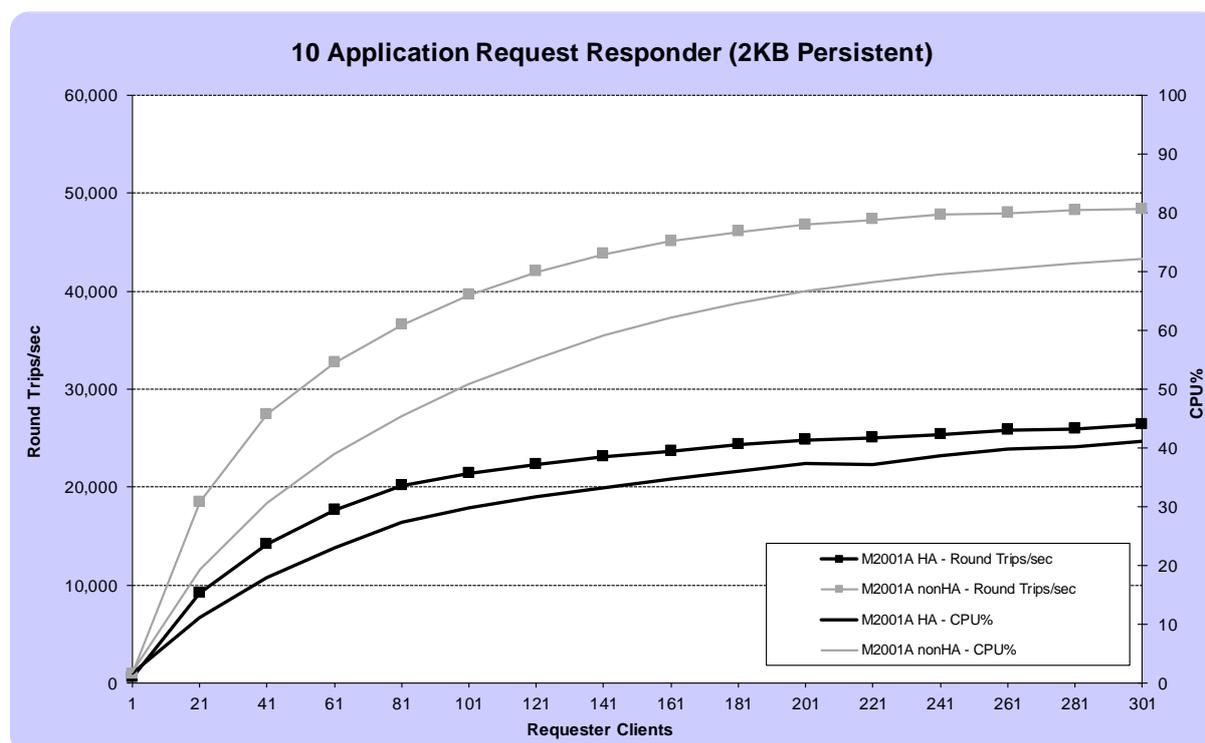


FIGURE 2 – PERFORMANCE RESULTS FOR 2KB PERSISTENT MESSAGING

Figure 2 shows that by enabling HA capability, the maximum throughput achieved in this single queue manager scenario with a 2K message size is approximately halved. There is a similar reduction in CPU utilisation, thus providing capacity to perform additional messaging on the appliance.

Test	M2001A HA				M2001A nonHA			
	Max Rate*	CPU%	Clients	Latency#	Max Rate*	CPU%	Clients	Latency#
10Q Request Responder (256b Persistent)	42,223	56.76	301	1.7	57,015	78.81	261	1.1
10Q Request Responder (2KB Persistent)	26,411	41.15	301	1.8	48,379	72.26	301	1.2
10Q Request Responder (20KB Persistent)	6,083	13.72	221	2.3	16,866	38.7	301	1.4
10Q Request Responder (200KB Persistent)	655	7.44	61	4.7	1,763	16.54	61	3.0

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 2 - PEAK RATES FOR PERSISTENT MESSAGING

4.3 Test Scenario HA2 – 10 applications per QM, 10 QM, Persistent

This test repeats the test C4 in the MQ Appliance performance report and is presented here with results from running tests against a standalone set of Queue Managers and also against a set of Queue Managers that are included in an HA group.

Results are presented for various numbers of requester threads distributed across the 10 Queue Managers who each host 10 pairs of queues (representing 10 applications per QM), 200 fixed responder threads (2 responders per request queue) will send the replies to the appropriate reply queue which are subsequently received by the originating requester threads, and the report will show the message rates achieved (in round trips/second) as the number of requesters is increased.

