IBM MQ Appliance TLS and Encryption Performance Report

Model: M2003

Version 1.0 - November 2025

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, November 2025.

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

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

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

This performance report at version 1.0 contains performance data based on the MQ Appliance models M2003A and M2003B. This report covers TLS, AMS and encrypted filesystem performance data and includes the following highlights:

- The M2003A provides encrypted performance with only a small (from 6%) reduction in throughput. See section 5.2
- The M2003A provides encrypted performance with only a small (from 1%) reduction in throughput with HA QM. See section 5.4
- The M2003A provides TLS encrypted transmission of message data using TLS12 and TLS13 ciphers. See section 7
- The M2003A provides AMS confidentiality protection of persistent messages whilst achieving a message rate larger than 50,000 round trips/second. See section 8
- Over 220,000 round trips/second peak messaging rate achieved in a Non-persistent messaging TLS scenario (~440,000 messages produced and ~440,000 messages consumed per second). See section 7
- Over 125,000 round trips/second peak messaging rate achieved in an HA enabled scenario with filesystem encryption (~250,000 messages produced and ~250,000 messages consumed per second). See section 5.4

The M2003 hardware components and how they compare to the previous model M2002 are shown below:

Model	M2002A	M2002B	M2003A	M2003B
CPU	2x12 Core HT	1x6 Core HT	2x16 Core HT (2.9GHz)	1x8 Core HT (2.9GHz)
RAM	192GB	192GB	256GB	256GB
Storage	6.4TB SSD	3.2TB SSD	6.4TB NVMe SSD	3.2TB NVMe SSD
IO Subsystem	RAID 10	RAID 10	RAID 10	RAID 10
Workload and replication network connectivity	8x1Gb 6x10Gb 4x40Gb	8x1Gb 6x10Gb 4x40Gb	8x1Gb 4x10Gb 4x40Gb 2x100Gb	8x1Gb 4x10Gb 4x40Gb 2x100Gb
Management	2x1Gb	2x1Gb	2x1Gb	2x1Gb
Chipset	Skylake	Skylake	Icelake	Icelake
RAID	SAS 12Gb/s 2GB cache	SAS 12Gb/s 2GB cache	PCIe 4 x8 16GB/s	PCIe 4 x8 16GB/s

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

There are local disks within the appliance to enable efficient persistent messaging by the local Queue Managers. The four 3.2TB NVMe SSD drives are configured in a RAID10 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 synchronously across the pair of appliances. Disaster Recovery (DR) may be achieved by the addition of a remote appliance to which QM data is distributed synchronously or asynchronously.

The MQ appliance can be purchased in two variants: M2003A and M2003B. There are two main differences for the M2003B as highlighted in the table above, reduced CPU capacity and reduced filesystem storage space.

As before, you can purchase an upgrade to convert an M2003B appliance to an M2003B+ appliance, which has the same capacity as an M2003A appliance.

Most of the tests use the M2003A variant of the MQ Appliance and this is the default hardware unless stated otherwise. Several tests were also conducted using the M2003B variant and provide comparative data points to the main testing to provide appropriate capacity planning information.

There are two modules that support 40Gb network connectivity with two ports available in each. There is a capacity limit of 40Gb per module. There is one module that supports 100Gb network connectivity with two ports available. There is a capacity limit of 100Gb per module. This report utilises 2 of the 100Gb links for workload traffic; and 1 40Gb port for replication traffic.

All 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.

A feature was released in MQ 9.2.5 which enables the Queue Manager (QM) to be created with an encrypted filesystem. This prevents access to the data stored in MQ queues (and the recovery log) in the event the appliance SSD disks are removed from the M2003 appliance. Storage used to retain QM configuration backup or diagnostic information can also be encrypted.

The version of the MQ Appliance as tested in this report is M2003A MQ 9.4.3 and where a comparison is made to the restricted appliance configuration, this uses the MQ Appliance M2003B MQ 9.4.3. If you plan on using an encrypted filesystem, we recommend using MQ 9.4.0.7(LTS), MQ 9.4.2(CD) or later, as large performance gains were achieved at this juncture.

4 Request/Responder Scenario

The scenario that will be used in this report reflects the most common, anticipated usage patterns for the MQ appliance and provides guidance for those customers performing capacity planning or migration activities.

Each test was initially conducted and graphs produced using a 2K (2048 byte) message size. Additional tests were also conducted using 256byte, 20K and 200K to provide further data for capacity planning and are found in the accompanying data table.

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 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.

The test scenario in Figure 1**Error! Reference source not found.** is a Request Responder scenario that simulates several 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.

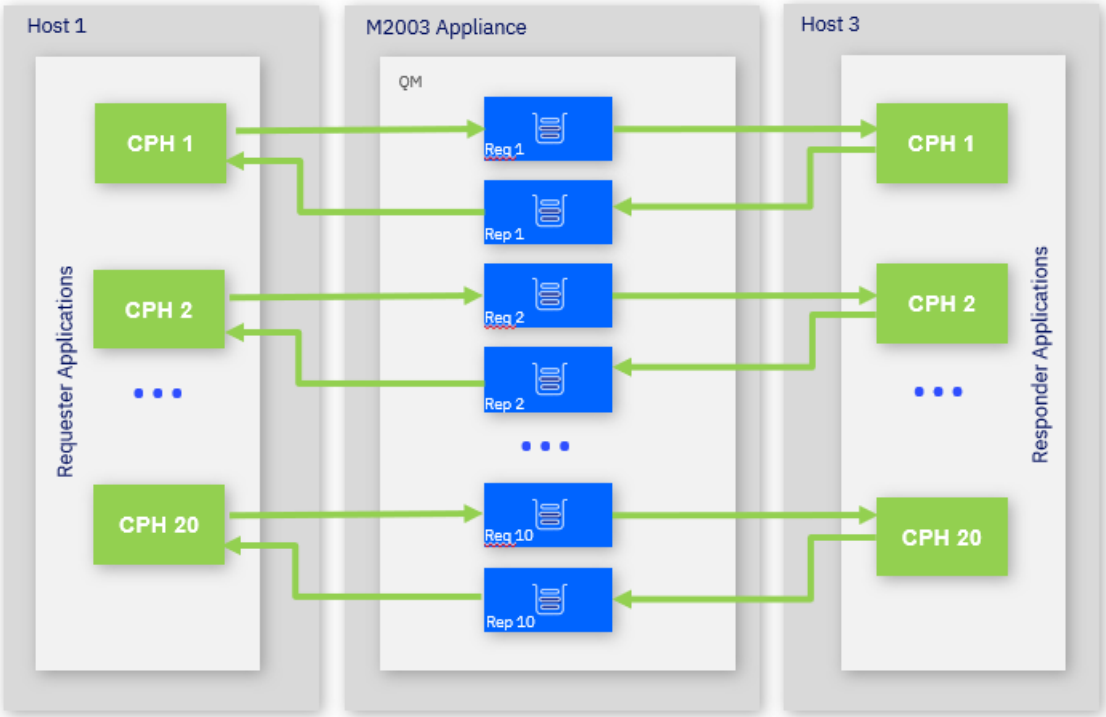

FIGURE 1 - REQUESTER-RESPONDER WITH REMOTE QUEUE MANAGER ON MQ APPLIANCE

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), 300 fixed responder threads (30 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.

