Cisco UCS C240 M6 Rack Server Disk I/O Characterization

October 2021

Contents

Executive summary

The Cisco Unified Computing System™ (Cisco UCS®) C240 M6 Rack Server is well-suited for a wide range of storage and input/output (I/O)-intensive applications such as big data analytics, databases, virtualization and high-performance computing in its 2-socket, 2-rack-unit (2RU) form factor.

The Cisco UCS C240 M6 server extends the capabilities of the Cisco UCS rack server portfolio with 3rd Gen Intel Xeon Scalable Processors supporting more than 43 percent more cores per socket and 33 percent more memory when compared with the previous generation. This server provides up to 40 percent more performance than the M5 generation for your most demanding applications.

You can deploy the Cisco UCS C-Series rack servers as standalone servers or as the part of Cisco UCS managed by Cisco Intersight™ technology, Cisco UCS Manager, or Intersight Managed Mode to take advantage of Cisco® standards-based unified computing innovations that can help reduce your Total Cost of Ownership (TCO) and increase your business agility.

These improvements deliver significant performance and efficiency gains that will improve your application performance. The Cisco UCS C240 M6 Rack Server delivers outstanding levels of expandability and performance.

This document outlines the I/O performance characteristics of the Cisco UCS C240 M6 Rack Server using the Cisco 12G SAS Modular RAID Controller with a 4 GB cache module and the Cisco 12G SAS Host Bus Adapter (HBA). Performance comparisons of solid-state disks (SSDs) with RAID and HBA controllers are presented. The goal of this document is to help customers make well-informed decisions in choosing the right internal disk types and configuring the right controller options to meet their individual I/O workload needs.

Performance data was obtained using the FIO measurement tool, with analysis based on the number of I/O operations per second (IOPS) for random I/O workloads and megabytes-per-second (MBps) throughput for sequential I/O workloads. From this analysis, specific recommendations are made for storage configuration parameters.

Many combinations of drive types and RAID levels are possible. For these characterization tests, performance evaluations were limited to small-form-factor (SFF) SSDs with RAID 0, 5, and 10 virtual drives and JBOD configurations.

Introduction

The widespread adoption of virtualization and data center consolidation technologies has profoundly affected the efficiency of the data center. Virtualization brings new challenges for the storage technology, requiring the multiplexing of distinct I/O workloads across a single I/O "pipe". From a storage perspective, this approach results in a sharp increase in random IOPS. For spinning media disks, random I/O operations are the most difficult to handle, requiring costly seek operations and rotations between microsecond transfers. Not only do the hard disks add a security factor, they also constitute the critical performance components in the server environment. Therefore, it is important to bundle the performance of these components through intelligent technology so that they do not cause a system bottleneck and they will compensate for any failure of an individual component. RAID technology offers a solution by arranging several hard disks in an array so that any disk failure can be accommodated.

According to conventional wisdom, data center I/O workloads are either random (many concurrent accesses to relatively small blocks of data) or sequential (a modest number of large sequential data transfers). Historically, random access has been associated with a transactional workload, which is the most common type of workload for an enterprise. Currently, data centers are dominated by random and sequential workloads resulting from the scale-out architecture requirements in the data center.

I/O challenges

The rise of technologies such as virtualization, cloud computing, and data consolidation poses new challenges for the data center and requires enhanced I/O requests. These enhanced requests lead to increased I/O performance requirements. They also require that data centers fully use available resources so that they can support the newest requirements of the data center and reduce the performance gap observed industrywide.

The following are the major factors leading to an I/O crisis:

- Increasing CPU usage and I/O: Multicore processors with virtualized server and desktop architectures increase processor usage, thereby increasing the I/O demand per server. In a virtualized data center, it is the I/O performance that limits the server consolidation ratio, not the CPU or memory.

- Randomization: Virtualization has the effect of multiplexing multiple logical workloads across a single physical I/O path. The greater the degree of virtualization achieved, the more random the physical I/O requests.

Scope of this document

For the I/O performance characterization tests, performance was evaluated using SSDs with configurations of RAID 0, 5, and 10 virtual drives to achieve maximum performance for SAS and SATA SSDs for the Cisco UCS C240 M6 Rack server. The Cisco UCS C240 M6SX server used for the I/O performance characterization tests supports up to 28 SSDs. The performance tests described here were performed with a RAID 0 28-disk configuration for SFF SSDs, JBOD 28-disk configuration with dual HBA adapters, and a combination of 24 RAID 0 SAS and SATA SSDs + 4 Nonvolatile Memory Express (NVMe) SSDs and 24 JBOD SAS and SATA SSDs + 4 NVMe SSDs.

Cisco UCS C240 M6 server models

Overview

The Cisco UCS C240 M6 SFF server extends the capabilities of the Cisco UCS System portfolio in a 2RU form factor with the addition of the gen3 Intel Xeon Scalable Processors (Ice Lake) and 16 dual-inline-memory-module (DIMM) slots per CPU for 3200-MHz double-data-rate 4 (DDR4) DIMMs with DIMM capacity points up to 256 GB. The maximum memory capacity for 2 CPUs follows:

- 8 TB (32 x 256 GB DDR4 DIMMs), or

- 12 TB (16 x 256 GB DDR4 DIMMs and 16 x 512 GB Intel Optane Persistent Memory Modules [PMEMs]).

There are several options to choose from:

Option 1:

- Up to 12 front SFF SAS/SATA HDDs or SSDs (optionally up to 4 of the drives can be NVMe)

- I/O-centric option, which provides up to 8 PCIe slots using all three rear risers

- Storage-centric option, which provides 6 PCIe slots and one rear riser with a total of up to 2 SFF drives (SAS/SATA or NVMe with 4 PCIe gen4 slots)

- Optional optical drive

Option 2:

- Up to 24 front SFF SAS/SATA HDDs or SSDs (optionally up to 4 of the slots can be NVMe)

- I/O-centric option, which provides up to 8 PCIe slots using all three rear risers

- Storage-centric option, which provides 3 PCIe slots using slots in one of the rear risers and two rear risers with a total of up to 4 SFF drives (SAS/SATA or NVMe with 4 PCIe gen4 slots)
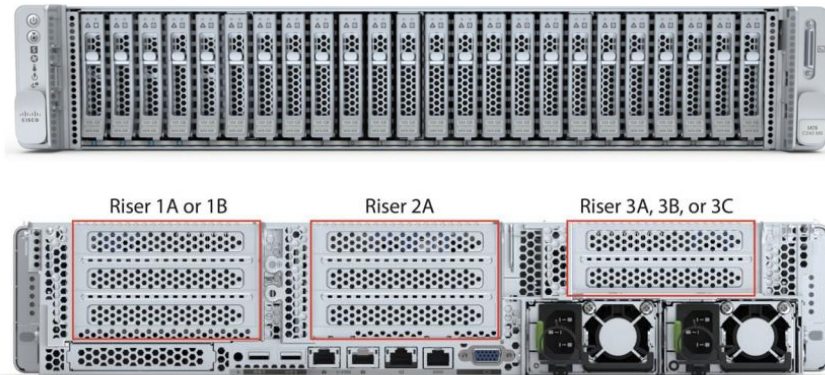
Option 3:

- Up to 12 front SFF NVMe-only drives

- I/O-centric option, which provides up to 6 PCIe slots using two rear risers

- Storage-centric option, which provides 3 PCIe slots using slots in one of the rear risers and up to 2 SFF drives (NVMe with 4 PCIe gen4 slots) using one of the rear risers

Option 4:

- Up to 24 front NVMe-only drives

- I/O-centric option, which provides up to 6 PCIe slots using two rear risers

- Storage-centric option, which provides 3 PCIE slots using slots in one of the rear risers and up to 2 SFF drives (NVMe with 4 PCIe gen4 slots) using one of the rear risers

The server provides one or two internal slots (depending on the server type) for the following:

- One slot for a SATA Interposer to control up to 8 SATA-only drives from the PCH (AHCI), or

- One slot for a Cisco 12G RAID controller with cache backup to control up to 28 SAS/SATA drives, or

- Two slots for Cisco 12G SAS pass-through HBAs; each HBA controls up to 16 SAS/SATA drives

- A mix of 24 SAS/SATA and NVMe front drives (up to 4 NVMe front drives) and optionally 4 SAS/SATA/NVMe rear drives

Figure 1.
Front view of server

A specifications sheet for the C240 M6 is available at:
https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m6-sff-specsheet.pdf.

SSDs and NVMe drives used

The SSDs selected for the tests were based on the following factors:

- 960 GB enterprise value 12G SAS SSD
- 960 GB enterprise value 6G SATA SSD
- 1.6 TB NVMe high-performance medium endurance NVMe SSD

Performance

When considering price, it is important to differentiate between enterprise performance SSDs and enterprise value SSDs. The differences between the two in performance, cost, reliability, and targeted applications are significant. Although it can be appealing to integrate SSDs with NAND flash technology in an enterprise storage solution to improve performance, the cost of doing so on a large scale may be prohibitive.

Virtual disk options

You can configure the following controller options with virtual disks to accelerate write and read performance and provide data integrity:

- RAID level
- Strip (block) size
- Access policy
- Disk cache policy
- I/O cache policy
- Read policy
- Write policy
- RAID levels

Table 1 summarizes the supported RAID levels and their characteristics.

Table 1.   RAID levels and characteristics

RAID level	Characteristics	Parity	Redundancy
RAID 0	Striping of 2 or more disks to achieve optimal performance	No	No
RAID 1	Data mirroring on 2 disks for redundancy with slight performance improvement	No	Yes
RAID 5	Data striping with distributed parity for improved fault tolerance	Yes	Yes
RAID 6	Data striping with dual parity with dual fault tolerance	Yes	Yes
RAID 10	Data mirroring and striping for redundancy and performance improvement	No	Yes
RAID 50	Block striping with distributed parity for high fault tolerance	Yes	Yes
RAID 60	Block striping with dual parity for performance improvement	Yes	Yes

To avoid poor write performance, full initialization is always recommended when creating a RAID 5 or RAID 6 virtual drive. Depending on the virtual disk size, the full initialization process can take a long time based on the drive capacity. Fast initialization is not recommended for RAID 5 and RAID 6 virtual disks.

Stripe size

Stripe size specifies the length of the data segments that the controller writes across multiple drives, not including parity drives. You can configure stripe size as 64, 128, 256, or 512 KB, or 1 MB. The default stripe size is 64 KB. The performance characterization tests in this paper used a stripe size of 64 KB. With random I/O workloads, no significant difference in I/O performance was observed by varying the stripe size. With sequential I/O workloads, performance gains are possible with a stripe size of 256 KB or larger; however, a stripe size of 64 KB was used for all the random and sequential workloads in the tests.

Strip (block) size versus stripe size

A virtual disk consists of two or more physical drives that are configured together through a RAID controller to appear as a single logical drive. To improve overall performance, RAID controllers break data into discrete chunks called strips that are distributed one after the other across the physical drives in a virtual disk. A stripe is the collection of one set of strips across the physical drives in a virtual disk.

Access policy

You can set access policy as follows:

- RW: Read and write access is permitted.
- Read Only: Read access is permitted, but write access is denied.
- Blocked**:** No access is permitted.

Disk cache policy

You can set disk cache policy as follows:

- Disabled: The disk cache is disabled. The drive sends a data transfer completion signal to the controller when the disk media has actually received all the data in a transaction. This process helps ensure data integrity if a power failure occurs.

- Enabled**:** The disk cache is enabled. The drive sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. However, the data has not actually been transferred to the disk media, so data may be permanently lost if a power failure occurs. Although disk caching can accelerate I/O performance, it is not recommended for enterprise deployments.

I/O cache policy

You can set I/O cache policy as follows:

- Direct: All read data is transferred directly to host memory, bypassing the RAID controller cache. Any read-ahead data is cached. All write data is transferred directly from host memory, bypassing the RAID controller cache if write-through cache mode is set. The direct policy is recommended for all configurations.

Read policy

You can set the read policy as follows:

- No read ahead (normal read): Only the requested data is read, and the controller does not read any data ahead.

- Always read ahead: The controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon.

Write policy

You can set the write policy as follows:

- Write through: Data is written directly to the disks. The controller sends a data transfer completion signal to the host when the drive subsystem has received all the data in a transaction.