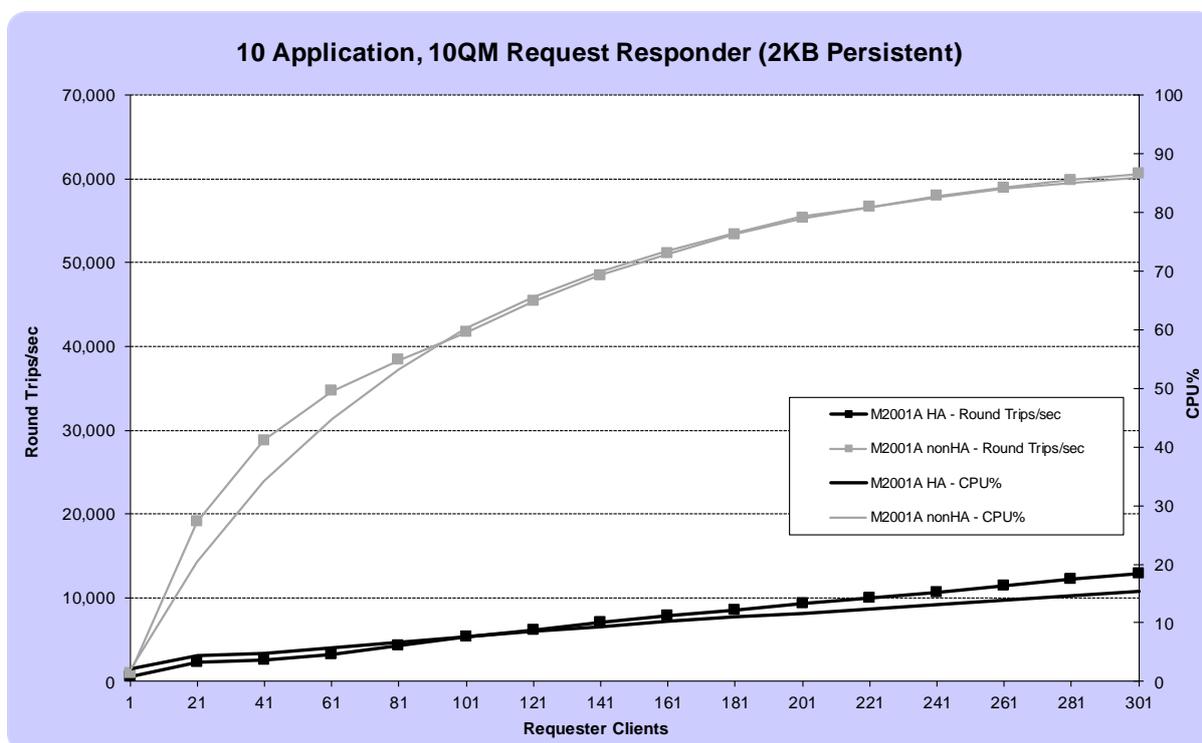


FIGURE 3 - PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING

Figure 3 shows that when we have multiple QM performing 2KB persistent messaging across a pair of HA appliances, the messaging rate is approximately 80% less than when distributed across a set of non HA Queue Managers. At larger message sizes, there is less impact on the maximum message rate, although more clients are required to be connected to achieve the maximum rate.

Test	M2001A HA				M2001A nonHA			
	Max Rate*	CPU%	Clients	Latency#	Max Rate*	CPU%	Clients	Latency#
10Q 10QM Request Responder (256b Persistent)	14,657	15.97	301	1.4	75,798	93.72	301	1.0
10Q 10QM Request Responder (2KB Persistent)	12,908	15.37	301	1.7	60,626	85.93	301	1.0
10Q 10QM Request Responder (20KB Persistent)	4,918	9.81	301	2.2	12,072	18.44	31	1.5
10Q 10QM Request Responder (200KB Persistent)	760	8.11	301	4.8	1,269	8.59	31	2.9

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 3 - PEAK RATES FOR 10QM PERSISTENT MESSAGING

4.4 How does HA perform over larger distances?

The previous section shows how the MQ appliance HA capability might perform if both appliances were located in the same data centre (i.e. 3m distance between the appliances). How would the HA performance differ if the pair of appliances were located a larger distance apart? Due to testing limitations, we need to simulate the latency that might be experienced as the distances between the appliances grows.

