Cisco UCS C220 M5 and C240 M5 Rack Server NVMe I/O Characterization

October 2018

Contents

Executive summary

This document summarizes the Non-Volatile Memory Express (NVMe) I/O performance characteristics of the Cisco UCS® C220 M5 and C240 M5 Rack Servers using NVMe solid-state disks (SSDs) and NVMe half-height half-length (HHHL) cards. Performance comparisons of various NVMe SSDs and NVMe HHHL cards are presented. The goal of this document is to help customers make well-informed decisions so that they can choose the right NVMe drives to meet their I/O workload needs.

Performance data was obtained using the Flexible I/O (Fio) measurement tool, with analysis based on the number of I/O operations per second (IOPS) for random I/O workloads, and megabytes per second (MBps) of throughput for sequential I/O workloads.

Introduction

The Cisco Unified Computing System™ (Cisco UCS) is a next-generation data center platform that unites computing, networking, storage, and virtualization resources in a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility.

As storage moves closer to the server, new opportunities for data center efficiency are arising. When applications that need greater storage performance achieve high availability by using cluster-capable file systems and other means for replicating data, major efficiencies can be gained.

Ultra-low-latency NVMe storage fully integrated into the Cisco UCS architecture enables servers to provide increased storage reliability and performance compared to spinning media. Organizations also gain the benefits of lower total cost of acquisition and lower TCO through reduced data center power and cooling needs, as well as lower cost per IOPS and lower wattage requirements per IOPS. Bringing storage inside the server on a high-performance NVMe tier can also reduce application licensing costs, making local flash storage a powerful solution for delivering more capabilities on a smaller budget. And all these benefits are more fully optimized on Cisco UCS than on any other server platform.

Cisco UCS implements local storage differently for a uniquely powerful experience. The Cisco UCS platform uses an advanced cooling methodology and zero-oversubscription CPU mapping to provide the highest levels of efficiency as well as best-in-class, consistent performance. Teams can manage hundreds of devices as easily as one with the Cisco® Integrated Management Controller (IMC) or Cisco UCS Manager. Customers can also choose the amount of storage necessary to meet their application needs: from 400 GB all the way up to 46 TB (for example, for a 2-rack-unit [2RU] server)

NVMe storage solutions

NVMe storage solutions offer the following main benefits:

- Reduced TCO: NVMe storage can be used to eliminate the need for SANs and network-attached storage (NAS) or to augment existing shared-array infrastructure. With significant performance improvements available in both cases, Cisco customers can reduce the amount of physical infrastructure they need to deploy, increase the number of virtual machines they can place on a single physical server, and improve overall system efficiency. These improvements provide savings in capital expenditures (CapEx) and operating expenses (OpEx), including reduced application licensing fees and savings related to space, cooling, and energy use.
- Strategic partnerships: Cisco tests a broad set of NVMe storage technologies and focuses on major vendors. With each partnership, devices are built exclusively in conjunction with Cisco engineering, so customers have the flexibility of a variety of endurance and capacity levels and the most relevant form factors, as well as the powerful management features and robust quality benefits that are unique to Cisco.

Overview

Cisco UCS NVMe storage offers the following main advantages:

- Manageability
  - Complete Cisco UCS Manager inventory, service-profile mapping, and firmware updates are supported for one or many Cisco NVMe storage devices.
  - Significantly reduce complexity compared to most competing solutions, which require manual command-line interface (CLI) processes for each PCIe-connected device.
- Performance
  - Get what you pay for with full CPU-to-PCIe lane connectivity to each storage device (zero oversubscription). For example, a Cisco server with six NVMe devices has 6 x 8 = 48 lanes of Generation 3 PCIe (PCIe 3.0) connectivity directly to the CPU. Not all server platforms provide this level of integration and thus offer lower performance than Cisco UCS rack servers.
  - Gain peace of mind with endurance and performance-state tracking, displayed in Cisco UCS Manager or Cisco IMC.
  - Power and cooling algorithms based on device system management bus (SMBus) integration provide exceptional power consumption efficiency and consistent performance in any data center, with server inlet temperatures of less than 35°C.
- Flexibility
  - Gain best-in-class capacity flexibility with the capability to choose from as little as 400 GB to as much as 77 TB of capacity in a 2RU server, with multiple NVMe device endurance level options, to meet the needs of any application.
  - Gain block storage that integrates transparently with Cisco UCS servers to immediately improve performance and relieve I/O bottlenecks.