- Write back: Data is first written to the controller cache memory, and then the acknowledgment is sent to the host. Data is written to the disks when the commit operation occurs at the controller cache. The controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction.

- Write back with battery backup unit (BBU): Battery backup is used to provide data integrity protection if a power failure occurs. Battery backup is always recommended for enterprise deployments.

Workload characterization

This section provides an overview of the specific access patterns used in the performance tests.

Table 2 lists the workload types tested.

Table 2.   Workload types

Workload type	RAID type	Access pattern type	Read:write (%)
OLTP	5	Random	70:30
Decision-support system (DSS), business intelligence, and video on demand (VoD)	5	Sequential	100:0
Database logging	10	Sequential	0:100
High-performance computing (HPC)	5	Random and sequential	50:50
Digital video surveillance	10	Sequential	10:90
Big data: Hadoop	0	Sequential	90:10
Apache Cassandra	0	Sequential	60:40
VDI: Boot process	5	Random	80:20
VDI: Steady state	5	Random	20:80

Tables 3 and 4 list the I/O mix ratios chosen for the sequential access and random access patterns, respectively.

Table 3.   I/O mix ratio for sequential access pattern

I/O mode	I/O mix ratio (read:write)	
Sequential	100:0	0:100
	RAID 0, 10	RAID 0, 10
	JBOD*	JBOD*

*JBOD is configured using HBA

Table 4.   I/O mix ratio for random access pattern

I/O mode	I/O mix ratio (read:write)			
Random	100:0	0:100	70:30	50:50
	RAID 0	RAID 0	RAID 0, 5	RAID 0, 5
	JBOD*	JBOD*	JBOD*	JBOD*

*JBOD is configured using HBA

Tables 5 and 6 show the recommended virtual drive configuration for SSDs and HDDs, respectively.

Table 5.   Recommended virtual drive configuration for SSDs

Access pattern	RAID level	Strip size	Disk cache policy	I/O cache policy	Read policy	Write policy
Random I/O	RAID 0	64 KB	Unchanged	Direct	No read ahead	Write through
Random I/O	RAID 5	64 KB	Unchanged	Direct	No read ahead	Write through
Sequential I/O	RAID 0	64 KB	Unchanged	Direct	No read ahead	Write through
Sequential I/O	RAID 10	64 KB	Unchanged	Direct	No read ahead	Write through

Note: Disk cache policy is set as 'Unchanged' as it cannot be modified for VDs created using SSD

Table 6.   Recommended virtual drive configuration for HDDs

Access pattern	RAID level	Strip size	Disk cache policy	I/O cache policy	Read policy	Write policy
Random I/O	RAID 0	256 KB	Disabled	Direct	Always read ahead	Write back good BBU
Random I/O	RAID 5	256 KB	Disabled	Direct	Always read ahead	Write back good BBU
Sequential I/O	RAID 0	256 KB	Disabled	Direct	Always read ahead	Write back good BBU
Sequential I/O	RAID 10	256 KB	Disabled	Direct	Always read ahead	Write back good BBU

Test configuration

The test configuration was as follows:

- 28 RAID 0 virtual drives were created with 28 disks for the SFF SSDs.

- 4 RAID 5 virtual drives were created with 7 disks in each virtual drive.

- 7 RAID 10 virtual drives were created with 4 disks in each virtual drive.

- The RAID configuration was tested with the Cisco M6 12G SAS RAID controller with SuperCap and a 4 GB cache module with 28 disks for the SFF SSDs.

- The JBOD configuration was tested with the Cisco 12G SAS HBA that supports up to 16 drives. The server is configured with dual HBA adapters, and 14 disks are configured on each HBA adapter (28 disks on two adapters).

- 4 NVMe SSDs were configured on front NVMe slots that are 4 PCIe lanes and directly managed by CPU.

- Random workload tests were performed using 4 and 8 KB block sizes for both SAS/SATA and NVMe SSDs.

- Sequential workload tests were performed using 256 KB and 1 MB block sizes for both SAS/SATA and NVMe SSDs.

Table 7 lists the recommended FIO settings.

Table 7.   Recommended FIO settings

Name	Value
FIO version	Fio-3.19
Filename	Device name on which FIO tests should run
Direct	For direct I/O, page cache is bypassed
Type of test	Random I/O or sequential I/O, read, write, or mix of read:write
Block size	I/O block size: 4, 8, or 256K, or 1M
I/O engine	FIO engine - libaio
I/O depth	Number of outstanding I/O instances
Number of jobs	Number of parallel threads to be run
Run time	Test run time
Name	Name for the test
Ramp-up time	Ramp-up time before the test starts
Time-based	To limit the run time of test

Note: The SSDs were tested with various combinations of outstanding I/O and number of jobs to get the best performance within acceptable response time.

SSD performance results

Performance data was obtained using the FIO measurement tool on RHEL, with analysis based on the IOPS rate for random I/O workloads and MBps throughput for sequential I/O workloads. From this analysis, specific recommendations can be made for storage configuration parameters.

The I/O performance test results capture the maximum IOPS and bandwidth achieved with the SATA SSDs and NVMe SSDs within the acceptable response time (latency) of 2 milliseconds (ms).

The RAID adapter (UCSC-RAID-M6SD) used in this testing uses Broadcom SAS3916 RAID-on-Chips (ROC) with Cub Expander internally. The performance specifications for this SAS3916 ROC is 3M IOPS for Random I/O when used with SAS SSDs. Please refer to Broadcom SAS3916 ROC product brief (https://docs.broadcom.com/doc/12397937) for more information.

Results showed that, for random I/O (reads) performance with SAS drives, 3+ million IOPS was achieved with 16 SAS drives populated on the server. However, given that the C240 M6SX server allows 28 disks populated in the server, the additional disks augment only extra storage capacity while achieving maximum performance, as captured in this document (Figure 2 later in this document).

Similarly, for random I/O (reads) performance with SATA drives, 1.53M IOPS was achieved with 28 SATA drives, a fully populated disk scenario. This result aligns with the IOPS expectation of the storage adapter design for the SATA interface (refer to Figure 8 later in this document).

For sequential I/O workloads (both with SAS and SATA drives), results showed that an adapter maximum of 14,000 MBps was achieved with 14 drives in the case of SAS and 28 drives in the case of SATA, as captured in this document (refer to Figure 2 for SAS and Figure 8 for SATA).

Note: These conclusions are not applicable for SAS HBA adaptors (UCSC-SAS-M6T) because the C240 M6SX is configured with two HBA adaptors and hence overall server performance matches with total drives multiplied by expected single-drive performance.

SAS SSD (960 GB enterprise value 12G SAS SSD) performance for 28-disk configuration

SAS SSD RAID 0 performance

Figure 2 shows the performance of the SSDs under test for RAID 0 configuration with a 100-percent random read and random write, and 70:30 percent random read:write access pattern. The graph shows the comparative performance values achieved for enterprise value SAS SSD drives to help customers understand the performance. The graph shows that the 960 GB enterprise value SAS drive provides performance of more than 3M IOPS with latency of 0.4 milliseconds for random read, performance of 2M IOPS with latency of 0.4 milliseconds for random write, and performance of 3.3M IOPS with latency of 0.5 milliseconds for random read:write 70:30 percent with a 4 KB block size. The result shows the maximum performance achievable from the C240 M6 RAID adapter. Latency is the time taken to complete a single I/O request from the viewpoint of the application.
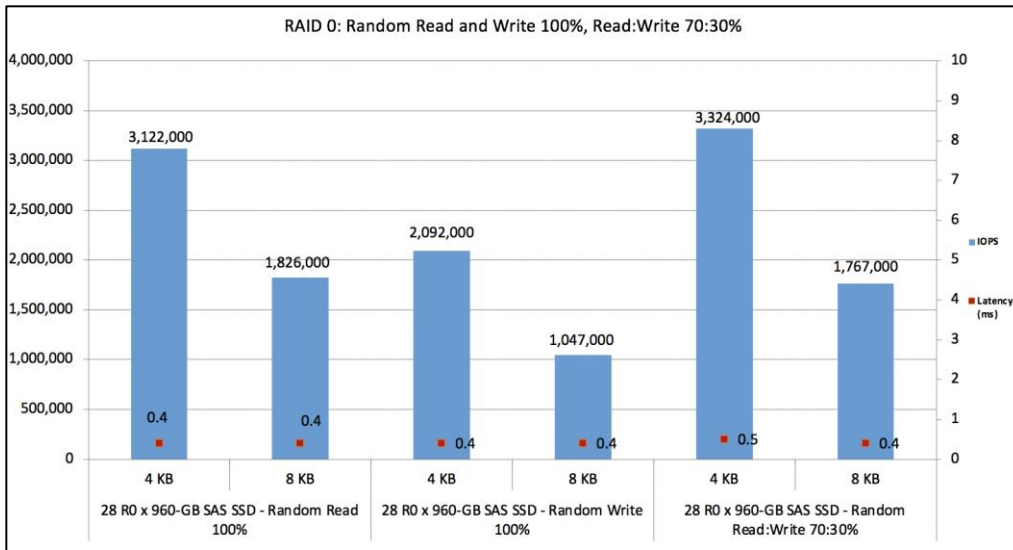
Figure 2.
Random read and write 100%; read:write 70:30%

Figure 3 shows the performance of the SSDs under test for the RAID 0 configuration with a 100-percent sequential read and write access pattern. The graph shows that the 960 GB enterprise value SAS SSD drive provides performance of more than 15,000 MBps for sequential read and more than 14,500 MBps for sequential write, with 256 KB and 1 MB block sizes within a latency of 2 milliseconds.
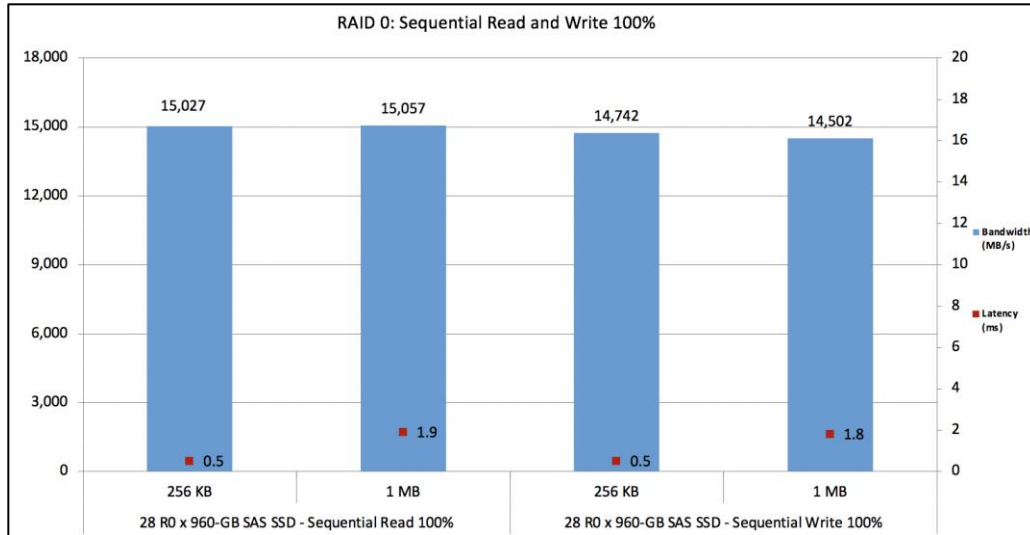
Figure 3.
Sequential read and write 100%

SAS SSD RAID 5 performance

Figure 4 shows the performance of the SSDs under test for the RAID 5 configuration with a 70:30- and 50:50-percent random read:write access pattern. The graph shows the comparative performance values achieved for enterprise value SAS SSD drives to help customers understand the performance for the RAID 5 configuration. The graph shows that the 960 GB enterprise value SAS drive provides performance of close to 800K IOPS with latency of 0.8 milliseconds for random read:write 70:30 percent and close to 480K IOPS with latency of 0.7 milliseconds for random read:write 50:50 percent with a 4 KB block size.