If the appliances are located 100Km apart, you might expect the smallest increase in packet transmission latency for each leg to be calculated as follows:

$$\text{distance} / \text{speed} = \text{time}$$

$$100,000\text{m} / 300,000,000\text{m/s}^a = 0.000333\text{s} = 333 \text{ microseconds}$$

There must also be an allowance for the refraction index of the cable

$$333 * 1.5 = 500 \text{ microseconds}$$

Switching hardware and non-linear cable routing will likely further increase the latency between the pair of HA appliances. It is currently advised to customers to site a pair of HA appliances so that the latency between the two appliances is no greater than 10ms and preferably within the same data center.

A delay can be inserted into the sending network layer of both appliances to simulate such latency and let us examine how this impacts the HA performance. The following chart repeats the test from section 4.3 and shows the effect of a 2ms round trip latency introduced into the network layer between the two HA appliances.

^a Assuming speed of light to be $3 \times 10^8 \text{m/s}$

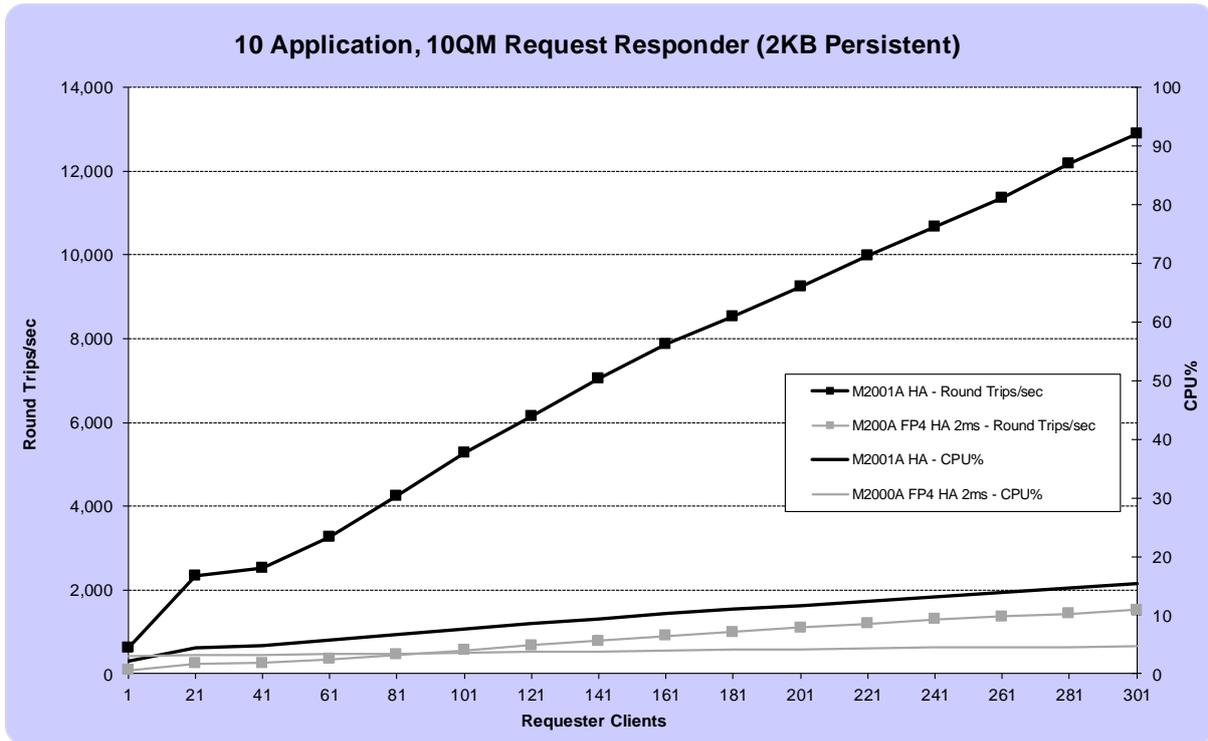


FIGURE 4 - PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING WITH/WITHOUT 2MS LATENCY

Figure 4 shows that an additional 2ms latency on the round trip time of the HA replication interface results in a ~88% reduction in performance than compared with the direct connection (no additional latency) between the appliances.

Test	M2001A HA				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 2ms Latency (256b Persistent)	1,519	2.27	301	9.8	10.4%
10Q 10QM Request Responder 2ms Latency (2KB Persistent)	1,494	2.43	301	11.7	11.6%
10Q 10QM Request Responder 2ms Latency (20KB Persistent)	1,281	2.98	301	12.4	26.1%
10Q 10QM Request Responder 2ms Latency (200KB Persistent)	555	5.75	301	14.9	73.0%

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 4 - PEAK RATES FOR 10QM PERSISTENT MESSAGING WITH 2MS SIMULATED LATENCY

The data in the following tables show additional data points with simulated latency delays of 1, 5 and 10ms.

