

Cisco UCS C240 M5 Rack Server Disk I/O Characterization



April 2018

Executive summary

This document outlines the I/O performance characteristics of the Cisco UCS® C240 M5 Rack Server using the Cisco® 12-Gbps modular RAID controller with a 4-GB cache module (UCSC-MRAID-M5HD) and Cisco 12-Gbps modular serial-attached Small Computer System Interface (SCSI; SAS) host bus adapter (HBA; UCSC-SAS-M5HD). Performance comparisons of various SAS solid-state disks (SSDs), hard-disk drives (HDDs), 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 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 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. 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 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.

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 an enterprise’s most common type of workload. 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 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 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 and HDDs with configurations of RAID 0, RAID 5, and RAID 10 virtual disks because most of the workloads targeted for Cisco UCS C240 M5 Rack Servers use these RAID levels. The Cisco UCS C240 M5SX server used for the I/O performance characterization tests supports up to 26 HDDs and SSDs. The performance tests described here were limited to 10-disk and 16-disk configurations for SFF SSDs and HDDs for RAID 0 and RAID 5, and 8-disk and 16-disk configurations for RAID 10 for SFF SSDs and HDDs.

Solution components

The tested solution used a Cisco UCS C240 M5SX Rack Server with SSDs and HDDs.

Cisco UCS C240 M5 Rack Server overview

The Cisco UCS C240 M5 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 M5 server provides:

- Support for a 2RU 2-socket server using Intel® Xeon® Scalable processors
- Support for 2666-MHz DDR4 DIMMs and 128-GB DIMMs
- Increased storage density 24 front-pluggable 2.5-inch SFF drive bays, or 12 front-pluggable 3.5-inch large-form-factor (LFF) drive bays and 2 rear 2.5-inch SFF drive bays
- Non-Volatile Memory Express (NVMe) PCI Express (PCIe) SSD support (for up to 2 drives on the standard chassis SKU or up to 10 drives on the NVMe-optimized SKU)
- Cisco 12-Gbps SAS RAID modular controller and Cisco 12-Gbps SAS HBA controller
- 2 Flexible Flash (FlexFlash) Secure Digital (SD) card slots or 2 modular M.2 SATA slots
- 10-Gbps embedded Intel x550 10GBASE-T LAN-on-motherboard (LOM) port
- 1 modular LOM (mLOM) slot
- 6 PCIe Generation 3 slots
- Up to 2 hot-pluggable redundant power supplies

The Cisco UCS C240 M5 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 M5 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 M5 Rack Servers
- CPU: 2 x 2.40-GHz Intel Xeon Gold 6148
- Memory: 12 x 16-GB (192-GB) DDR4

- Cisco UCS Virtual Interface Card (VIC) 1387 mLOM 10-Gbps Enhanced Small Form-Factor Pluggable (SFP+)
- Cisco 12-Gbps modular RAID controller with 4-GB cache module
- Cisco 12-Gbps modular SAS HBA

Cisco UCS C240 M5 server models

The Cisco UCS C240 M5 server can be configured in four different models to match specific customer environments. The Cisco UCS C240 M5 can be used as a standalone server or as part of Cisco UCS, which unifies computing, networking, management, virtualization, and storage access into a single integrated architecture that enables end-to-end server visibility, management, and control in both bare-metal and virtualized environments.


The Cisco UCS C240 M5 server includes a dedicated internal mLOM slot for installation of a Cisco VIC or third-party network interface card (NIC), without consuming a PCI slot, in addition to two 10GBASE-T Intel x550 embedded (on the motherboard) LOM ports.

Cisco UCS C240 M5 servers are broadly categorized into SFF and LFF models as follows:

- Cisco UCS C240 M5 Rack Server (SFF model)
 - UCSC-C240-M5SX
 - 24 SFF front-facing SAS and SATA HDDs or SAS and SATA SSDs
 - Optionally, up to 2 front-facing SFF NVMe PCIe SSDs (replacing the SAS and SATA drives)
 - Optionally, up to 2 rear-facing SFF SAS and SATA HDDs or SSDs, or up to 2 rear-facing SFF NVMe PCIe SSDs
 - UCSC-C240-M5SN
 - Up to 8 front-facing SFF NVMe PCIe SSDs only (replacing the SAS and SATA drives)
 - 16 front-facing SFF SAS and SATA HDDs or SAS and SATA SSDs; drives occupy slots 9 to 24
 - Optionally, up to two rear-facing SFF NVMe PCIe SSDs (must be NVMe only); rear-facing NVMe drives are connected from Riser 2
 - UCSC-C240-M5S
 - 8 front-facing SFF SAS and SATA HDDs or SSDs
 - Optionally, up to 2 front-facing NVMe PCIe SSDs (replacing the SAS and SATA drives)
 - Optionally, up to 2 rear-facing SFF SAS/SATA HDDs or SSDs, or up to 2 rear-facing SFF NVMe PCIe SSDs
 - Optionally, 1 front-facing DVD drive
- Cisco UCS C240 M5 Rack Server (LFF model)
 - LFF drives with 12-drive backplane; the server can hold up to:
 - 12 LFF 3.5-inch front-facing SAS and SATA HDDs or SAS and SATA SSDs
 - Optionally, up to 2 front-facing SFF NVMe PCIe SSDs (replacing the SAS and SATA drives)
 - Optionally, up to 2 SFF 2.5-inch, rear-facing SAS and SATA HDDs or SSDs, or up to 2 rear-facing SFF NVMe PCIe SSDs


The Cisco UCS C240 M5SX SFF server (Figure 1) extends the capabilities of the Cisco UCS portfolio in a 2RU form factor with the addition of Intel Xeon Scalable processors, 24 DIMM slots for 2666-MHz DDR4 DIMMs and capacity points of up to 128 GB, up to 6 PCIe 3.0 slots, and up to 26 internal SFF drives. The C240 M5 SFF server also includes one dedicated internal slot for a 12-Gbps SAS storage controller card.