Each RAID 5 virtual drive was carved out using 7 disks, and a total of 4 virtual drives were created using 28 disks. The performance data shown in the graph is cumulative data by running tests on all four virtual drives in parallel.
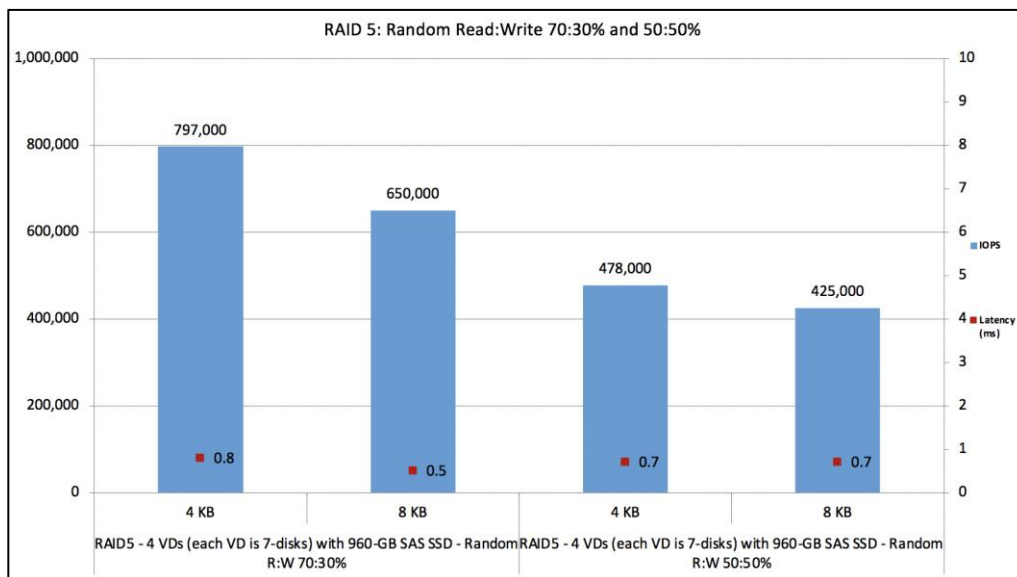
Figure 4.
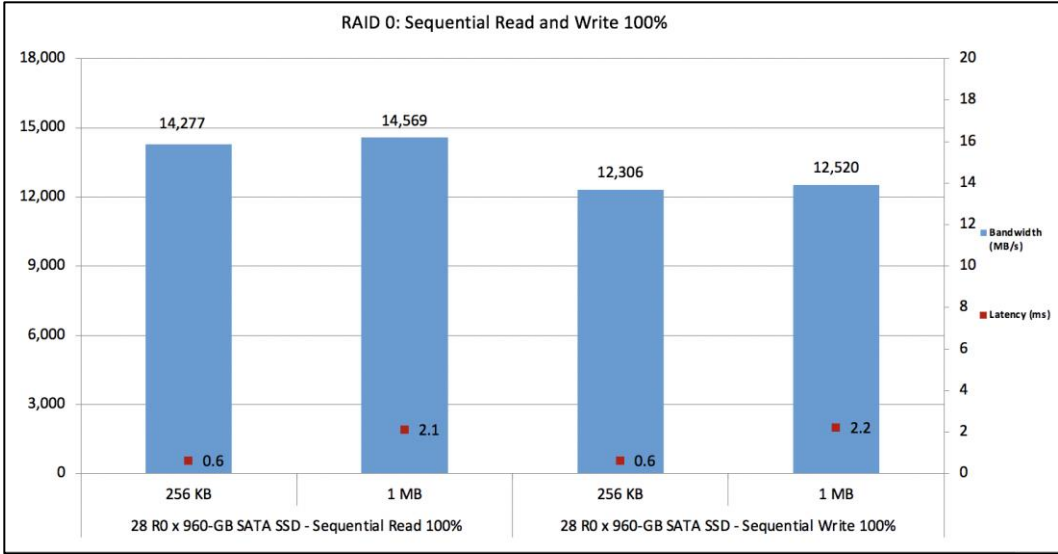Random read:write 70:30% and 50:50%

SAS SSD RAID 10 performance

Figure 5 shows the performance of the SSDs under test for the RAID 10 configuration with a 100-percent sequential read and write access pattern. The graph shows the comparative performance values achieved for enterprise value SAS SSD drives to help customers understand the performance for the RAID 10 configuration. The graph shows that the 960 GB enterprise value SAS drive provides performance of close to 15,000 MBps with latency of 0.7 milliseconds for sequential read and close to 7,400 MBps with latency of 0.9 milliseconds for sequential write with 256 KB and a 1 MB block size.

Each RAID 10 virtual drive was carved out using 4 disks with 2 spans of 2 disks each, and a total of 7 virtual drives were created using 28 disks. The performance data shown in the graph is cumulative data by running tests on all 7 virtual drives in parallel.
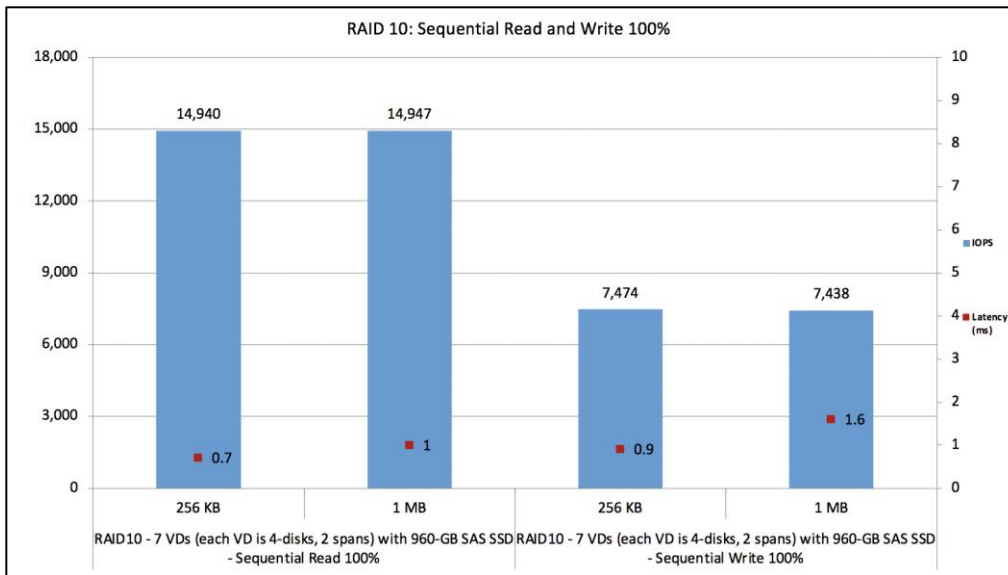
Figure 5.
Sequential read and write 100%

SAS SSD HBA performance

The C240 M6SX server supports the dual 12G SAS HBA adapters. Each adapter manages 14 disks on a 28-disk system. This configuration helps to scale performance from individual adapters, based on the performance specification of the drives. This dual adapter configuration is for performance, not for any fault tolerance between adapters.

Figure 6 shows the performance of the SSDs under test for the JBOD configuration with the HBA adapter for the 100-percent random read and random write and 70:30-percent random read:write access patterns. The graph shows that the 960 GB enterprise value SAS drive provides performance of close to 6M IOPS with latency of 0.2 milliseconds for random read, 2M IOPS with latency of 0.3 milliseconds for random write, and 5.4M IOPS with latency of 0.5 milliseconds for random read:write 70:30 percent with a 4 KB block size.
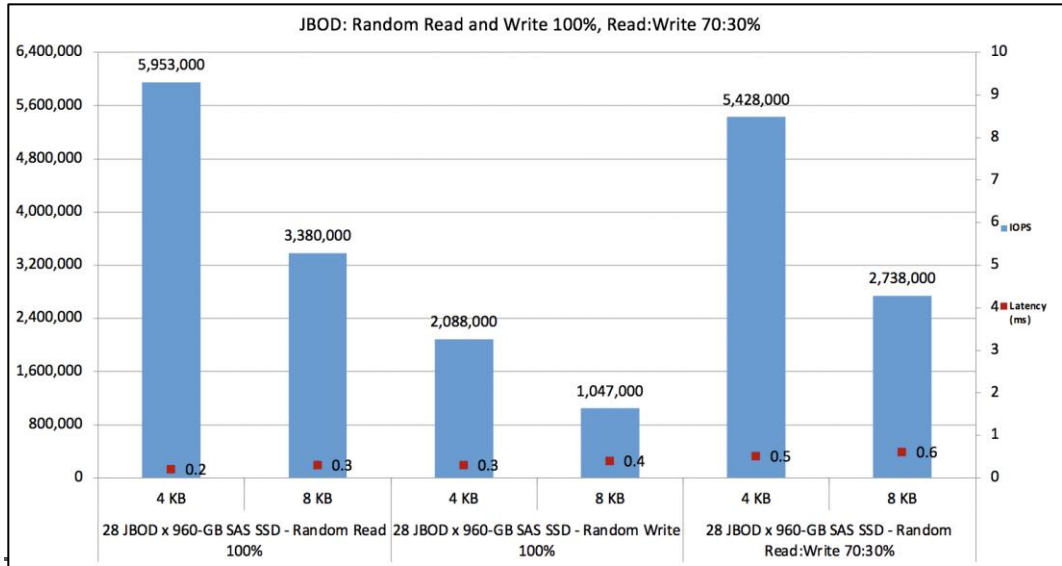
Figure 6.
Random read and write 100%, read:write 70:30%

Figure 7 shows the performance of the SSDs under test for the JBOD configuration with a 100-percent sequential read and write access pattern. The graph shows that the 960 GB enterprise value SAS SSD drive provides performance of more than 29,000 MBps for sequential read and write with 256 KB and 1 MB block sizes within latency of 2 milliseconds.
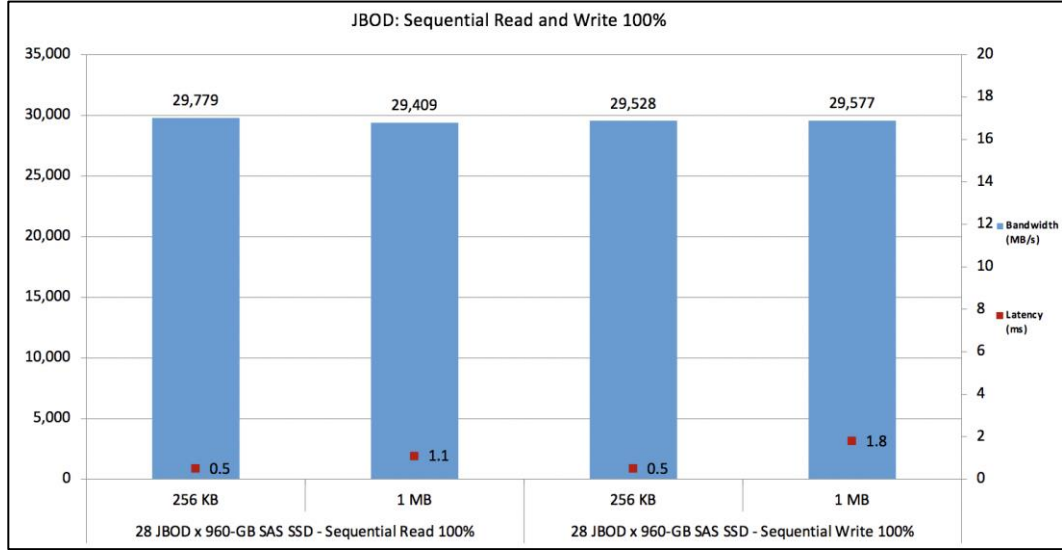
Figure 7.
Sequential read and write 100%

SATA SSD (960 GB enterprise value 6G SATA SSD) performance for 28-disk configuration

