White Paper

Cisco C240 M4 Rack Server Disk I/O Characterization

December 2016

Contents

Executive Summary

This document outlines the I/O performance characteristics of the Cisco UCS® C240 M4 Rack Server using the Cisco® 12-Gbps SAS modular RAID controller (UCSC-MRAID12G). Performance comparisons of various serial-attached Small Computer System Interface (SCSI; SAS) solid-state disk (SSD) drives, hard-disk drives (HDD), Redundant Array of Independent Disk (RAID) configurations, and controller options 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 and RAID levels to meet their individual I/O workload needs.

Performance data was obtained using the Iometer 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 and HDDs, and large form-factor (LFF) HDDs with configurations of RAID-0, RAID-5, and RAID-10 virtual disks.

Introduction

The widespread adoption of virtualization and data center consolidation technologies has had a profound impact on the efficiency of the data center. Virtualization brings with it added 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. The hard disks not only add a security factor but also are 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 also so 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 hard disk failure can be accommodated.

Conventional wisdom holds that 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 an enterprise's most common type of workload. Currently, data centers are dominated by random and sequential workloads, brought in by 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 data centers to fully utilize 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 utilization and increasing I/O:** Multicore processors with virtualized server and desktop architectures increase processor utilization, increasing the I/O demand per server. In a virtualized data center, it is the I/O performance that limits server consolidation ratio, not CPU or memory.
- **Randomization:** Virtualization has the effect of multiplexing multiple logical workloads across a single physical I/O path. The greater the virtualization achieved, the more random the physical I/O requests.

Scope of This Document

For the I/O performance characterization tests, performance evaluations are done with SSDs and HDDs with configurations of RAID-0, RAID-5, and RAID-10 virtual disks because most of the workloads targeted for Cisco UCS C240 M4 servers use these RAID levels. The Cisco UCS C240 M4S2 server used for the I/O performance characterization tests supports up to 16 HDDs and SSDs with SAS expanders, and the Cisco UCS C240 M4L server supports up to 12 LFF HDDs used for I/O performance characterization. The performance tests described here were limited to an 8-disk configuration for SSDs and SFF HDDs, and a 12-disk configuration for LFF HDDs.

Disk I/O performance for the Cisco UCS C220 M4 was tested and validated. These results are comparable to those for the Cisco UCS C240 M4 because the hardware, firmware, and disk drives remain the same.

Infrastructure Components

Hardware Components

Cisco UCS C240 M4 Overview

The Cisco UCS C240 M4 Rack Server is an enterprise-class server in a 2-rack-unit (2RU) form factor. It is designed to deliver exceptional performance, expandability, and efficiency for storage and I/O-intensive infrastructure workloads. These workloads include big data analytics, virtualization, and graphics-intensive and bare-metal applications.

The Cisco UCS C240 M4 server provides:

- Dual Intel® Xeon® processor E5-2600 v4 CPUs for improved performance suitable for nearly all 2-socket applications
- Next-generation double-data-rate 4 (DDR4) memory and 12-Gbps SAS throughput
- Innovative Cisco UCS virtual interface card (VIC) support in PCI Express (PCIe) or modular LAN-on-motherboard (mLOM) form factor
- Graphics-intensive experiences to more virtual users with support for the latest NVIDIA graphics processing units (GPUs)

The Cisco UCS C240 M4 server also offers exceptional reliability, availability, and serviceability (RAS) features, including:

- Tool-free CPU insertion
- Easy-to-use latching lid
- Hot-swappable and hot-pluggable components
- Redundant Cisco Flexible Flash (FlexFlash) Secure Digital (SD) cards

The Cisco UCS C240 M4 server can be deployed as a standalone device or as part of a managed Cisco Unified Computing System™ (Cisco UCS) environment. Cisco UCS unifies computing, networking, management, virtualization, and storage access into a single integrated architecture that can enable end-to-end server visibility, management, and control in both bare-metal and virtualized environments.

With a Cisco UCS managed deployment, the Cisco UCS C240 M4 takes advantage of our standards-based unified computing innovations to significantly reduce customers' total cost of ownership (TCO) and increase business agility.

Server Specifications

The server specifications are as follows:

- Cisco UCS C240 M4 Rack Servers
- CPU: 2 x 2.20-GHz Intel Xeon processors E5-2699 v4
- Memory: 16 x 16 GB (256 GB) DDR4
- Cisco UCS VIC 1227 mLOM 10-Gbps SFP+
- Cisco 12-Gbps SAS modular RAID controller (2-GB* cache module)
- Cisco 12-Gbps modular SAS host bus adapter (HBA)

  * All tests used a 2-GB cache module on the RAID controller with SSDs and HDDs to handle mixed environments.

Models of Cisco UCS C240 M4 Servers

The Cisco UCS C240 M4 server can be configured in four different models to match specific customer environments.

The Cisco UCS C240 M4SX provides outstanding storage expandability and performance with up to 24 x 2.5-inch 12-Gbps SFF HDDs or SSDs plus two optional internal 2.4-inch boot drives (Figure 1).

Figure 1. Cisco UCS C240 M4SX

The Cisco UCS C240 M4S2 balances cost and expandability with up to 16 x 2.5-inch 12-Gbps SFF HDDs or SSDs (Figure 2).

Figure 2. Cisco UCS C240 M4S2

The Cisco UCS C240 M4S is the cost-effective choice with up to 8 x 2.5-inch 12-Gbps SFF HDDs or SSDs (Figure 3).

Figure 3. Cisco UCS C240 M4S

The Cisco UCS C240 M4L is the choice for outstanding storage capability with up to 12 x 3.5-inch 12-Gbps LFF HDDs plus two optional internal 2.5-inch boot drives (Figure 4).

Figure 4. Cisco UCS C240 M4L

For details about how to configure specific models, please refer to the appropriate specification sheets:

- Small form-factor models (either HDD or SSD): http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m4-sff-spec-sheet.pdf
- Large form-factor model (HDD) http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m4-lff-spec-sheet.pdf

Table 1 provides part numbers for the SFF and LFF server models.

Table 1. Part Numbers of the Cisco UCS C240 M4 High-Density SFF Rack Base Server Models

Part Number	Description
UCSC-C240-M4SX	Cisco UCS C240 M4 SFF base server, with 2.5-inch 24-drive backplane and SAS expander
UCSC-C240-M4S2	Cisco UCS C240 M4 SFF base server, with 2.5-inch 16-drive backplane and SAS expander
UCSC-C240-M4S	Cisco UCS C240 M4 SFF base server, with 2.5-inch 8-drive backplane and no SAS expander
UCSC-C240-M4L	Cisco UCS C240 M4 LFF base server, with 3.5-inch 12-drive backplane and SAS expander

The performance testing described in this document uses the Cisco UCS C240 M4S2 server, which supports 16 HDDs and SSDs with SAS expanders, and the Cisco UCS C240 M4L server, which supports 12 LFF HDDs. In the disk I/O characterization described here, tests were limited to 8 SFF SSDs and SFF HDDs and 12 LFF HDDs for all RAID configurations.

Server Specifications

Table 2 summarizes the server specifications.

Table 2. Server Specifications

Item	Specifications
Chassis	2RU
Processors	Either 1 or 2 Intel Xeon processor E5-2600 v4 product family CPUs
Chip set	Intel C610 series
Memory	Up to 24 DDR4 dual in-line memory modules (DIMMs) with speeds up to 2400 MHz
PCIe slots	Up to 6 PCIe Generation 3 slots (4 full-height and full-length, 4 network connectivity status indicator (NCSI) capable and VIC ready, and 2 GPU ready)
Hard drives	Up to 24 SFF drives or 12 LFF drives, plus 2 optional internal SATA boot drives
Embedded network interface card (NIC)	Two 1-Gbps Intel i350-based Gigabit Ethernet ports
mLOM	Slot can flexibly accommodate 1-, 10-, or 40-Gbps adapters
RAID controller	• Cisco 12-Gbps SAS modular RAID controller (UCSC-MRAID12G) for internal drives • Cisco UCS 12-Gbps SAS modular HBA • Cisco 9300-8E 12-Gbps SAS HBA for external drives • Embedded software RAID (entry RAID solution) for up to 4 SATA drives

Storage Components

The Cisco C240 M4 SFF and LFF server (Figure 5) supports the following RAID controllers:

- **Embedded RAID (Intel Platform Controller Hub [PCH] SATA):** This onboard controller supports 8 internal front-facing SATA drives in two groups of four drives. This type of controller supports default RAID with RAID-0, RAID-1, and RAID-10; software RAID-5 key upgrade with RAID-0, RAID-1, RAID-5, and RAID-10; and JBOD mode. For more information about how to enable the embedded RAID option, see http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/C240M4/install/C240M4/raid.html.
- **Cisco UCS 12-Gbps SAS modular RAID controller:** This PCIe card supports 8, 12, 16, and 24 internal HDDs and SSDs. The Cisco UCS C240 M4S2 SFF model used in the tests supports 16 internal drives, and the Cisco UCS C240 M4L supports 12 internal drives with a SAS expander and provides internal SAS connectivity. The controller can be ordered with modular flash-based write cache (FBWC) options: 1, 2, or 4 GB. This controller supports all RAID levels and JBOD mode.
- **Cisco UCS 12-Gbps SAS modular HBA:** This PCIe card supports 8, 12, 16, and 24 internal HDDs and SSDs. The Cisco UCS C240 M4S2 SFF model used in the tests supports 16 internal drives, and the Cisco UCS C240 M4L supports 12 internal drives with a SAS expander and provides internal SAS connectivity. This type of controller supports JBOD mode only.
- **Cisco UCS SAS 9300-8e HBA:** This HBA supports all Cisco UCS C240 M4 server models with 8 external SAS and SATA ports, controlling up to 1024 non-RAID external drives.

Figure 5. Cisco UCS C240 M4 SFF Motherboard with RAID Options

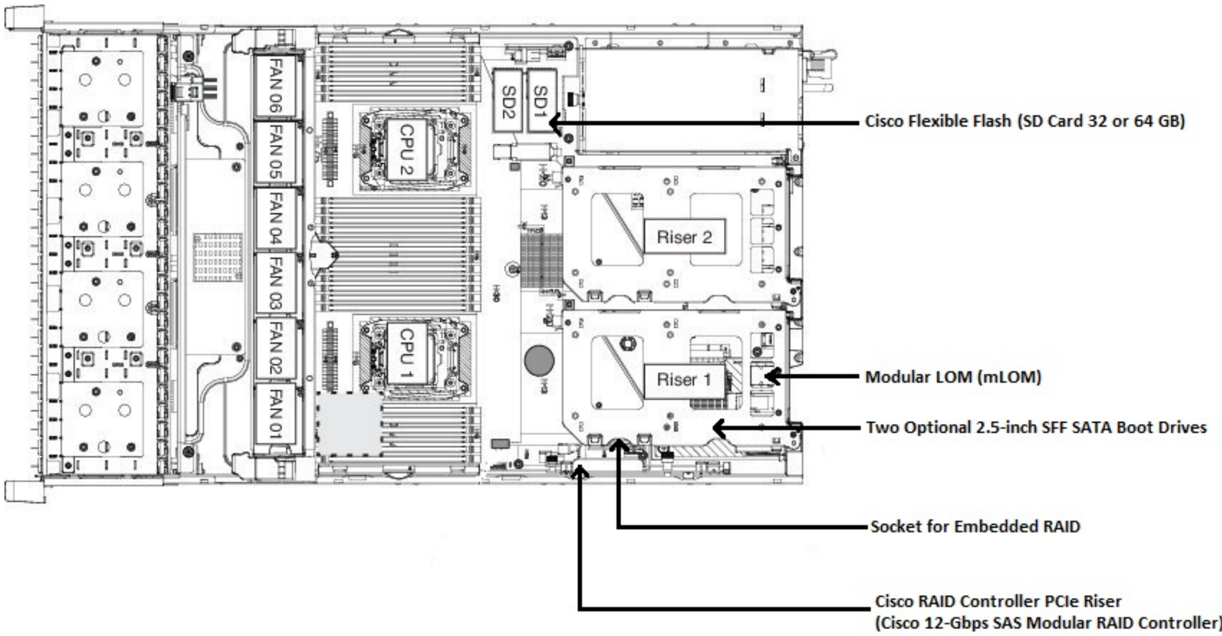

Hard-Disk Drives Versus Solid-State Drives

The choice of HDDs or SSDs is critical for enterprise customers and involves considerations of performance, reliability, price, and capacity. Part of the challenge is the sheer amount of today's data. The huge growth in data is threatening traditional computing infrastructures based on HDDs. However, the problem isn't simply growth; it is also the speed at which applications operate.

The mechanical nature of HDDs in high-I/O environments is the problem. Deployment of very fast SSDs is the increasingly popular solution to this problem.

For performance, without question, SSDs are faster than HDDs.

HDDs have an unavoidable overhead because they physically scan the disk for read and write operations. In an HDD array, I/O read and write requests are directed to physical disk locations. In response, the platter spins, and the disk-drive heads seek the location to write or read the I/O request. Latency from noncontiguous write locations multiplies the seek-time problem.

SSDs have no physical tracks or sectors and no mechanical movement. Thus, SSDs can reach memory addresses much more quickly than HDD heads can physically move. Because SSDs have no moving parts, there is no mechanical seek time or latency to overcome.

Even the fastest 15,000-rpm HDDs may not keep pace with SSDs in a high-demand I/O environment. Parallel disks, caching, and additional memory certainly help, but the inherent physical disadvantages have limited the capability of HDDs to keep pace with today's seemingly limitless data growth.

Choosing Between HDDs and SSDs

Customers should consider both performance and price when choosing between SSDs and HDDs. SSDs offer significant benefits for some workloads. Customer applications with the most random data requirements will gain the greatest benefit from SSDs compared to HDDs.

Even for sequential workloads, SSDs can offer increased I/O performance compared to HDDs. However, the performance improvement may not justify their additional cost for sequential operations. Therefore, Cisco recommends HDDs for predominantly sequential I/O applications.

For typical random workloads, SSDs offer tremendous performance improvements with less need for concern about reliability and write endurance and wear-out. Performance improves further as applications become more parallel and use the full capabilities of SSDs with tiering software or caching applications. The performance improvements gained from SSDs can provide a strong justification for their additional cost for random operations. Therefore, Cisco recommends SSDs for random I/O applications.

Solid-State Disk Drives

SSDs have no physical tracks or sectors and no mechanical movement. Thus, SSDs can reach logical block addresses (LBA) much more quickly than HDD heads can physically move. Because SSDs have no moving parts, there is no mechanical seek time or latency to overcome.

Even the fastest 15,000-rpm HDDs cannot keep pace with SSDs in a high-demand I/O environment. Parallel disks, caching, and additional memory certainly help, but the inherent physical disadvantages have limited the ability of HDDs to keep pace with today's seemingly limitless data growth.

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

- Low cost with higher-capacity SSD drive (480-GB SATA with 6-Gbps)
- High speed with higher-capacity SSD drive (3.8-TB SATA with 6-Gbps)
- High speed and high endurance with high-capacity SSD (1.6-TB SAS with 12-Gbps)

To meet the requirements of different application environments, Cisco offers both enterprise performance SSD drives and enterprise value SSD drives.

SSD Types Used in I/O Characterization

Comparing Enterprise Performance SSDs and Enterprise Value SSDs

As SSD technology evolves, manufacturers are improving their reliability processes. With maturity, reliability will become a smaller differentiating factor in the choice between SSDs and HDDs. With the availability of latest SSDs at higher capacities, combined with better performance and lower cost, they are increasingly becoming an obvious storage choice for enterprises.

To meet the requirements of different application environments, Cisco offers both enterprise performance SSDs and enterprise value SSDs. They all deliver superior performance compared to HDDs; however, enterprise performance SSDs support higher read-write workloads and have a longer expected service life. Enterprise value SSDs provide relatively large storage capacities at lower cost, but they do not have the endurance of the enterprise performance SSDs.

Enterprise performance SSDs provide high endurance and support up to 10 full-drive write operations per day. These SSDs are targeted at write-centric I/O applications such as caching, online transaction processing (OLTP), data warehousing, and virtual desktop infrastructure (VDI).

Enterprise value SSDs provide low endurance and support up to 1 full-drive write operation per day. These SSDs are targeted at read-centric I/O applications such as OS boot, streaming media, and collaboration.

Reliability

Cisco uses several different technologies and design requirements to help ensure that our SSDs can meet the reliability and endurance demands of server storage.

Reliability depends on many factors, including use, physical environment, application I/O demands, vendor, and mean time before failure (MTBF).

In challenging environments, the physical reliability of SSDs is clearly better than that of HDDs given SSDs' lack of mechanical parts. SSDs can survive cold and heat, drops, and extreme gravitational forces (G-forces). However, these extreme conditions are not a factor in typical data centers. Although SSDs have no moving heads or spindles, they have their own unique stress points and failures in components such as transistors and capacitors. As an SSD ages, its performance slows. The SSD controller must read, modify, erase, and write increasing amounts of data. Eventually, memory cells wear out.

Some common SSD points of failure include:

- **Bit errors:** Random data bits may be stored in cells.
- **Flying or shorn write operations:** Correct write content may be written in the wrong location, or write operations may be truncated due to power loss.
- **Unserializability:** Write operations may be recorded in the wrong order.
- **Firmware:** Firmware may fail, become corrupted, or upgrade improperly.
- **Electronic failures:** Even though SSDs have no moving parts, physical components such as chips and transistors may fail.
- **Power outages:** SSDs are subject to damaged file integrity if they are reading or writing during power interruptions. Enterprise performance SSD drives support greater read-write workloads and have a longer expected service life.

Performance

When considering price, it is important to differentiate between enterprise performance SSDs and enterprise value SSDs. It is important to recognize the significant differences between the two in performance, cost, reliability, and targeted applications. 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.

When price is measured on a per-gigabyte basis (US$/GB), SSDs are significantly more expensive than HDDs. Even when price is measured in terms of bandwidth per Gbps (US$/Gbps), enterprise performance SSDs remain more expensive.

In addition to the price of individual drives, you should consider TCO. The higher performance of SSDs may allow I/O demands to be met with a lower number of SSDs than HDDs, providing a TCO advantage.

Capacity

HDDs provide the highest capacity and storage density, up to 2TB in a 2.5” SFF form factor or 10TB in a 3.5” LFF form factor. Storage requirements may outweigh performance depending how much data must be retained in “online” accessibility or “hot” storage.

Virtual Disk Options

The following controller options can be configured 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 3 summarizes the supported RAID levels and their characteristics.

Table 3. 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 and 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 with performance improvement	Yes	Yes

RAID-0, RAID-5, and RAID-10 were used for these performance characterization tests.

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. In this mode, the controller is fully utilized to perform the initialization, and it blocks any I/O operations. 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. Stripe size can be configured as 64, 128, 256, or 512 KB or 1 MB. The default stripe size is 64 KB. A stripe size of 64 KB was used for these performance characterization tests. 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.

Comparing Strip (Block) Size and 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. Stripe size is not configured. It is a product of the strip size, the number of physical drives in the virtual disk, and the RAID level.

Access Policy

- **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

- **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 in the event of a power failure.
- **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 in the event of a power failure. Although disk caching can accelerate I/O performance, it is not recommended for enterprise deployments.

I/O Cache Policy

- **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.
- **Cached:** All read and write data passes through the controller cache memory on its way to or from host memory. Subsequent read requests for the same data can then be addressed from the controller cache. Note that “cached I/O” refers to the caching of read data, and “read ahead” refers to the caching of speculative future read data.

Read Policy

- **No Read Ahead (Normal Read):** Only the requested data is read, and the controller does not read ahead any data.
- **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

- **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 in the event of a power failure. 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 4 lists the workload types tested.

Table 4. 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
Virtual desktop infrastructure (VDI): Boot process	5	Random	80:20
Virtual desktop infrastructure (VDI): Steady state	5	Random	20:80

Tables 5 and 6 list the I/O mix ratios chosen for the sequential access and random access patterns, respectively.

Table 5. I/O Mix Ratio for Sequential Access Pattern

I/O Mode	I/O Mix Ratio (Read:Write)		
Sequential	100:0	0:100	
	RAID-0	RAID-0	RAID-10

Table 6. 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-5	RAID-5

Tables 7 and 8 list the recommended virtual drive configuration parameters for deployment of SSDs and HDDs.