Scope of this document

For the NVMe I/O performance characterization tests, performance was evaluated using NVMe SSDs and NVMe HHHL cards for random and sequential access patterns for Cisco UCS C220 M5SX and C240 M5SX servers. The Cisco C220 M5SX server supports two NVMe SSDs connected directly by PCIe x4 lanes to the CPU. The Cisco UCS C240 M5SX supports four NVMe SSDs, in which two SSDs are front facing and two SSDs are rear facing, connected directly by PCIe x4 lanes to the CPU.

The NVMe performance tests used the following configurations:

- For the SSD form factor:
  - 2 disks
  - 4 disks
- For the HHHL card form factor:
  - 4 HHHL cards (on the Cisco UCS C240 M5SX)

Hardware RAID options are not available for NVMe drives, so they are not considered in this document.

Solution components

The tested solution used Cisco UCS C220 M5SX and C240 M5SX Rack Servers for 2-disk, 4-disk, and 4-HHHL cards configurations.

Cisco UCS C220 M5 and C240 M5 Rack Server overview

See the following data sheets for more information about the C220 M5 and C240 M5 servers:

- https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-739281.html
- https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-739279.html

Server specifications

The server specifications are as follows:

- Cisco UCS C220 M5SX and C240 M5SX Rack Servers
- CPU: 2 x 2.60-GHz Intel® Xeon® Gold 6126 processors
- Memory: 24 x 16-GB (384-GB) DDR4
- Cisco UCS Virtual Interface Card (VIC) 1387 modular LAN-on-motherboard (mLOM) 10-Gbps Enhanced Small Form-Factor Pluggable (SFP+)

Cisco UCS C220 M5 server models

The Cisco UCS C220 M5 server can be configured in three different models to match specific customer environments. The Cisco UCS C220 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 C220 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 C220 M5 servers are broadly categorized into small-for-factor (SFF) and large-form-factor (LFF) models as follows:

- Cisco UCS C220 M5 Rack Server (SFF model)
  - UCSC-C220-M5SX
    - 10 SFF front-facing SAS and SATA hard-disk drives (HDDs) or SAS and SATA SSDs
    - Optionally, up to 2 front-facing SFF NVMe PCIe SSDs (replacing the SAS and SATA drives)
  - UCSC-C220-M5SN
    - Up to 10 front-facing SFF NVMe PCIe SSDs only (replacing the SAS and SATA drives)
- Cisco UCS-C220-M5L Rack Server (LFF model)
  - 4 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)

The Cisco UCS C220 M5SX SFF server (Figure 1) extends the capabilities of the Cisco UCS portfolio in a 1RU 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 2 PCIe 3.0 slots, and up to 10 internal SFF drives. The C220 M5SX server also includes one dedicated internal slot for a 12-Gbps SAS storage controller card.

Figure 1. Cisco UCS C220 M5SX

The Cisco UCS C220 M5SN SFF server (Figure 2) extends the capabilities of the Cisco UCS portfolio in a 1RU 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 2 PCIe 3.0 slots, and up to 10 internal SFF NVMe drives. The C220 M5SN server is 10-drive NVMe-only system.

Figure 2. Cisco UCS C220 M5SN

The Cisco UCS C220 M5 LFF server (Figure 3) extends the capabilities of the Cisco UCS portfolio in a 1RU 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 2 PCIe 3.0 slots, and up to 4 front-facing internal LFF drives. The C220 M5 LFF server also includes one dedicated internal slot for a 12-Gbps SAS storage controller card. Figure 3 shows the Cisco UCS C220 M5L model.