Figure 1. Cisco UCS C240 M5SX



The Cisco UCS C240 M5 LFF server (Figure 2) extends the capabilities of the Cisco UCS portfolio in a 2RU form factor with the addition of Intel Xeon Scalable processors, 24 DIMM slots for 2666-MHz DDR4 DIMMs and capacity points of up to 128 GB, up to 6 PCIe 3.0 slots, and up to 12 front-facing internal LFF drives. The C240 M5 LFF server also includes one dedicated internal slot for a 12-Gbps SAS storage controller card.

Figure 2. Cisco UCS C240 M5L



For details about configuring a specific model of C240 M5 server, please refer to the appropriate specification sheets:

- SFF models (either HDD or SSD): <https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m5-sff-specsheet.pdf>
- LFF model (HDD): <https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m5-lff-specsheet.pdf>

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

Table 1. Part numbers for Cisco UCS C240 M5 high-density SFF and LFF base rack server models

Part number	Description
UCSC-C240-M5SX	Cisco UCS C240 M5: 24 SFF front drives with option for 2 SFF rear drives with no CPU, memory, HDD, PCIe cards, or power supply
UCSC-C240-M5SN	Cisco UCS C240 M5: 10 SFF NVMe (8 front and 2 rear) and 16 (front) SATA and SATA drives with no CPU, memory, HDD, PCIe cards, or power supply
UCSC-C240-M5S	Cisco UCS C240 M5: 8 SFF front drives with option for 2 SFF rear drives with no CPU, memory, HDD, PCIe cards, or power supply
UCSC-C240-M5L	Cisco UCS C240 M5: LFF with no CPU, memory, HDD, SSD, PCIe cards, tool-free rail kit, or power supply, with 12-drive backplane

The performance testing described in this document uses the Cisco UCS C240 M5SX server, which supports 26 HDDs and SSDs with SAS expanders.

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.

SSDs have no physical tracks or sectors and no mechanical movement. Thus, SSDs can reach memory addresses (logical block addresses [LBAs]) 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 strong justification for their additional cost for random operations. Therefore, Cisco recommends SSDs for random I/O applications.

SSDs and HDDs used in the tests

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

- Low cost with high-capacity SSD (1.9-GB enterprise value SAS at 12 Gbps)
- High speed with 10x endurance SSD (400-GB enterprise performance SAS at 12 Gbps)
- High-capacity SATA SSD (7.6-TB enterprise value SATA at 6 Gbps)
- High-speed SAS HDD (600-GB 15,000-rpm 12-Gbps SAS HDDs)
- High-capacity SAS HDD (1.2-TB 10,000-rpm 12-Gbps SAS HDDs)

To meet the requirements of different application environments, Cisco offers both enterprise performance (EP) SSDs and enterprise value (EV) SSDs.

SSD types used in I/O characterization

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 the latest SSDs with higher capacities, plus better performance and lower cost, SSDs are increasingly becoming an obvious storage choice for enterprises.

Enterprise performance SSDs versus enterprise value SSDs

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 one 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 SSDs 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 2 TB in a 2.5-inch SFF model or 10 TB in a 3.5-inch LFF model. 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 2 summarizes the supported RAID levels and their characteristics.

Table 2. 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

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.

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

Access policy can be set 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

Disk cache policy can be set 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 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

I/O cache policy can be set 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.
- **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

Read policy can be set as follows:

- **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 policy can be set 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 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 3 lists the workload types tested.

Table 3. 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 4 and 5 list the I/O mix ratios chosen for the sequential access and random access patterns, respectively.

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

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

Table 5. I/O mix ratio for random access pattern

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

Tables 6 and 7 list the recommended virtual drive configuration parameters for deployment of SSDs and HDDs.

Table 6. 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

Table 7. 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:

- Ten RAID 0 virtual drives were created with 10 disks for the SFF SSDs and SFF HDDs.
- Sixteen RAID 0 virtual drives were created with 16 disks for the SFF SSDs and SFF HDDs.
- One RAID 5 virtual drive was created with 10 and 16 disks for the SFF SSDs and SFF HDDs.
- One RAID 10 virtual drive was created with 10 and 16 disks for the SFF SSDs and SFF HDDs.
- The RAID configuration was tested with the Cisco 12-Gbps modular RAID controller with a 4-GB cache module with 10 and 16 disks for the SFF SSDs and SFF HDDs.
- The JBOD configuration was tested with the Cisco UCS 12-Gbps SAS modular HBA with 10 and 16 disks for the SFF SSDs and SFF HDDs.
- Random workload tests were performed using 4- and 8-KB block sizes for all SSDs and HDDs.
- Sequential workload tests were performed using a 256-KB block size for all SSDs and HDDs.

Table 8 lists the recommended lometer settings.

Table 8. Recommended lometer settings

Name	Value
lometer version	Release 1.1.0
Run time	30 minutes per access specifications
Ramp-up time	10 seconds for random I/O; 20 seconds for sequential I/O
Record results	All
Number of workers	10 and 16 workers for random I/O (equal to the number of SSDs/HDDs); 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 and HDDs 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 lometer 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.