Table 7. Recommended Virtual Drive Configuration for SSDs

Workload	RAID Level	Strip Size	Disk Cache Policy	I/O Cache Policy	Read Policy	Write Policy
Random I/O	RAID-0	64 KB	Disabled	Direct	No Read Ahead	Write Through
Random I/O	RAID-5	64 KB	Disabled	Direct	No Read Ahead	Write Through
Sequential I/O	RAID-0	64 KB	Disabled	Direct	No Read Ahead	Write Through
Sequential I/O	RAID-10	64 KB	Disabled	Direct	No Read Ahead	Write Through

Table 8. Recommended Virtual Drive Configuration for HDDs

Workload	RAID Level	Strip Size	Disk Cache Policy	I/O Cache Policy	Read Policy	Write Policy
Random I/O	RAID-0	64 KB	Disabled	Direct	No Read Ahead	Write Through
Random I/O	RAID-5	64 KB	Disabled	Direct	No Read Ahead	Write Back Good BBU

Sequential I/O	RAID-0	64 KB	Disabled	Direct	Always Read Ahead	Write Through
Sequential I/O	RAID-10	64 KB	Disabled	Direct	Always Read Ahead	Write Through

Test Configuration

The test configuration was as follows:

- One virtual drive was created with eight disks for the SFF SSD and SFF HDD.
- One virtual drive was created with 12 disks for the LFF HDD.
- RAID-0, RAID-5, and RAID-10 configurations were tested with the SFF SSD, SFF HDD, and LFF HDD.
- JBOD configuration was tested with the Cisco UCS 12-Gbps SAS modular HBA pass-through controller with 8 disks for the SFF SSD and SFF HDD and 12 disks for the LFF HDD.

Table 9 lists the recommended Iometer settings.

Table 9. Recommended Iometer Settings

Name	Value
Iometer version	1.1.0
Run time	30 minutes per access specifications
Ramp-up time	20 seconds for random I/O; 60 seconds for sequential I/O
Record results	All
Number of workers	8 workers for random I/O (equal to the number of SSDs); 1 worker for sequential I/O
Write I/O data pattern	Repeating bytes
Transfer delay	1 I/O operation
Align I/O on	Request size boundaries
Reply size	No reply

Note: The SSDs were tested with various numbers of outstanding I/O operations to get the best performance within an acceptable response time.

SSD Performance Results

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

Figures 6 through 13 were prepared from Iometer measurement data. They illustrate the I/O performance of 480-GB enterprise performance and enterprise value, 3.8-TB enterprise value, and 1.6-TB enterprise performance and enterprise value SSDs used in these comparison tests. The server specifications and BIOS settings used in these performance characterization tests are detailed in the appendix, “Test Environment.”

The I/O performance test results capture the maximum read IOPS and bandwidth achieved with the SSDs within the acceptable response time (latency) of 2 milliseconds (ms). However, the SSDs under test are capable of a much higher IOPS rate and much greater bandwidth with increased latency.

SSD RAID-0 Performance

Figure 6 shows the performance of the SSDs under test with a RAID-0 configuration with a 100 percent random read access pattern. The graph shows the comparitive performance values achieved for enterprise performance (EP) drives and enterprise value (EV) drives to help customers understand the performance trade-off when choosing an SSD type. The graph shows that the 1.6-TB enterprise performance drive with 10x endurance provides about 100 percent greater performance at the 4-KB block size and 86 percent greater performance at the 8-KB block size compared to the 1.6-TB enterprise value drive. Latency is the time taken to complete a single I/O request from the application's viewpoint.

Figure 6. Random Read 100%

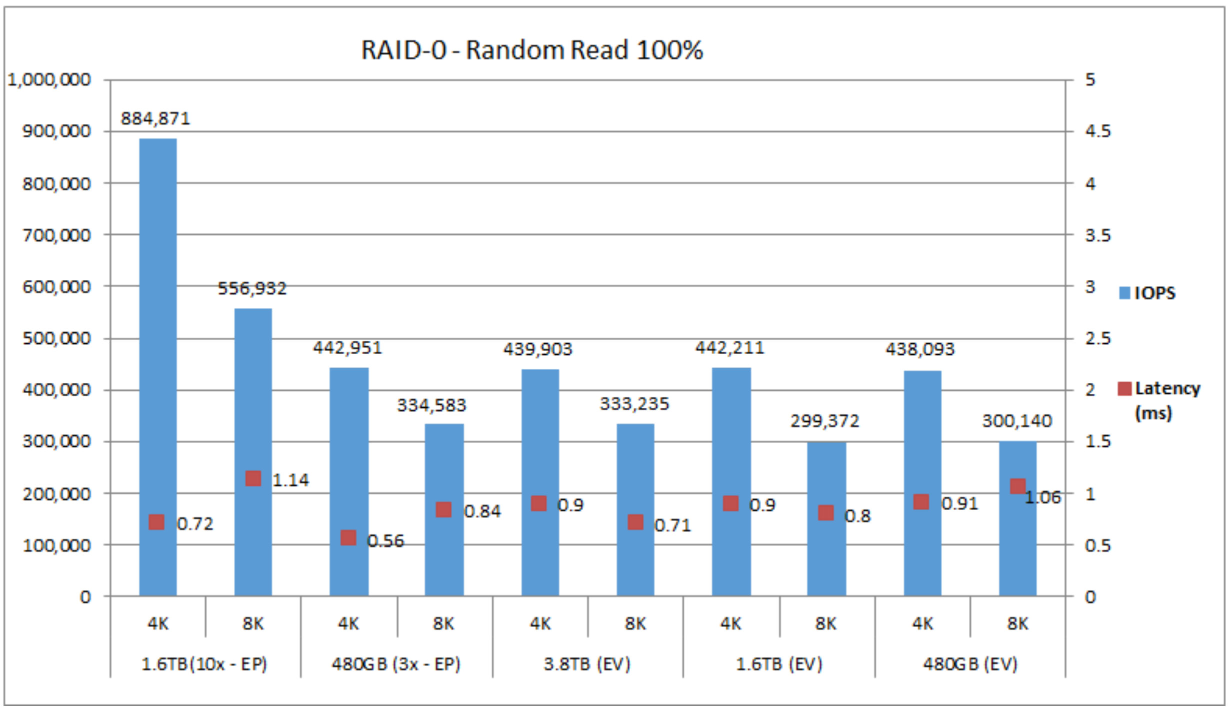

Figure 7 shows the performance of SSDs under test for a RAID-0 configuration with a 100 percent random write access pattern. The numbers in the graph affirm that for all SSDs the IOPS rate for the 4-KB block size is well over 100,000, and for the 8-KB block size it is over 50,000.

Figure 7. Random Write 100%

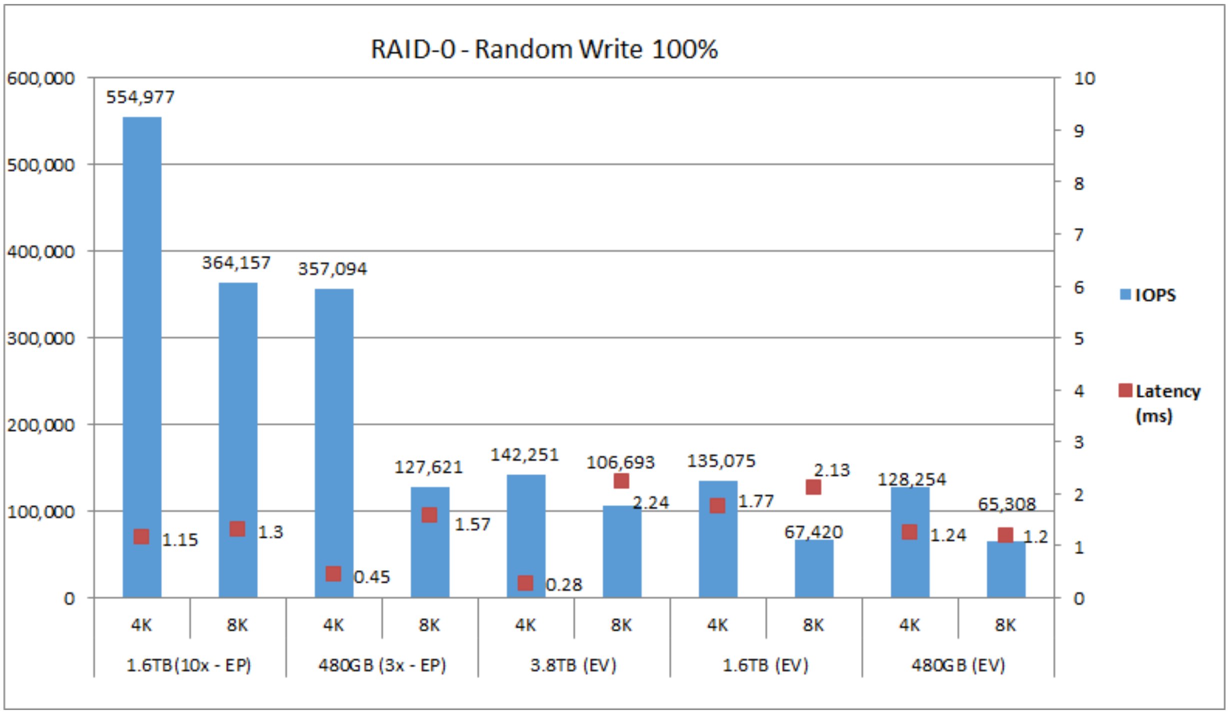

Figure 8 shows the performance of SSDs under test for a RAID-0 configuration with a random 70 percent read and 30 percent write access pattern. The numbers in the graph affirm that for all SSDs the IOPS rate for the 4-KB block size is well over 200,000, and for the 8-KB block size it is over 100,000.

Figure 8. Random Read:Write 70%:30%

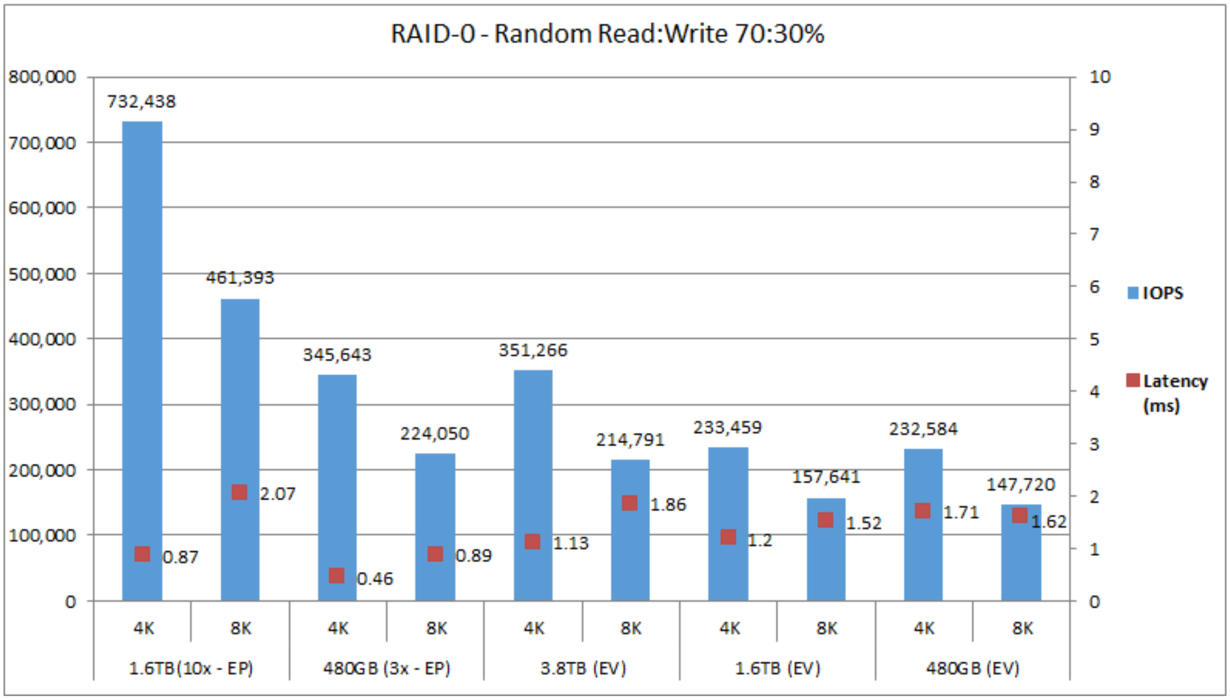

Figure 9 shows the performance of SSDs under test for a RAID-0 configuration with a 100 percent sequential read access pattern. The graph shows that the 1.6-TB enterprise performance drive with 10x endurance provides 100 percent greater performance at 256-KB and 1-MB block sizes compared to the 1.6-TB enterprise value drive. The SSDs under test are capable of much greater bandwidth with greater latency.

Figure 9. Sequential Read 100%

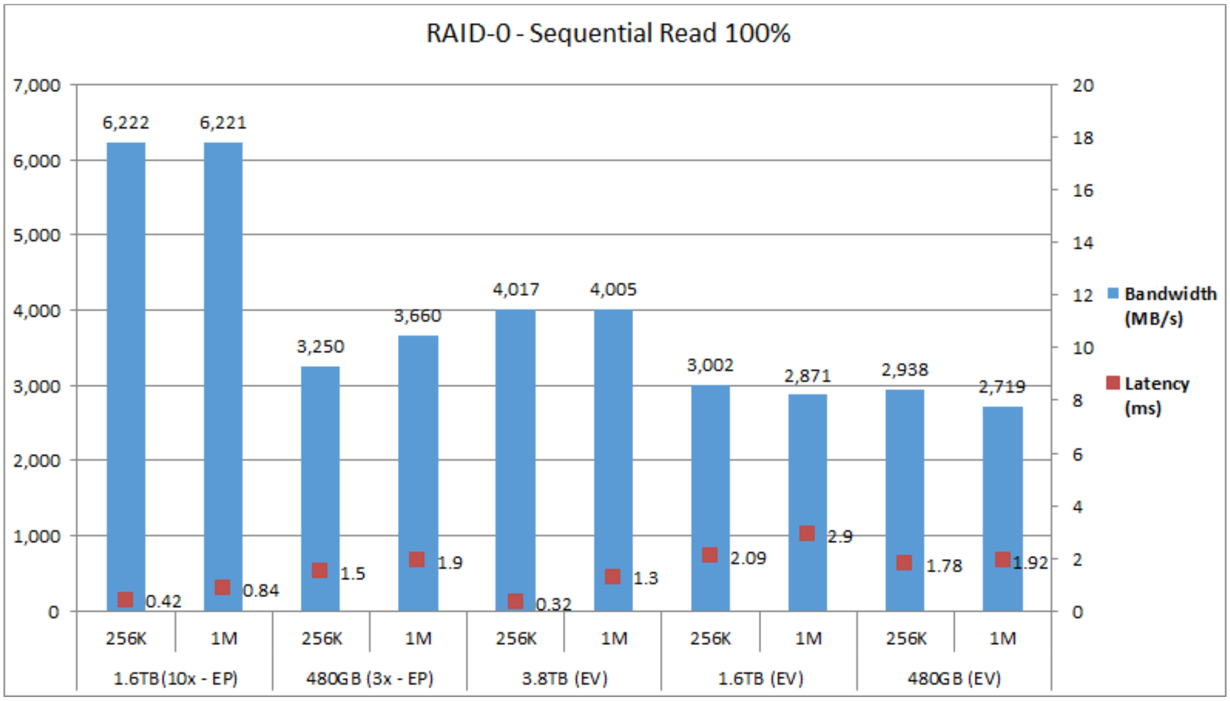

Figure 10 shows the performance of SSDs under test for a RAID-0 configuration with a 100 percent sequential write access pattern. The numbers in the graph affirm that the bandwidth is well over 1500 MBps for both the 256-KB and 1-MB block sizes for all the SSDs. Even though the bandwidth is optimal for the 1.6-TB enterprise value SSD, the latency is greater, which is expected for a sequential write access pattern.

Figure 10. Sequential Write 100%

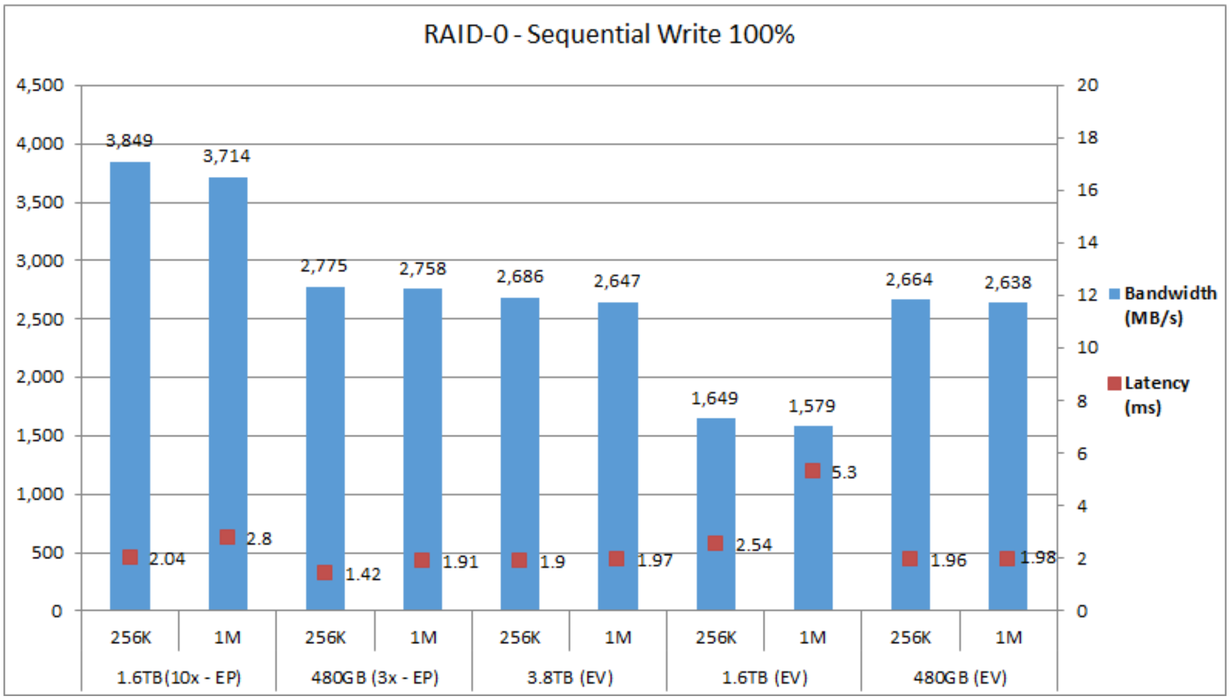

SSD RAID-5 Performance

Figure 11 shows the performance of SSDs under test for a RAID-5 configuration with a random 70 percent read and 30 percent write access pattern. The graph shows that the 1.6-TB enterprise performance SSD with 10x endurance provides 8 percent greater performance at the 4-KB block size and 100 percent greater performance at the 8-KB block size compared to the 1.6-TB enterprise value SSD. A write-back cache is not used for SSDs because they perform better with write policy set to Write Through (this cache typically is used in HDDs to enhance write performance).