Figure 3. Cisco UCS C220 M5L

For information about configuring a specific model of C220 M5 server, please refer to the appropriate specification sheet:

- SFF models (either HDD or SSD): https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m5-sff-specsheet.pdf
- LFF model (HDD): https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m5-lff-specsheet.pdf

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

Table 1. Part numbers for Cisco UCS C220 M5 high-density SFF and LFF base rack server models

Part number	Description
UCSC-C220-M5SX	Cisco UCS C220 M5: 10 SFF front drives
UCSC-C220-M5SN	Cisco UCS C220 M5: 10 SFF NVMe drives
UCSC-C220-M5L	Cisco UCS C220 M5: 4 LFF front drives

The performance testing described in this document uses the Cisco UCS C220 M5SX server, which supports 10 SFF SSDs and HDDs with SAS expanders.

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 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 2 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 and 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 4) 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 4. Cisco UCS C240 M5SX

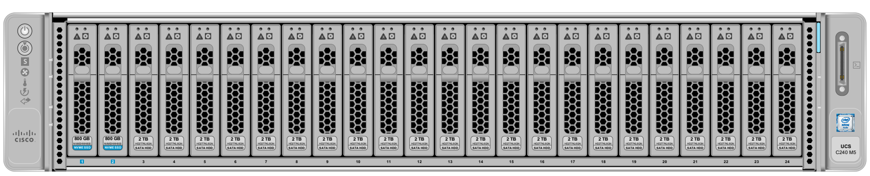

The Cisco UCS C240 M5SN server (Figure 5) extends the capabilities of the Cisco UCS portfolio in a 2RU form factor with the addition of Intel Xeon Scalable processors, 24 DIMM slots for 2666-MHz DDR4 DIMMs and capacity points of up to 128 GB, up to 6 PCIe 3.0 slots, and up to 10 SFF NVMe drives (8 front and 2 rear) and 16 front-facing internal SFF drives. The C240 M5SN server also includes one dedicated internal slot for a 12-Gbps SAS storage controller card.

Figure 5. Cisco UCS C240 M5SN

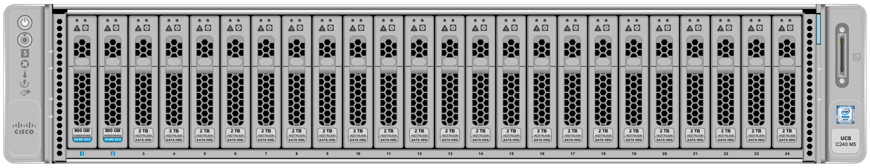

The Cisco UCS C240 M5 LFF server (Figure 6) extends the capabilities of the Cisco UCS portfolio in a 2RU form factor with the addition of Intel Xeon Scalable processors, 24 DIMM slots for 2666-MHz DDR4 DIMMs and capacity points of up to 128 GB, up to 6 PCIe 3.0 slots, and up to 12 front-facing internal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

Figure 6. Cisco UCS C240 M5L

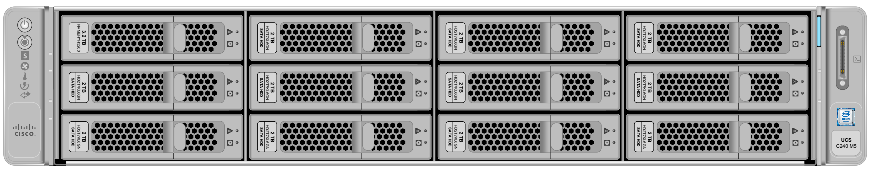

For information about configuring a specific model of C240 M5 server, please refer to the appropriate specification sheet:

- 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 2 provides part numbers for the SFF and LFF server models.