Test	M2001A HA				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 1ms Latency (256b Persistent)	2,763	3.73	301	5.7	18.9%
10Q 10QM Request Responder 1ms Latency (2KB Persistent)	2,699	3.87	301	6.7	20.9%
10Q 10QM Request Responder 1ms Latency (20KB Persistent)	2,046	4.47	301	7.2	41.6%
10Q 10QM Request Responder 1ms Latency (200KB Persistent)	634	6.68	301	9.7	83.5%

*Round trips/sec

#Single thread round trip latency (ms)

Test	M2001A HA				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 5ms Latency (256b Persistent)	674	0.9	301	21.9	4.6%
10Q 10QM Request Responder 5ms Latency (2KB Persistent)	668	1.42	301	26.5	5.2%
10Q 10QM Request Responder 5ms Latency (20KB Persistent)	633	1.8	301	27.9	12.9%
10Q 10QM Request Responder 5ms Latency (200KB Persistent)	522	6.07	301	30.5	68.7%

*Round trips/sec

#Single thread round trip latency (ms)

Test	M2001A HA				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 10ms Latency (256b Persistent)	419	0.43	301	42.2	2.9%
10Q 10QM Request Responder 10ms Latency (2KB Persistent)	424	0.79	301	50.8	3.3%
10Q 10QM Request Responder 10ms Latency (20KB Persistent)	429	0.57	301	53.6	8.7%
10Q 10QM Request Responder 10ms Latency (200KB Persistent)	386	5.48	301	56.0	50.8%

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 5 - PEAK RATES FOR 10QM PERSISTENT MESSAGING WITH 1, 5 AND 10MS SIMULATED LATENCY

4.5 How long does it take before the QM on the standby appliance becomes available?

There are a number of factors to consider when evaluating the time it takes for a client to failover to the second appliance:

1. Detection by the MQ HA group that the main appliance is no longer responding
2. Making the QM on second appliance available, this involves replaying the transaction log to bring the queue files up to date (this will depend on the current persistent messaging rate, the time since last log checkpoint and the queue depth)
3. Time for client to notice disconnection and reconnect to second appliance

In our testing where a manual failover scenario is triggered, we are only examining the time spent in (2), because triggering a manual failover doesn't use the appliance heartbeating(1) functionality to detect loss of communication between the pair of appliances, and additionally all clients are informed to reconnect so do not have to wait for (3).

Our testing client was configured to attempt to reconnect to the QM on the second appliance when any errors were received whilst communicating with the QM on the main appliance. The location of the active QM was switched to the standby appliance by executing the **sethapreferred** command at the standby appliance. The time taken before a connection to the QM on the second appliance was recorded and can be seen in Table below.

M2001A HA		
Message size	Queue Depth	Failover time (s)
N/A	0	6
2K	500,000	11
20K	100,000	15
200K	25,000	12

TABLE 6 - FAILOVER TIMES FOR SELECTED MSG SIZES AND QUEUE DEPTHS

5 DR Scenarios

From Fixpack 4, the ability to configure a Queue Manager on a recovery appliance to perform a Disaster Recovery (DR) role is provided. This configuration allows the Queue Manager on the recovery appliance to resume work should an outage occur that results in the main appliance becoming unavailable.

When DR was introduced, you could choose to have a queue manager as part of a disaster recovery configuration, or it could be configured to be a member of a high availability group, but not both; this restriction has now been removed and you can now configure an HA and DR enabled queue manager. The performance of this configuration will be examined in section 6, whilst in this section we will look at the standalone performance of DR.

The Queue Manager data is replicated asynchronously to the recovery appliance, which can result in messaging data loss (up to a maximum of 4MB per QM is held in the TCP send buffer) should the main appliance become unavailable. The Queue Manager at the recovery appliance must be manually started before it can start accepting connections from clients.

To illustrate the cost of enabling the DR infrastructure, tests will be performed on two of the scenarios featured in the base MQ Appliance performance report.

- 1) Request Responder 1QM Persistent (Test C2)
- 2) Request Responder 10QM Persistent (Test C4)

Each test will be conducted with both a standalone QM and a QM configured with a remote DR appliance, so that the cost of the asynchronous replication can be evaluated.