SATA SSD RAID 0 performance

Figure 8 shows the performance of the SSDs under test for the RAID 0 configuration with a 100-percent random read and random write and 70:30-percent random read:write access pattern. The graph shows the comparative performance values achieved for enterprise value SATA SSD drives to help customers understand the performance for the RAID 0 configuration. The graph shows that the 960 GB enterprise value SATA drive provides performance of more than 1.5M IOPS with latency of 0.6 milliseconds for random read, close to 800K IOPS with latency of 0.5 milliseconds for random write, and 1.5M IOPS for random read:write, 70:30 percent with a 4 KB block size. Latency is the time taken to complete a single I/O request from the viewpoint of the application.
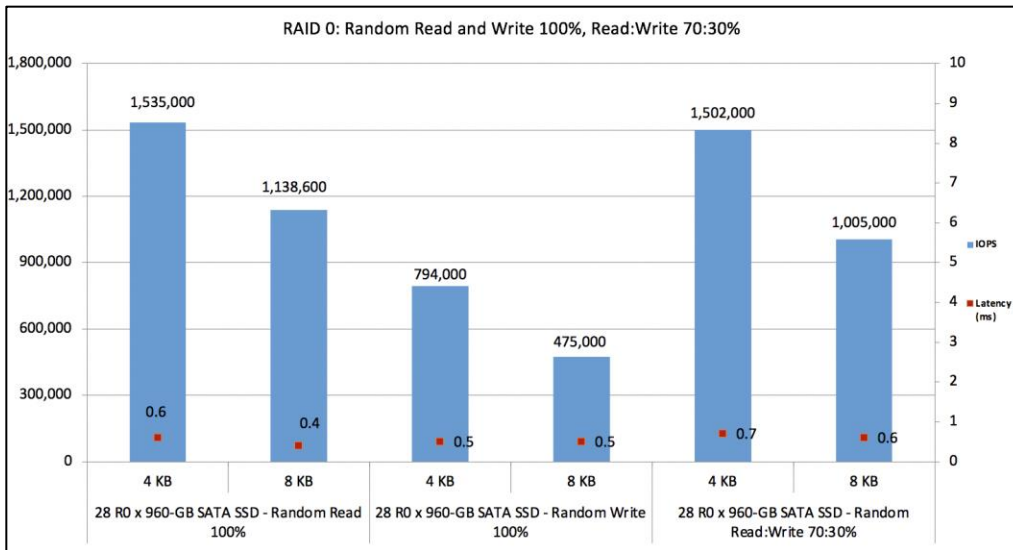
Figure 8.
Random read and write 100%, read:write 70:30%

Figure 9 shows the performance of the SSDs under test for the RAID 0 configuration with a 100-percent sequential read and write access pattern. The graph shows that the 960 GB enterprise value SATA SSD drive provides performance of more than 14,000 MBps with latency of 0.6 milliseconds for sequential read and more than 12,000 MBps for sequential write with 256 KB and 1 MB block sizes.
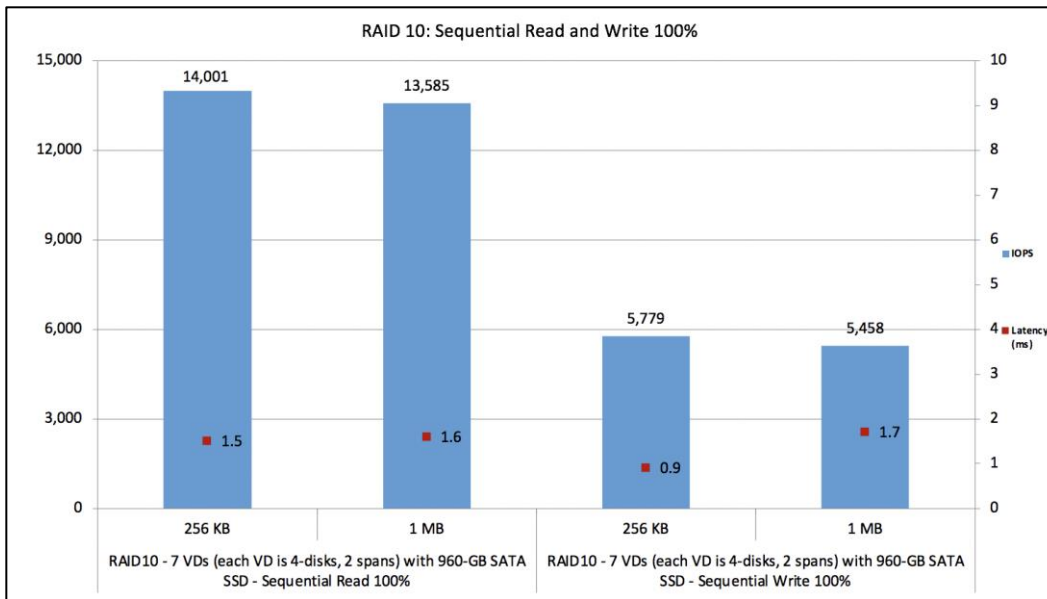
Figure 9.
Sequential read and write 100%

SATA SSD RAID 5 performance

Figure 10 shows the performance of the SSDs under test for the RAID 5 configuration with 70:30- and 50:50-percent random read:write access patterns. The graph shows the comparative performance values achieved for enterprise value SATA SSD drives to help customers understand the performance for the RAID 5 configuration. The graph shows that the 960 GB SATA drive provides performance of close to 760K IOPS with latency of 0.8 milliseconds for random read:write 70:30 percent and 470K IOPS with latency of 0.8 milliseconds for random read:write 50:50 percent with a 4 KB block size.

Each RAID 5 virtual drive was carved out using 7 disks, and a total of 4 virtual drives were created using 28 disks. The performance data shown in the graph is cumulative data by running tests on all 4 virtual drives in parallel.
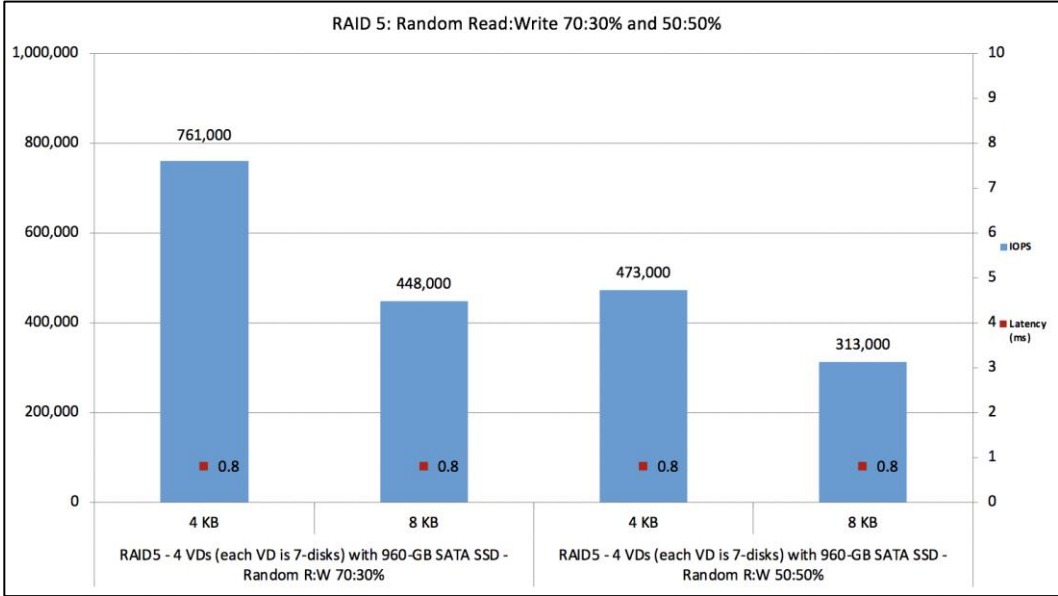
Figure 10.
Random read:write 70:30% and 50:50%

SATA SSD RAID 10 performance

Figure 11 shows the performance of the SSDs under test for the RAID 10 configuration with a 100-percent sequential read and write access pattern. The graph shows the comparative performance values achieved for enterprise value SATA SSD drives to help customers understand the performance for the RAID 10 configuration. The graph shows that the 960 GB SATA drive provides performance of close to 14,000 MBps with latency of 1.5 milliseconds for sequential read and close to 5,700 MBps with latency of 0.9 milliseconds for sequential write with a 256 KB block size.

Each RAID 10 virtual drive was carved out using 4 disks with 2 spans of 2 disks each, and a total of 7 virtual drives were created using 28 disks. The performance data shown in the graph is cumulative data by running tests on all 7 virtual drives in parallel.

Figure 11.
Sequential read and write 100%

SATA SSD HBA performance

The C240 M6SX server supports dual 12G SAS HBA adapters. Each adapter manages 14 disks on a 28-disk system. This configuration helps in scaling of performance from individual adapters, based on the performance specification of drives. This dual adapter configuration is for performance, not for any fault tolerance between adapters.

Figure 12 shows the performance of the SSDs under test for the JBOD configuration with a 100-percent random read and random write and 70:30-percent random read:write access pattern. The graph shows that the 960 GB enterprise value SATA drive provides performance of more than 2.1M IOPS with latency of 0.4 milliseconds for random read, 810K IOPS with latency of 0.5 milliseconds for random write, and 1.8M IOPS for random read:write 70:30 percent with a 4 KB block size. Latency is the time taken to complete a single I/O request from the viewpoint of the application.
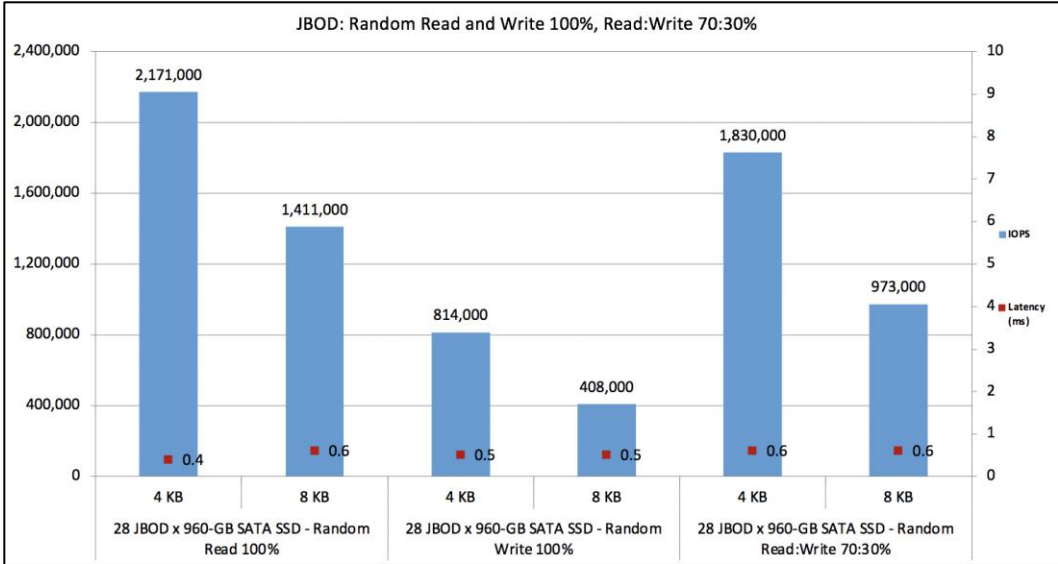
Figure 12.
Random read and write 100%, read:write 70:30%