The recommendations reflect the I/O performance of 400-GB 12-Gbps SAS enterprise performance with 10x endurance, 1.9-TB 12-Gbps SAS enterprise value SSDs and 7.6-TB 6-Gbps SATA 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 higher latency.

Note: All the performance metrics presented in this document have been revalidated with BIOS microcode updates from Intel with the fix for Spectre Variant 2 issue.

SSD RAID 0 performance for 10-disk configuration

Figure 3 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 comparative performance values achieved for enterprise performance drives and enterprise value drives to help customers understand the performance trade-off when choosing a SSD type. The graph shows that the 1.9-TB enterprise value drive provides better performance at the 4-KB block size. Latency is the time taken to complete a single I/O request from the application’s viewpoint.

Figure 3. Random read 100%




Figure 4 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent random write access pattern. The numbers in the graph show that for all SSDs the IOPS rate for the 4- and 8-KB block sizes is well over 100,000, and that the 400-GB 12-Gbps SAS SSD provides better performance.

Figure 4. Random write 100%

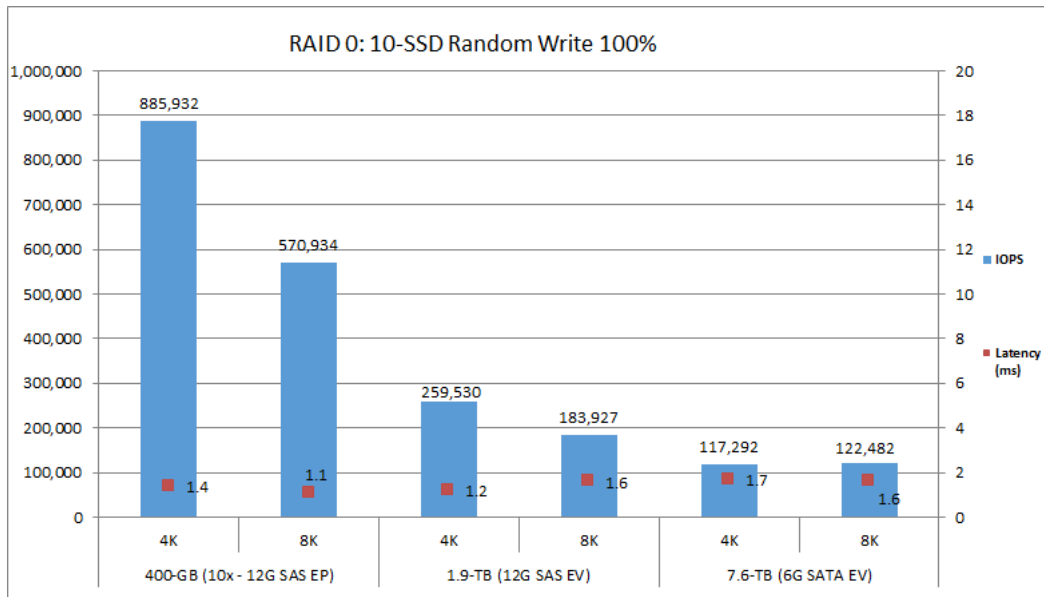


Figure 5 shows the performance of the SSDs under test for a RAID 0 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph affirm that for all SSDs the IOPS rate for the 4- and 8-KB block sizes is well over 300,000.

Figure 5. Random read:write 70%:30%




Figure 6 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent sequential read access pattern. The graph shows that the 12-Gbps SSDs (400-GB enterprise performance drives and the 1.9-TB enterprise value drives) show bandwidth more than 7000 MBps approximately.

Figure 6. Sequential read 100%





Figure 7 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent sequential write access pattern. The 12-Gbps SSDs have shown bandwidth (6800MBps) slightly less than the bandwidth seen with sequential read access (~7000 MBps).

Figure 7. Sequential write 100%



SSD RAID 5 performance for 10-disk configuration

Figure 8 shows the performance of the SSDs under test for a RAID 5 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given SSDs the IOPS rate is well over 150,000.

Figure 8. Random read:write 70%:30%





Figure 9 shows the performance of the SSDs under test for a RAID 5 configuration with a 50 percent random read and 50 percent random write access pattern. The graph shows that for the given SSDs the IOPS rate is well over 100,000.


Figure 9. Random read:write 50%:50%



SSD RAID 10 performance for 10-disk configuration

Figure 10 shows the performance of the SSDs under test for a RAID 10 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that the bandwidth achieved is over 3000 MBps with 12-Gbps SAS SSDs.

Figure 10. Sequential write 100%



SSD RAID 0 performance for 16-disk configuration

The test results prove that 16-disk configuration has higher performance than 10-disk configuration (increased IOPS and bandwidth) which is expected (given 6 more additional disks) for all access patterns.

Figure 11 shows the performance of the SSDs under test with a RAID 0 16-disk configuration with a 100 percent random read access pattern. The graph shows the comparative 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 400-GB 12-Gbps SAS drive provides better performance and can achieve 1.4 million IOPS at the 4-KB block size.

Figure 11. Random read 100%




Figure 12 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent random write access pattern. The graph shows maximum IOPS (1.2 Million) is achieved for 400GB 12-Gbps SAS SSD for 4-KB block size.

Figure 12. Random write 100%




Figure 13 shows the performance of the SSDs under test for a RAID 0 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for all SSDs the IOPS rate for the 4- and 8-KB block sizes is well over 450,000, and that the 12-Gbps SAS SSDs can achieve a rate of more than 1 million IOPS for 4-KB block size.