Table 2. 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
UCSC-C240-M5SN	Cisco UCS C240 M5: 10 SFF NVMe (8 front and 2 rear) and 16 (front) SAS and SATA drives
UCSC-C240-M5S	Cisco UCS C240 M5: 8 SFF front drives with option for 2 SFF rear drives
UCSC-C240-M5L	Cisco UCS C240 M5: 12 LFF front drives with option for 2 SFF rear drives

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

NVMe SFF SSDs and HHHL cards used in the tests

The following NVMe SFF SSDs and HHHL cards were selected for the tests:

- 800-GB SN200 NVMe High-Performance High-Endurance SSD (UCSC-NVMEHW-H800)
- 1.6-TB SN200 NVMe High-Performance High-Endurance SSD (UCSC-NVMEHW-H1600)
- 7.7-TB SN200 NVMe High-Performance High-Endurance SSD (UCSC-NVMEHW-H7680)
- 1-TB P4500 NVMe High-Performance Value-Endurance SSD (UCSC-NVMEHW-I1000)
- 1.6-TB SN260 NVMe Extreme-Performance High-Endurance HHHL card (UCSC-F-H16003)

These models were used because they have high-performance and high-endurance characteristics but differ by capacity, vendor, and form factor.

Workload characterization

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

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

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

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

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

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

Test configuration

The test configuration was as follows:

- 2 NVMe SFF SSDs on C220 M5SX server
- 4 NVMe SFF SSDs on C240 M5SX server (2 front-facing SFF NVMe PCIe SSDs and 2 rear-facing SFF NVMe PCIe SSDs)
- 4 NVMe HHHL PCIe cards on C240 M5SX server
- Random workload tests performed using 4- and 8-KB block sizes for all NVMe SSDs and NVMe HHHL cards
- Sequential workload tests performed using a 256-KB and 1-MB block size for all NVMe SSDs and NVMe HHHL cards

Table 5 lists the recommended Fio settings.

Table 5. Recommended Fio settings

Name	Value
Fio	Release 3.6
Run time	60 minutes per access specifications
Ramp-up time	10 seconds for random I/O and sequential I/O
Record results	All
Number of workers	16 workers for random I/O; 1 worker for sequential I/O
I/O engine	Libaio
Operating system	Red Hat Enterprise Linux (RHEL) 7.5 with latest kernel updates and patches

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

NVMe SSD performance results

Performance data was obtained using the Fio 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 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 rate and bandwidth achieved with the NVMe SSDs and NVMe HHHL cards within the acceptable response time (latency) of 2 milliseconds (ms). However, the NVMe drives under test are capable of a much higher IOPS rate and much greater bandwidth with greater latency.

Note: All the performance metrics presented in this document have been tested with the latest kernel updates for the RHEL 7.5 operating system and BIOS microcode updates from Intel with the solution for the Spectre and Meltdown issue.

NVMe SSD performance for two- and four-disk configurations

The NVMe two-disk SFF SSD configuration was tested on the Cisco UCS C220 M5SX server, which is front facing (slots 1 and 2) and managed by the x4 PCIe lane.

The NVMe four-disk SFF SSD configuration was tested on the Cisco UCS C240 M5SX server, in which two disks are front-facing (slots 1 and 2) and two disks are rear facing and both are managed by x4 PCIe lanes.

Figure 7 shows the performance of the NVMe SSDs under test with a two-disk configuration for a 100 percent random read access pattern. The graph shows the comparative performance values achieved to help customers understand the performance trade-off when choosing an NVMe SSD type for front-facing SFF slots. SN200 NVMe SSD performance is more than 1.6 million IOPS with a 4-KB block size, and almost the same for 800-GB and 1.6- and 7.7-TB SSDs.