For the 10QM tests, there are 10 QM with 10 applications per QM (again using 10 pairs of queues). There are still 300 overall responder threads, but as we now have 100 pairs of queues, we have 3 responders per request queue.

5 Filesystem encryption performance

It has been noted that more IBM MQ customers often require message payloads to be encrypted while in transit and at rest to comply with various security mandates. Before the release of this feature, customers who wished to ensure their data was encrypted at rest had to used AMS (Advanced Message Security).

Customers can now use TLS to protect the message data in transit and an encrypted QM filesystem to protect that data when at rest in the appliance. Note that the use of an encrypted filesystem does not require that all messaging is TLS enabled.

It is also possible to enable an encrypted fs for an HA and/or DR QM as well as a standalone QM. The passphrase used to encrypt the filesystem will be propagated to the HA partner automatically, but will need to be manually added to any DR recovery appliance. The impact on HA performance will be examined in the following chapter.

The following sections show the impact of utilising an encrypted filesystem for Persistent messaging across several different scenarios.

5.1 Single QM

The graph below shows the results from the single QM test using a 2KB message size on an M2003A appliance:

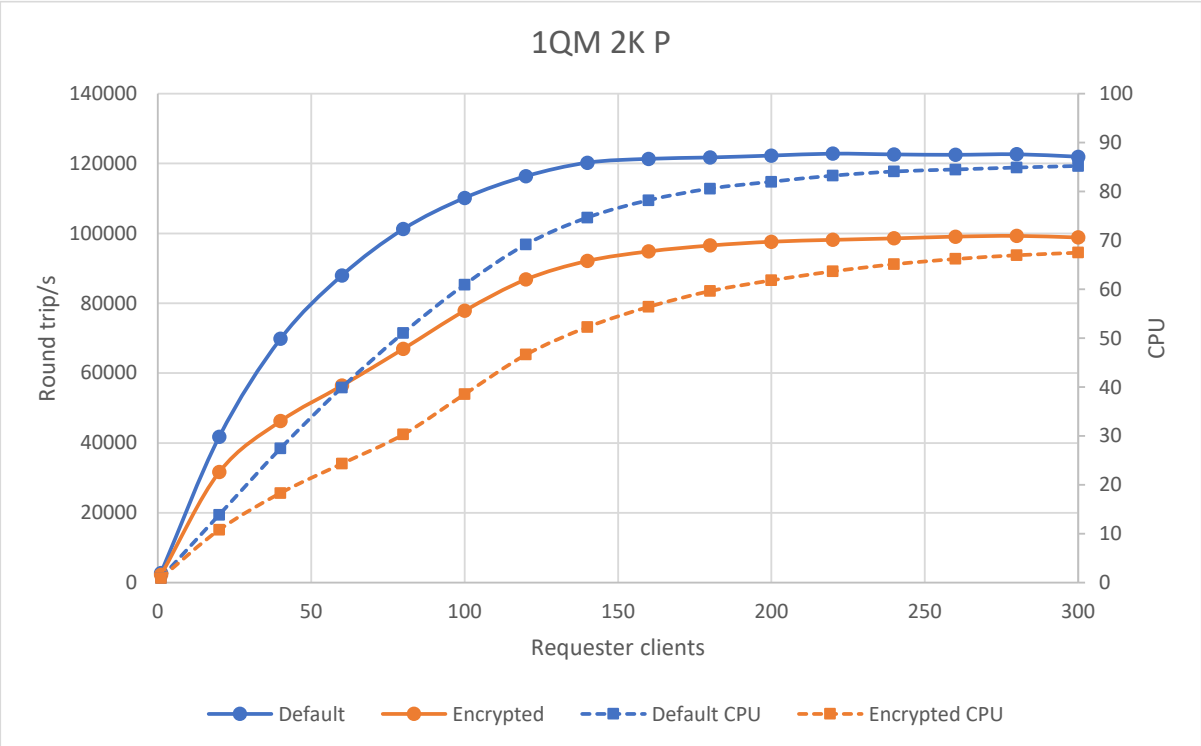

FIGURE 2 - PERFORMANCE RESULTS FOR 2KB, 1QM PERSISTENT MESSAGING, M2003A

The impact of enabling encryption increases latency of a single requester thread sending and receiving 2KB messages by 26%. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 19%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Single thread latency increase	Minimum impact on throughput (round trips/s)
256b	12%	-11%
2K	26%	-19%
20K	19%	-31%
200K	34%	-27%

TABLE 1 - IMPACT ON LATENCY AND THROUGHPUT AT VARIOUS MESSAGE SIZES FOR A SINGLE QM ON M2003A

The following graph shows the results from the single QM test using a 2KB message size on an M2003B appliance:

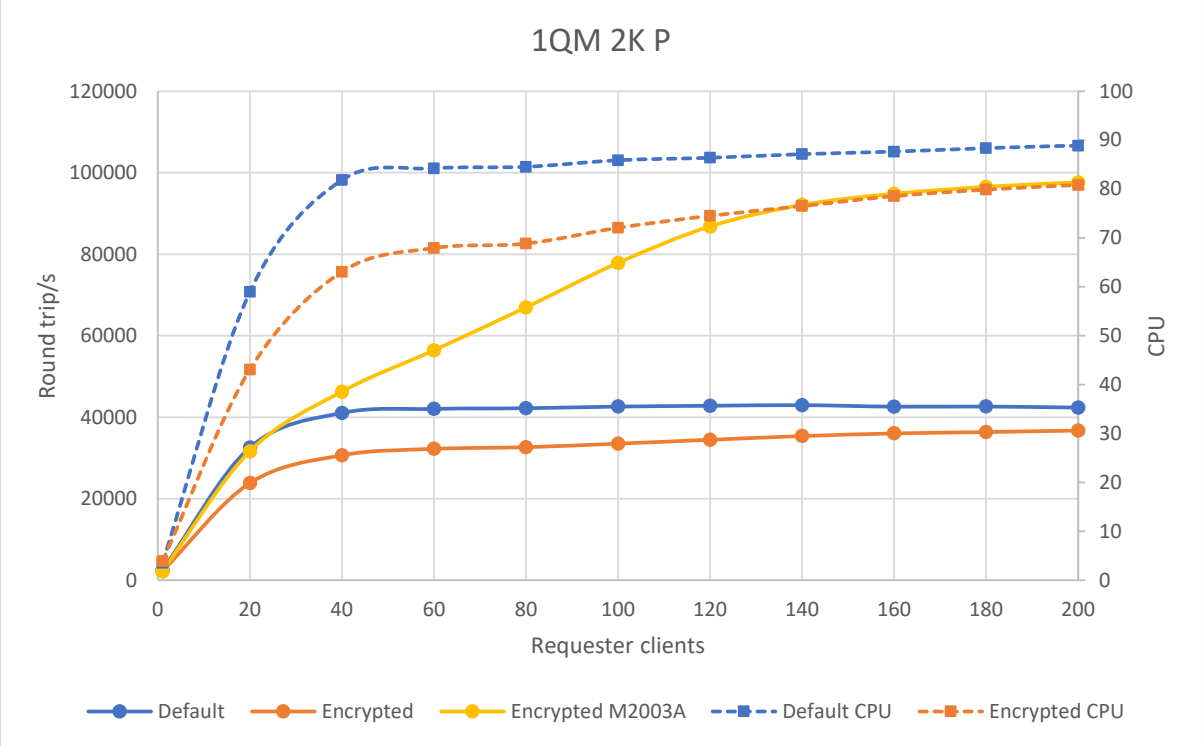

FIGURE 3 - PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING, M2003B VS M2003A

The impact of enabling encryption increases latency of a single requester thread sending and receiving 2KB messages by 12%. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 12%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Single thread latency increase	Minimum impact on throughput (round trips/s)
256b	17%	-7%
2K	12%	-12%
20K	41%	-37%
200K	28%	-30%

TABLE 2 - IMPACT ON LATENCY AND THROUGHPUT AT VARIOUS MESSAGE SIZES FOR A SINGLE QM ON M2003B

5.2 Multiple QM

The graph below shows the results from the multiple QM test using a 2KB message size on a M2003A appliance:

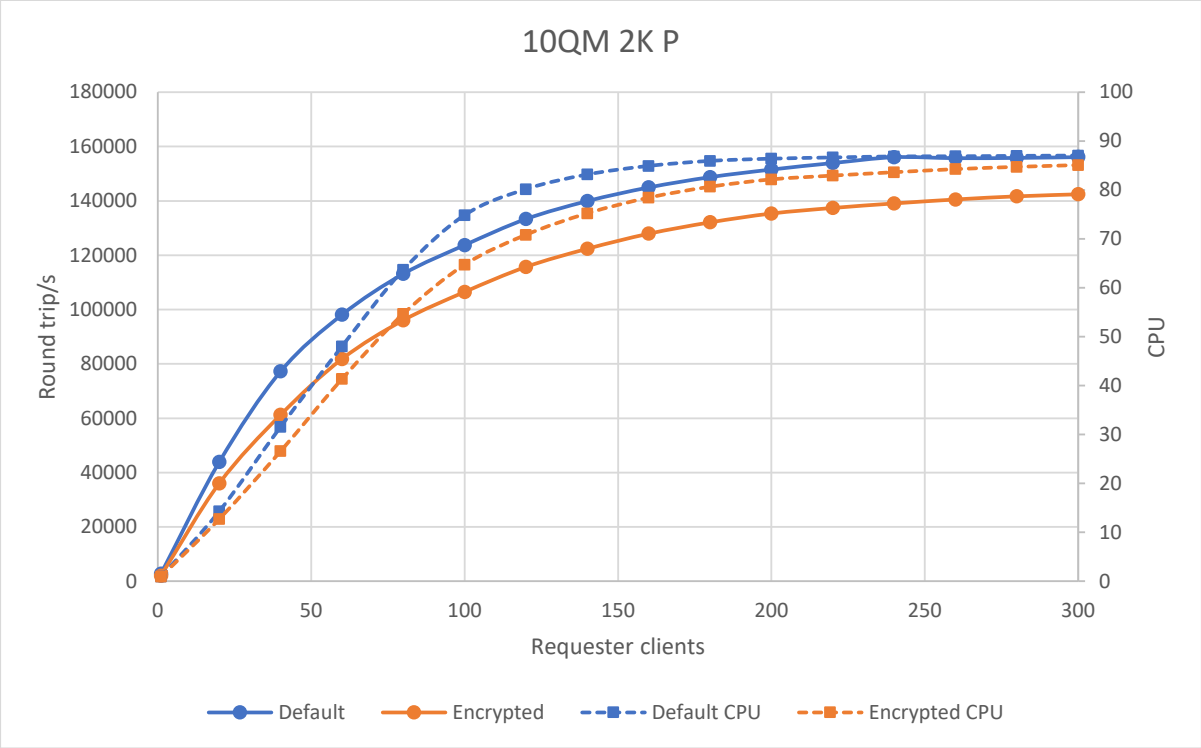

FIGURE 4 - PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING, M2003A

The impact of enabling encryption has a reduced effect on multiple QM because there are multiple threads processing the encryption and recovery log writes. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 9%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Minimum impact on throughput (round trips/s)
256b	-6%
2K	-9%
20K	-23%
200K	-4%

TABLE 3 - IMPACT ON THROUGHPUT AT VARIOUS MESSAGE SIZES FOR MULTIPLE QM ON M2003A

The graph below shows the results from the multiple QM test using a 2KB message size on a M2003B appliance:

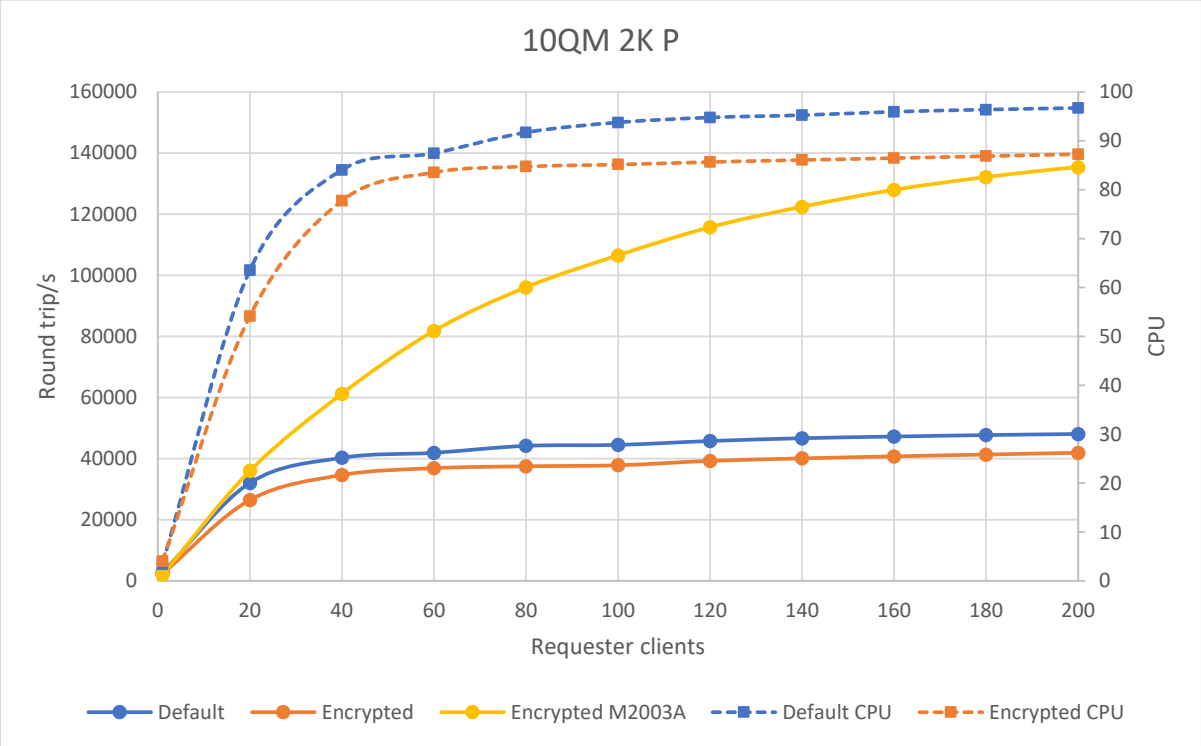

FIGURE 5 – PERFORMANCE RESULTS FOR 2KB, 10QM PERSISTENT MESSAGING, M2003B VS M2003A

The impact of enabling encryption has a reduced effect on multiple QM because there are multiple threads processing the encryption and recovery log writes. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 12%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Minimum impact on throughput (round trips/s)
256b	-9%
2K	-12%
20K	-23%
200K	-32%

TABLE 4 - IMPACT ON THROUGHPUT AT VARIOUS MESSAGE SIZES FOR MULTIPLE QM ON M2003B

5.3 Single HA QM

The graph below shows the results from the single HA QM test using a 2KB message size on a M2003A appliance:

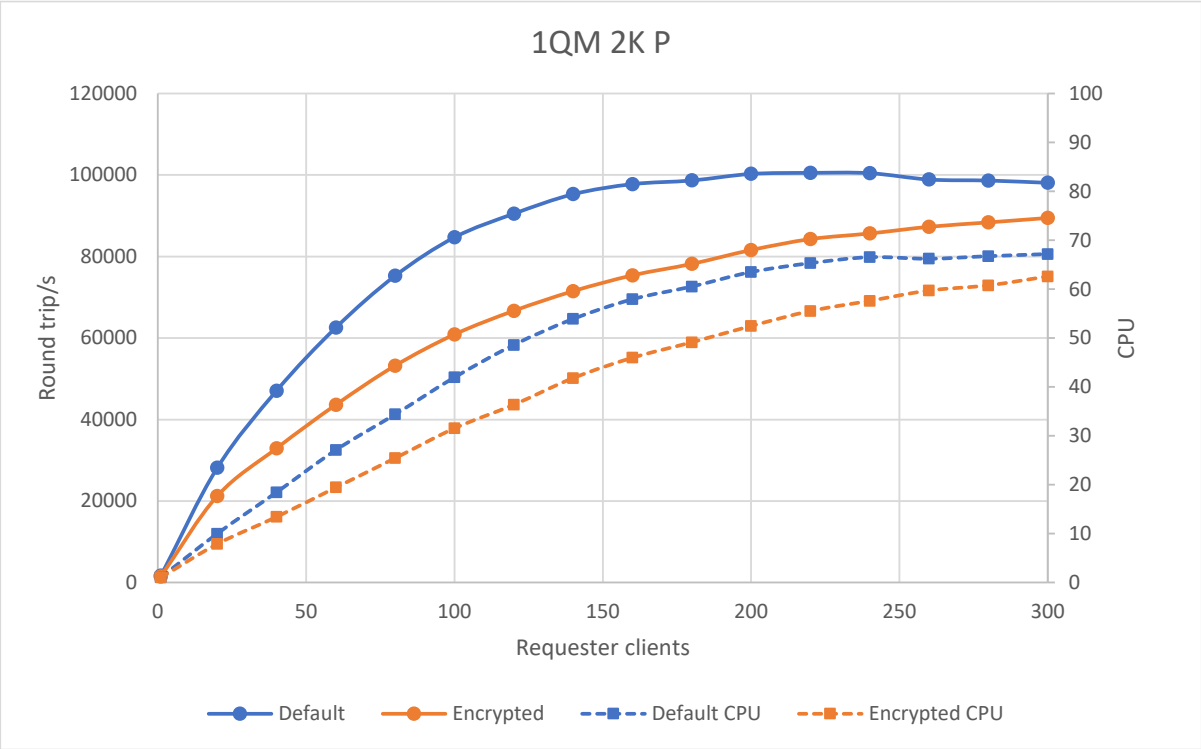

FIGURE 6 - PERFORMANCE RESULTS FOR 2KB, 1QM HA PERSISTENT MESSAGING, M2003A

The impact of enabling encryption increases latency of a single requester thread sending and receiving 2KB messages by 11%. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 9%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Single thread latency increase	Minimum impact on throughput (round trips/s)
256b	6%	-12%
2K	11%	-9%
20K	10%	-4%
200K	15%	-1%

TABLE 5 - IMPACT ON LATENCY AND THROUGHPUT AT VARIOUS MESSAGE SIZES FOR A SINGLE HA QM ON M2003A

The graph below shows the results from the single HA QM test using a 2KB message size on a M2003B appliance:

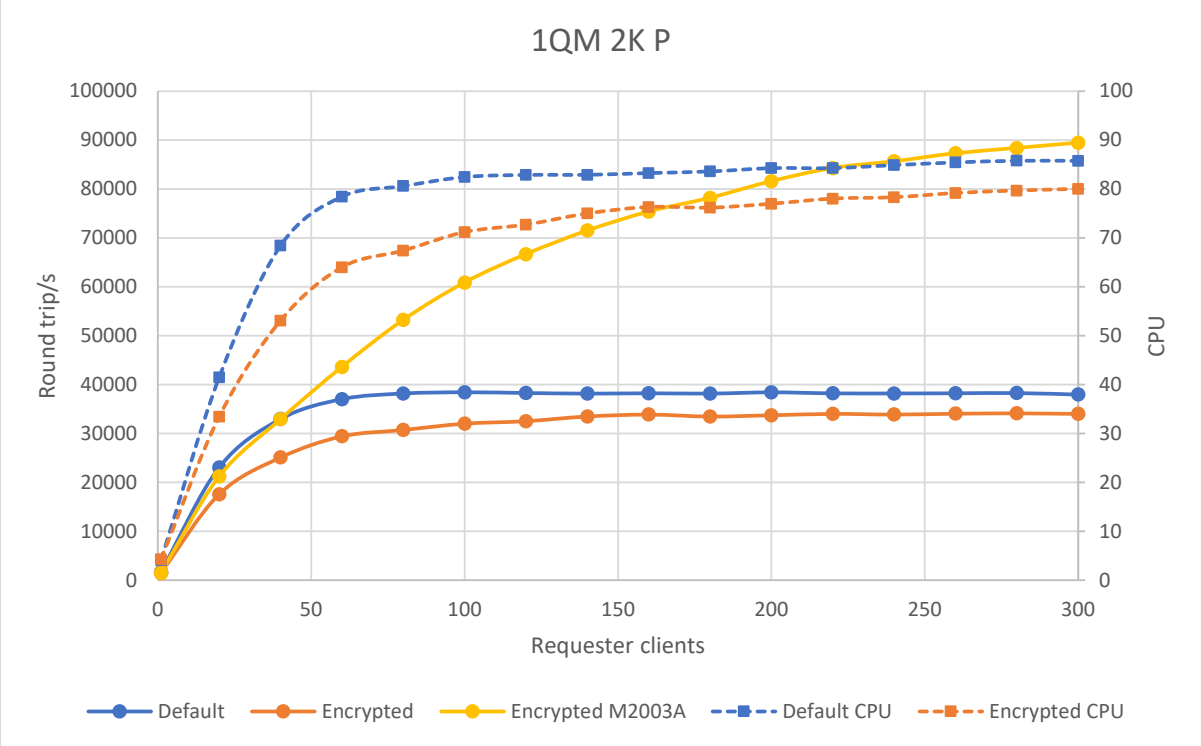

FIGURE 7 - PERFORMANCE RESULTS FOR 2KB, 1QM HA PERSISTENT MESSAGING, M2003B VS M2003A

The impact of enabling encryption increases latency of a single requester thread sending and receiving 2KB messages by 11%. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 10%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Single thread latency increase	Minimum impact on throughput (round trips/s)
256b	11%	-10%
2K	11%	-10%
20K	18%	-18%
200K	22%	-10%

TABLE 6 - IMPACT ON LATENCY AND THROUGHPUT AT VARIOUS MESSAGE SIZES FOR A SINGLE HA QM ON M2003B

5.4 Multiple HA QM

The graph below shows the results from the multiple HA QM test using a 2KB message size on M2003A appliances:

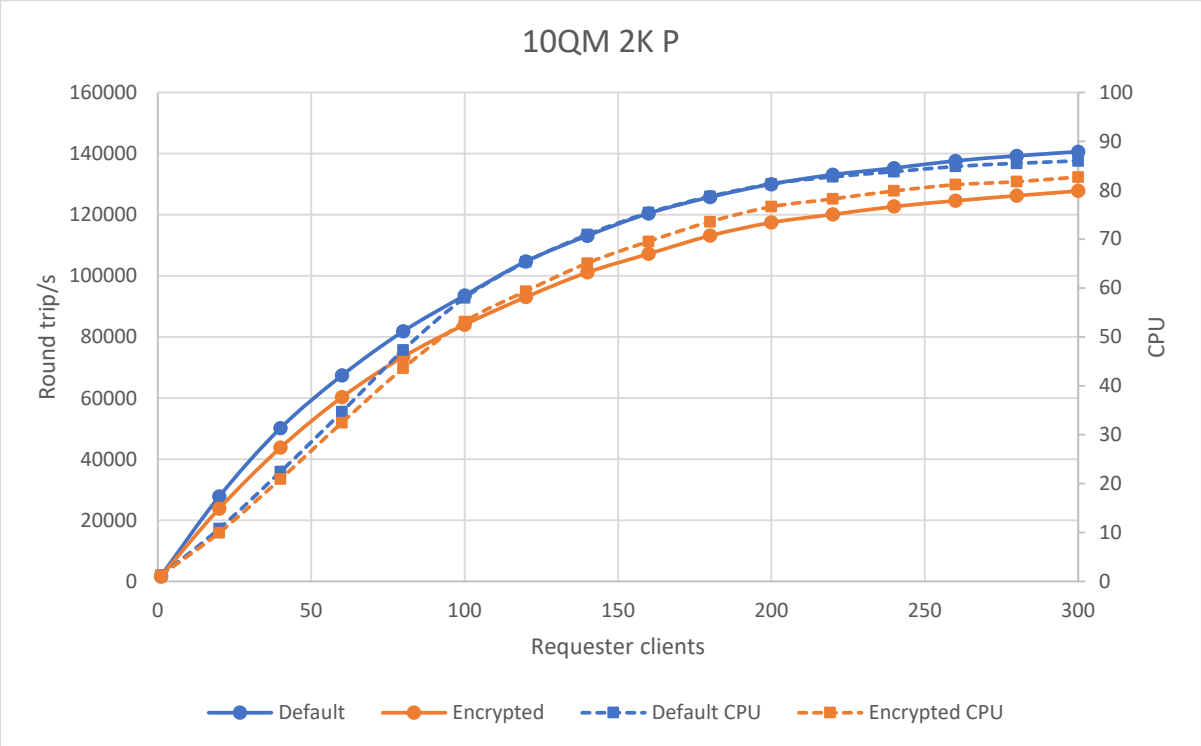

FIGURE 8 - PERFORMANCE RESULTS FOR 2KB, 10QM HA PERSISTENT MESSAGING, M2003A

The impact of enabling encryption has a reduced effect on multiple QM because there are multiple threads processing the encryption and recovery log writes. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 9%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Minimum impact on throughput (round trips/s)
256b	-6%
2K	-9%
20K	-18%
200K	-2%

TABLE 7 - IMPACT ON THROUGHPUT AT VARIOUS MESSAGE SIZES FOR MULTIPLE HA QM ON M2003A

The graph below shows the results from the multiple HA QM test using a 2KB message size on M2003B appliances:

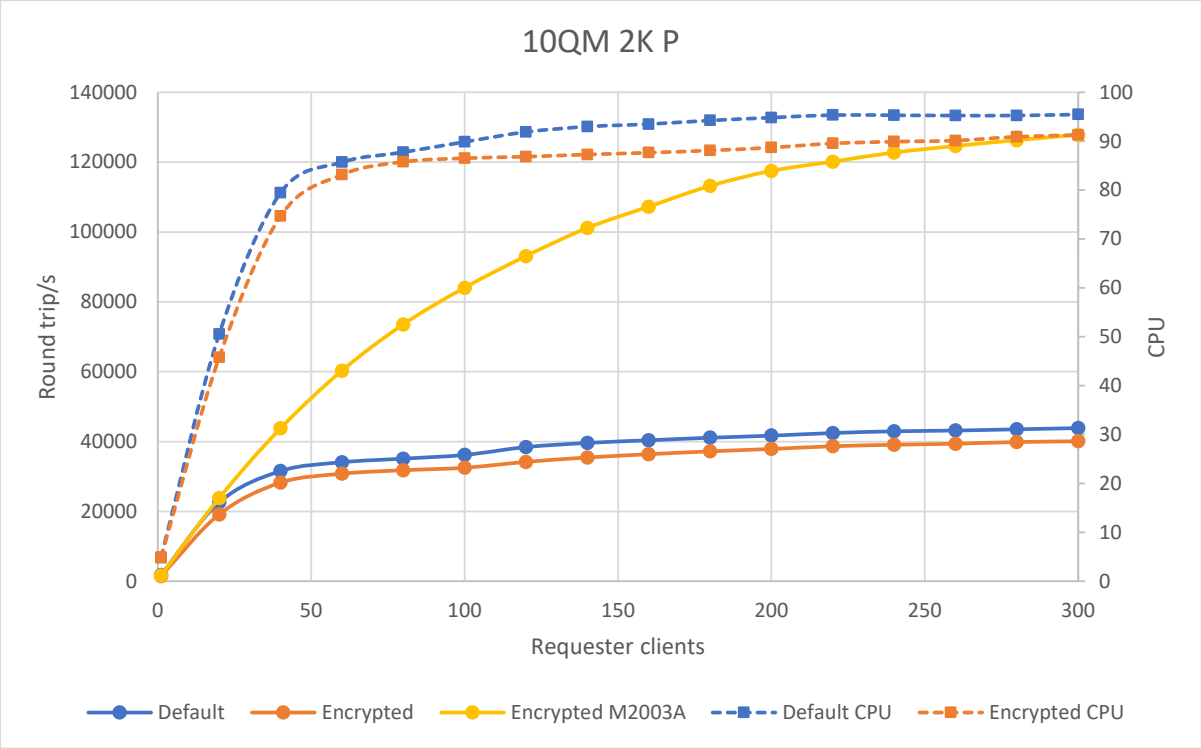

FIGURE 9 - PERFORMANCE RESULTS FOR 2KB, 10QM HA PERSISTENT MESSAGING, M2003B VS M2003A

The impact of enabling encryption has a reduced effect on multiple QM because there are multiple threads processing the encryption and recovery log writes. The peak throughput achieved across a varying number of requester threads is reduced by a minimum of 8%.

The following table contains the datapoints for the other message sizes in this scenario:

Message Size	Minimum impact on throughput (round trips/s)

256b	-7%
2K	-8%
20K	-19%
200K	-12%

TABLE 8 - IMPACT ON THROUGHPUT AT VARIOUS MESSAGE SIZES FOR MULTIPLE HA QM ON M2003B

6 Filesystem encryption conclusions

The filesystem encryption functionality offers protection for your data at rest within the MQ appliance. There is a small increase in CPU which reflects the cost of encrypting the message payload before persisting that data to storage. There is a similar cost when decrypting the message data after retrieval from storage. Note that in many scenarios, the QM (and the OS) optimize message retrieval by avoiding reading (and therefore decryption) from the IO subsystem.

Peak throughput on a single Non HA QM on M2003A will be impacted by a minimum of 19% as the latency of writing data is increased for a 2KB message size. Using multiple Non HA QM helps mitigate the impact of this increase in latency, resulting in a minimum reduction of 9% of peak throughput for this scenario.

Peak throughput on a single HA QM on M2003A will be impacted by a minimum of 9% as the latency of writing data is increased for a 2KB message size. Using multiple HA QM helps mitigate the impact of this increase in latency, although the minimum reduction of of peak throughput remains 9% for this scenario.

The data from the M2003B appliance has also been included to help guide you to which model is most appropriate for your deployment.

7 TLS

This section illustrates the cost of enabling TLS communication between the clients and the QM. We will use scenarios C1 and C2 from Section 6 of MPA5 and apply two of the strongest TLS 1.2 CipherSpecs and one strong TLS 1.3 CipherSpec to compare their performance.

The following TLS 1.2 CipherSpecs were tested (all utilise 256bit encryption and are FIPS compliant):

- ECDHE_ECDSA_AES_256_CBC_SHA384
- ECDHE_ECDSA_AES_256_GCM_SHA384 (Suite B compliant)
- ECDHE_RSA_AES_256_CBC_SHA384
- ECDHE_RSA_AES_256_GCM_SHA384

Results for the suite B compliant CipherSpec (ECDHE_ECDSA_AES_256_GCM_SHA384), along with a CBC based CipherSpec (ECDHE_RSA_AES_256_CBC_SHA384) and a FIPS compliant TLS 1.3 CipherSpec (TLS_AES_256_GCM_SHA384) are plotted below. As will be seen in the tables below, the remaining tested CipherSpecs exhibited a performance profile similar to one of these.

Queue Manager authentication is used to setup the TLS conversation.