Figure 13 shows the performance of the SSDs under test for the JBOD configuration with a 100-percent sequential read and write access pattern. The graph shows that the 960 GB enterprise value SATA SSD drive provides performance of more than 15,000 MBps with latency of 0.9 milliseconds for sequential read and more than 13,000 MBps for sequential write with the 256 KB and 1 MB block sizes.
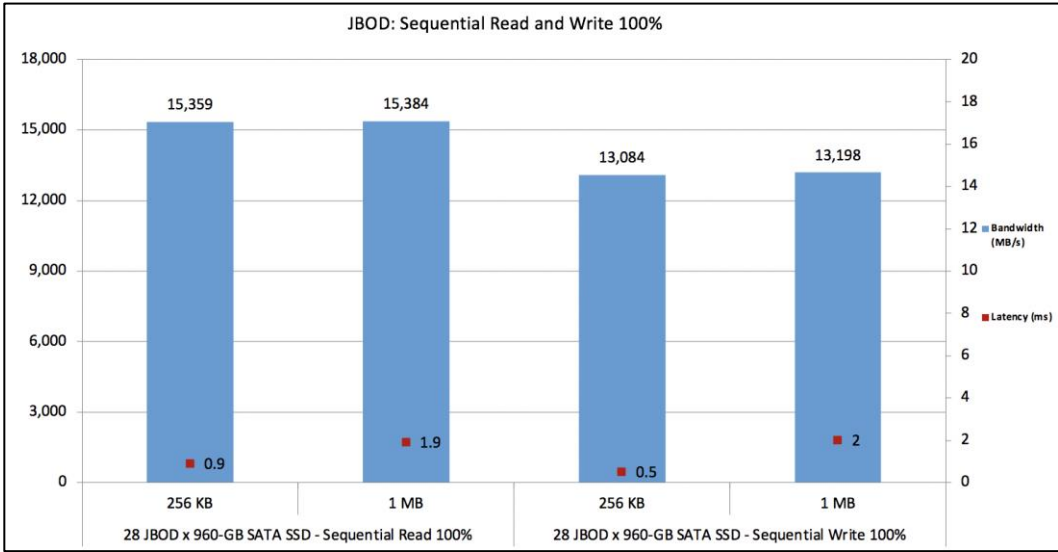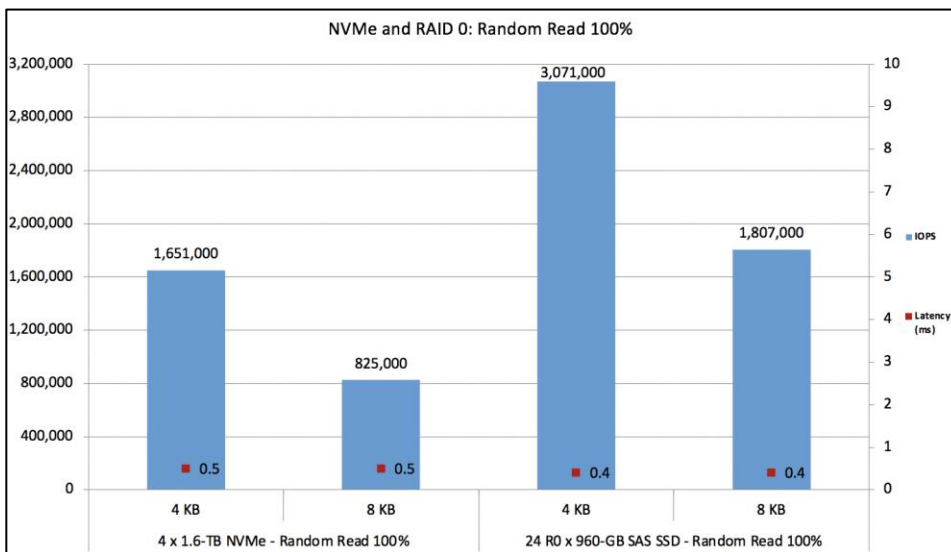

Figure 13.
Sequential read and write 100%

NVMe and SSD combined performance results

The C240 M6SX server supports up to 8 NVMe SSDs, 4 drives in the first four 4 slots and 4 drives in the rear slots. These NVMe drives are CPU-managed PCIe with 4 lanes. This section shows the performance achieved on the C240 M6SX server with a combination of 4 NVMe SSDs and 24 SAS or 24 SATA SSDs managed by either the RAID adapter or HBA adapter configuration. This information gives an idea of what can be the overall server performance with a combination of NVMe and SSDs. The combined performance can be even higher if all 4 slots are configured with high-performing NVMe SSDs. The configurations with 8 NVMe drives populated is expected to increase overall performance further (not in this scope of the current version of the document).

For the performance data updated in the following section of this document, we have considered NVMe SSDs populated only on the front 4 slots.

NVMe and SAS SSD combined performance with RAID

Figure 14 shows the performance of the 4 NVMe SSDs (1.6 TB Intel P5600 NVMe high-performance medium endurance) on the first 4 front slots and 24 SAS SSDs (960 GB enterprise value 12G SAS SSDs) on the remaining slots under test for a 100-percent random read access pattern. The graph shows the performance values achieved with a combination of NVMe and SAS SSDs to help customers understand the performance trade-off when choosing this configuration. SAS SSDs are configured as single-drive RAID 0 volumes (24 R0 volumes). The graph shows that the NVMe SSD provides performance of more than 1.6M IOPS with latency of 0.4 milliseconds and the R0 SAS SSD 3M IOPS with latency of 0.4 milliseconds for random read with a 4 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.

Figure 14.
Random read 100%

Figure 15 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on remaining slots under test for a 100-percent random write access pattern. The graph shows that the 1.6 TB NVMe SSD provides performance of 530K IOPS with latency of 0.5 milliseconds and the R0 SAS SSD close to 1.8M IOPS with latency of 0.5 milliseconds for random write with a 4 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.
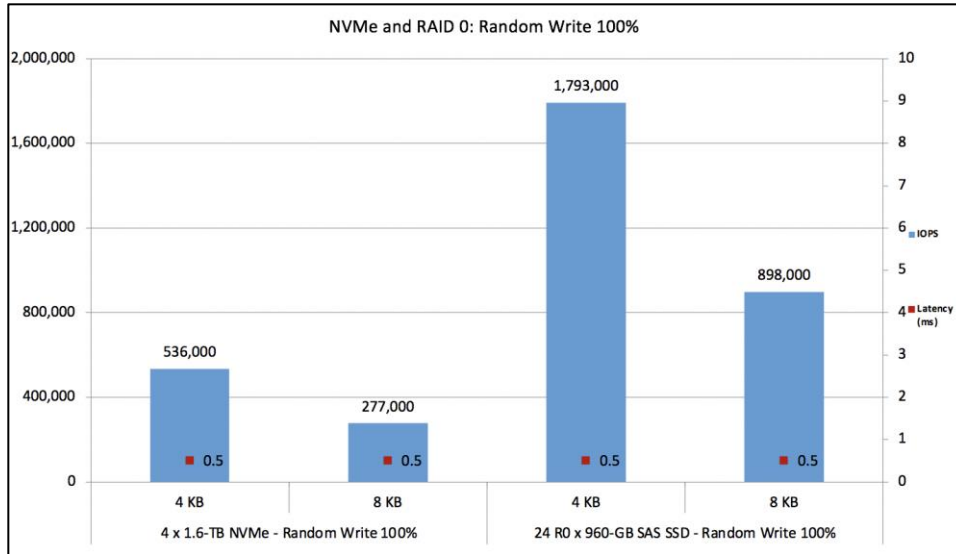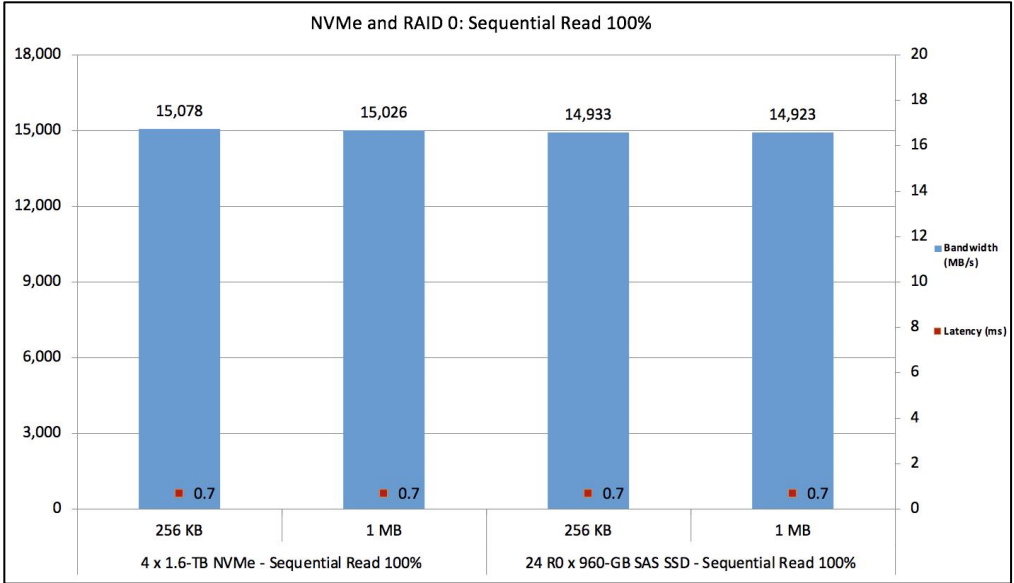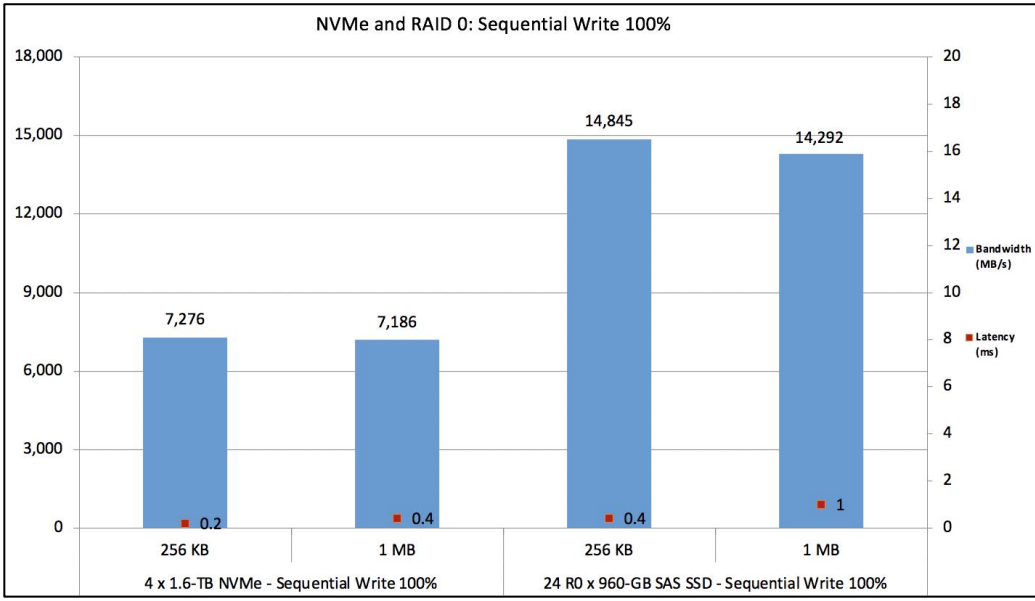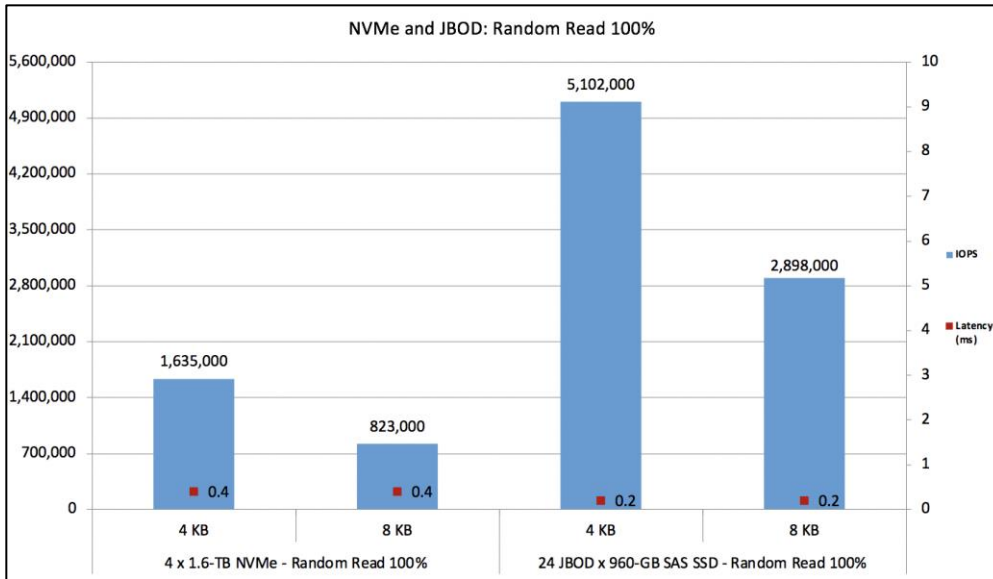
Figure 15.
Random write 100%