Figure 11. Random Read:Write 70%:30%

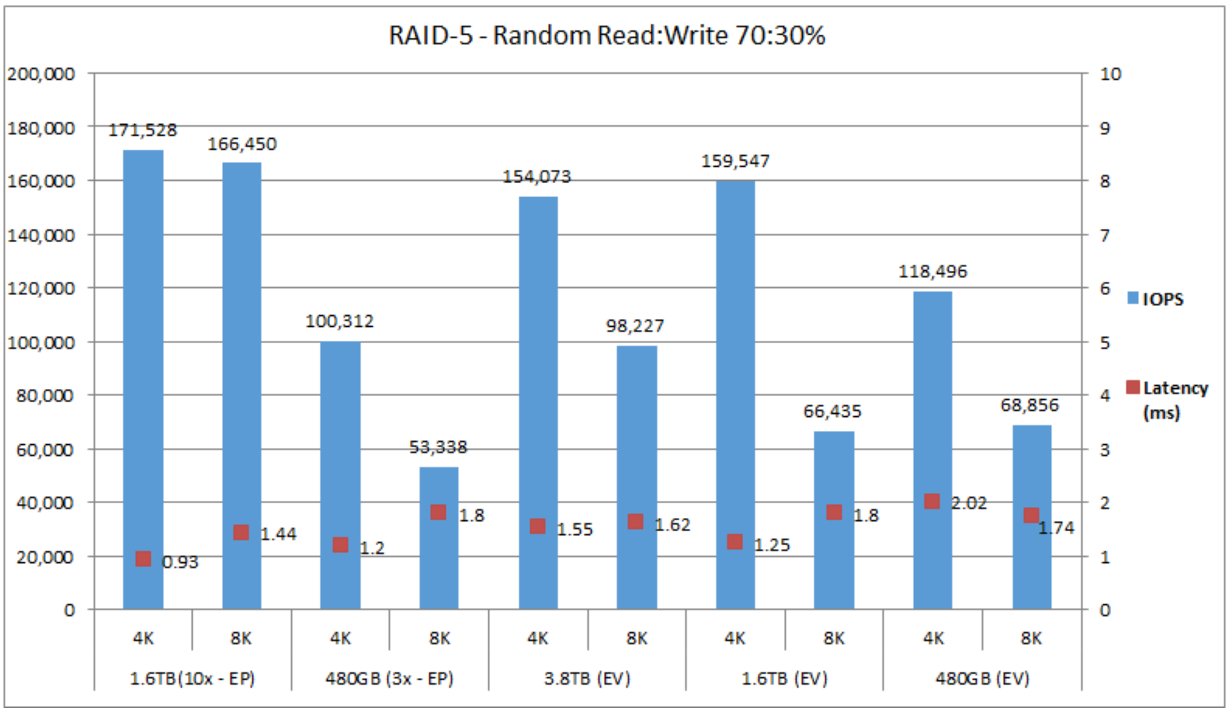

Figure 12 shows the performance of SSDs under test for a RAID-5 configuration with a random 50 percent read and 50 percent write access pattern. The graph shows that the 1.6-TB enterprise performance SSD with 10x endurance provides 19 percent greater performance at the 4-KB block size and 15 percent percent greater performance at the 8-KB block size compared to the 1.6-TB enterprise value SSD.

Figure 12. Random Read:Write 50%:50%

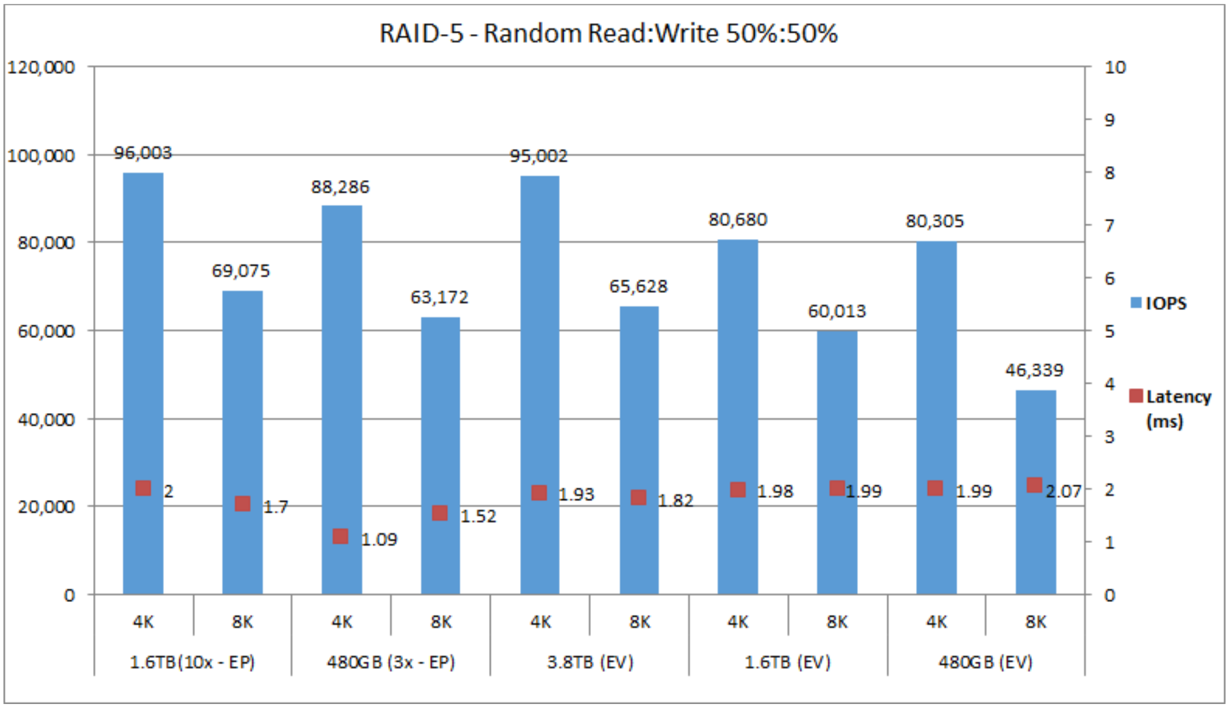

SSD RAID-10 Performance

Figure 13 shows the performance of SSDs under test for a RAID-10 configuration with a 100 percent sequential write access pattern. The numbers in the graph affirm that the bandwidth achieved is over 1000 MBps for both the 4-KB and 8-KB block sizes for all the SSDs. The SSDs under test are capable of a much greater IOPS rate with greater latency.

Figure 13. Sequential Write 100%

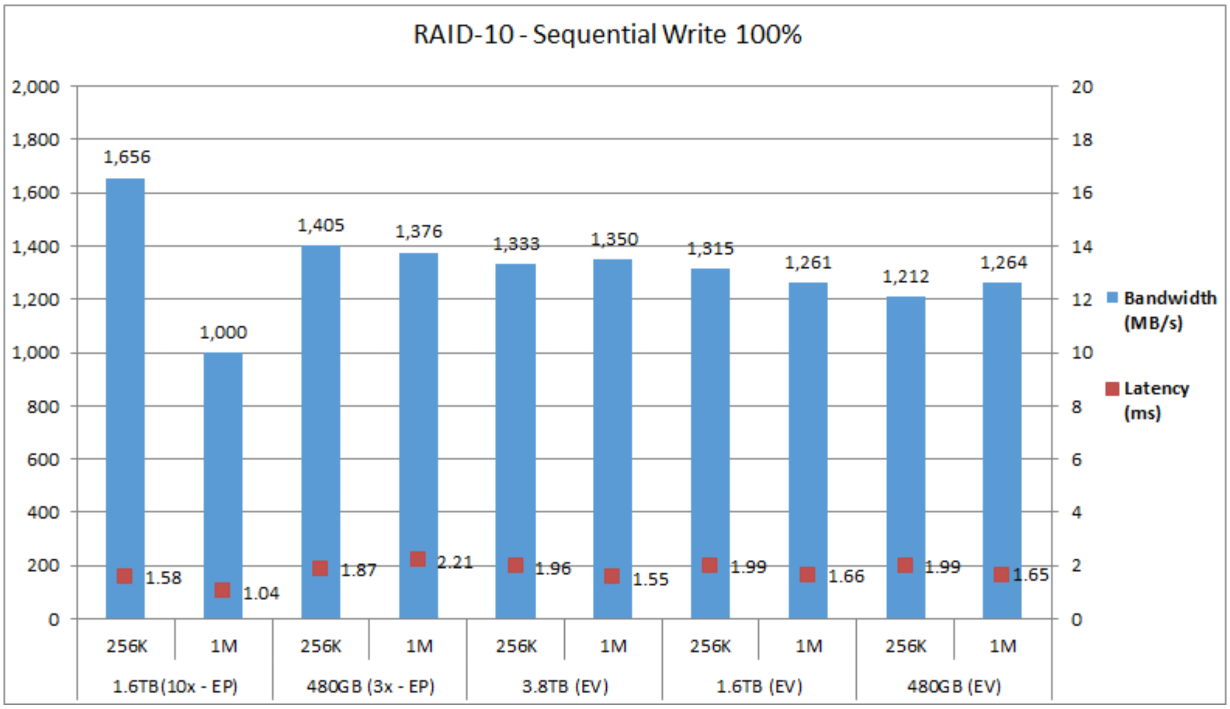

SSD JBOD Performance

The I/O performance characterization results include JBOD pass-through controller configuration results because some software-defined storage solutions use JBOD storage. JBOD storage provides better control over individual disks. It is robust and inexpensive and allows you to use all the disk space (no RAID-only pass-through). Availability is assured through replication. JBOD setups are becoming popular in scale-out environments because they can provide pooled storage effectively.

The tests used the Cisco 12-Gbps modular SAS HBA (UCSC-SAS12GHBA) for JBOD. The JBOD performance tests used eight SSDs. Figures 14 through 18 show the results.

Figure 14 shows the performance of SSDs under test for a JBOD configuration with a 100 percent random read access pattern. The graph shows the comparitive performance values achieved for enterprise performance drives and enterprise value drives to help customers understand the performance trade-off when choosing an SSD type. The graph shows that the 1.6-TB enterprise performance drive with 10x endurance provides 100 percent greater performance at the 4-KB and 8-KB block sizes compared to the 1.6-TB enterprise value drive.

Figure 14. Random Read 100%

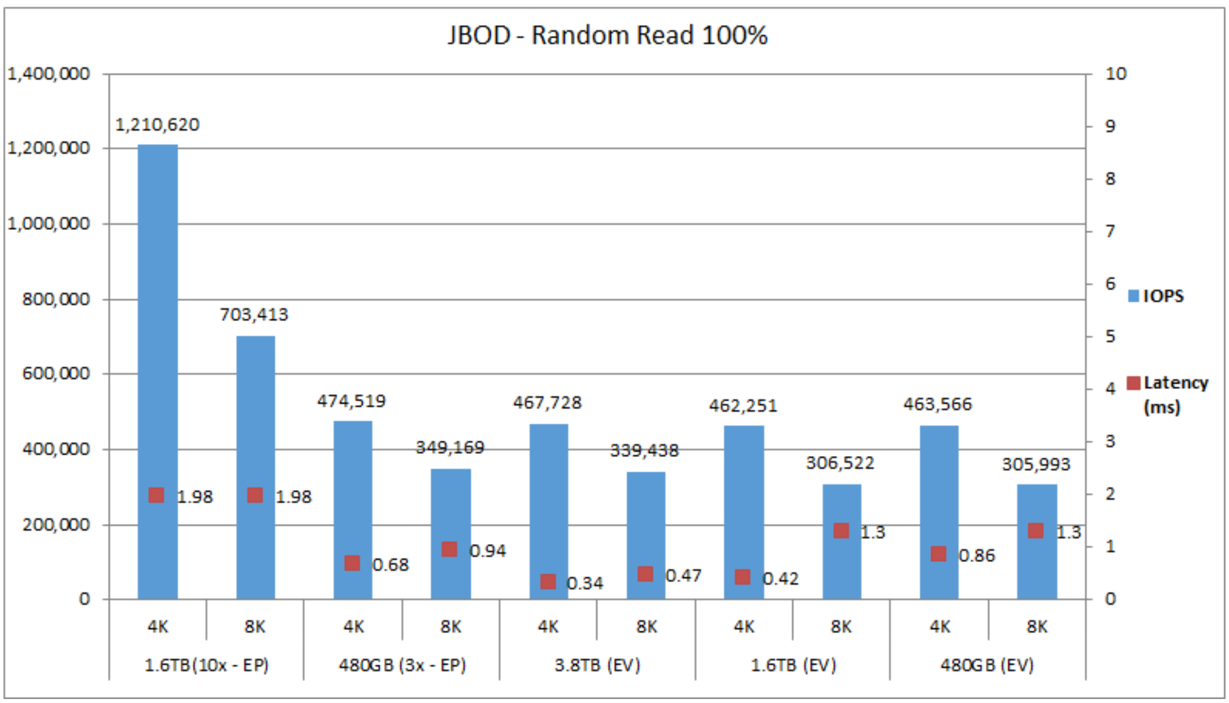

Figure 15 shows the performance of SSDs under test for a JBOD configuration with a 100 percent random write access pattern. The numbers in the graph affirm that for all the SSDs the IOPS rate for the 4-KB block size is well over 100,000, and for the 8-KB block size it is well over 50,000.

Figure 15. Random Write 100%

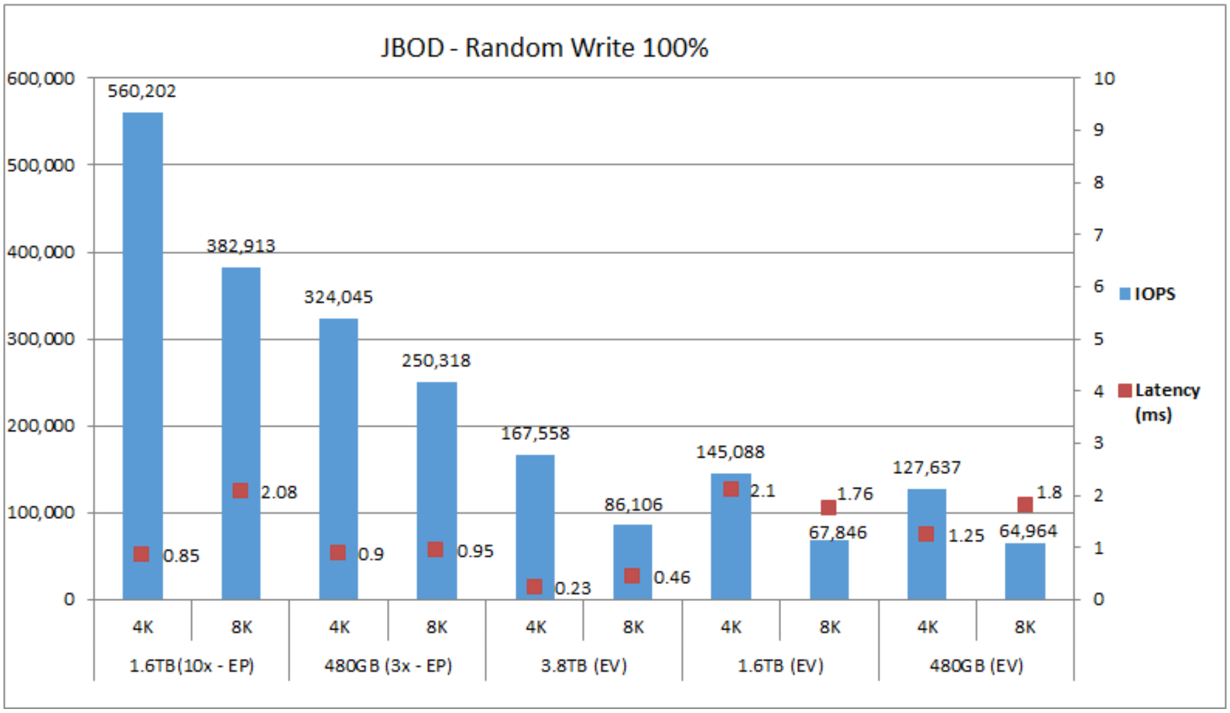

Figure 16 shows the performance of SSDs under test for a JBOD configuration with a random 70 percent read and 30 percent write access pattern. The numbers in the graph affirm that for all the SSDs the IOPS rate for the 4-KB block size is well over 200,000, and for the 8-KB block size it is over 100,000.

Figure 16. Random Read:Write 70%:30%

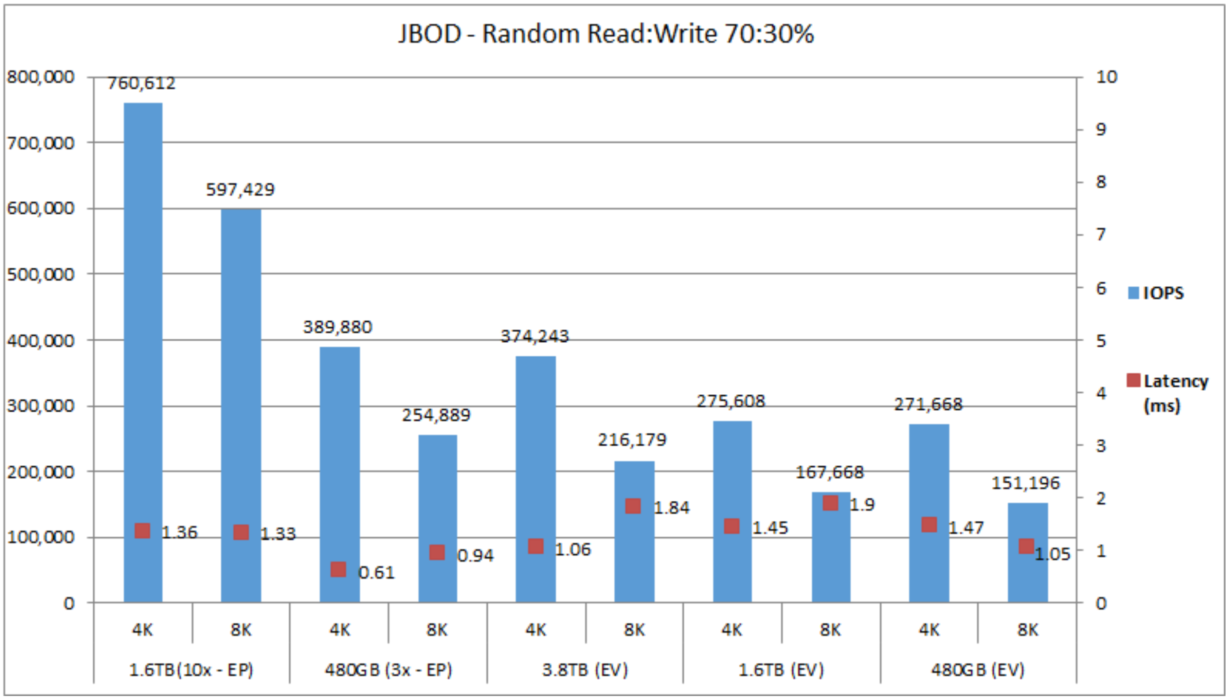

Figure 17 shows the performance of SSDs under test for a JBOD configuration with a 100 percent sequential read access pattern. The graph shows that the 1.6-TB enterprise performance drive with 10x endurance provides 54 percent greater performance at the 256-KB and 1-MB block sizes than the 1.6-TB enterprise value drive.

Figure 17. Sequential Read 100%

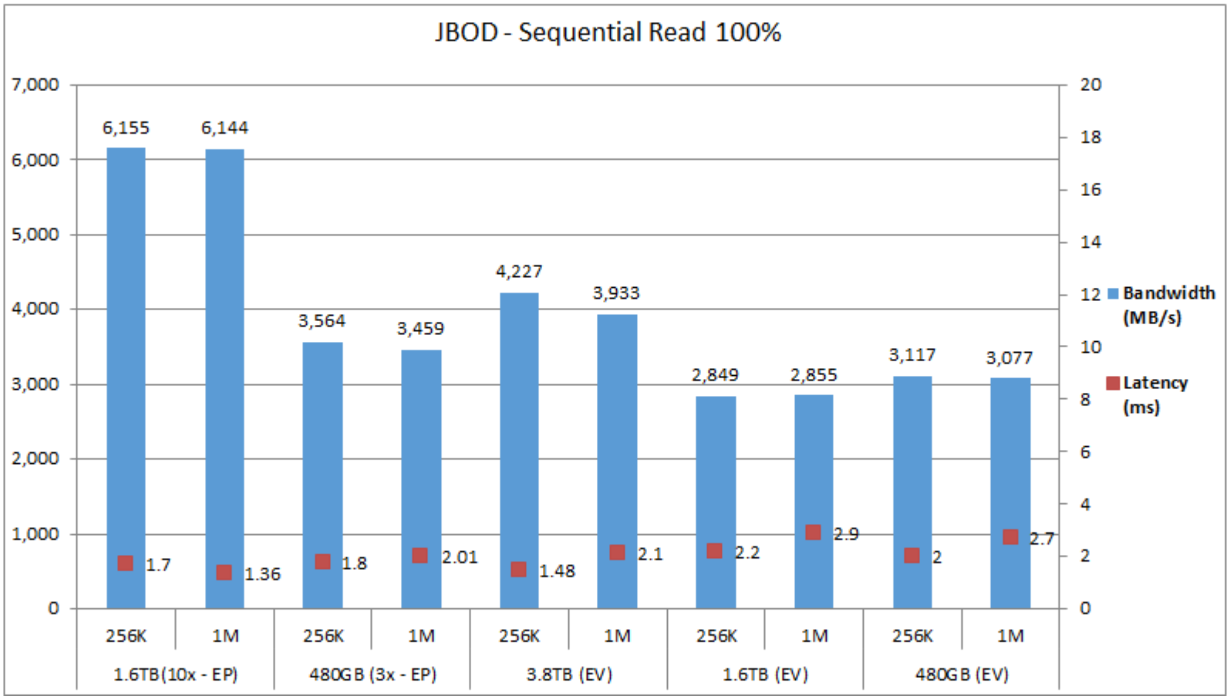

Figure 18 shows the performance of SSDs under test for a JBOD configuration with a 100 percent sequential write access pattern. The numbers in the graph affirm that the bandwidth is well over 1900 MBps for both the 256-KB and 1-MB block sizes for all the SSDs. Even though the bandwidth is optimal for the 1.6-TB enterprise value SSD, the latency is greater, which is expected for a sequential write access pattern.