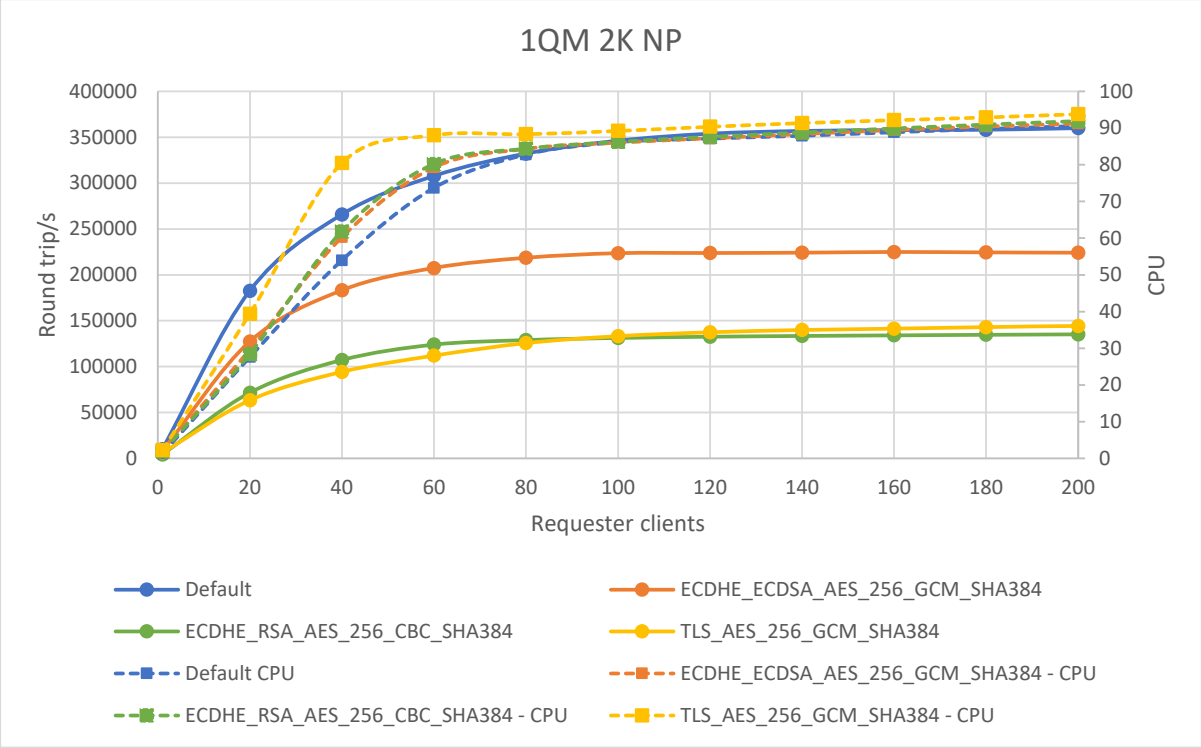

FIGURE 10 - PERFORMANCE RESULTS FOR 2KB NON-PERSISTENT MESSAGING WITH TLS

The graph in Figure 10 illustrates that ECDHE_ECDSA_AES_256_GCM_SHA384 was the best performing TLS12 CipherSpec, whilst the TLS12 CBC and TLS13 Ciphers perform similarly.

The table below shows the peak throughput for all the referenced CipherSpecs. It should be noted that all TLS13 CipherSpecs performed similarly.

Cipher	**TLS Version**	**Peak Throughput**	**Compared with No TLS**
None	None	361981	
ECDHE_ECDSA_AES_256_GCM_SHA384	1.2	224892	-38%
ECDHE_ECDSA_AES_256_CBC_SHA384	1.2	136017	-62%
ECDHE_RSA_AES_256_CBC_SHA384	1.2	131537	-64%
ECDHE_RSA_AES_256_GCM_SHA384	1.2	213682	-41%
TLS_AES_256_GCM_SHA384	1.3	144724	-60%

TABLE 9 - PEAK RATES FOR 2KB NON-PERSISTENT MESSAGING WITH TLS

The following chart illustrates the impact of TLS encryption on Persistent messaging scenarios:

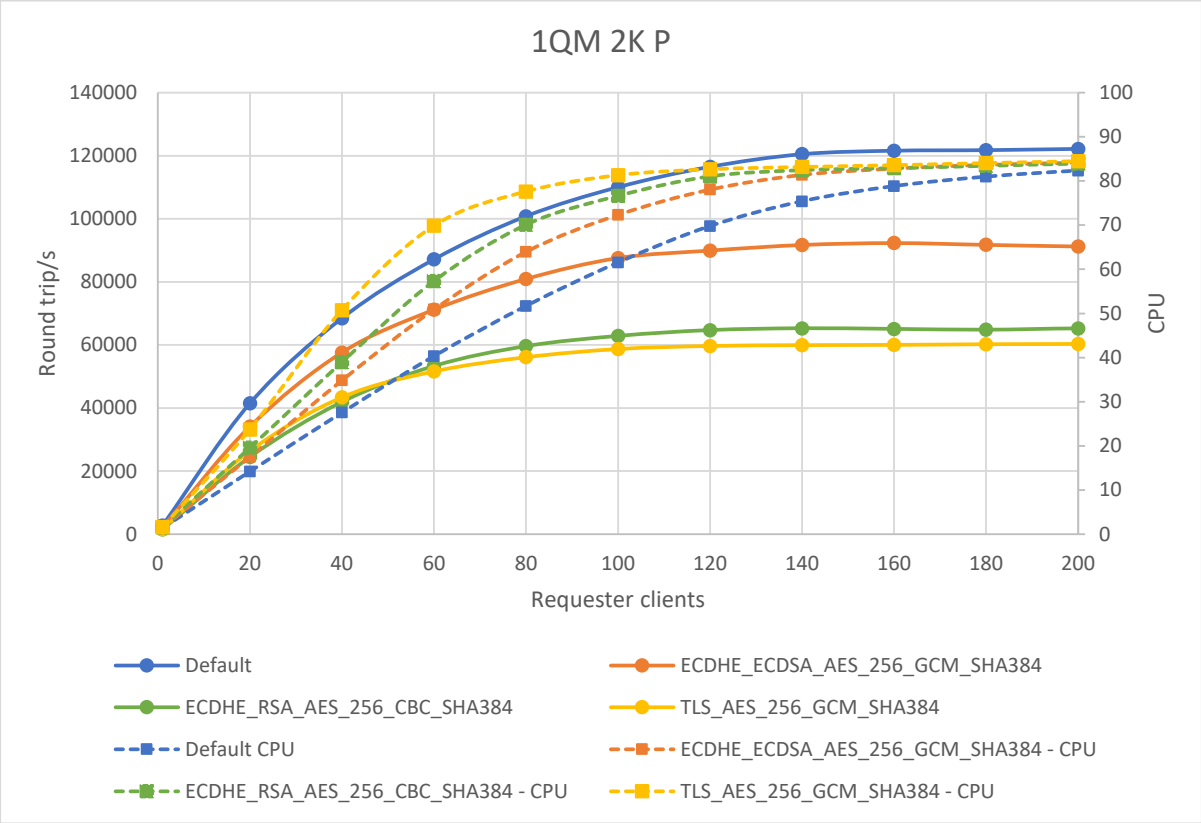

FIGURE 11 - PERFORMANCE RESULTS FOR 2KB PERSISTENT MESSAGING WITH TLS

The graph in Figure 11 illustrates that ECDHE_ECDSA_AES_256_GCM_SHA384 was the best performing TLS12 CipherSpec, whilst the TLS12 CBC and TLS13 Ciphers perform similarly. TLS encryption has less impact on Persistent scenarios than Non-Persistent scenarios.

The table below shows the peak throughput for all the referenced CipherSpecs. It should be noted that all TLS13 CipherSpecs performed similarly.

Cipher	TLS Version	Peak Throughput	Compared with No TLS
None	None	122452	
ECDHE_ECDSA_AES_256_GCM_SHA384	1.2	92327	-25%
ECDHE_ECDSA_AES_256_CBC_SHA384	1.2	65273	-47%
TLS_AES_256_GCM_SHA384	1.3	60620	-50%

TABLE 10 - PEAK RATES FOR 2KB PERSISTENT MESSAGING WITH TLS

8 AMS

This section illustrates the cost of enabling AMS to protect the message contents in transit between the clients and the QM and at rest at the QM. We will use the scenario C2 from Section 6 of MPA5, and compare Integrity, Privacy and Confidentiality mode with the Non-AMS performance.