Figure 13. Random read:write 70%:30%




Figure 14 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent sequential read access pattern. The graph shows that with the 16-drive configuration all SSDs (both 12-Gbps SAS and 6-Gbps SATA) perform better (~7000 MBps) compared to Figure 6: Sequential Read 100%: 10 drive configuration.

Figure 14. Sequential read 100%





Figure 15 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that the bandwidth is well over 4000 MBps for all SSDs.

Figure 15. Sequential write 100%



SSD RAID 5 performance for 16-disk configuration

Figure 16 shows the performance of the SSDs under test for a RAID 5 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that the IOPS rate is well over 400,000 for 12-Gbps SAS SSDs and above 200,000 for 6-Gbps SATA SSD.

Figure 16. Random read:write 70%:30%





Figure 17 shows the performance of the SSDs under test for a RAID 5 configuration with a 50 percent random read and 50 percent random write access pattern. The graph shows that for the given SSDs the IOPS rate for the 4- and 8-KB block sizes is well over 300,000 for the 12-Gbps SSDs.


Figure 17. Random read:write 50%:50%



SSD RAID 10 performance for 16-disk configuration

Figure 18 shows the performance of the 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 3800 MBps for 12-Gbps SAS SSDs.

Figure 18. Sequential write 100%



SSD JBOD performance

The I/O performance characterization tests include JBOD pass-through controller configuration results because some software-defined storage solutions use JBOD storage. JBOD storage provides better control over individual disks. This storage type is robust and inexpensive and allows you to use all the disk space (it has no RAID requirements; all storage is 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-SAS-M5HD) for JBOD. The JBOD performance tests used 10 and 16 disks for the SSD and HDD configurations.

SSD JBOD performance for 10-disk configuration

Figure 19 shows the performance of SSDs under test for a JBOD configuration with a 100 percent random read access pattern. The graph shows the comparative performance values achieved for enterprise performance drives and enterprise value drives to help customers understand the performance trade-off when choosing an SSD type for a JBOD configuration. The graph shows that the 12-Gbps SAS SSDs at the 4-KB block size can achieve well over 1.3 million IOPS.

Figure 19. Random read 100%




Figure 20 shows the performance of SSDs under test for a JBOD configuration with a 100 percent random write access pattern. The numbers in the graph shows that for 400GB 12-Gbps SAS SSD the IOPS rate for the 4-KB block size is well over 900,000.

Figure 20. Random write 100%

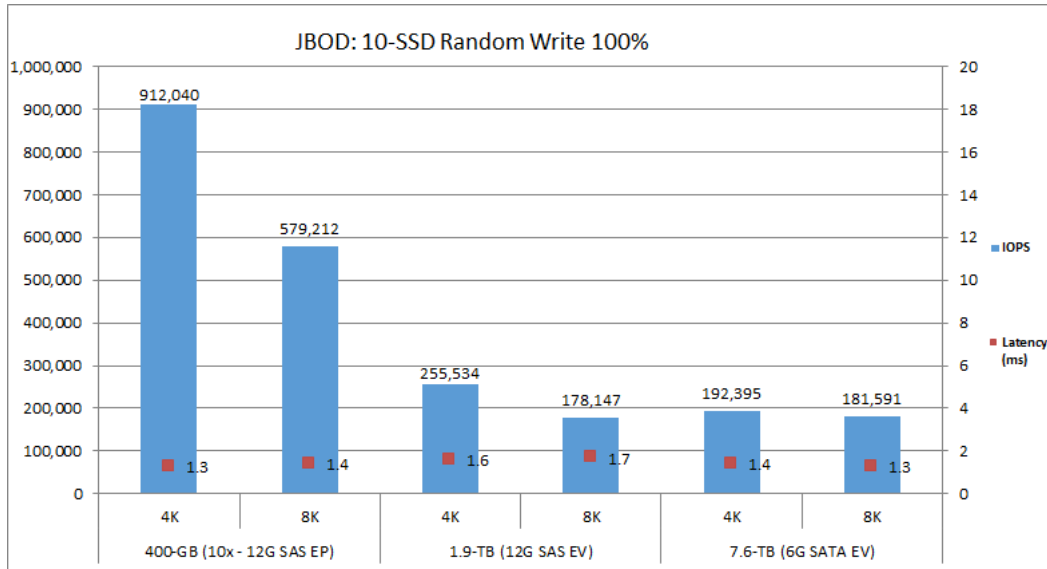


Figure 21 shows the performance of the SSDs under test for a JBOD configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given configuration the IOPS rate for the 4-KB block size is 1.2 million for the 400-GB 12-Gbps SAS SSD.

Figure 21. Random read:write 70%:30%




Figure 22 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent sequential read access pattern. The graph shows that the 12-Gbps enterprise value SAS SSDs provide performance of 7000 MBps.

Figure 22. Sequential read 100%





Figure 23 shows the performance of the SSDs under test for a RAID 0 configuration with a 100 percent sequential write access pattern. The graph shows that the 12-Gbps enterprise value SAS SSDs provide bandwidth around ~7000 MBps.

Figure 23. Sequential write 100%



SSD JBOD performance for 16-disk configuration

The test results prove that 16-disk configuration has higher performance than 10-disk configuration (increased IOPS and bandwidth) which is expected (6 more additional disks) for all access patterns.

Figure 24 shows the performance of the SSDs under test for a JBOD configuration with a 100 percent random read access pattern. The graph shows that with a 16-SSD configuration and 4-KB block size for 12-Gbps SAS SSD, the system can achieve well over 1.5 million IOPS with latency within 2 ms.