Figure 7. NVMe two-SSD configuration: Random read 100 percent

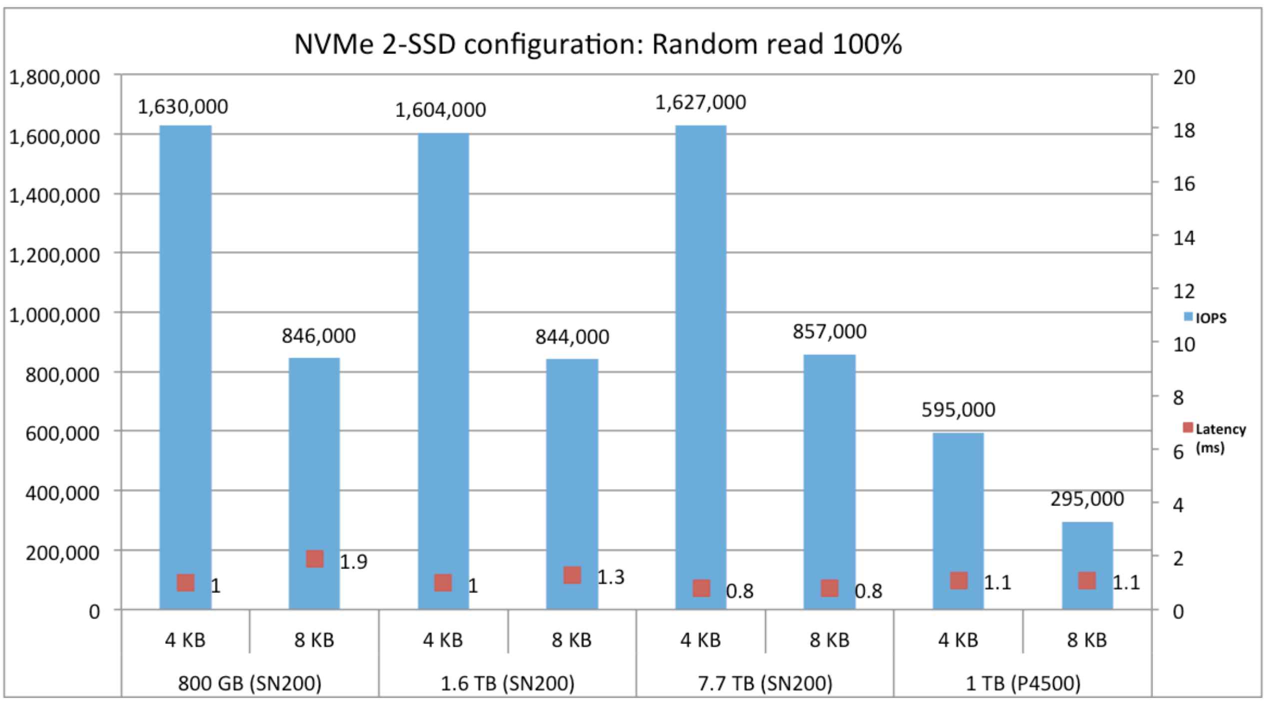

Figure 8 shows the performance of the NVMe SSDs under test with a two-disk configuration for a 100 percent random write access pattern. The graph shows that performance of the 1.6-TB SN200 NVMe SSD is higher than that for other NVMe SSDs.