Figure 18. Sequential Write 100%

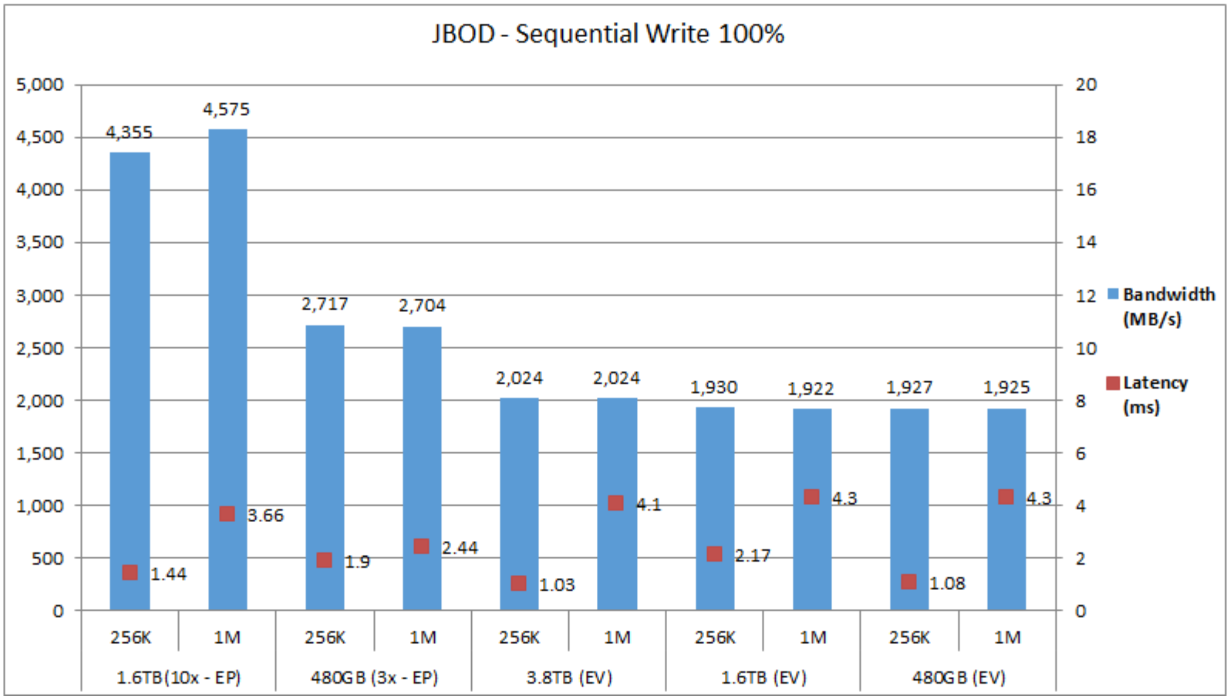

HDD Performance Results

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

Figures 19 through 48 were prepared from Iometer measurement data. They illustrate the I/O performance of 300- and 600-GB and 1.2-TB SFF HDDs, and of 4-, 6-, and 10-TB LFF HDDs. The server specifications and BIOS settings used in these performance characterization tests are detailed in the appendix, "Test Environment."

The I/O performance test results capture the maximum IOPS and bandwidth achieved with the HDDs within the acceptable response time (latency) of 20 ms and maximum achievable IOPS rate and bandwidth with higher latency.

HDD RAID-0 Performance

LFF 12-HDD RAID-0 Performance

Figure 19 shows the performance of LFF HDDs under test for a RAID-0 configuration with a 100 percent random read access pattern. The graph shows the comparitive performance values achieved for LFF HDDs to help customers understand the performance trade-off when choosing an LFF HDD type. The graph shows that performance for the 4-, 6-, and 10-TB LFF HDDs is almost same for latency within 20 ms, and that drives perform better with increased latency for 4- and 8-KB block sizes.

Figure 19. Random Read 100%

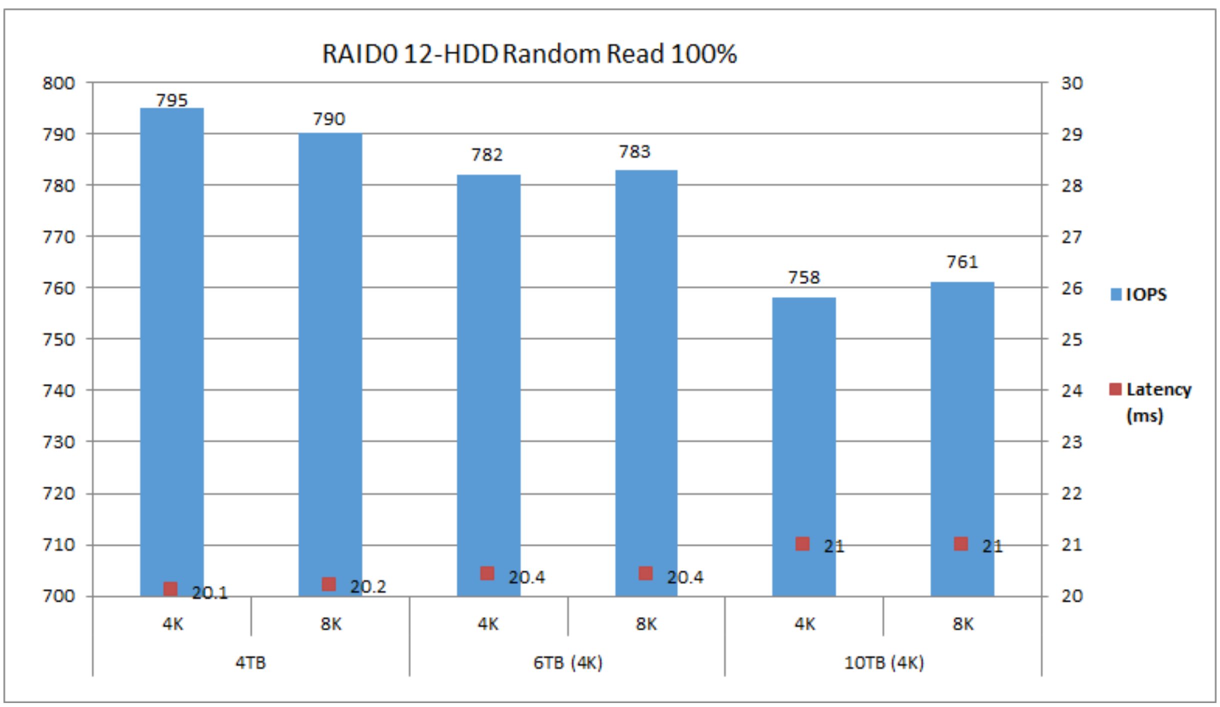

Figure 20 shows the performance of LFF HDDs under test for a RAID-0 configuration with a 100 percent random write access pattern. The numbers in the graph affirm that the IOPS rate for the 4- and 8-KB block sizes for 6- and 10-TB (4-KB format) LFF HDDs is better for random write operations than it is for 4-TB LFF HDDs (512-KB format).

Figure 20. Random Write 100%

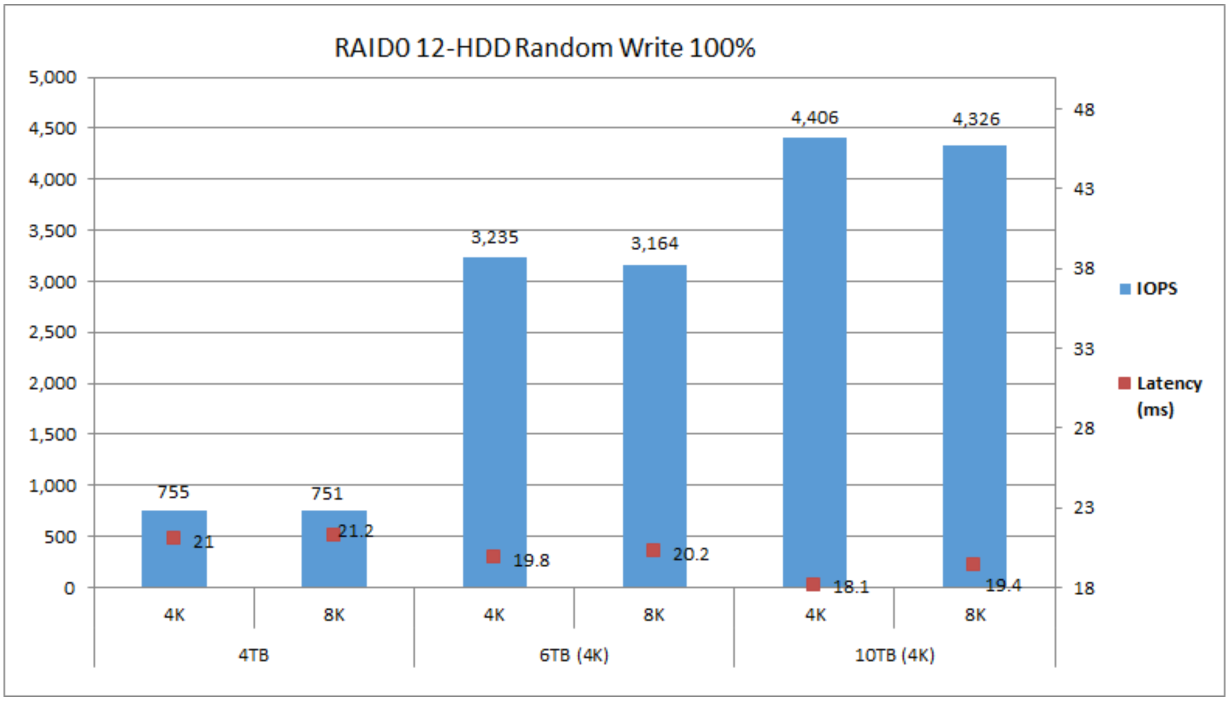

Figure 21 shows the performance of LFF HDDs under test for a RAID-0 configuration with a random 70 percent read and 30 percent write access pattern. The numbers in the graph affirm that performance for the 4-KB format LFF HDDs is better within a latency of 20 ms.

Figure 21. Random Read:Write 70:30%

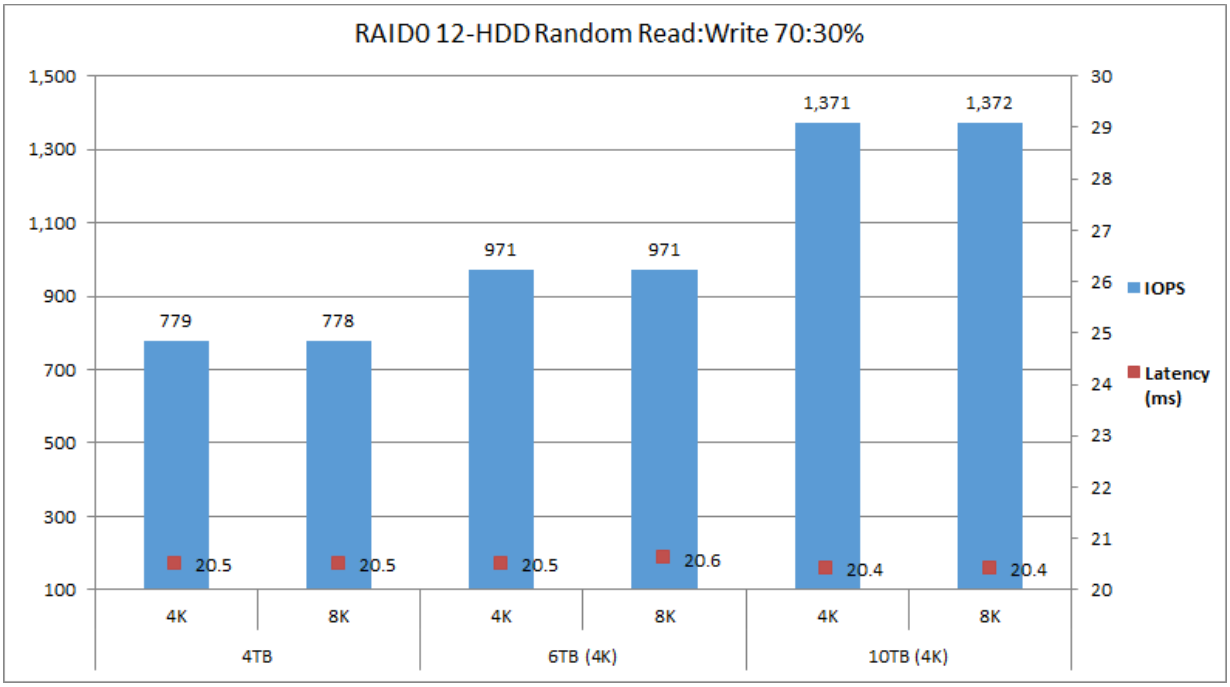

Figure 22 shows the performance of LFF HDDs under test for a RAID-0 configuration with a 100 percent sequential read access pattern. The graph shows that the 10-TB drive provides greater performance at the 256-KB and 1-MB block sizes than do the 4- and 6-TB LFF HDDs.

Figure 22. Sequential Read 100%

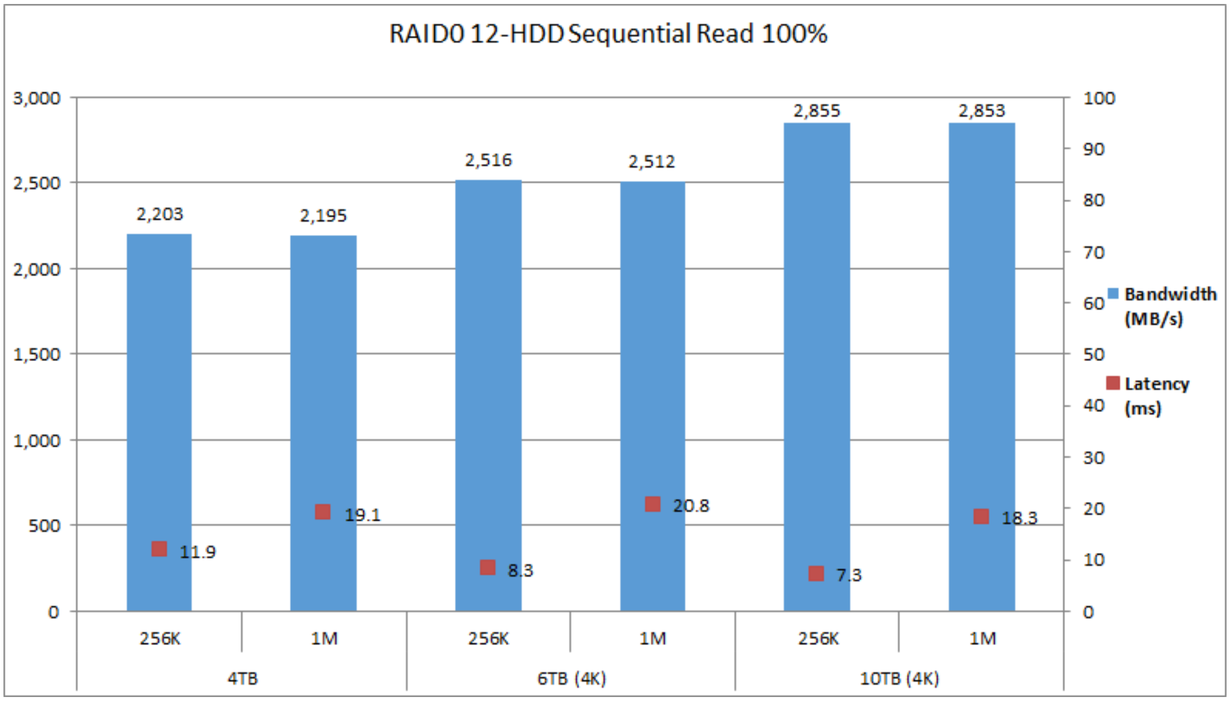

Figure 23 shows the performance of LFF HDDs under test for a RAID-0 configuration with a 100 percent sequential write access pattern. The graph shows that the 10-TB drive provides greater performance at the 256-KB and 1-MB block sizes than do the 4- and 6-TB LFF HDDs.

Figure 23. Sequential Write 100%

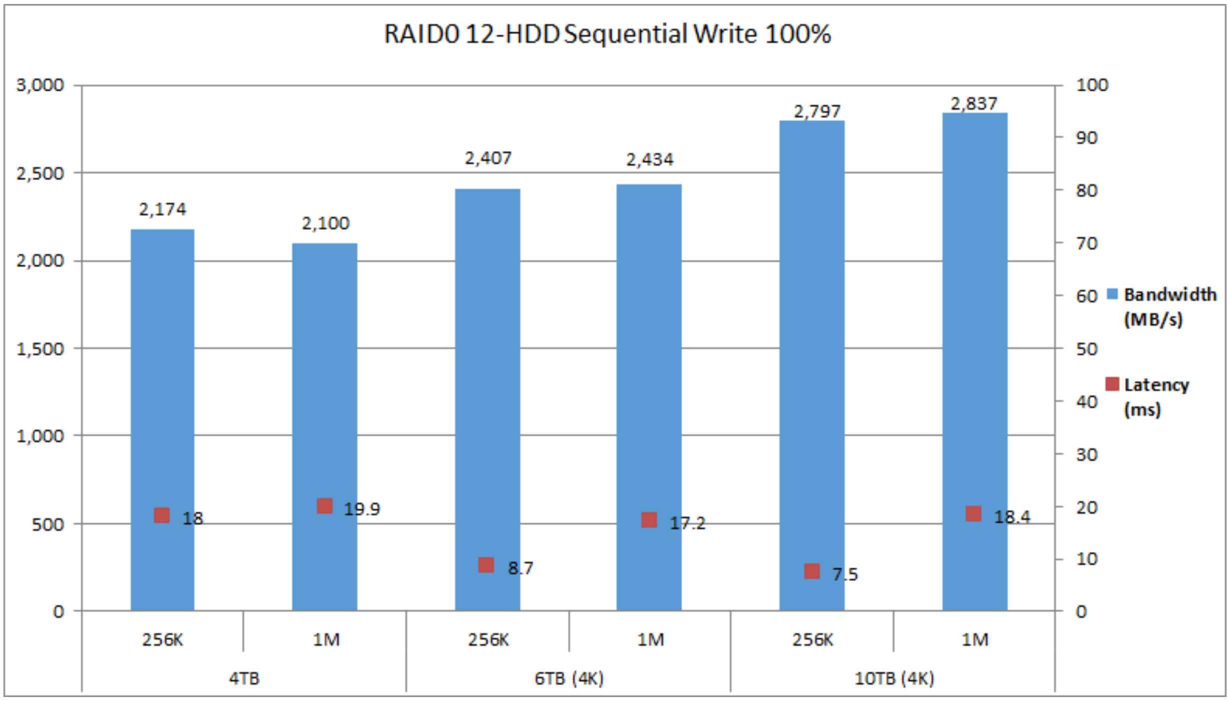

SFF 8-HDD RAID-0 Performance

Figure 24 shows the performance of the SFF HDDs under test with a RAID-0 configuration with a 100 percent random read access pattern. The graph shows the comparative performance values achieved for 10,000- and 15,000-rpm SFF HDDs to help customers understand the performance trade-off when choosing an HDD type. The graph shows that the 600-GB 15,000-rpm HDD can perform better than the 10,000-rpm drives.