Figure 16 shows the performance of the 4 NVMe SSDs on first 4 slots and 24 SAS SSDs on remaining slots under test for a 70:30-percent random read:write access pattern. The graph shows that the NVMe SSD provides performance of 930K IOPS with latency of 0.5 milliseconds and the R0 SAS SSD provides performance of 3.2M IOPS with latency of 0.4 milliseconds for random read:write 70:30 with a 4 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.
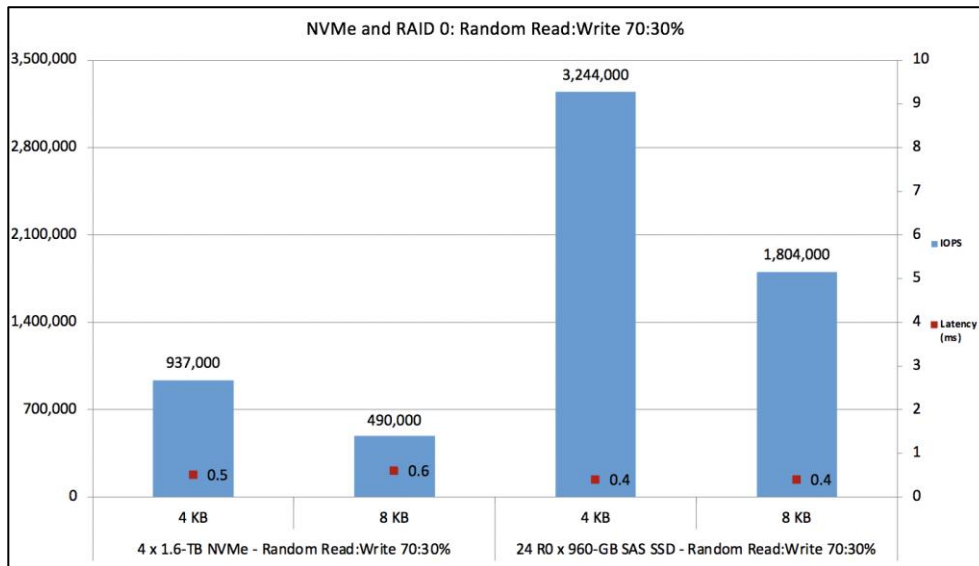
Figure 16.
Random read:write 70:30%

Figure 17 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 100-percent sequential read access pattern. The graph shows that the NVMe SSD provides performance of 15,000 MBps with latency of 0.7 milliseconds and the R0 SAS SSD close to 15,000 MBps with latency of 0.7 milliseconds for sequential read with a 256 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.
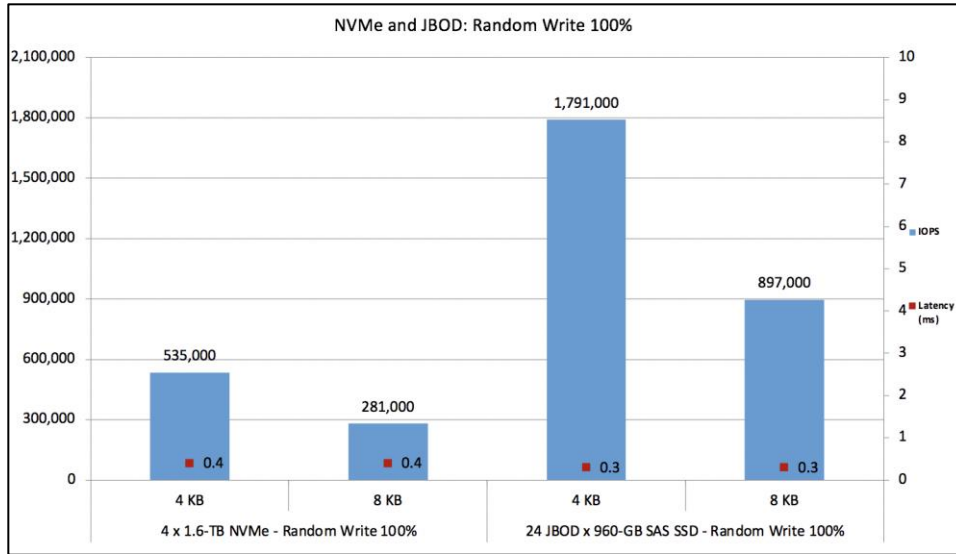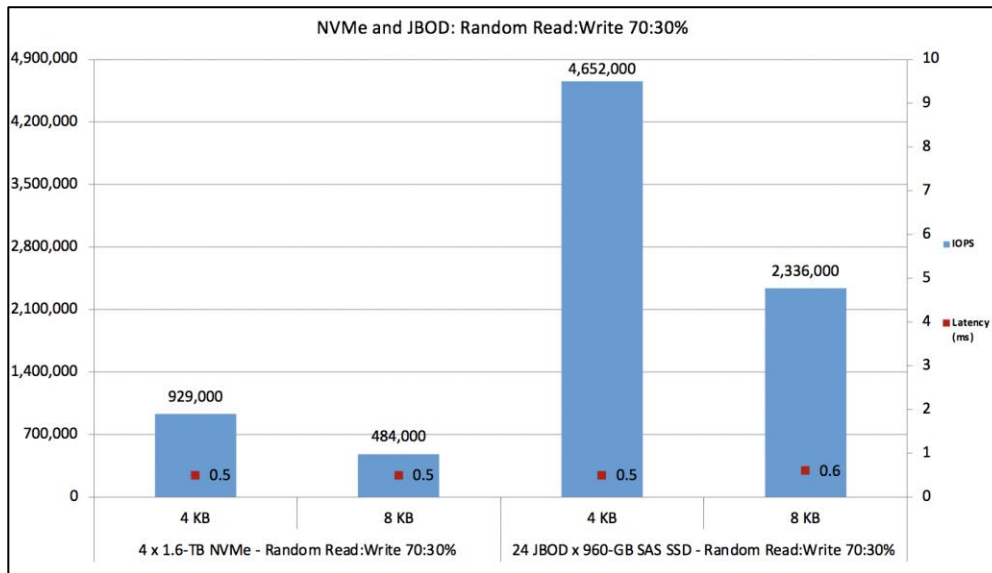
Figure 17.
Sequential read 100%

Figure 18 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 100-percent sequential write access pattern. The graph shows that the NVMe SSD provides performance of 7,200 MBps with latency of 0.2 milliseconds and the R0 SAS SSD 14,800 MBps with latency of 0.4 milliseconds for sequential write with a 256 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.

Figure 18.
Sequential write 100%
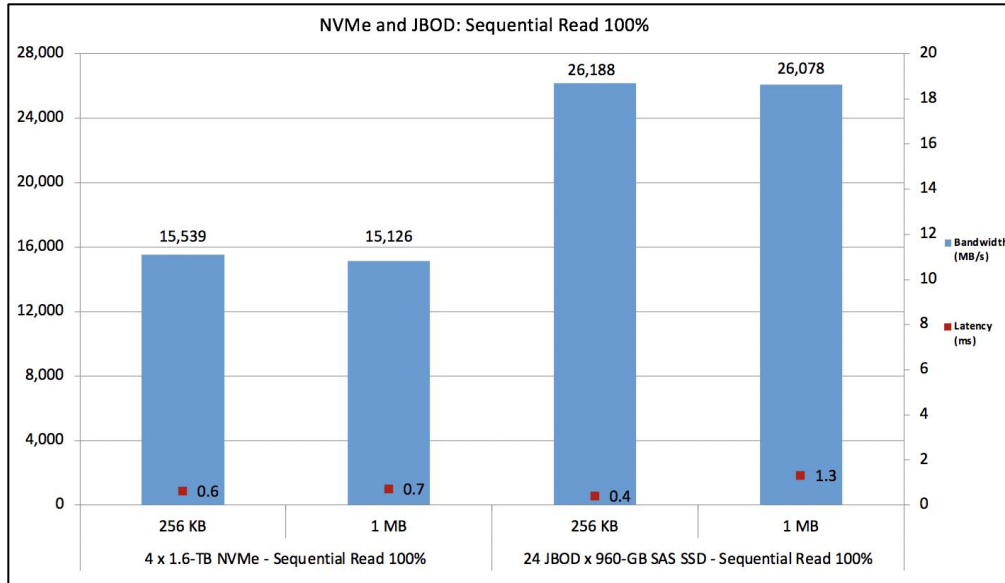
NVMe and SAS SSD combined performance with HBA

Figure 19 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 100-percent random read access pattern. The graph shows the performance values achieved with a combination of NVMe and SAS SSDs to help customers understand the performance trade-off when choosing this configuration. SAS SSDs are configured as JBOD managed by an HBA adapter. The graph shows that the 1.6 TB NVMe SSD provides performance of more than 1.6M IOPS with latency of 0.4 milliseconds and the JBOD SAS SSD 5.1M IOPS with latency of 0.2 milliseconds for random read with a 4 KB block size. Tests on the NVMe and SAS SSDs were run in parallel. SAS SSDs are managed by dual HBA adapters.
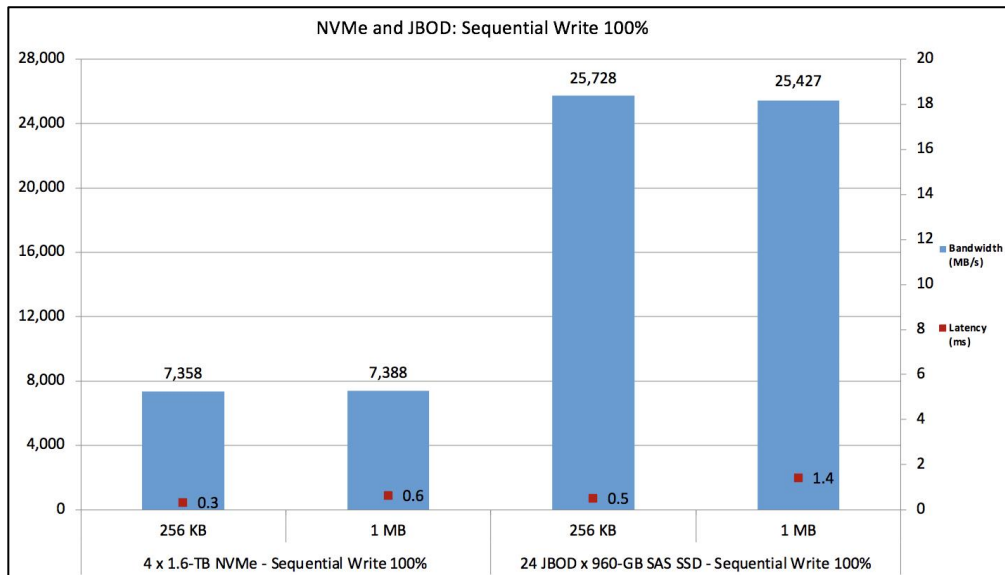

Figure 19.
Random read 100%

Figure 20 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 100-percent random write access pattern. The graph shows that the 1.6 TB NVMe SSD provides performance of 530K IOPS with latency of 0.4 milliseconds and the JBOD SAS SSD close to 1.8 M IOPS with latency of 0.3 milliseconds for random write with a 4 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.
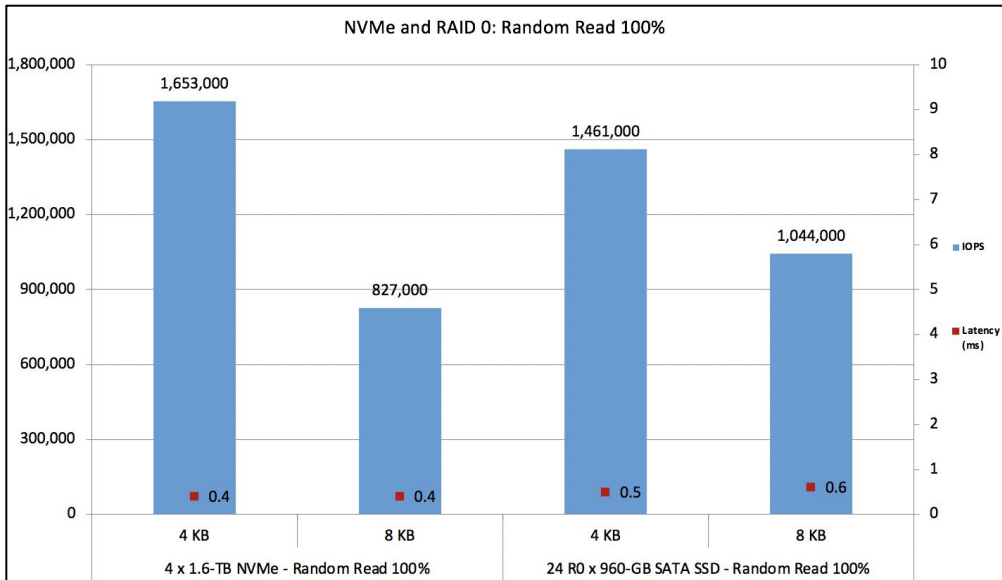
Figure 20.
Random write 100%

Figure 21 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 70:30-percent random read:write access pattern. The graph shows that the NVMe SSD provides performance of 930K IOPS with latency of 0.5 milliseconds and the JBOD SAS SSD 4.6M IOPS with latency of 0.5 milliseconds for random read:write 70:30 with a 4 KB block size. Tests on the NVMe and SAS SSDs were run in parallel.