The M2001A and M2001B appliances are supplied with 4x10Gb Ethernet Network links and 8x1Gb Ethernet network links. When the appliances are configured for Disaster Recovery(DR), 1x10Gb link is reserved for use by the appliance, leaving a total of 38Gb for customer workloads. When the appliances are configured for HA and DR, 2x1Gb links and 2x10Gb links are reserved for use by appliance, leaving a total of 26Gb for customer workloads. In Non-DR mode, all 48Gb connectivity can be utilised for workload traffic. There are a further two separate 1Gb links that are explicitly reserved for appliance administration. This section utilises the following connections:

Appliance A	Appliance B	Notes
eth13	eth13	Connected directly between appliances with 1Gb copper patch cable. Used in section 6 only
eth17	eth17	Connected directly between appliances with 1Gb copper patch cable. Used in section 6 only

eth20	eth20	Connected directly between appliances with 10Gb fibre cable in this section for DR. Connected via switch to Appliance C in section 6
eth21	eth21	Connected directly between HA appliances with 10Gb copper cable. Used in section 6 only
eth22		Workload driven via this interface
eth23		Workload driven via this interface

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 section is M2001A 8.0.0.5+APAR IT 15823 (which we will refer to subsequently as M2001A).

5.1 Test Scenario DR1 – 10 Applications per QM, 1 QM, Persistent

This test is identical to the test in Section 4.1 and is presented here with results from running tests against a standalone QM and also against a QM that is configured for Disaster Recovery (although the recovery appliance is located 3m from the main appliance).

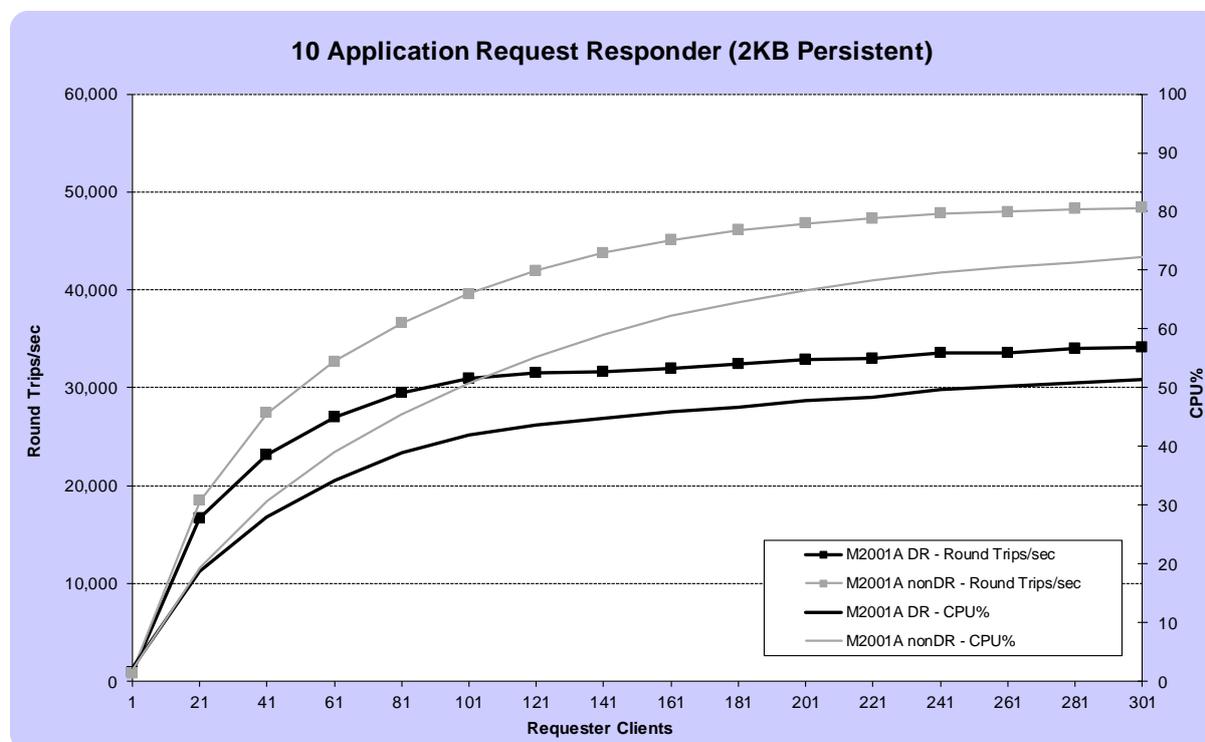


FIGURE 5 – PERFORMANCE RESULTS FOR 2KB PERSISTENT MESSAGING

Figure 5 shows that by enabling DR capability, the maximum throughput achieved with a 2K message size on a single Queue Manager is reduced by up to 30%. There is a similar reduction in CPU utilisation, thus providing capacity to perform additional messaging on the appliance.