The certificate key size used is 1024 bytes and the key reuse limit in Confidentiality mode was set to 32. The symmetric key encryption algorithm used was AES256. The cryptographic hash function used was SHA512.

The default certificate key size changed from 1024 to 2048 in MQ 9.1.4 and has an impact on AMS performance; the performance of Confidentiality mode with the increased key size is also shown.

Note that client CPU (rather than server CPU) is featured on the graph below as that shows the increase in computation performed by the clients in encrypting and decrypting the AMS protected messages.

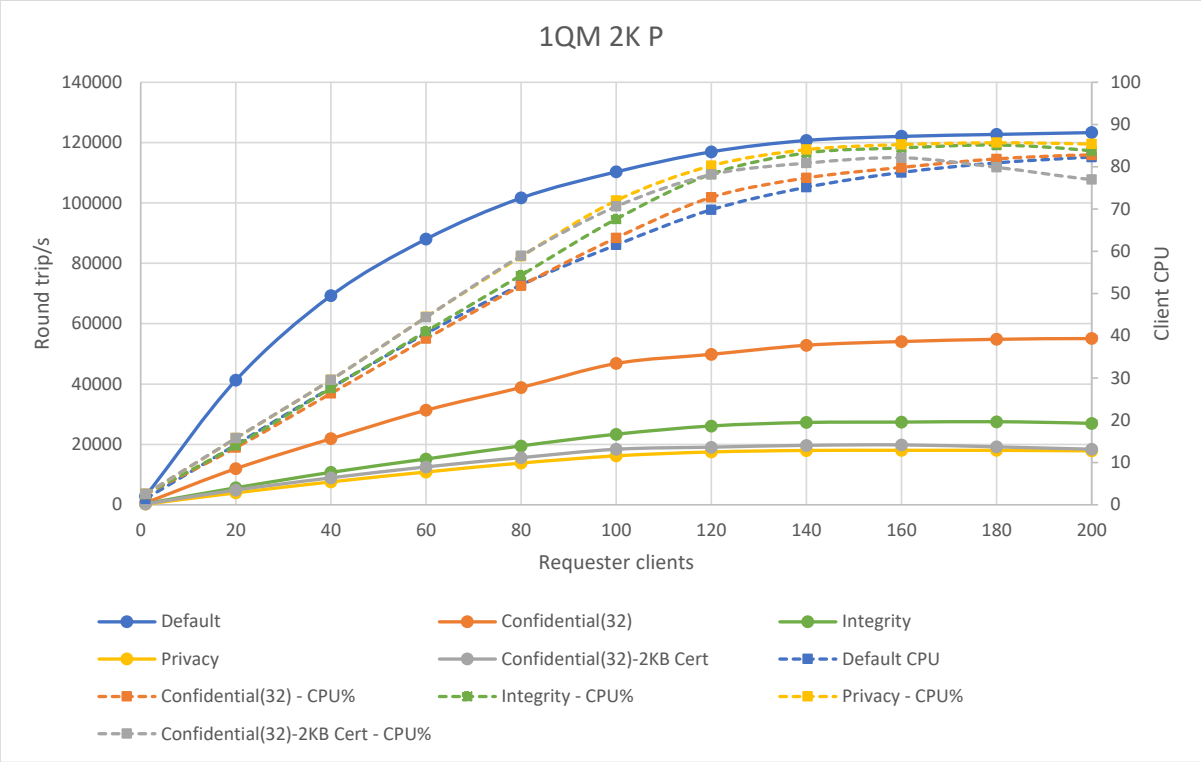

FIGURE 12 - PERFORMANCE RESULTS FOR 2KB PERSISTENT MESSAGING WITH AMS

The graph in Figure 12 shows that the performance of AMS confidentiality is almost half of the Non-AMS performance. The cost of the larger certificate key (resulting in more complex cryptographic calculations) reduces the performance by over 80%.

The peak throughput achieved for the AMS Confidentiality(32) measurement at 200 clients was over 55,000 round trip/s. The request/responder scenario utilizes a request and a reply queue, so for each round trip, 2 message put and 2 message get operations take place. For a single put/single get scenario, the peak performance that you might obtain in the same environment is over 110,000 msg/sec.

9 Appendix A – Client machine specification

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

Category	Value
Machine	Lenovo ThinkSystem SR630 V2
OS	Red Hat Enterprise Linux Server 8.10
CPU	2x16 (3.1Ghz)
RAM	256GB RAM
Network	10/100Gb Ethernet
Disks	2x 3TB NVMe SSD in RAID-0 array
RAID	Linux mdraid MQ Logs hosted on RAID-0 partition

10 Appendix B – QM Configuration

The following commands and expect scripts were used to create the standalone Queue Managers for this report:

```
crtmqm -lp 64 -lf 16384 -h 5000 -fs 16 PERF0
setmqini -m PERF0 -s TuningParameters -k DefaultPQBufferSize -v 10485760
setmqini -m PERF0 -s TuningParameters -k DefaultQBufferSize -v 10485760


proc configureQM { QMname QMport QMqueues } {
    send "runmqsc $QMname\n"
    send "define listener(L1) trptype(tcp) port($QMport) control(qmgr)\n"
    send "start listener(L1)\n"
    send "alter channel(SYSTEM.DEF.SVRCONN) chltype(SVRCONN) sharecnv(1) maxmsgl(104857600)\n"
    send "alter qmgr maxmsgl(104857600)\n"
    send "alter qlocal(system.default.local.queue) maxmsgl(104857600)\n"
    send "alter qmodel(system.default.model.queue) maxmsgl(104857600)\n"
    send "alter qmodel(system.jms.model.queue) maxmsgl(104857600)\n"
    send "alter qmodel(system.jms.tempq.model) maxmsgl(104857600)\n"
    send "alter qlocal(system.dead.letter.queue) maxmsgl(104857600)\n"
    send "define channel(SYSTEM.ADMIN.SVRCONN) chltype(SVRCONN)\n"
    send "alter qmgr chlauth(disabled)\n"
    send "alter authinfo(SYSTEM.DEFAULT.AUTHINFO.IDPWOS) authtype(IDPWOS) chckclnt(OPTIONAL)\n"
    send "refresh security type(CONNAUTH)\n"
    send "define qlocal(queue) maxdepth(5000) replace\n"
    send "define qlocal(request) maxdepth(5000) replace\n"
    send "define qlocal(reply) maxdepth(5000) replace\n"
    for {set j 0} {$j <= $QMqueues} {incr j 1} {
        send "define qlocal(request$j) maxdepth(5000) replace\n"
        send "define qlocal(reply$j) maxdepth(5000) replace\n"
    }
    send "end\n"
}
```