Figure 21.
Random read:write 70:30%

Figure 22 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 100-percent sequential read access pattern. The graph shows that the NVMe SSD provides performance of 15,000 MBps with latency of 0.6 milliseconds and the JBOD SAS SSD close to 26,000 MBps with latency of 0.4 milliseconds for sequential read with 256 KB and 1 MB block sizes. Tests on the NVMe and SAS SSDs were run in parallel.
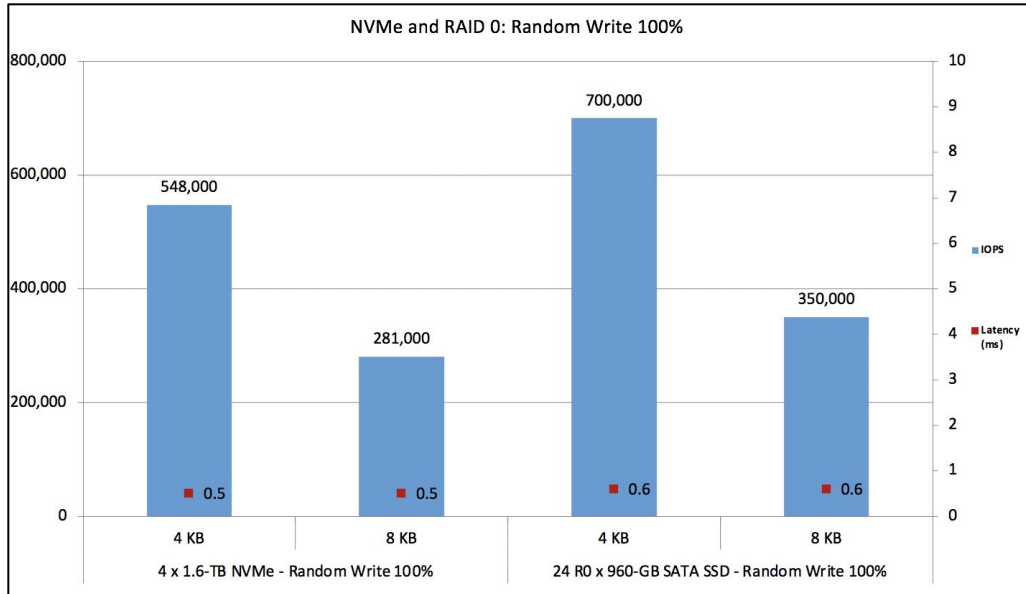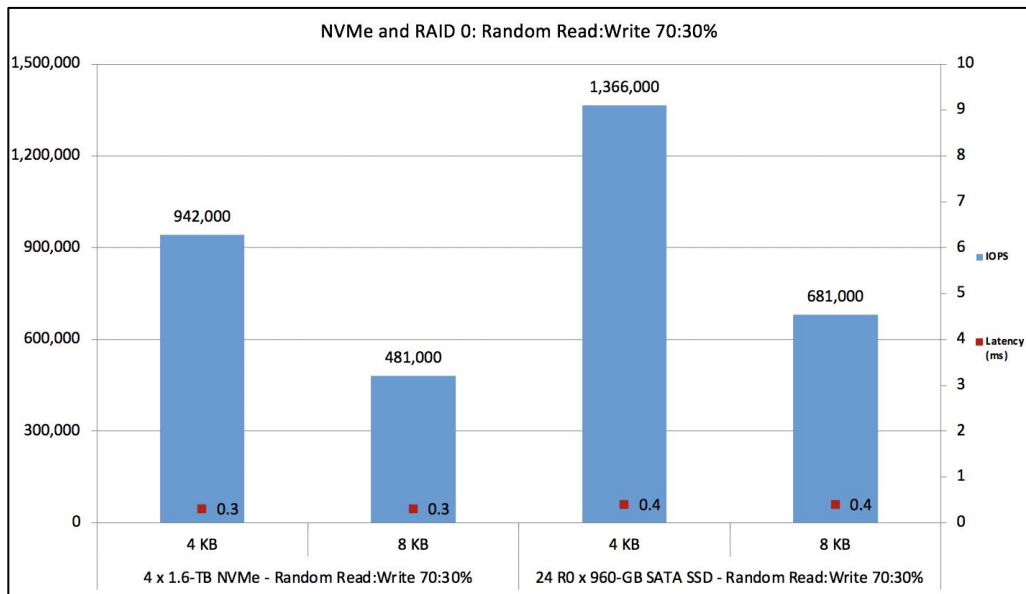
Figure 22.
Sequential read 100%

Figure 23 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SAS SSDs on the remaining slots under test for a 100-percent sequential write access pattern. The graph shows that the NVMe SSD provides performance of 7,300 MBps with latency of 0.3 milliseconds and the JBOD SAS SSD 25,000 MBps with latency of 0.5 milliseconds for sequential write with 256 KB and 1 MB block sizes. Tests on the NVMe and SAS SSDs were run in parallel.

Figure 23.
Sequential write 100%

NVMe and SATA SSD combined performance with RAID

Figure 24 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent random read access pattern. SATA SSDs are configured as single-drive RAID 0 volumes (24 R0 volumes). The graph shows that the 1.6 TB NVMe SSD provides performance of more than 1.6M IOPS with latency of 0.4 milliseconds and the R0 SATA SSD 1.46 M IOPS with latency of 0.5 milliseconds for random read with a 4 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.

Figure 24.
Random read 100%

Figure 25 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent random write access pattern. The graph shows that the 1.6 TB NVMe SSD provides performance of 540K IOPS with latency of 0.5 milliseconds and the R0 SATA SSD close to 700K IOPS with latency of 0.6 milliseconds for random write with a 4 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.
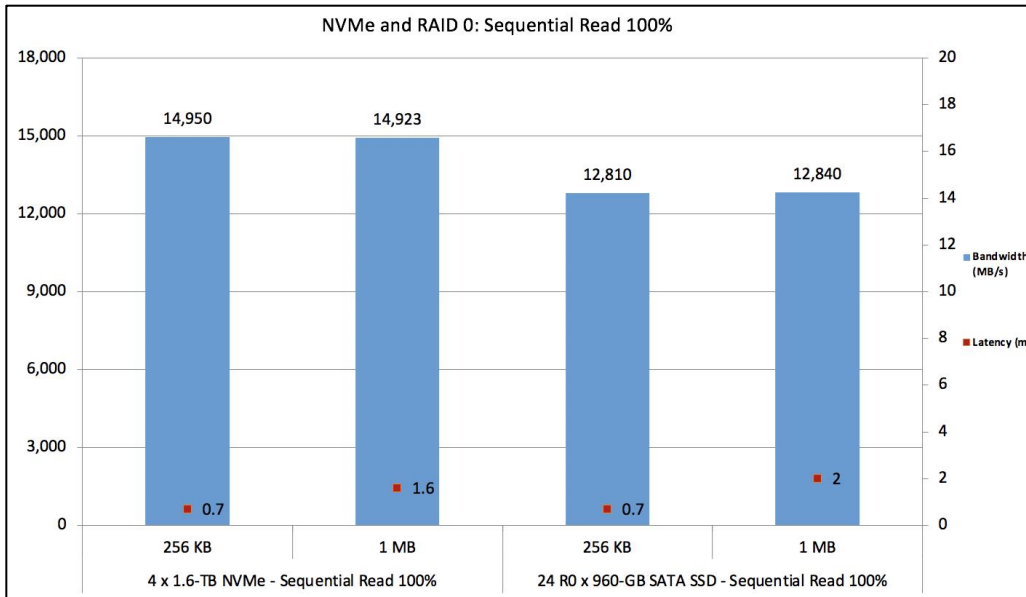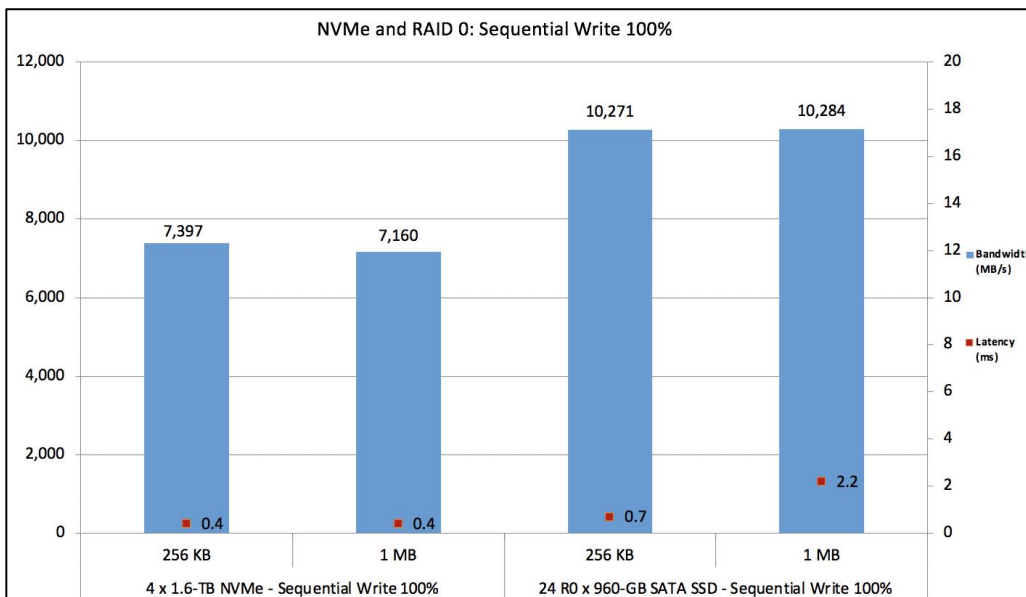
Figure 25.
Random write 100%

Figure 26 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 70:30-percent random read:write access pattern. The graph shows that the NVMe SSD provides performance of 940K IOPS with latency of 0.3 milliseconds and the R0 SATA SSD 1.3M IOPS with latency of 0.4 milliseconds for random read:write 70:30 with a 4 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.

Figure 26.
Random read:write 70:30%

Figure 27 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent sequential read access pattern. The graph shows that the NVMe SSD provides performance of close to 15,000 MBps with latency of 0.7 milliseconds and the R0

SATA SSD 12,800 MBps with latency of 0.7 milliseconds for sequential read with a 256 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.