Test	M2001A DR				M2001A nonDR			
	Max Rate*	CPU%	Clients	Latency#	Max Rate*	CPU%	Clients	Latency#
10Q Request Responder (256b Persistent)	52,948	73.19	301	1.0	57,015	78.81	261	1.1
10Q Request Responder (2KB Persistent)	34,164	51.32	301	1.1	48,379	72.26	301	1.2
10Q Request Responder (20KB Persistent)	7,243	17.73	241	1.4	16,866	38.7	301	1.4
10Q Request Responder (200KB Persistent)	793	8.88	51	3.5	1,763	16.54	61	3.0

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 7 - PEAK RATES FOR PERSISTENT MESSAGING

5.2 Test Scenario DR2 – 10 Applications per QM, 10 QM, Persistent

This test is identical to the test in Section 4.3 and is presented here with results from running tests against a standalone QM and also against a QM that is configured for Disaster Recovery (although the recovery appliance is located 3m from the main appliance).

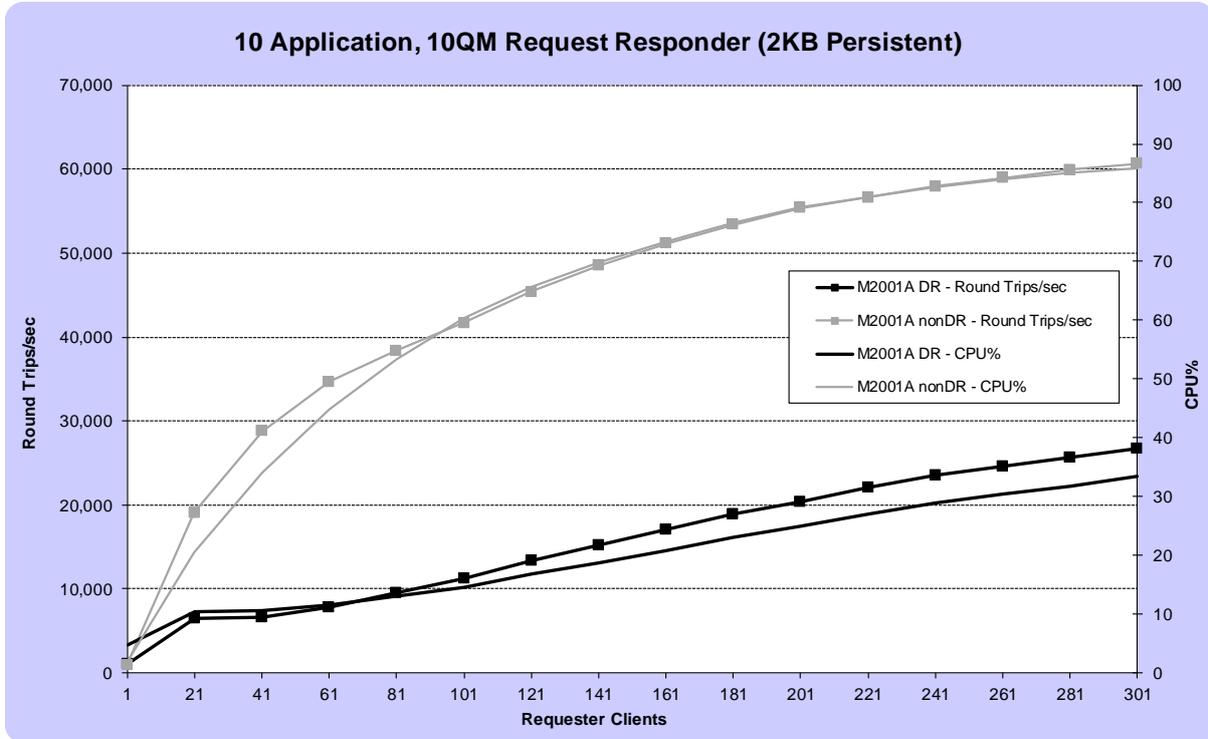


FIGURE 6 - PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING

Figure 6 shows that when we have multiple QM performing 2KB persistent messaging across a pair of DR appliances, the peak messaging rate is less than half than when distributed across a set of non DR Queue Managers. Increasing the numbers of clients further or utilising message batching would narrow the difference in behaviour shown.