Figure 24. Random read 100%




Figure 25 shows the performance of the SSDs under test for a JBOD configuration with a 100 percent random write access pattern. The numbers in the graph affirm that for the 400-GB 12-Gbps SAS SSDs, the IOPS rate for the 4-KB block size is well over 1.2 million.

Figure 25. Random write 100%




Figure 26 shows the performance of the SSDs under test for a JBOD configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given configuration the IOPS rate for the 4-KB block size is 1.1 million for 12-Gbps SAS SSDs.

Figure 26. Random read:write 70%:30%




Figure 27 shows the performance of the SSDs under test for a JBOD configuration with a 100 percent sequential read access pattern. The graph shows that the SSDs considered for this testing can perform around 7000 MBps.

Figure 27. Sequential read 100%

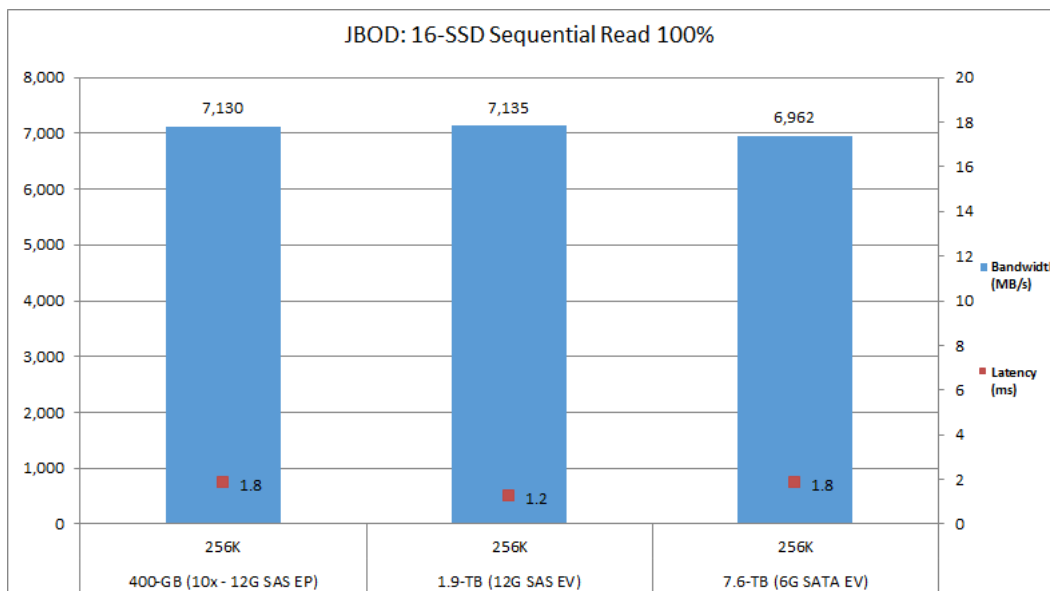
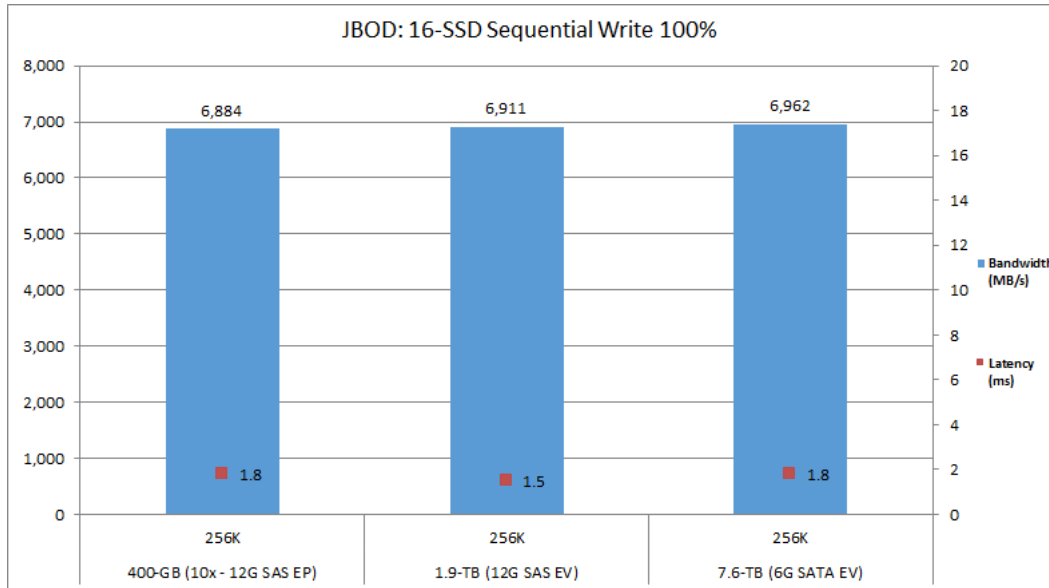


Figure 28 shows the performance of the SSDs under test for a JBOD configuration with a 100 percent sequential write access pattern. The numbers in the graph show that the bandwidth is well over 6800 MBps for the SSDs considered in this testing.

Figure 28. Sequential write 100%



HDD performance results

Figures 29 through 54 were prepared from Iometer measurement data. They illustrate the I/O performance of 600-GB and 1.2-TB SFF 12-Gbps SAS 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. These drives are capable of delivering more IOPS and bandwidth with higher queue depths and latency.