Figure 24. Random Read 100%

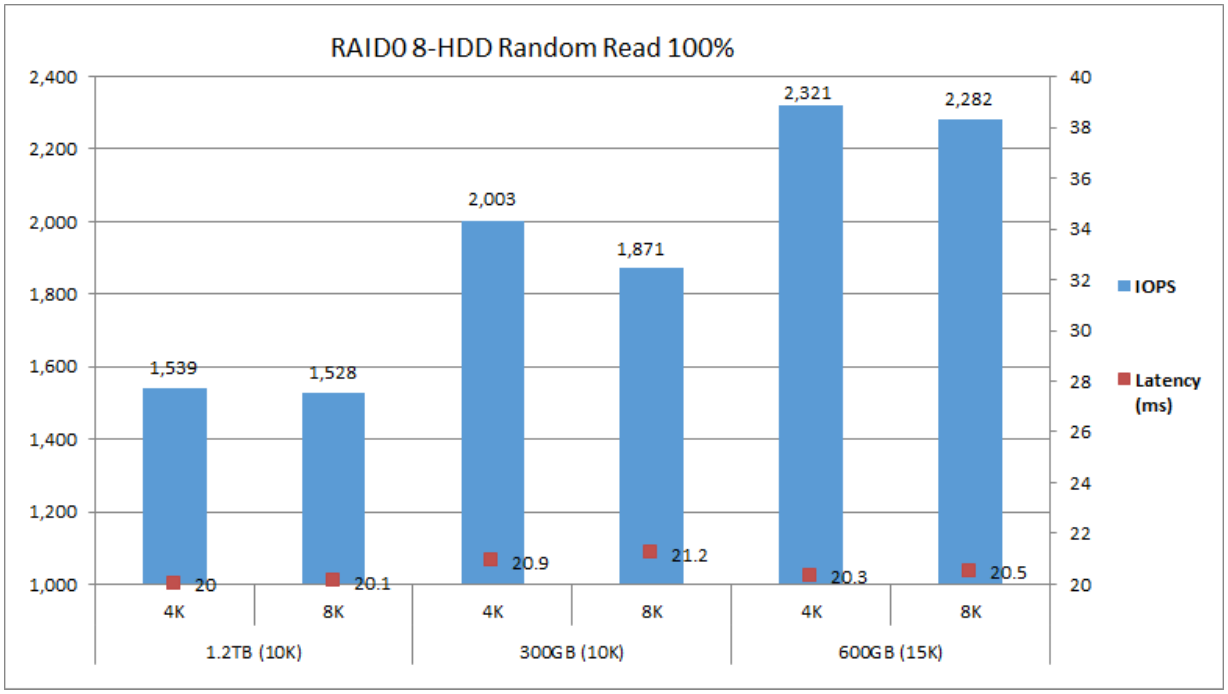

Figure 25 shows the performance of the SFF HDDs under test with a RAID-0 configuration with a 100 percent random write access pattern. The graph shows that the 600-GB 15,000-rpm HDD can perform better than the 10,000-rpm drives for the 4- and 8-KB block sizes.

Figure 25. Random Write 100%

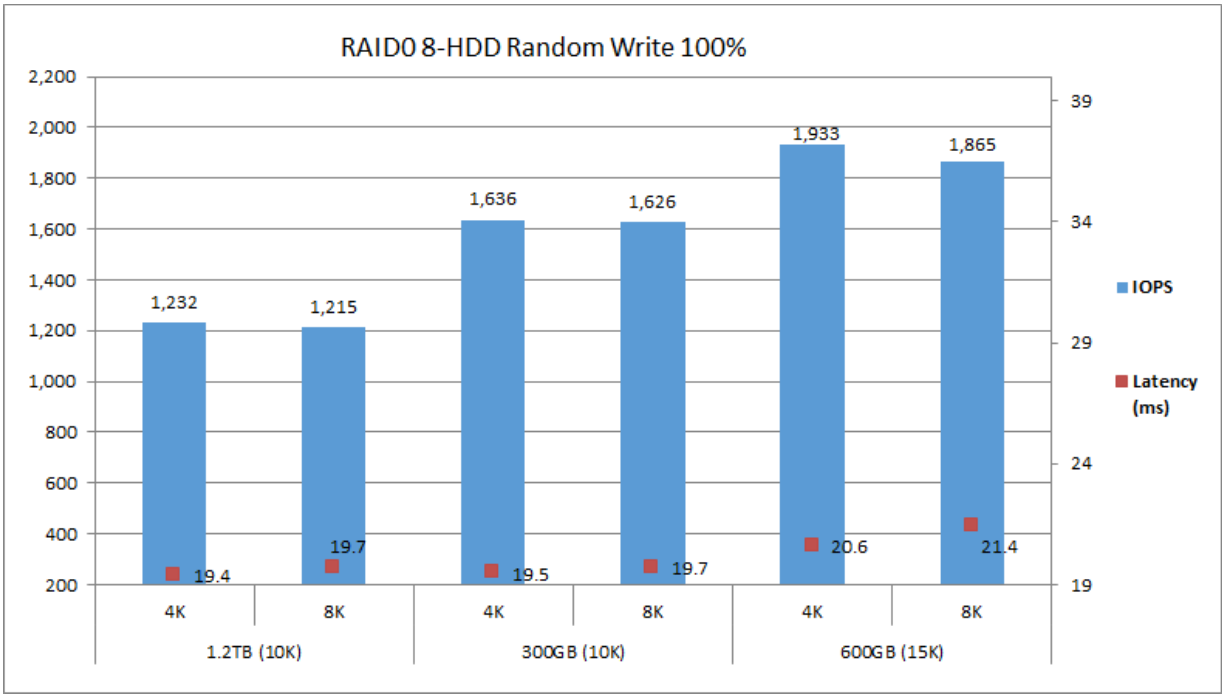

Figure 26 shows the performance of LFF HDDs under test for a RAID-0 configuration with a random 70 percent read and 30 percent write access pattern. The graph affirms that the 600-GB 15,000-rpm HDD can perform better than the 10,000-rpm drives for the 4- and 8-KB block sizes.

Figure 26. Random Read:Write 70:30%

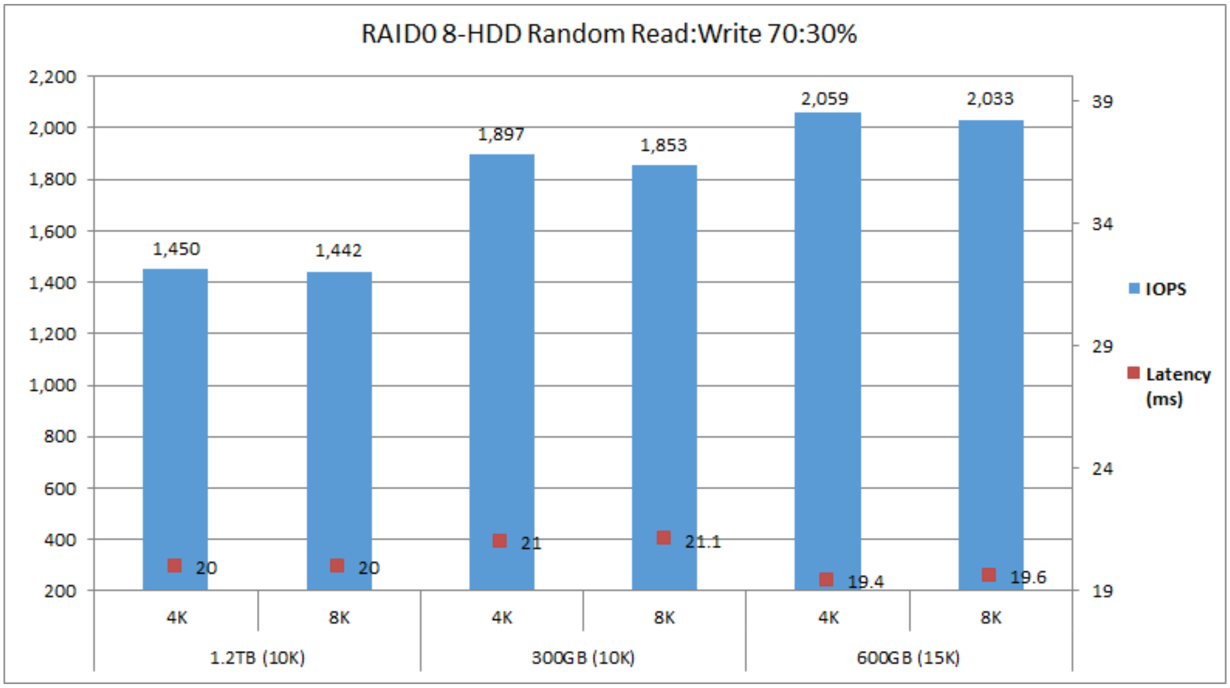

Figure 27 shows the performance of SFF HDDs under test for a RAID-0 configuration with a 100 percent sequential read access pattern. The graph shows that the 600-GB 15,000-rpm drive performs better for the 256-KB and 1-MB block sizes than do the 1.2-TB and 300-GB 10,000-rpm drives. The SFF HDDs under test can perform better with greater latency.

Figure 27. Sequential Read 100%

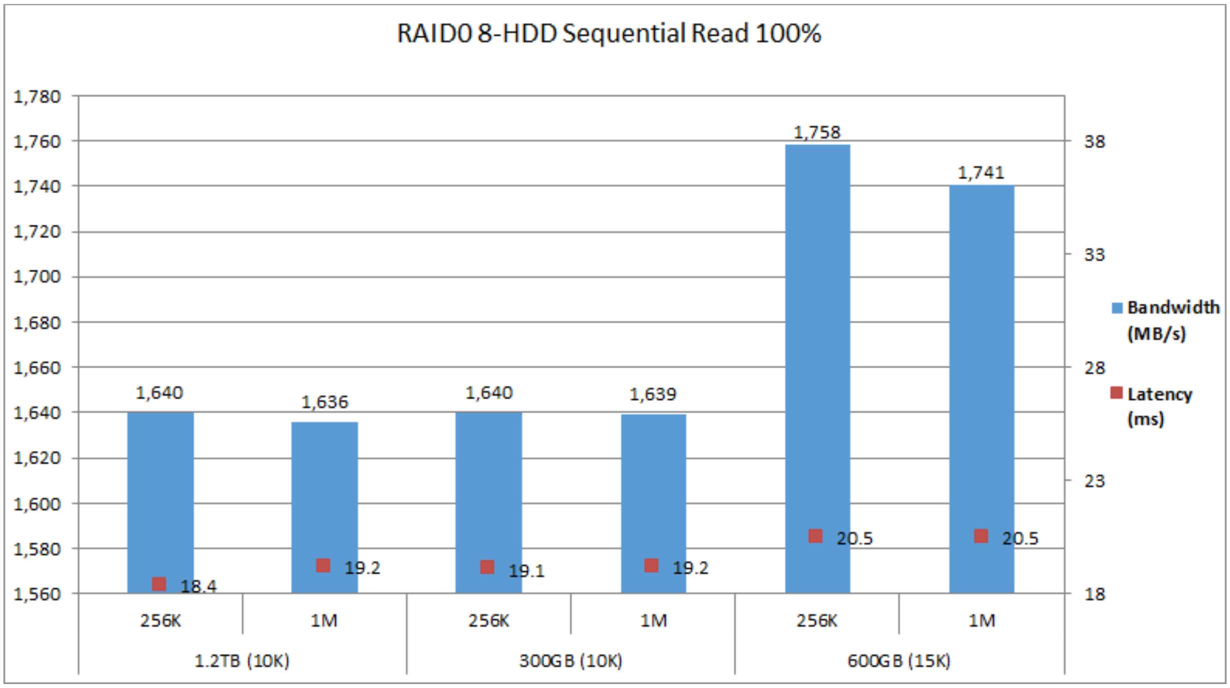

Figure 28 shows the performance of SFF HDDs under test for a RAID-0 configuration with a 100 percent sequential write access pattern. The numbers in the graph affirm that the bandwidth is well over 1600 MBps for both the 256-KB and 1-MB block sizes for all the SFF HDDs.

Figure 28. Sequential Write 100%

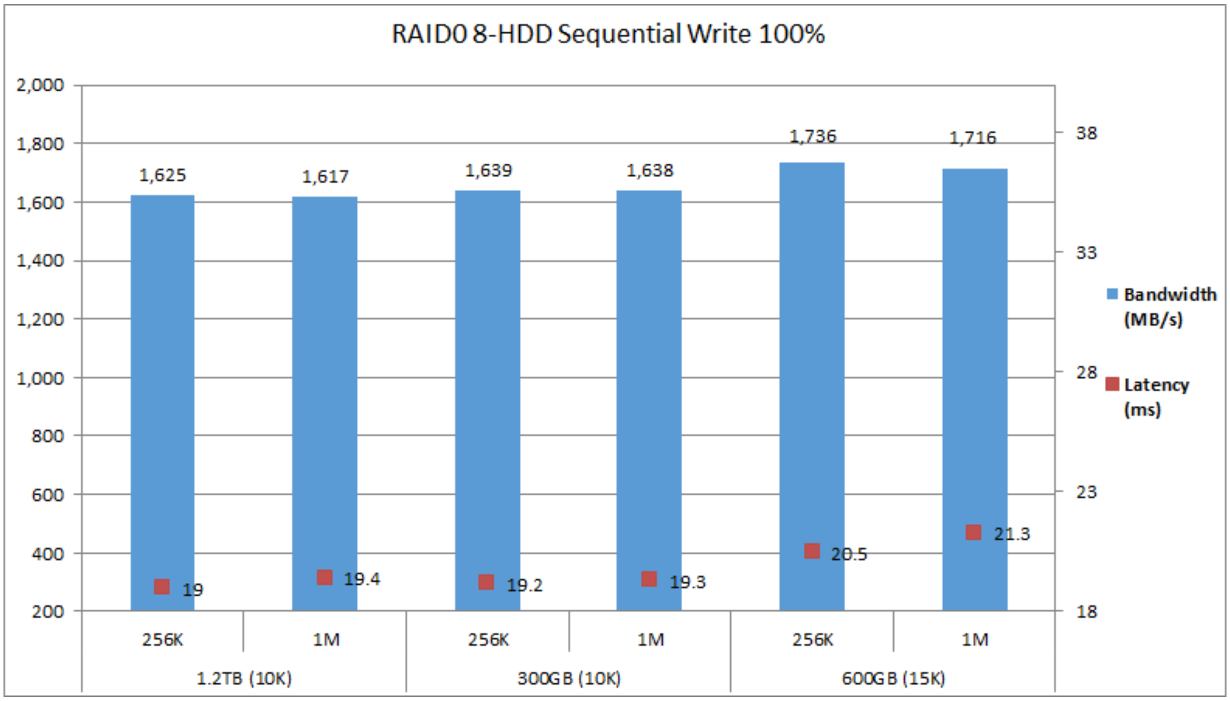

Maximum RAID-0 Performance

Maximum performance with a RAID-0 configuration was tested with three disk types: 15,000-rpm SFF HDD, 10,000-rpm SFF HDD, and 7200-rpm LFF HDD. This test was performed to achieve the best performance from these drives without concern for latency. The SFF HDD was tested with a RAID-0 8-disk configuration, and the LFF HDD was tested with a RAID-0 12-disk configuration.

Figure 29 shows the maximum performance of the 15,000- and 10,000-rpm SFF HDDs and 7200-rpm LFF HDD under test with a RAID-0 configuration with a 100 percent random read access pattern. The graph shows the comparative performance values achieved for each drive type with the 4- and 8-KB block sizes. Maximum performance for SFF HDDs was tested with an 8-disk configuration, and for LFF HDDs it was tested with a 12-disk configuration.

Figure 29. Random Read 100%

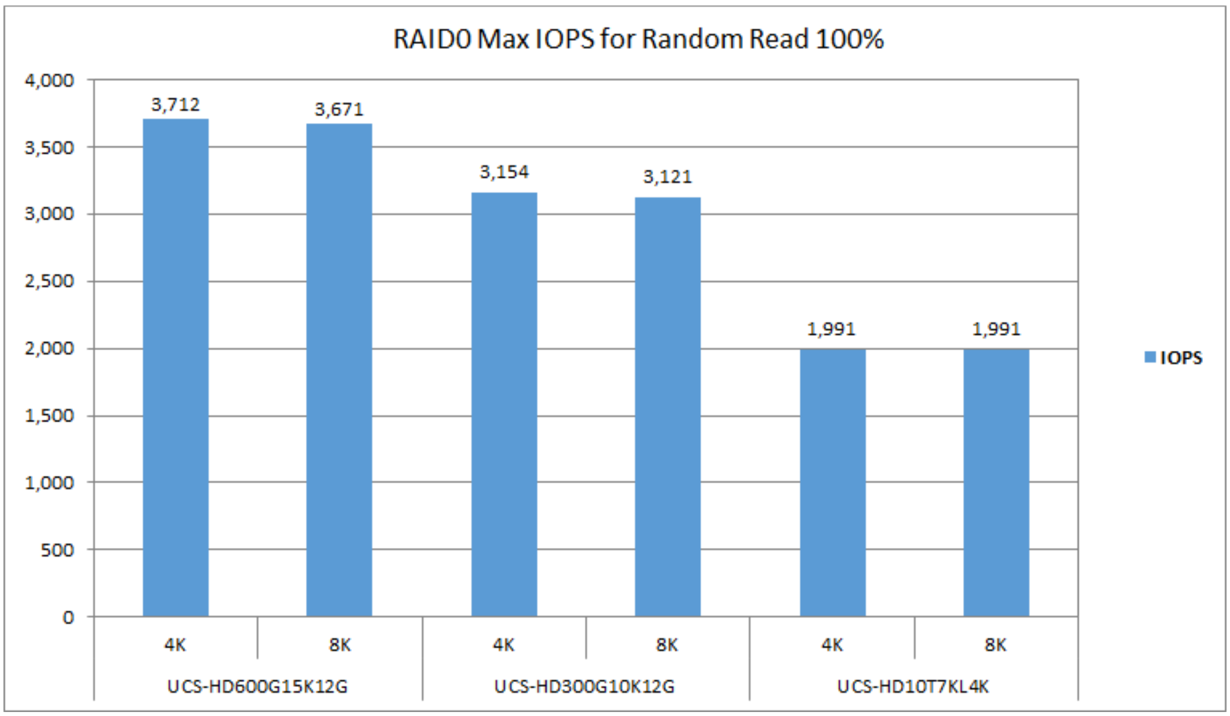

Figure 30 shows the maximum performance of the 15,000- and 10,000-rpm SFF HDDs and 7200-rpm LFF HDD under test with a RAID-0 configuration with a 100 percent random write access pattern. The graph shows the comparative performance values achieved for each drive type with the 4- and 8-KB block sizes.