Test	M2001A DR				M2001A nonDR			
	Max Rate*	CPU%	Clients	Latency#	Max Rate*	CPU%	Clients	Latency#
10Q 10QM Request Responder (256b Persistent)	36,941	42.8	301	0.8	75,798	93.72	301	1.0
10Q 10QM Request Responder (2KB Persistent)	26,682	33.37	301	0.9	60,626	85.93	301	1.0
10Q 10QM Request Responder (20KB Persistent)	6,768	13.6	301	1.2	12,072	18.44	31	1.5
10Q 10QM Request Responder (200KB Persistent)	719	7.97	31	3.2	1,269	8.59	31	2.9

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 8 - PEAK RATES FOR 10QM PERSISTENT MESSAGING

5.3 How does DR perform over larger distances?

DR configuration usually requires the pair of appliances to be situated a large distance apart so that any particular event that might affect one appliance would be hoped not to affect the second appliance.

The data in the following tables show the results from the test scenario featured in the previous section but with additional data points using simulated latency delays of 10, 20, 50 and 100ms. A comparison against the DR scenario in which the MQ Appliances are directly connected is also included.

Test	M2001A DR				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 10ms Latency (256b Persistent)	40,065	49.03	301	0.9	108.5%
10Q 10QM Request Responder 10ms Latency (2KB Persistent)	26,454	40.81	301	0.9	99.1%
10Q 10QM Request Responder 10ms Latency (20KB Persistent)	6,498	13.66	281	1.2	96.0%
10Q 10QM Request Responder 10ms Latency (200KB Persistent)	652	7.34	41	3.0	90.7%

*Round trips/sec

#Single thread round trip latency (ms)

Test	M2001A DR				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 20ms Latency (256b Persistent)	41,139	51.53	301	0.7	111.4%
10Q 10QM Request Responder 20ms Latency (2KB Persistent)	25,036	32.8	261	0.9	93.8%
10Q 10QM Request Responder 20ms Latency (20KB Persistent)	5,780	13.26	301	1.3	85.4%
10Q 10QM Request Responder 20ms Latency (200KB Persistent)	640	6.86	31	3.0	88.9%

*Round trips/sec

#Single thread round trip latency (ms)

Test	M2001A DR				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 50ms Latency (256b Persistent)	40,018	49.83	301	0.8	108.3%
10Q 10QM Request Responder 50ms Latency (2KB Persistent)	23,343	31.78	301	1.0	87.5%
10Q 10QM Request Responder 50ms Latency (20KB Persistent)	5,311	13.38	301	1.3	78.5%
10Q 10QM Request Responder 50ms Latency (200KB Persistent)	568	7.08	41	7.4	79.0%

*Round trips/sec

#Single thread round trip latency (ms)

Test	M2001A DR				vs Direct
	Max Rate*	CPU%	Clients	Latency#	
10Q 10QM Request Responder 100ms Latency (256b Persistent)	11,434	10.21	301	1.0	31.0%
10Q 10QM Request Responder 100ms Latency (2KB Persistent)	10,278	15.32	301	1.1	38.5%
10Q 10QM Request Responder 100ms Latency (20KB Persistent)	4,363	6.61	261	2.2	64.5%
10Q 10QM Request Responder 100ms Latency (200KB Persistent)	424	3.96	41	15.1	58.9%

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 9 - PEAK RATES FOR 10QM PERSISTENT MESSAGING WITH 10, 20, 50 AND 100MS SIMULATED LATENCY

6 HA and DR Scenarios

From Fixpack 5, you can now configure an HA and DR enabled queue manager. The data that is asynchronously replicated for disaster recovery is sent from the currently active instance of the HA pair.

This configuration allows the Queue Manager on the DR recovery appliance to resume work should an outage occur that results in both the appliances in the HA group becoming unavailable.

The performance of these scenarios is very close to that which has been measured in the HA scenarios featured in sections 4.2 and 4.3; the throughput is very similar and there is a small increase in CPU that accounts for the additional work performing the asynchronous replication.