SFF HDD RAID 0 performance for 10-disk configuration

Figure 29 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random read access pattern. The numbers in the graph affirm that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 3000 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

SFF HDD RAID 0 performance for 10-disk configuration

Figure 29 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random read access pattern. The numbers in the graph affirm that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 2000 for a 10-HDD configuration, and that the IOPS rate scales as the number of drives increases for each configuration.

Figure 29. Random read 100%




Figure 30 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 8000 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 30. Random write 100%




Figure 31 shows the performance of the HDDs under test for a RAID 0 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 3500 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 31. Random read:write 70:30%




Figure 32 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent sequential read access pattern. The numbers in the graph show that for the given configuration of HDDs, the bandwidth is well over 3000 MBps for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 32. Sequential read 100%





Figure 33 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is well over 2200 MBps.

Figure 33. Sequential write 100%



SFF HDD RAID 5 performance for 10-disk configuration

Figure 34 shows the performance of the HDDs under test for a RAID 5 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for 600-GB 12-Gbps 15,000-rpm SAS HDDs the IOPS rate for the 4-KB block size is well over 2000.

Figure 34. Random read:write 70:30%





Figure 35 shows the performance of the HDDs under test for a RAID 5 configuration with a 50 percent random read and 50 percent random write access pattern. The numbers in the graph show that the IOPS rate for the 4- and 8-KB block sizes is well over 1000 for given configuration.


Figure 35. Random read:write 50:50%



SFF HDD RAID 10 performance for 10-disk configuration

Figure 36 shows the performance of the HDDs under test for a RAID 10 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that the bandwidth is around ~1500 MBps for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 36. Sequential write 100%



SFF HDD RAID 0 performance for 16-disk configuration

These tests are done for RAID 0: 16-disk configuration, hence the IOPS/Bandwidth show higher numbers compared to RAID 0: 10 Disk configuration test results (Figure 29 - Figure 33)

Figure 37 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random read access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 3400.

Figure 37. Random read 100%




Figure 38 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 10,000 for both HDDs.

Figure 38. Random write 100%




Figure 39 shows the performance of the HDDs under test for a RAID 0 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 6000 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 39. Random read:write 70:30%




Figure 40 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent sequential read access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is around ~5000 MBps for 600-GB 12-Gbps 15,000-rpm SAS HDDs, and that the performance scales as the number of HDDs increases in the given configuration.

Figure 40. Sequential read 100%





Figure 41 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is well over 4800 MBps for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 41. Sequential write 100%



SFF HDD RAID 5 performance for 16-disk configuration

Figure 42 shows the performance of HDDs under test for a RAID 5 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that the IOPS rate for the 4- and 8-KB block sizes is well over 3000 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 42. Random read:write 70:30%





Figure 43 shows the performance of the HDDs under test for a RAID 5 configuration with a 50 percent random read and 50 percent random write access pattern. The numbers in the graph show that the IOPS rate for the 4- and 8-KB block sizes is well over 2500 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.


Figure 43. Random read:write 50:50%



SFF HDD RAID 10 performance for 16-disk configuration

Figure 44 shows the performance of the HDDs under test for a RAID 10 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that the bandwidth is well over 2300 MBps for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 44. Sequential write 100%



HDD JBOD performance

SFF HDD JBOD performance for 10-disk configuration

Figure 45 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random read access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 3000 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 45. Random read 100%




Figure 46 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 6000 for a 10-HDD configuration.

Figure 46. Random write 100%




Figure 47 shows the performance of the HDDs under test for a RAID 0 configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 3600 for 600-GB 12-Gbps 15,000-rpm SAS HDDs.

Figure 47. Random read:write 70:30%




Figure 48 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent sequential read access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is well over 2000 MBps.

Figure 48. Sequential read 100%





Figure 49 shows the performance of the HDDs under test for a RAID 0 configuration with a 100 percent sequential write access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is well over 2200 MBps.

Figure 49. Sequential write 100%



SFF HDD JBOD performance for 16-disk configuration

Figure 50 shows the performance of the HDDs under test for a JBOD configuration with a 100 percent random read access pattern. The numbers in the graph show that for the given configuration of HDDs, the IOPS rate for the 4- and 8-KB block sizes is well over 3400.

Figure 50. Random read 100%




Figure 51 shows the performance of the HDDs under test for a JBOD configuration with a 100 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4-KB block size is 11,200.

Figure 51. Random write 100%




Figure 52 shows the performance of the HDDs under test for a JBOD configuration with a 70 percent random read and 30 percent random write access pattern. The numbers in the graph show that for the given configuration of HDDs the IOPS rate for the 4- and 8-KB block sizes is well over 3300.

Figure 52. Random read:write 70:30%




Figure 53 shows the performance of the HDDs under test for a JBOD configuration with a 100 percent sequential read access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is well over 3500 MBps.

Figure 53. Sequential read 100%





Figure 54 shows the performance of the HDDs under test for a JBOD configuration with a 100 percent sequential write access pattern. The numbers in the graph show that for the given configuration of HDDs the bandwidth is well over 3500 MBps.

Figure 54. Sequential write 100%



For more information

For additional information, see <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-739279.html>.

Appendix: Test environment

Table 9 lists the details of the server under test.