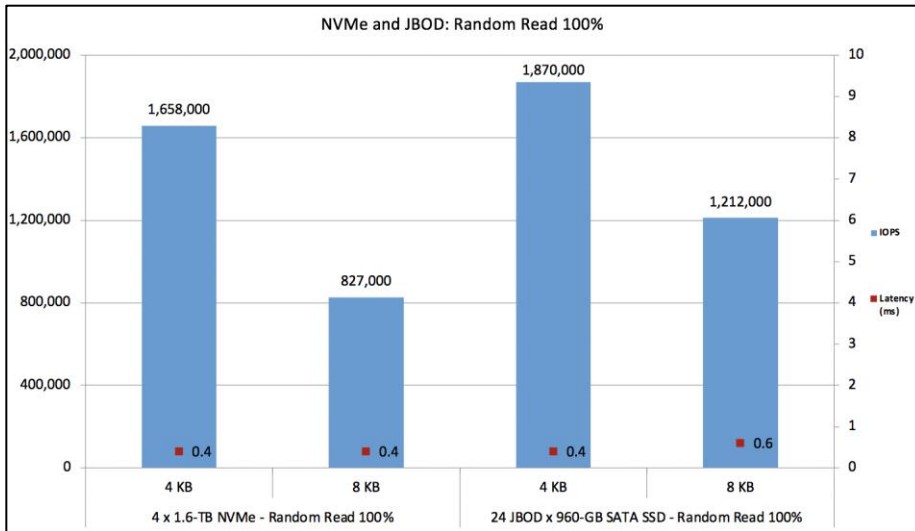**Figure 27.**
Sequential read 100%

Figure 28 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent sequential write access pattern. The graph shows that the NVMe SSD provides performance of 7,300 MBps with latency of 0.4 milliseconds and the R0 SATA SSD 10,200 MBps with latency of 0.7 milliseconds for sequential write with a 256 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.
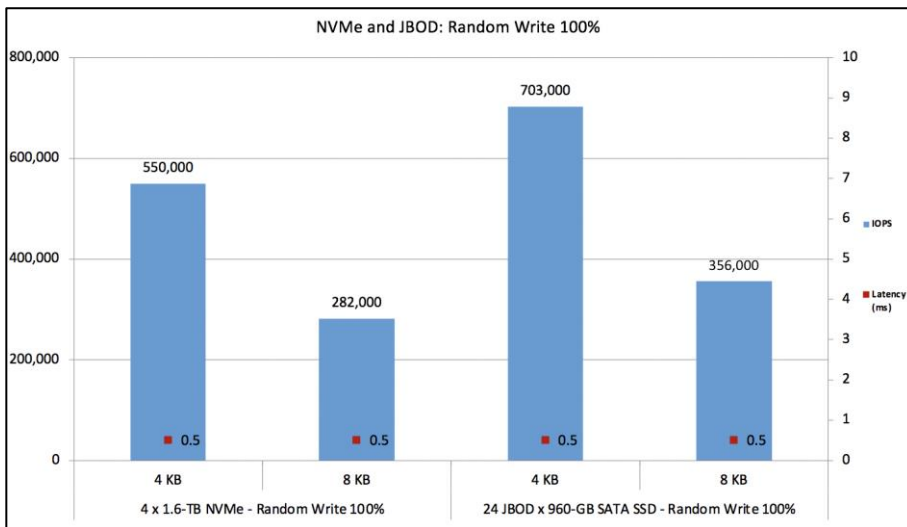

Figure 28.
Sequential write 100%

NVMe and SATA SSD combined performance with HBA

Figure 29 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent random read access pattern. SATA SSDs are configured as JBOD managed by an HBA adapter. The graph shows that the 1.6 TB NVMe SSD provides performance of more than 1.6M IOPS with latency of 0.4 milliseconds and the JBOD SATA SSD 1.8M IOPS with latency of 0.4 milliseconds for random read with a 4 KB block size. Tests on the NVMe and SATA SSDs were run in parallel. SATA SSDs are managed by dual HBA adapters.
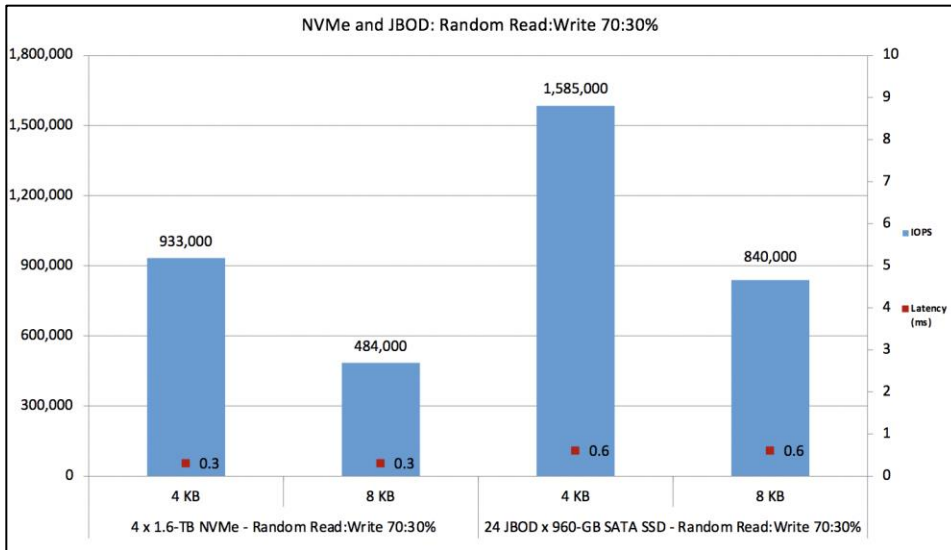

Figure 29.
Random read 100%

Figure 30 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent random write access pattern. The graph shows that the 1.6 TB NVMe SSD provides performance of 550K IOPS with latency of 0.5 milliseconds and the JBOD SATA SSD close to 700K IOPS with latency of 0.5 milliseconds for random write with a 4 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.
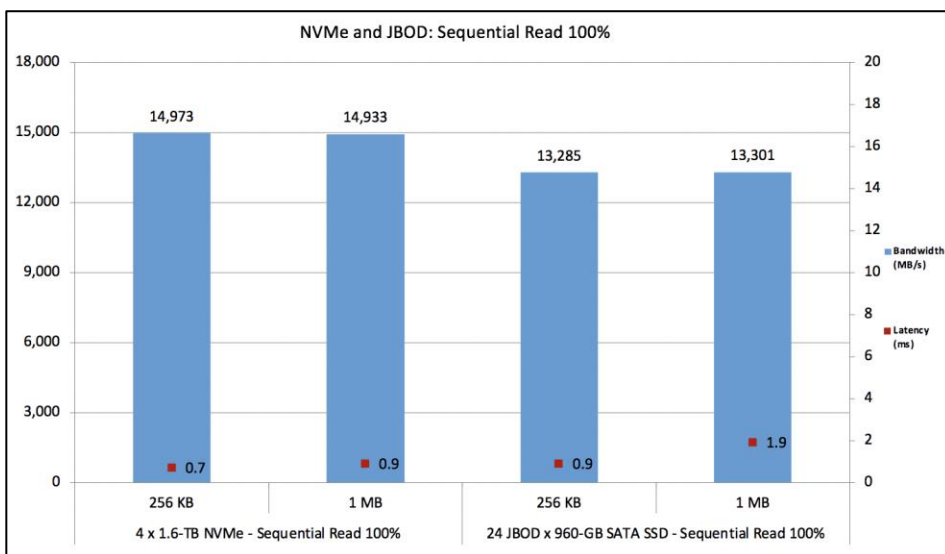

Figure 30.
Random write 100%

Figure 31 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 70:30-percent random read:write access pattern. The graph shows that the NVMe SSD provides performance of 930K IOPS with latency of 0.3 milliseconds and the JBOD SATA SSD close to 1.6M IOPS with latency of 0.6 milliseconds for random read:write 70:30 with a 4 KB block size. Tests on the NVMe and SATA SSDs were run in parallel.
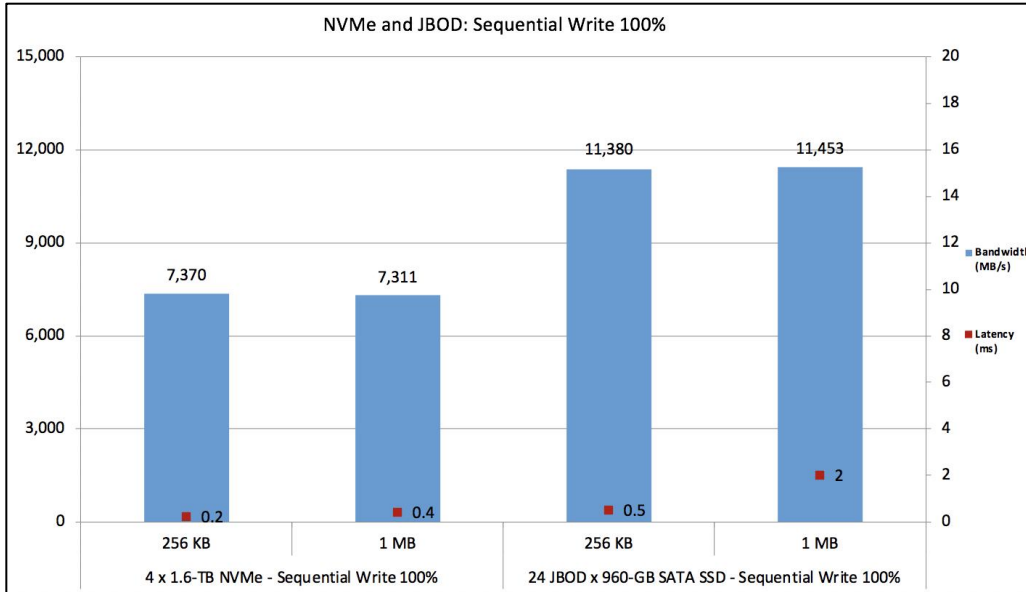

Figure 31.
Random read:write 70:30%

Figure 32 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent sequential read access pattern. The graph shows that the NVMe SSD provides performance of close to 15,000 MBps with latency of 0.7 milliseconds and the JBOD SATA SSD close to 13,200 MBps with latency of 0.9 milliseconds for sequential read with 256 KB and 1 MB block sizes. Tests on the NVMe and SATA SSDs were run in parallel.

Figure 32.
Sequential read 100%

Figure 33 shows the performance of the 4 NVMe SSDs on the first 4 slots and 24 SATA SSDs on the remaining slots under test for a 100-percent sequential write access pattern. The graph shows that the NVMe SSD provides performance of 7,300 MBps with latency of 0.2 milliseconds and the JBOD SATA SSD 11,300 MBps with latency of 0.5 milliseconds for sequential write with 256 KB and 1 MB block sizes. Tests on the NVMe and SATA SSDs were run in parallel.

Figure 33.
Sequential write 100%

For more information

For additional information, refer to: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/at-a-glance-c45-2377656.html.

For information on BIOS tunings for different workloads, refer BIOS tuning guide: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.pdf

For information on FIO tool, refer to: https://fio.readthedocs.io/en/latest/fio_doc.html

Appendix: Test environment

Table 8 lists the details of the server under test.

Table 8.   Server properties

Name	Value
Product names	Cisco UCS C240 M6SX
CPUs	CPU: Two 2.00-GHz Intel Xeon Gold 6330
Number of cores	28
Number of threads	56
Total memory	512 GB
Memory DIMMs (16)	32 GB x 16 DIMMs
Memory speed	2933 MHz
Network controller	Cisco LOM X550-TX; Two 10 Gbps interfaces
VIC adapter	Cisco UCS VIC 1477 mLOM 40/100 Gbps Small Form-Factor Pluggable Plus (SFP+)
RAID controllers	• Cisco M6 12G SAS RAID Controller with 4 GB flash-backed write cache (FBWC; UCSC-RAID-M6SD) • Cisco 12G SAS HBA (UCSC-SAS-M6T)
SFF SSDs	• 960 GB 2.5-inch enterprise value 12G SAS SSD (UCS-SD960GK1X-EV) • 960 GB 2.5-inch enterprise value 6G SATA SSD (UCS-SD960G61X-EV)
SFF NVMe SSDs	• 1.6 TB 2.5-inch Intel P5600 NVMe High Perf Medium Endurance (UCSC-NVMEI4-I1600)

Table 9 lists the server BIOS settings applied for disk I/O Testing

Table 9.   Server BIOS settings

Name	Value
BIOS version	Release 4.2.1c
Cisco Integrated Management Controller (Cisco IMC) version	4.2(1a)
Cores enabled	All
Hyper-Threading (All)	Enable
Hardware prefetcher	Enable
Adjacent-cache-line prefetcher	Enable

Name	Value
DCU streamer	Enable
DCU IP prefetcher	Enable
NUMA	Enable
Memory refresh enable	1x Refresh
Energy-efficient turbo	Enable
Turbo mode	Enable
EPP profile	Performance
CPU C6 report	Enable
Package C state	C0/C1 state
Power Performance Tuning	OS controls EPB
Workload configuration	I/O sensitive

Note: Rest of the BIOS settings are platform default values

Printed in USA

221220   09/21