Test	M2001A HA and DR			
	Max Rate*	CPU%	Clients	Latency#
10Q Request Responder (256b Persistent)	42,114	56.37	301	1.5
10Q Request Responder (2KB Persistent)	27,269	42.51	301	1.7
10Q Request Responder (20KB Persistent)	6,001	17.8	301	1.9
10Q Request Responder (200KB Persistent)	643	10.43	41	4.4

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 10 - PEAK RATES FOR PERSISTENT MESSAGING, HA AND DR

Test	M2001A HA and DR			
	Max Rate*	CPU%	Clients	Latency#
10Q 10QM Request Responder (256b Persistent)	15,155	17.03	301	1.3
10Q 10QM Request Responder (2KB Persistent)	12,947	16.73	301	1.6
10Q 10QM Request Responder (20KB Persistent)	4,604	10.63	301	1.9
10Q 10QM Request Responder (200KB Persistent)	717	9.77	301	4.8

*Round trips/sec

#Single thread round trip latency (ms)

TABLE 11 - PEAK RATES FOR 10QM PERSISTENT MESSAGING, HA AND DR

7 Frequently Asked Questions

Can I run any other HA solution in conjunction with the appliance?

No, the HA support by placing a pair of appliances into an MQ Appliance HA group is the only supported method of HA on the MQ Appliance.

What HA cables come with an MQ Appliance?

From FP3, MQ appliances should come supplied with 2x1Gb 10m copper cable and 1x10Gb 5m copper cable.

Can I use a fibre cable for 10Gb HA replication?

Yes, a fibre cable (with correct transceivers) can be used in place of the 10Gb copper cable.

Does the cable medium affect performance?

Copper and fibre cables can be used for the 10Gb HA workload connection. The majority of results in this report were collected using a 3m copper cable. To compare the performance of using a 3m fibre and a 3m copper cable (with appropriate transceivers), the single QM tests featured in section 4.1 were re-run using a 3m copper cable. There was no discernible change in performance between the two cables over that distance.

Which connections do I need to make to setup the HA group?

Interface eth13 on appliance A needs to be connected to eth13 on appliance B with a 1Gb copper cable. Interface eth17 on appliance A needs to be connected to eth17 on appliance B with a 1Gb copper cable. Interface eth21 on appliance A needs to be connected to eth21 on appliance B with a 10Gb cable.

Which connections do I need to make to setup the DR configuration?

Interface eth20 on appliance A needs to be connected to eth20 on appliance B with a 10Gb cable. The DR configuration doesn't use heartbeating between the appliances, so eth13 and eth17 are not reserved for DR use and can be used for appliance workload.

After starting a QM in an HA group on appliance A, how do I move it to appliance B?

From the mqcli on appliance B, run ***sethappreferred QM*** after which you be able to see it now running on appliance B.

My QM was created on appliance A, an outage caused the QM to be started on appliance B, how do I avoid the QM moving back to appliance A when it becomes available?

When the QM was created, a correlation was made between the QM and the appliance it was created on; where possible the QM will run on that appliance. To prevent the QM from reverting to appliance A, run **clearhapreferred QM** on appliance A or run **sethapreferred QM** on appliance B.

Is there a maximum number of HA QM supported on an appliance?

When using the M2000A/B up to 16 HA QM are supported within an HA group and each appliance can only be a member of one HA group. If you have Queue Managers active on the second appliance (with their standby QM on the first appliance), these are all part of the same HA group and count towards the limit of 16. You can however, have many more non HA QM configured and active.

If you are using the M2001, the restriction of 16 HA QM has been removed, although due to the resources required to support the QM processing and synchronous replication, we don't advise using more than 32 HA QM on a pair of appliances.

How long does it take to detect an appliance in an HA group is no longer active?

Heartbeats are issued every second between the HA appliances, and after 5 failures it is presumed the other appliance is no longer active.

Can I run HA and DR functionality on my pair of appliances?

At MQ Appliance Fixpack 4, HA **OR** DR capability can be selected. You cannot configure a Queue Manager within an HA group and provide DR capability (to the same secondary appliance or a separate MQ Appliance). From Fixpack 5 onwards, you can now run QM in both HA and DR configurations.

How do I start messaging processing at the DR recovery appliance?

When asynchronous replication is being performed, the secondary Queue Manager at the recovery appliance is in stopped state. You will need to run the **makedrprimary -m QM** and **strmqm QM** commands to start messaging at the recovery appliance.

How does HA and DR work together?

In a three appliance topology where you wish to have an HA pair of appliances and a remote DR appliance, DR data is transferred from the Primary appliance (of the HA pair) to the DR appliance. So that the DR appliance doesn't need to know when the HA pair fails over, a floating IP address is assigned and is always associated with the current HA primary appliance on Ethernet port eth20.

8 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 480GB SSD
RAID	ServeRAID M5210 (4GB Flash RAID cache)