Table 9. Server properties

Name	Value
Product names	Cisco UCS C240 M5SX
CPUs	CPU: 2 x 2.40-GHz Intel Xeon Gold 6148
Number of cores	40
Number of threads	80
Total memory	192 GB
Memory DIMMs (16)	16 GB x 12 DIMM channels
Memory speed	2666 MHz
Network controller	Cisco LOM X550-T2; 2 x 10-Gbps interfaces
VIC adapter	Cisco UCS VIC 1387 mLOM 10-Gbps SFP+
RAID controllers	<ul style="list-style-type: none"> • Cisco 12-Gbps modular RAID controller with 4-GB flash-backed write cache (FBWC; UCSC-RAID-M5HD) • Cisco 12-Gbps modular SAS HBA (UCSC-SAS-M5HD)
SFF SSDs	<ul style="list-style-type: none"> • 400-GB 2.5-inch enterprise value 12-Gbps SAS (UCS-SD400G12TX-EP) • 1.9-TB 2.5-inch enterprise value 12-Gbps SAS (UCS-SD19TB121X-EV) • 7.6-TB 2.5-inch enterprise value 6-Gbps SATA (UCS-SD76TM1X-EV)
SFF HDDs	<ul style="list-style-type: none"> • 1.2-TB 12-Gbps SAS 10,000-rpm SFF HDD (UCS-HD12TB10K12N) • 600-GB 12-Gbps SAS 15,000-rpm SFF HDD (UCS-HD600G15K12N)

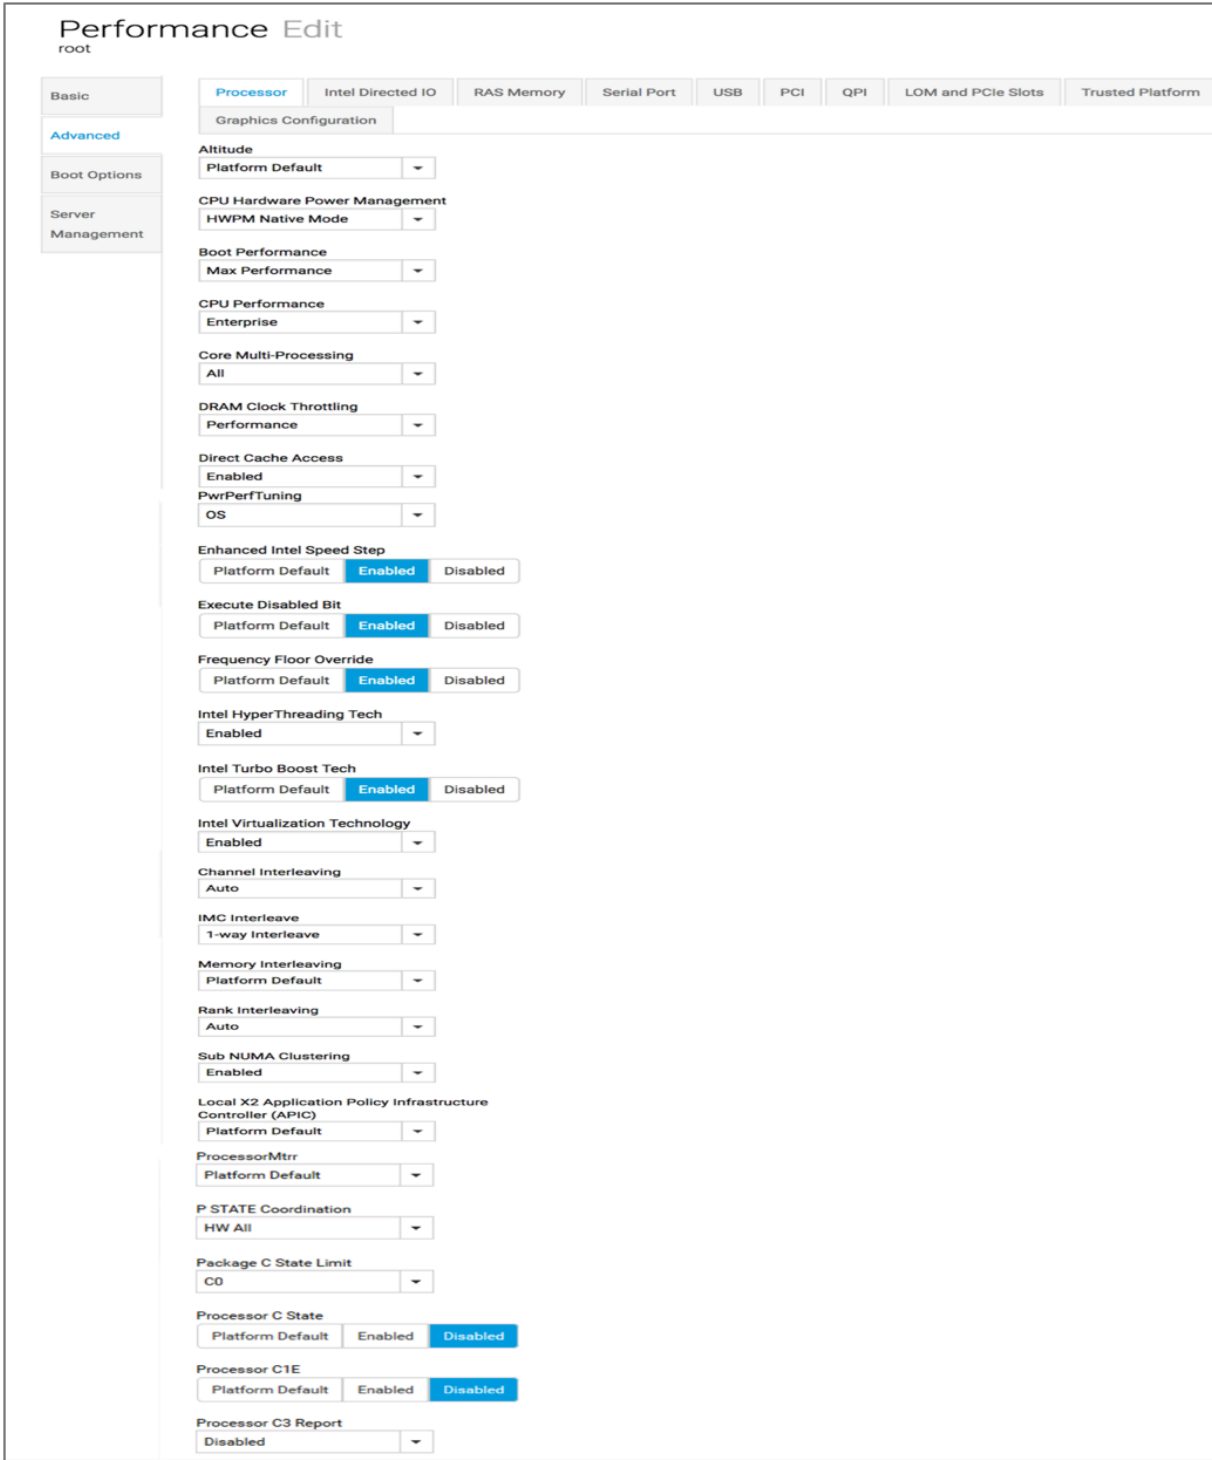
Table 10 lists the recommended server BIOS settings.