Figure 8. NVMe two-SSD configuration: Random write 100 percent

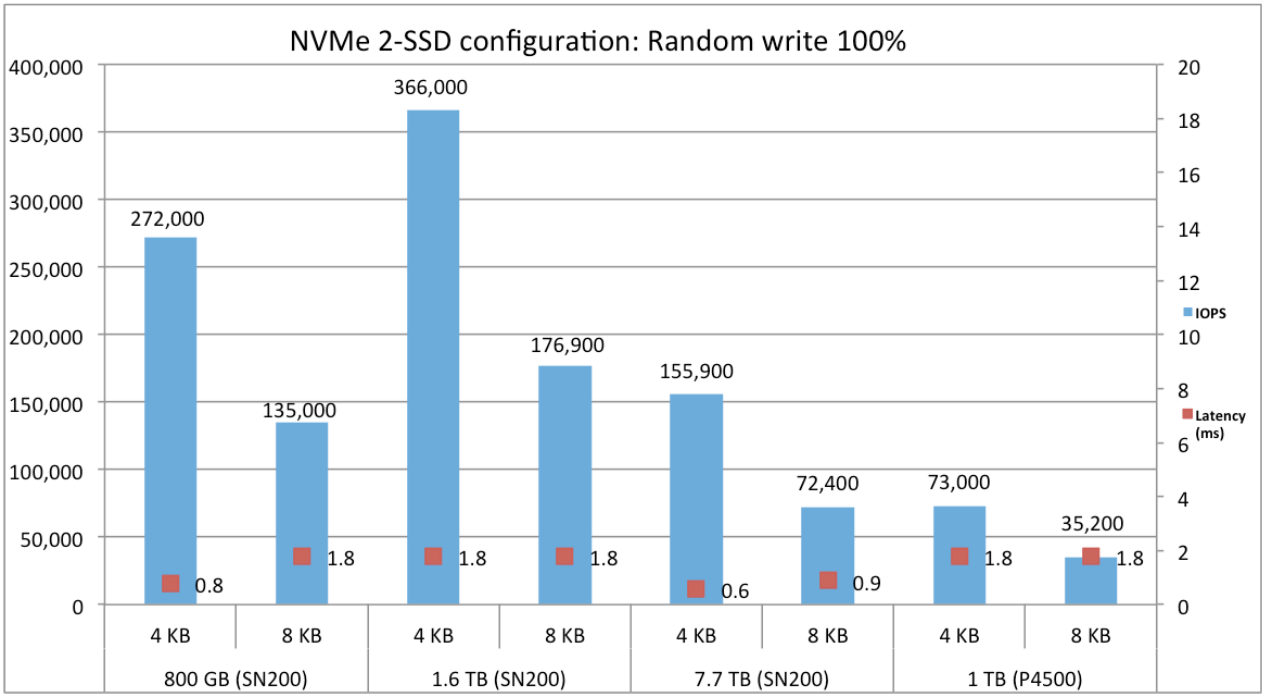

Figure 9 shows the performance of the NVMe SSDs under test with a two-disk configuration for a random read:write 70:30 percent access pattern. The performance of the 1.6-TB SN200 NVMe SSD is higher than that for other NVMe SSDs.