Figure 30. Random Write 100%

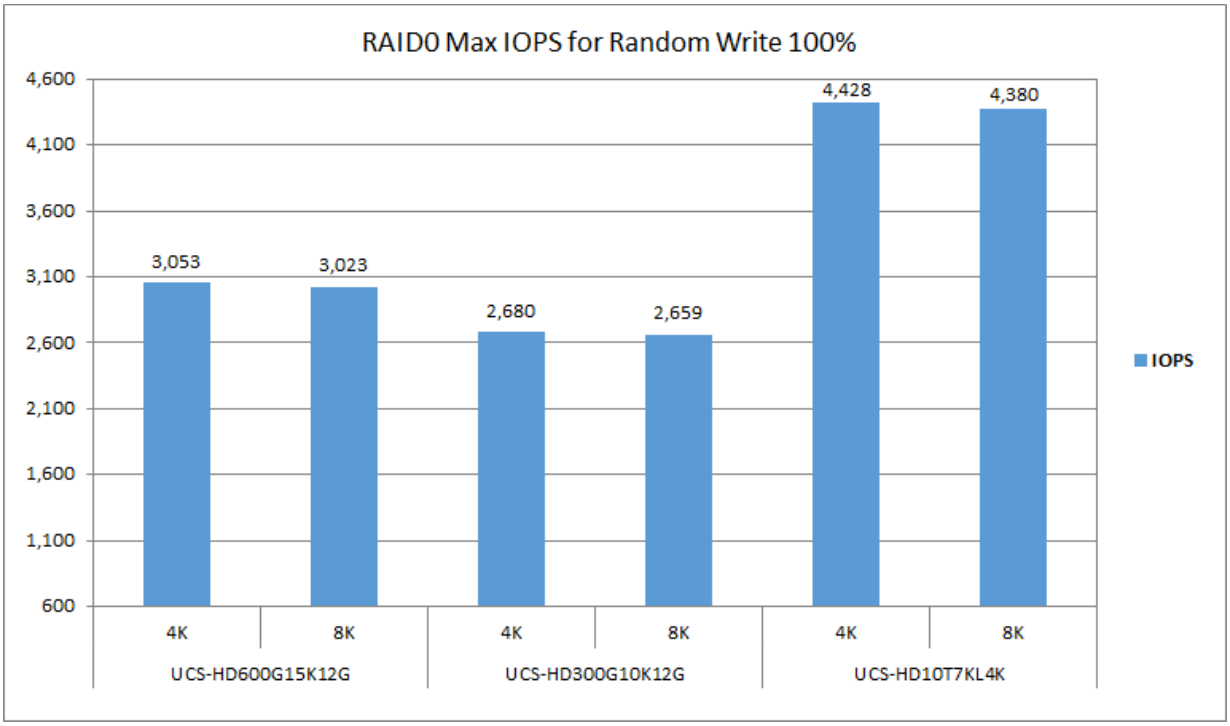

Figure 31 shows the maximum performance of the 15,000- and 10,000-rpm SFF HDDs and 7200-rpm LFF HDD under test with a RAID-0 configuration with a 100 percent sequential read access pattern.

Figure 31. Sequential Read 100%

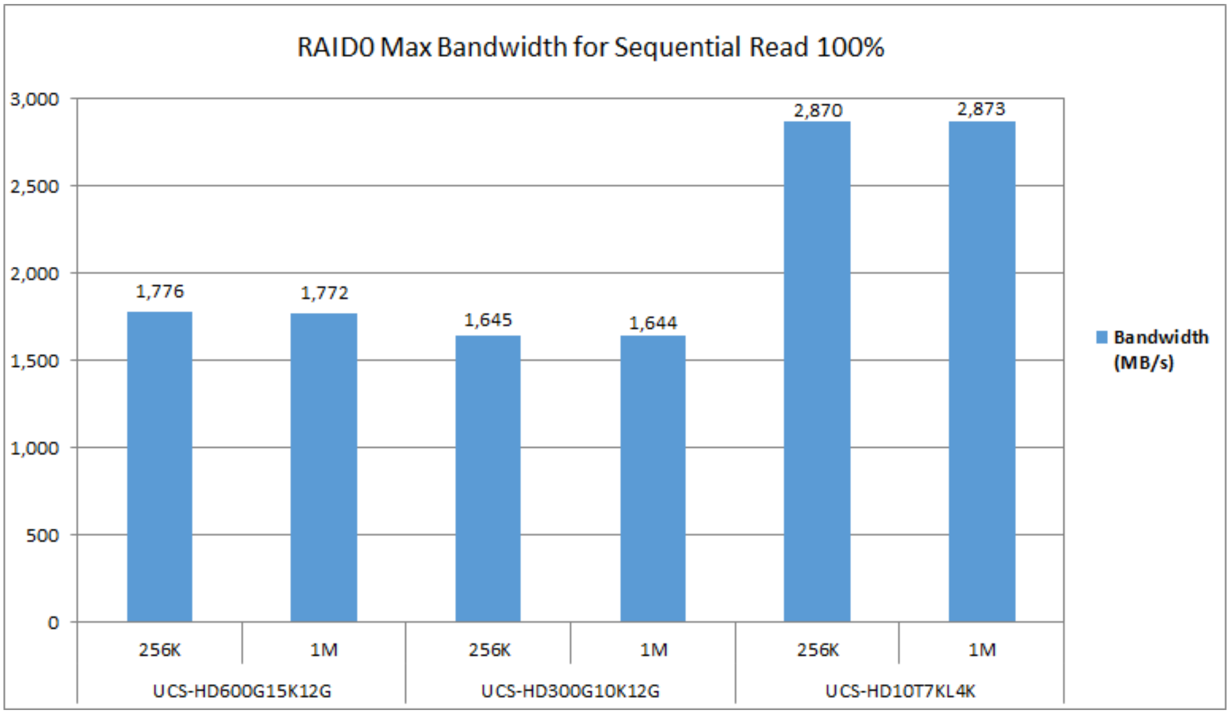

Figure 32 shows the maximum performance of the 15,000- and 10,000-rpm SFF HDDs and 7200-rpm LFF HDD under test with a RAID-0 configuration with a 100 percent sequential write access pattern.

Figure 32. Sequential Write 100%

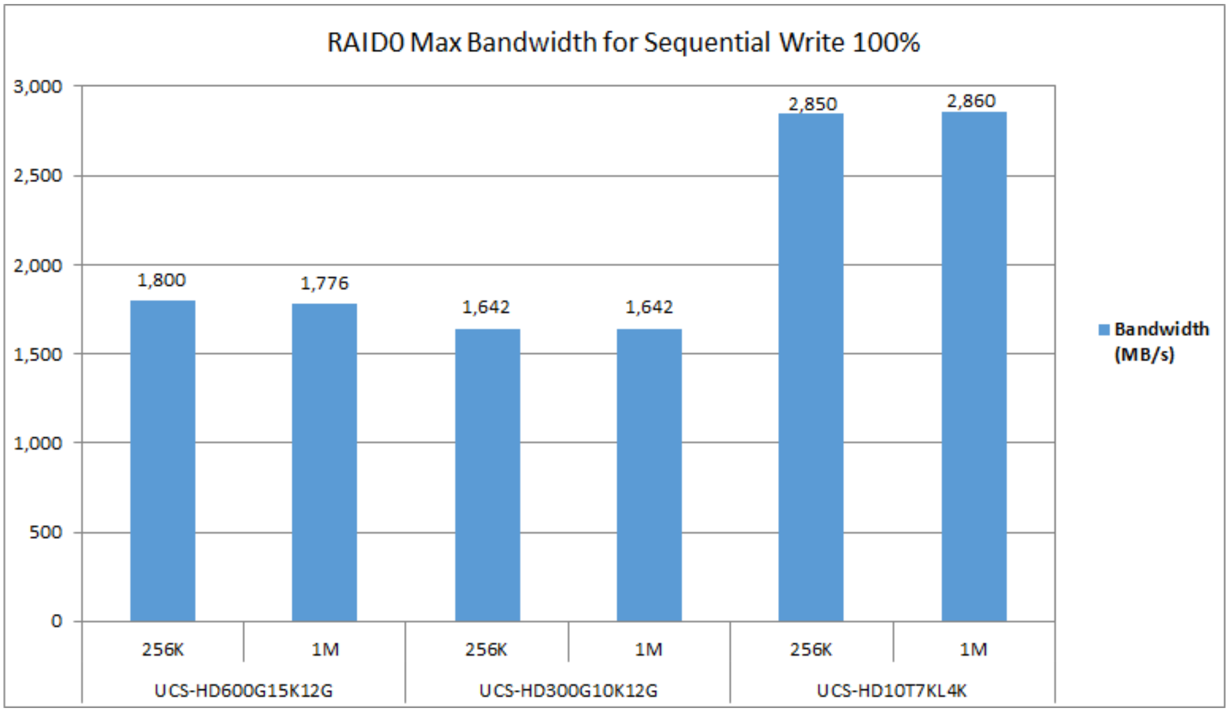

HDD RAID-5 Performance

LFF 12-HDD RAID-5 Performance

Figure 33 shows the performance of the LFF HDDs under test for a RAID-5 configuration with a random 70 percent read and 30 percent write access pattern. The graph affirms that the 10-TB LFF HDD can perform better for the 4- and 8-KB block sizes for the 12-disk configuration.

Figure 33. Random Read:Write 70:30%

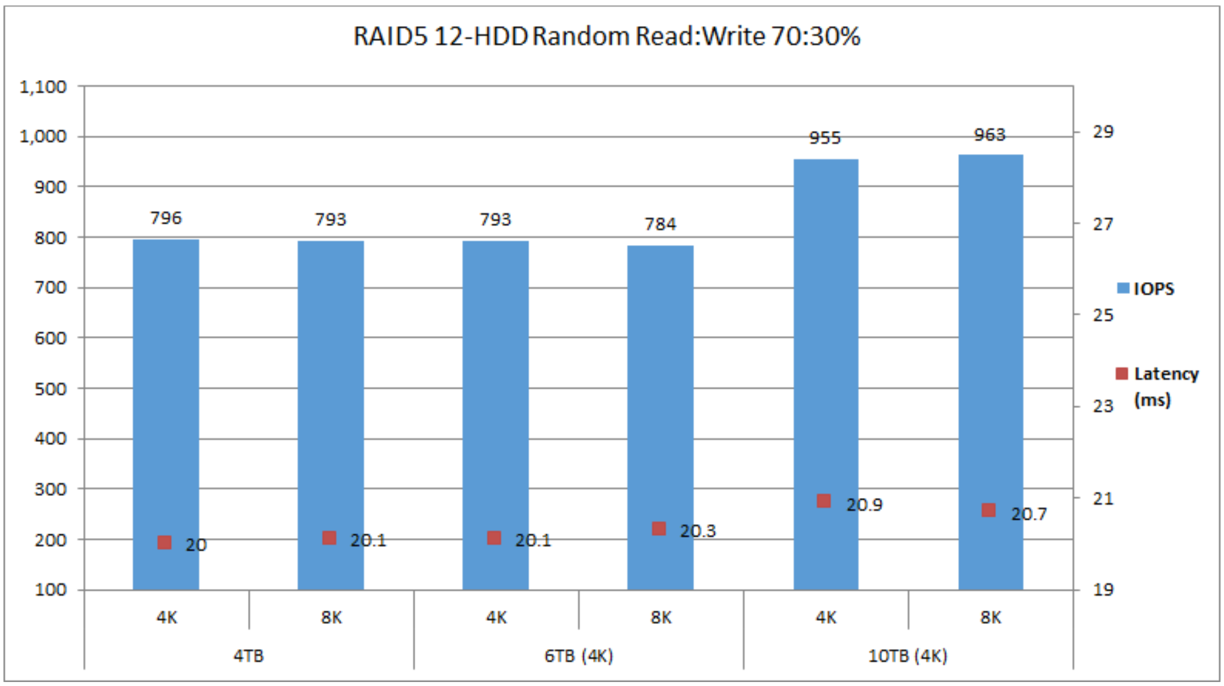

Figure 34 shows the performance of the LFF HDDs under test for a RAID-5 configuration with a random 50 percent read and 50 percent write access pattern. The graph affirms that the 10-TB LFF HDD can perform better for the 4- and 8-KB block sizes for the 12-disk configuration.

Figure 34. Random Read:Write 50:50%

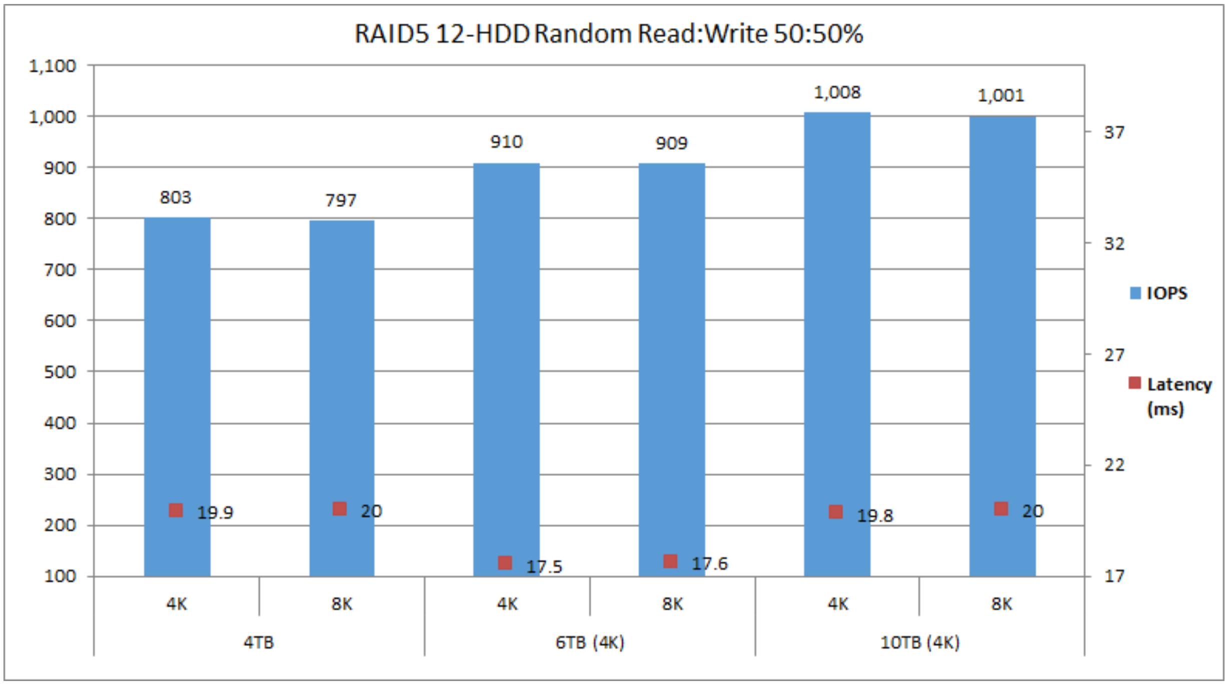

SFF 8-HDD RAID-5 Performance

Figure 35 shows the performance of the SFF HDDs under test for a RAID-5 configuration with a random 70 percent read and 30 percent write access pattern. The graph affirms that the 600-GB 15,000-rpm SFF HDD can perform better for the 4- and 8-KB block sizes for the 8-disk configuration.

Figure 35. Random Read:Write 70:30%

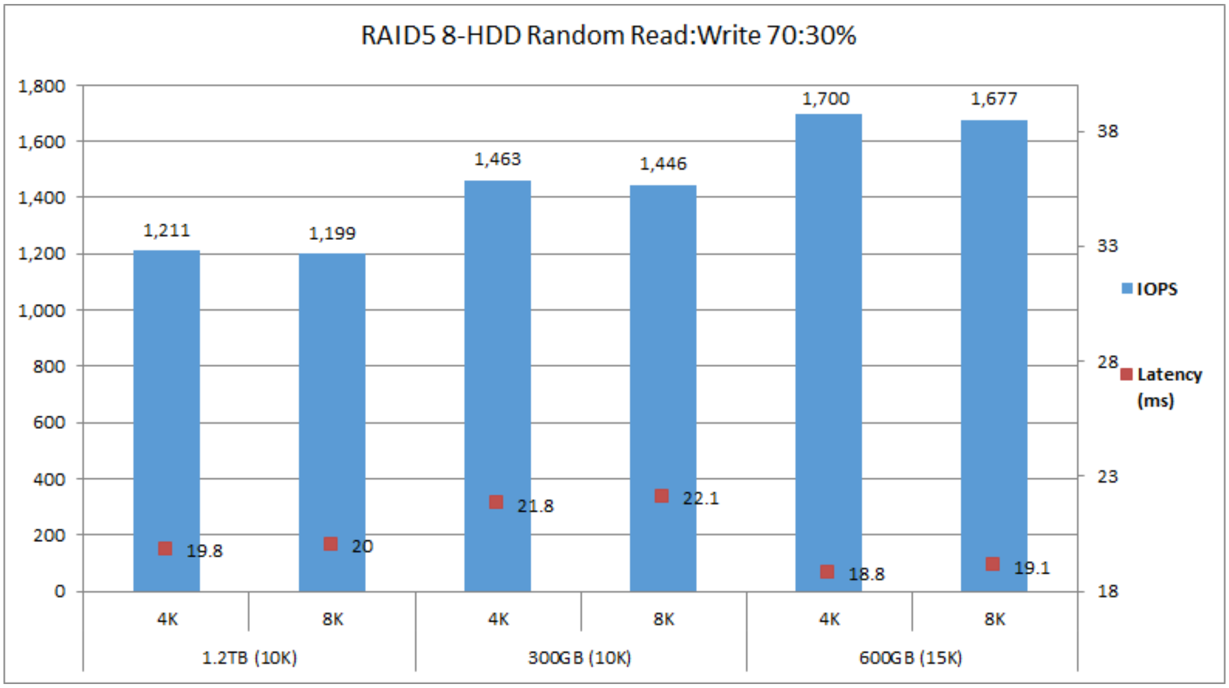

Figure 36 shows the performance of the SFF HDDs under test for a RAID-5 configuration with a random 50 percent read and 50 percent write access pattern. The graph affirms that the 600-GB 15,000-rpm SFF HDD can perform better for the 4- and 8-KB block sizes for the 8-disk configuration.

Figure 36. Random Read:Write 50:50%

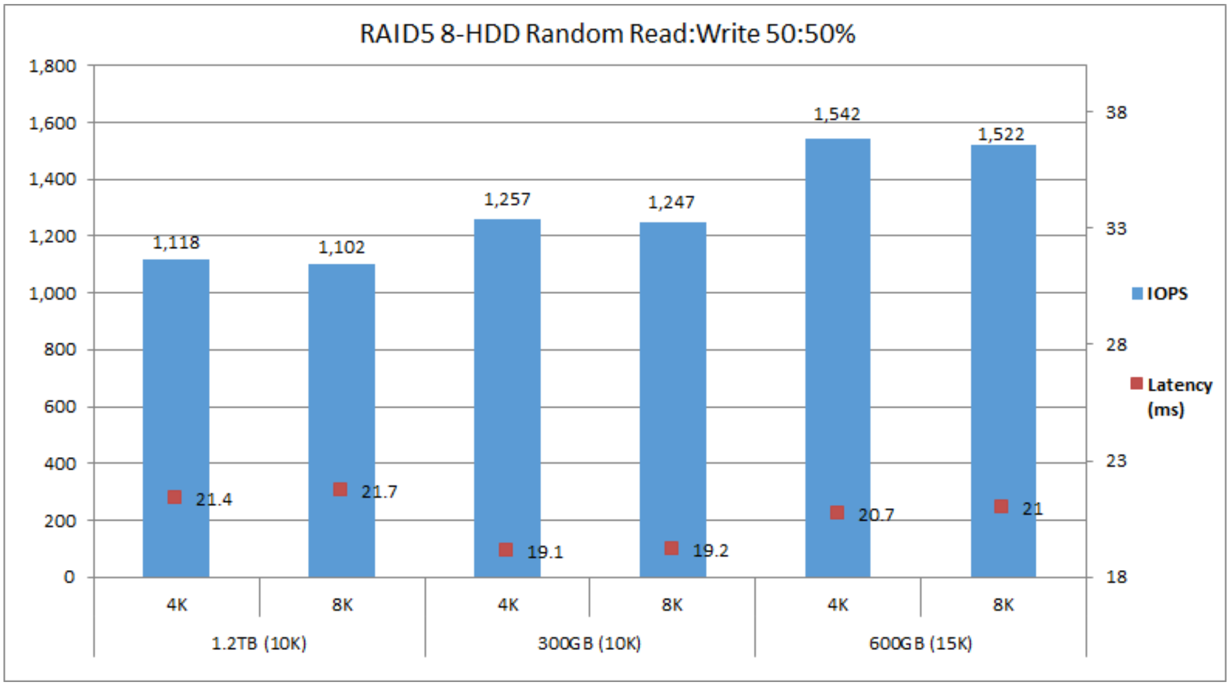

HDD RAID-10 Performance

LFF 12-HDD RAID-10 Performance

Figure 37 shows the performance of the LFF HDDs under test for a RAID-10 configuration with a sequential 100 percent write access pattern. The graph affirms that the 10-TB LFF HDD can perform better with bandwidth greater than 1400 MBps for the 4- and 8-KB block sizes with the 12-disk configuration.