Table 10. BIOS settings for standalone rack server

Name	Value
BIOS version	Release C240M5.3.1.3d
CIMC version	3.1(3a)
Cores Enabled	All
Hyper-Threading (All)	Enable
Execute disable bit	Enable
Intel (R) (VT)	Enable
Hardware Prefetcher	Enable
Adjacent-cache-line Prefetcher	Enable
DCU Streamer	Enable
DCU IP Prefetcher	Enable
Extended APIC	Disable
NUMA	Enable
SNC	Enable
IMC Interleaving	1-way-Interleave
Mirror mode	Disable
Patrol Scrub	Disable

Name	Value
Intel VT for Directed I/O (VT-d)	Enable
Interrupt Remapping	Enable
PassThrough DMA	Disable
ATS	Enable
Posted Interrupt	Enable
Coherency Support	Disable
SpeedStep (Pstates)	Enable
EIST PSD Function	HW_ALL
Boot Performance Mode	Max Performance
Energy Efficient Turbo	Disable
Turbo mode	Enable
Hardware P-States	Native mode
EPP profile	Balanced Performance
Autonomous Core	Disable
CPU C6 Report	Disable
Enhanced Halt State (C1E)	Disable
OS ACPI Cx	ACPI C2
Package C State	C0/C1 state
Power Performance Tuning	OS controls EPB
PECI PCB EPB	OS controls EPB
Workload configuration	Balanced

Figure 55. BIOS setting for UCS-managed rack server



Performance Edit
root

Basic | **Processor** | Intel Directed IO | RAS Memory | Serial Port | USB | PCI | QPI | LOM and PCIe Slots | Trusted Platform

Graphics Configuration

Altitude
Platform Default

CPU Hardware Power Management
HWPM Native Mode

Boot Performance
Max Performance

CPU Performance
Enterprise

Core Multi-Processing
All

DRAM Clock Throttling
Performance

Direct Cache Access
Enabled

PwrPerfTuning
OS

Enhanced Intel Speed Step
Platform Default **Enabled** Disabled

Execute Disabled Bit
Platform Default **Enabled** Disabled

Frequency Floor Override
Platform Default **Enabled** Disabled

Intel HyperThreading Tech
Enabled

Intel Turbo Boost Tech
Platform Default **Enabled** Disabled

Intel Virtualization Technology
Enabled

Channel Interleaving
Auto

IMC Interleave
1-way interleave

Memory Interleaving
Platform Default

Rank Interleaving
Auto

Sub NUMA Clustering
Enabled

Local X2 Application Policy Infrastructure Controller (APIC)
Platform Default

ProcessorMtrr
Platform Default

P STATE Coordination
HW All

Package C State Limit
C0

Processor C State
Platform Default **Enabled** Disabled

Processor C1E
Platform Default **Enabled** Disabled

Processor C3 Report
Disabled

Processor C6 Report
 Platform Default Enabled Disabled

Processor C7 Report
 Disabled

Processor CMCI
 Platform Default Enabled Disabled

Power Technology
 Performance

Energy Performance
 Performance

Adjacent Cache Line Prefetcher
 Platform Default Enabled Disabled

Data Cache Unit (DCU) IP Prefetcher
 Platform Default Enabled Disabled

Data Cache Unit (DCU) IP Prefetcher
 Platform Default Enabled Disabled

Hardware Prefetcher
 Platform Default Enabled Disabled

UPI Prefetcher
 Platform Default Enabled Disabled

LLC Prefetcher
 Platform Default Enabled Disabled

XPT Prefetcher
 Platform Default Enabled Disabled

Demand Scrub
 Platform Default Enabled Disabled

Patrol Scrub
 Platform Default Enabled Disabled

Workload Configuration
 Balanced

Basic Processor Intel Directed IO **RAS Memory** Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform

Advanced Graphics Configuration

Boot Options

Server Management

DDR3 Voltage Selection
 Platform Default

DRAM Refresh Rate
 Platform Default

LvDDRMode
 Performance Mode

Mirroring Mode
 Platform Default

NUMA optimized
 Platform Default Enabled Disabled

SelectMemory RAS Configuration
 Maximum Performance

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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)