Figure 9. NVMe two-SSD configuration: Random read:write 70:30 percent

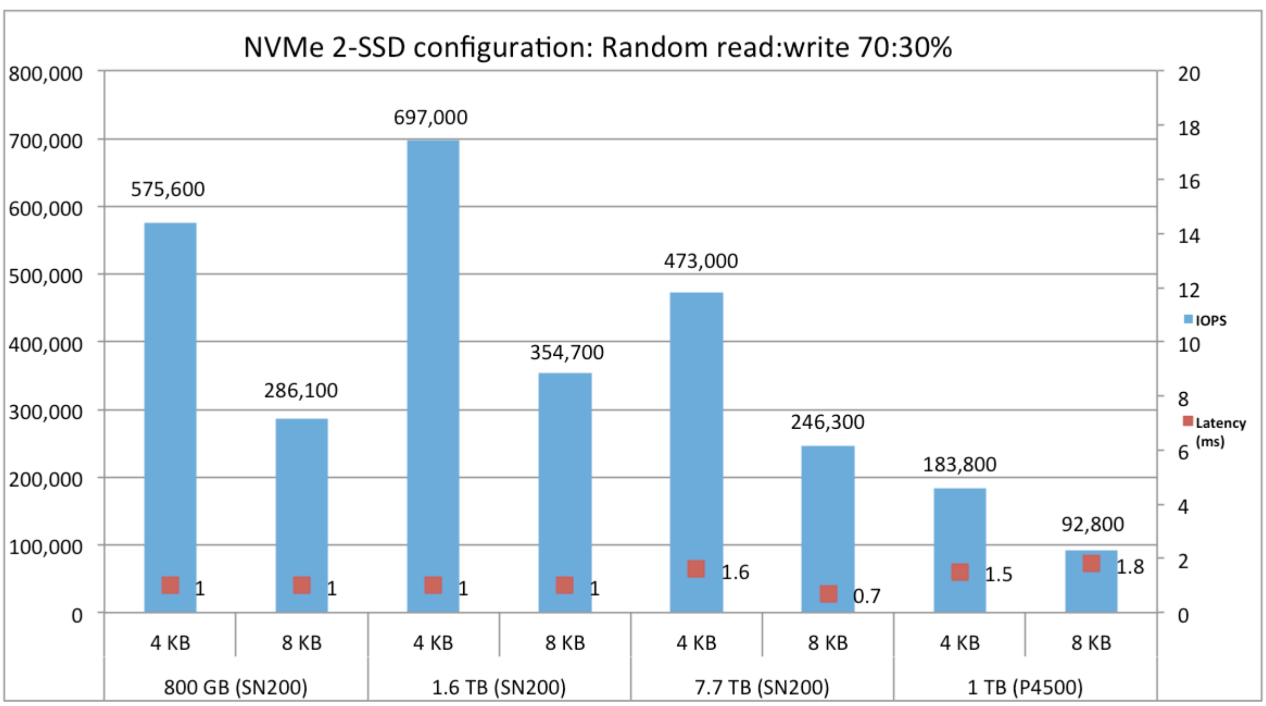

Figure 10 shows the performance of the NVMe SSDs under test with a two-disk configuration for a 100 percent sequential read access pattern. The graph shows that performance is more than 7000 MBps for SN200 NVMe SSDs with a 256-KB block size.

Figure 10. NVMe two-SSD configuration: Sequential read 100 percent

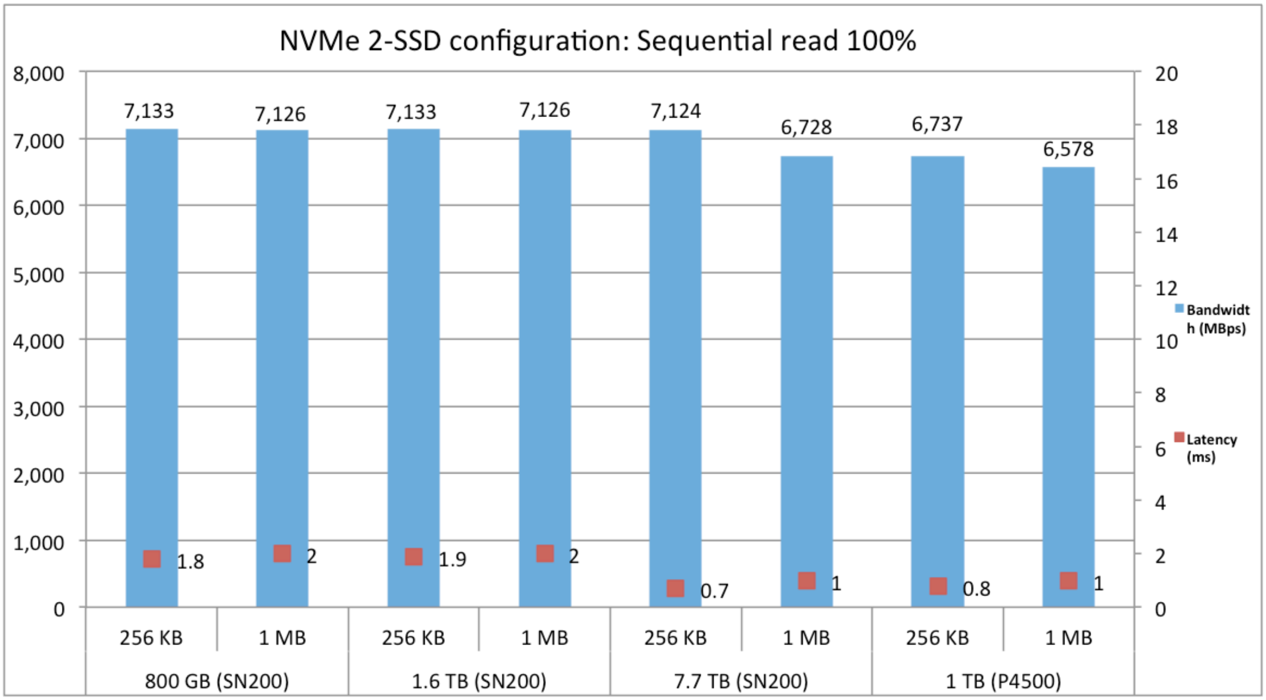

Figure 11 shows the performance of the NVMe SSDs under test with a two-disk configuration for a 100 percent sequential write access pattern. The performance for 1.6- and 7.7-TB SN200 NVMe SSDs is more than 4400 MBps for the 256-KB and 1-MB block sizes.