Figure 37. Sequential Write 100%

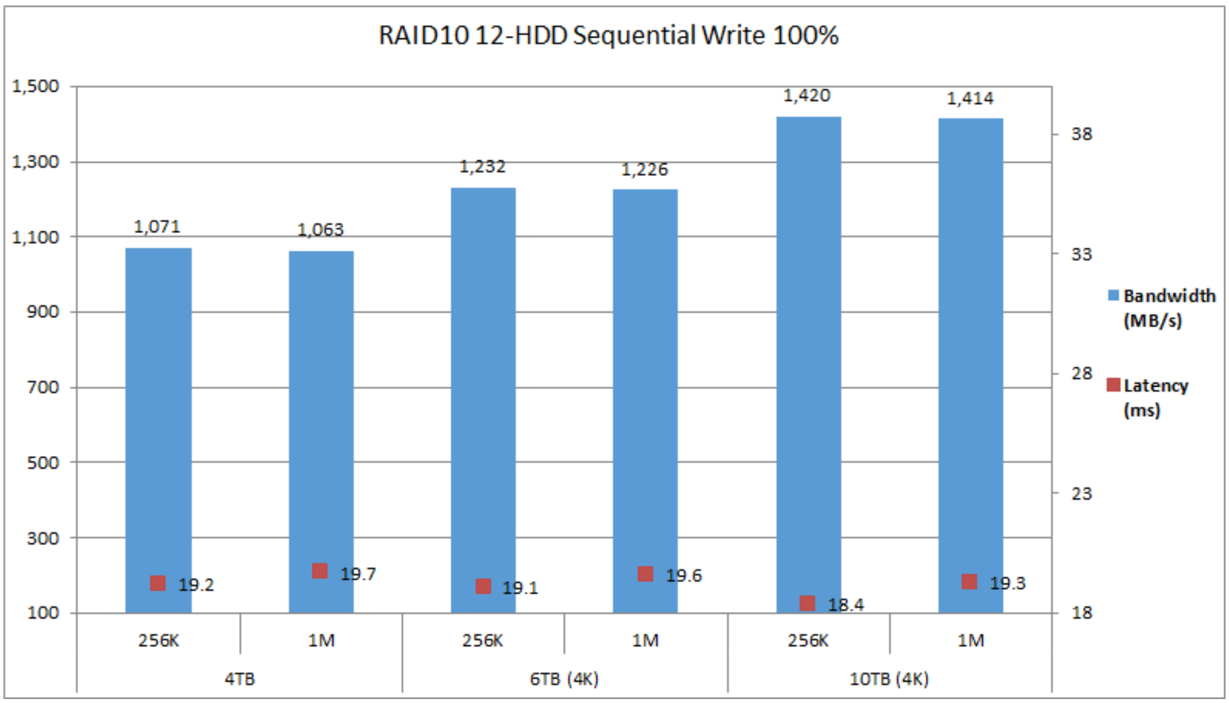

SFF 8-HDD RAID-10 Performance

Figure 38 shows the performance of the SFF HDDs under test for a RAID-10 configuration with a sequential 100 percent write access pattern. The graph affirms that the 600-GB 15,000-rpm SFF HDD can perform better for the 4- and 8-KB block sizes for the 8-disk configuration.

Figure 38. Sequential Write 100%

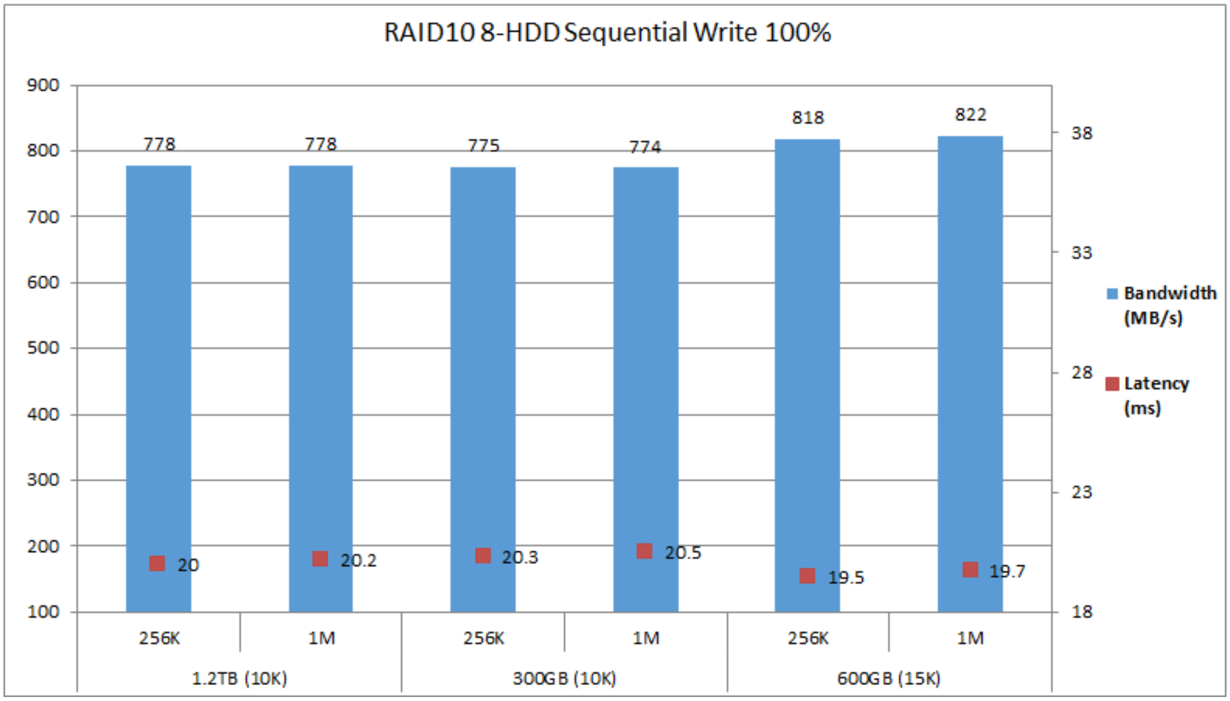

HDD JBOD Performance

LFF 12-HDD JBOD Performance

Figure 39 shows the performance of the LFF HDDs under test for a JBOD configuration with a random read access pattern. The graph shows that the performance for the 4- and 8-KB block sizes is more then 1000 IOPS within the acceptable latency of 20 ms.

Figure 39. Random Read 100%

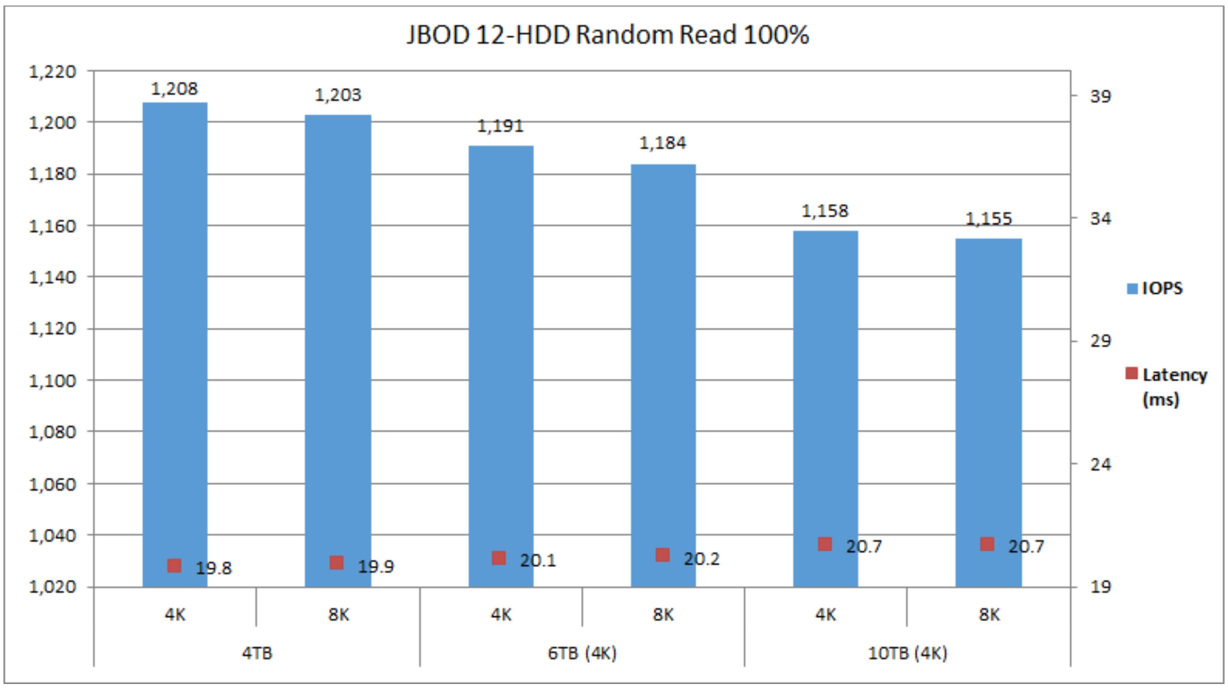

Figure 40 shows the performance of the LFF HDDs under test for a JBOD configuration with a random write access pattern. The graph shows that the 10-TB HDD achieves greater performance at the 4- and 8-KB block sizes than do the 4- and 6-TB LFF HDDs.

Figure 40. Random Write 100%

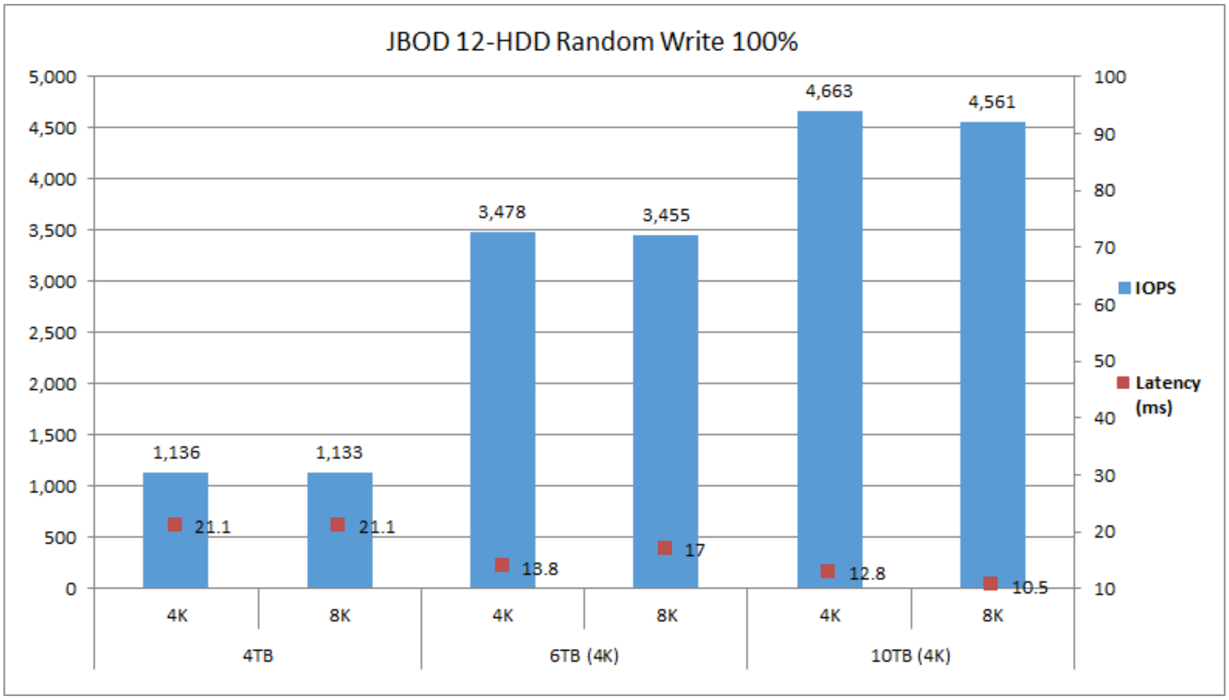

Figure 41 shows the performance of the LFF HDDs under test for a JBOD configuration with a random 70 percent read and 30 percent write access pattern. The graph shows that the 10-TB HDD achieves greater performance at the 4- and 8-KB block sizes than do the 4- and 6-TB LFF HDDs.

Figure 41. Random Read:Write 70:30%

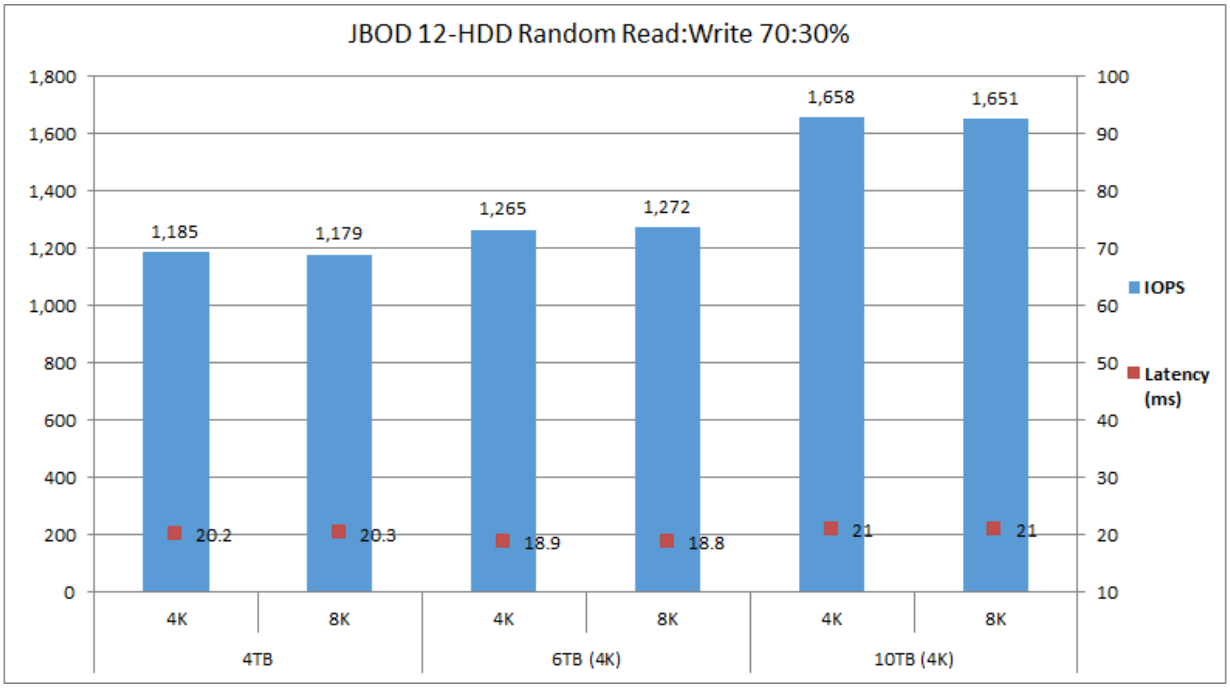

Figure 42 shows the performance of the LFF HDDs under test for a JBOD configuration with a 100 percent sequential read access pattern. The graph shows that the 10-TB HDD achieves greater performance at the 256-KB and 1-MB block sizes than do the 4- and 6-TB LFF HDDs.

Figure 42. Sequential Read 100%

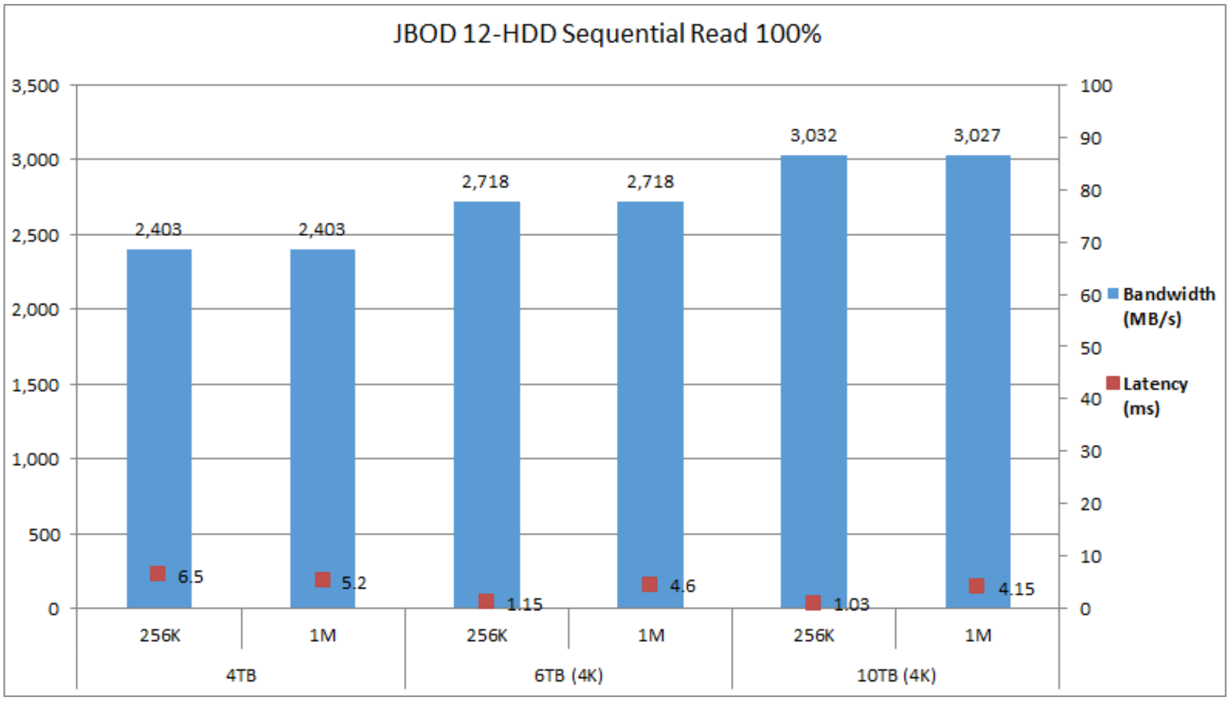

Figure 43 shows the performance of the LFF HDDs under test for a JBOD configuration with a 100 percent sequential write access pattern. The graph shows that the 10-TB HDD achieves greater performance at the 256-KB and 1-MB block sizes than do the 4- and 6-TB LFF HDDs.

Figure 43. Sequential Write 100%

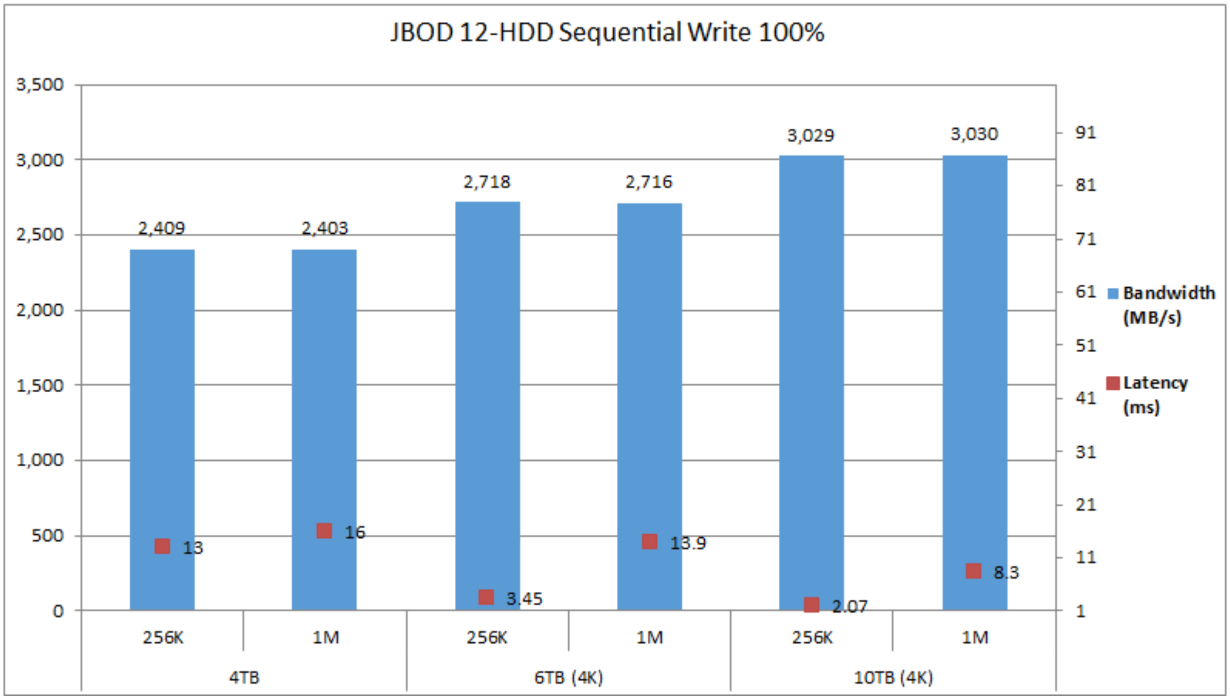

SFF 8-HDD JBOD Performance

Figure 44 shows the performance of the SFF HDDs under test for a JBOD configuration with a 100 percent random read access pattern. The graph shows the comparitive performance values achieved for SFF HDDs to help customers understand the performance trade-off when choosing a hard drive type. The graph shows that the 600-GB SFF HDD at 15,000 rpm achieves better performance with an acceptable latency of 20 ms. The testing also showed that the drives perform better with increased latency for the 4- and 8-KB block sizes.

Figure 44. Random Read 100%

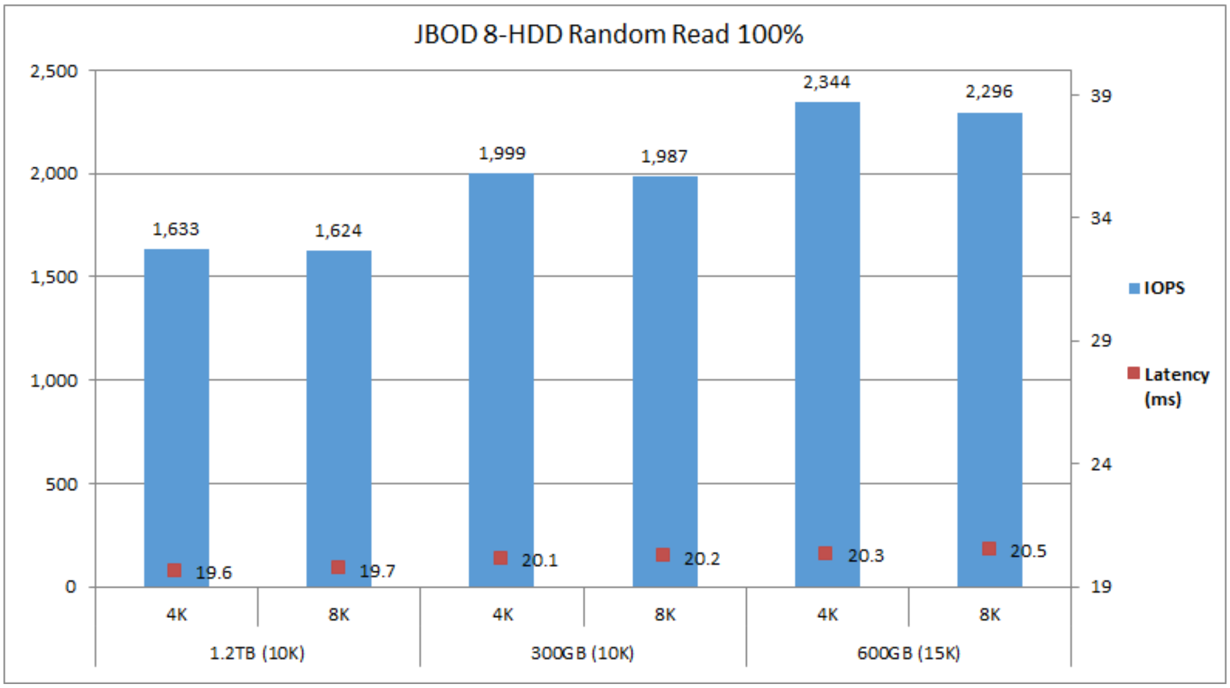

Figure 45 shows the performance of the SFF HDDs under test for a JBOD configuration with a 100 percent random write access pattern. The numbers in the graph affirm that the IOPS rate for the 4- and 8-KB block sizes for the 600-GB HDD at 15,000 rpm is better than that for the 10,000-rpm HDDs.

Figure 45. Random Write 100%

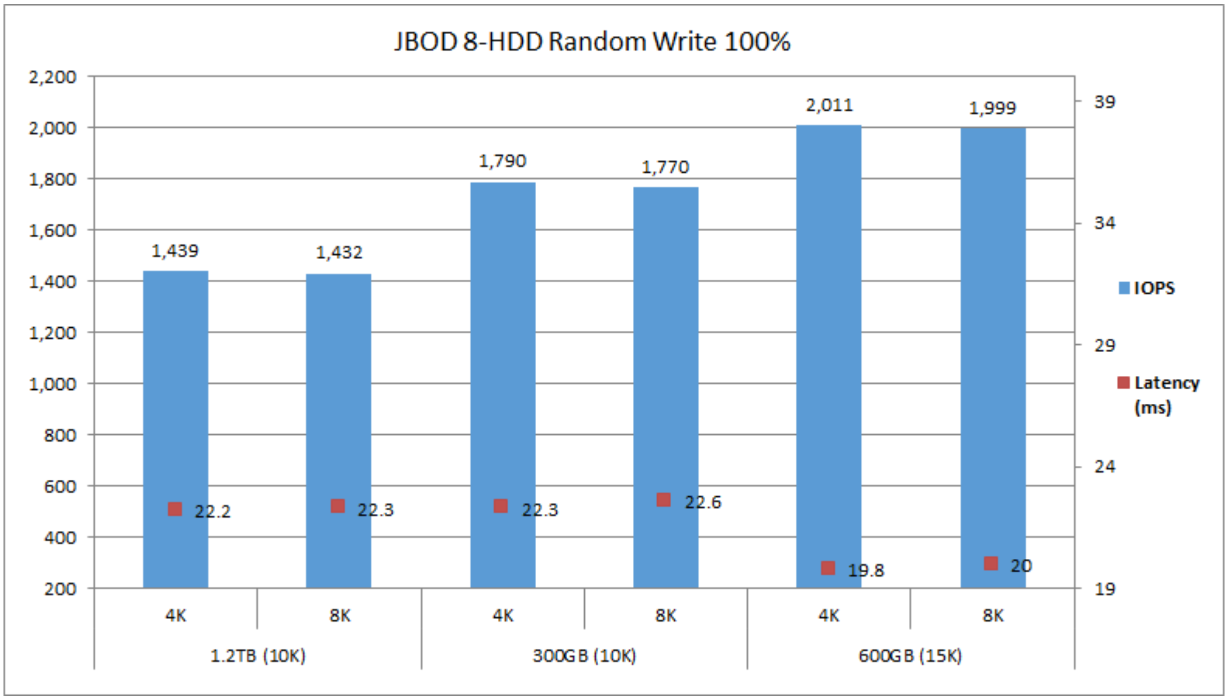

Figure 46 shows the performance of the SFF HDDs under test for a JBOD configuration with a random 70 percent read and 30 percent write access pattern. The numbers in the graph affirm that the IOPS rate for the 4- and 8-KB block sizes for the 600-GB HDD at 15,000 rpm is better than that for the 10,000-rpm HDDs.

Figure 46. Random Read:Write 70:30%

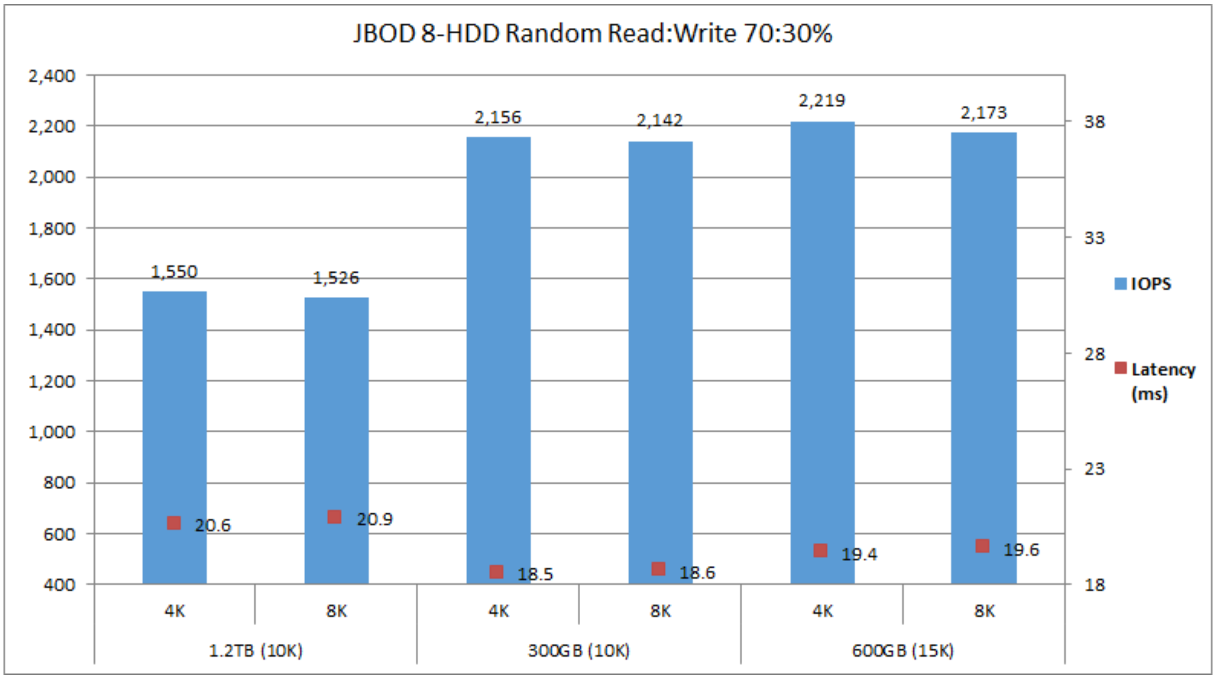

Figure 47 shows the performance of the SFF HDDs under test for a JBOD configuration with a 100 percent sequential read access pattern. The graph shows that the 600-GB drive with 15,000 rpm provides greater performance at the 256-KB and 1-MB block sizes than do the 1.2-TB and 300-GB drives with 10,000 rpm.

Figure 47. Sequential Read 100%

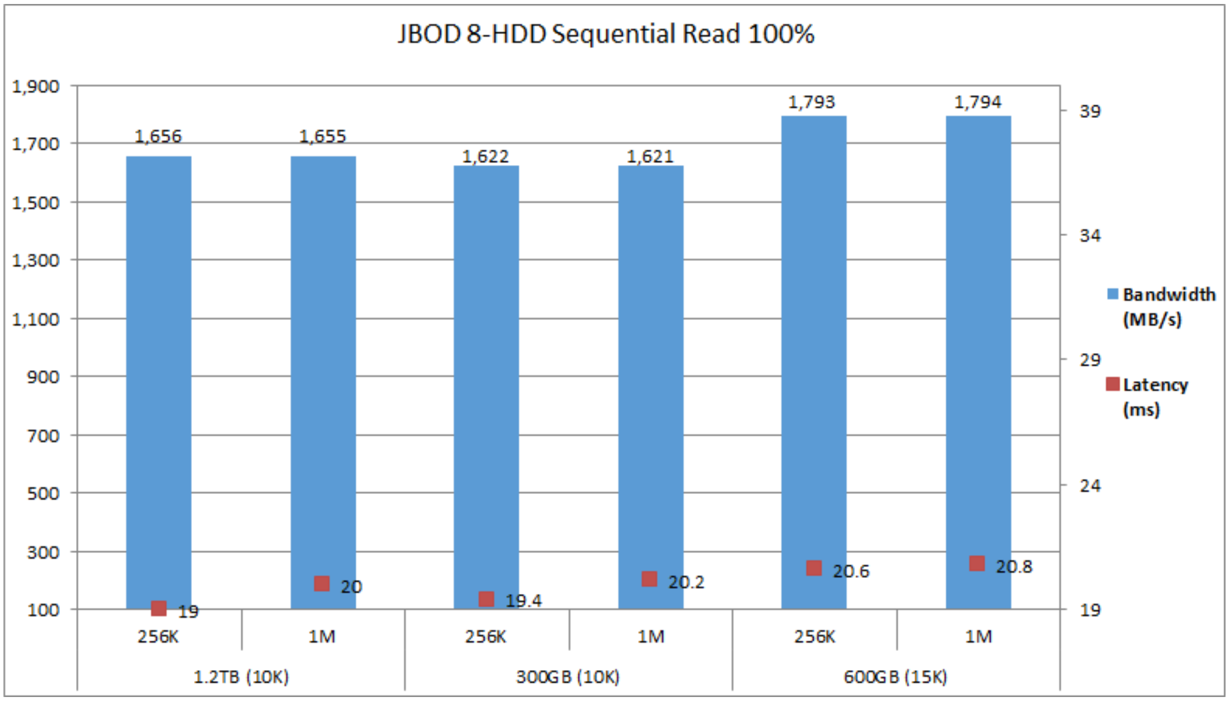

Figure 48 shows the performance of the SFF HDDs under test for a JBOD configuration with a 100 percent sequential write access pattern. The graph shows that the 600-GB drive with 15,000 rpm provides greater performance at the 256-KB and 1-MB block sizes than do the 1.2-TB and 300-GB drives with 10,000 rpm.

Figure 48. Sequential Write 100%

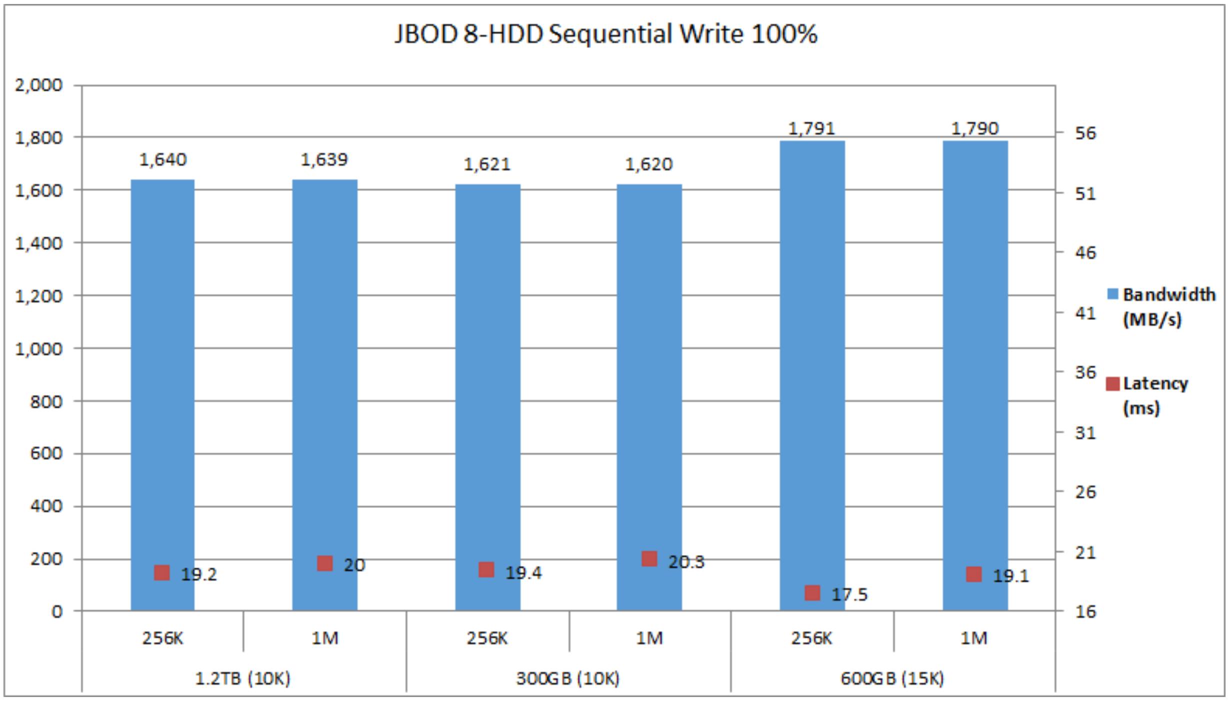

For More Information

- http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c240-m4-rack-server/index.html
- http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/C240M4/install/C240M4/raid.html
- http://www.lsi.com/downloads/Public/Direct Assets/LSI/Benchmark_Tips.pdf

Appendix: Test Environment

Table 10 lists the details of the server under test.

Table 10. Server Properties

Name	Value
Product names	Cisco UCS C240 M4SX and C240 M4L
CPUs (2)	Intel Xeon processor CPU E5-2699 v4 at 2.20 GHz
Number of cores	48
Number of threads	96
Total memory	256 GB
Memory DIMMs (16)	16 GB at 2 DIMMs per channel (DPCs)
Memory speed	2400 MHz
Network controller	Intel I350 1-Gbps network controller
VIC adapter	Cisco UCS VIC 1227 mLOM 10-Gbps SFP+ pluggable
Storage controller	Intel I350 mLOM 1-Gbps network controller
RAID controllers	• Cisco 12-Gbps SAS modular RAID controller (UCSC-MRAID12G) and 1-GB FBWC (UCSC-MRAID12G-1GB) • Cisco 12-Gbps modular SAS HBA (UCSC-SAS12GHBA)
SSDs	• 1.6-TB 2.5-inch enterprise performance 12-Gbps SAS SSD, 10x endurance (UCS-SD16TB12S4-EP) • 480-GB 2.5-inch enterprise performance 6-Gbps SATA SSD, 3x endurance (UCS-SD480G12S3-EP) • 3.8-TB 2.5-inch enterprise value 6-Gbps SATA SSD (UCS-SD38TBKS4-EV) • 1.6-TB 2.5-inch enterprise value 12-Gbps SAS (UCS-SD16TBKS4-EV) • 480-GB 2.5-inch enterprise value 6-Gbps SATA SSD (UCS-SD480GBKS4-EV)
SFF HDDs	• 300-GB 12-Gbps SAS 10,000-rpm SFF HDD (UCS-HD300G10K12G) • 600-GB 12-Gbps SAS 15,000-rpm SFF HDD (UCS-HD600G15K12G) • 1.2-TB 12-Gbps SAS 10,000-rpm SFF HDD (UCS-HD12TB10K12G)
LFF HDDs	• 4-TB 12-Gbps SAS 7200-rpm LFF HDD (UCS-HD4T7KL12G) • 6-TB 12-Gbps SAS 72-rpm LFF HDD (UCS-HD6T7KL) • 10-TB 12-Gbps SAS 7200-rpm LFF HDD (UCS-HD10T7KL)

Table 11 lists the recommended server BIOS settings.

Table 11. Server BIOS Settings

Name	Value
BIOS version	C240M4.2.0.13d.0.0812161132
Intel Hyper-Threading Technology	Enabled
Number of enabled cores	All
Execute disable bit	Enabled
Intel Virtualization Technology (VT)	Enabled
Intel VT for Directed I/O (VT-d)	Enabled
Intel interrupt remapping	Disabled
Intel pass-through direct memory access (DMA)	Disabled
Intel VT-d coherence support	Disabled
Intel VT-d Address Translation Services (ATS) support	Enabled
CPU performance	Enterprise
Hardware prefetcher	Enabled
Adjacent-cache-line prefetcher	Enabled

Data cache unit (DCU) streamer prefetch	Enabled
DCU IP prefetcher	Enabled
Direct cache access support	Auto
Power technology	Custom
Enhanced Intel SpeedStep Technology	Enabled
Intel Turbo Boost Technology	Enabled
Processor C3 report	Disabled
Processor C6 report	Enabled
Processor power state C1 enhanced	Enabled
P-state coordination	Hardware all
Energy performance tuning	BIOS
Energy performance	Balanced performance
Package C-state limit	C0/C1
Extended Cisco Application Policy Infrastructure Controller (APIC)	XAPIC
Workload configuration	I/O sensitive
CPU hardware power management (HWPM)	Disabled
CPU autonomous C-state	Disabled
Processor corrected machine check interrupt (CMCI)	Disabled
Select memory RAS	Maximum performance
Non-uniform memory access (NUMA)	Enabled
Channel interleaving	Auto
Rank interleaving	Auto
Patrol scrub	Disabled
Demand scrub	Disabled
Altitude	300 meters
Intel QuickPath Interconnect (QPI) link frequency select	Auto
QPI snoop mode	Home directory snoop

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-738090-00 12/16