Figure 11. NVMe two-SSD configuration: Sequential write 100 percent

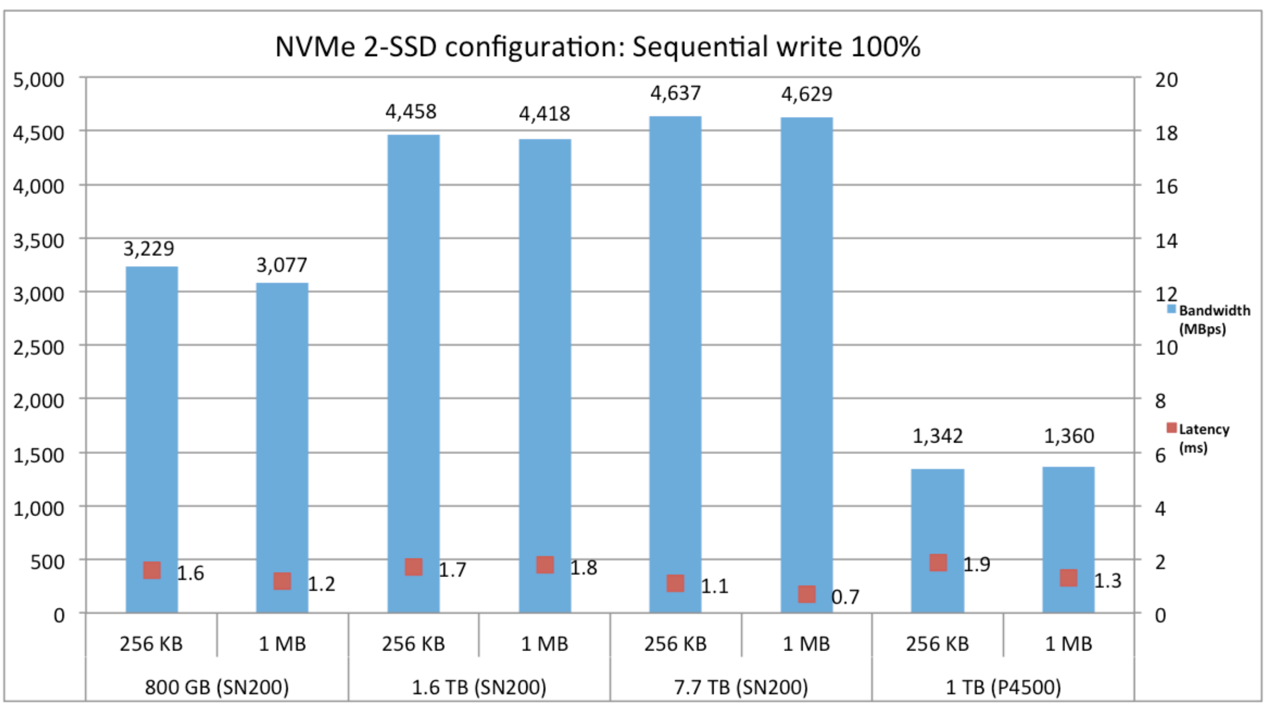

Figure 12 shows the performance of the NVMe SSDs under test with a four-disk configuration for a 100 percent random read access pattern. The graph shows the comparative performance values achieved to help customers understand the performance trade-off when choosing an NVMe SSD type for front-facing and rear-facing SFF NVMe PCIe SSDs. SN200 NVMe SSD performance for 800-GB and 1.6- and 7.7-TB SSDs is more than 3.2 million IOPS for a 4-KB block size.

Figure 12. NVMe four-SSD configuration: Random read 100 percent

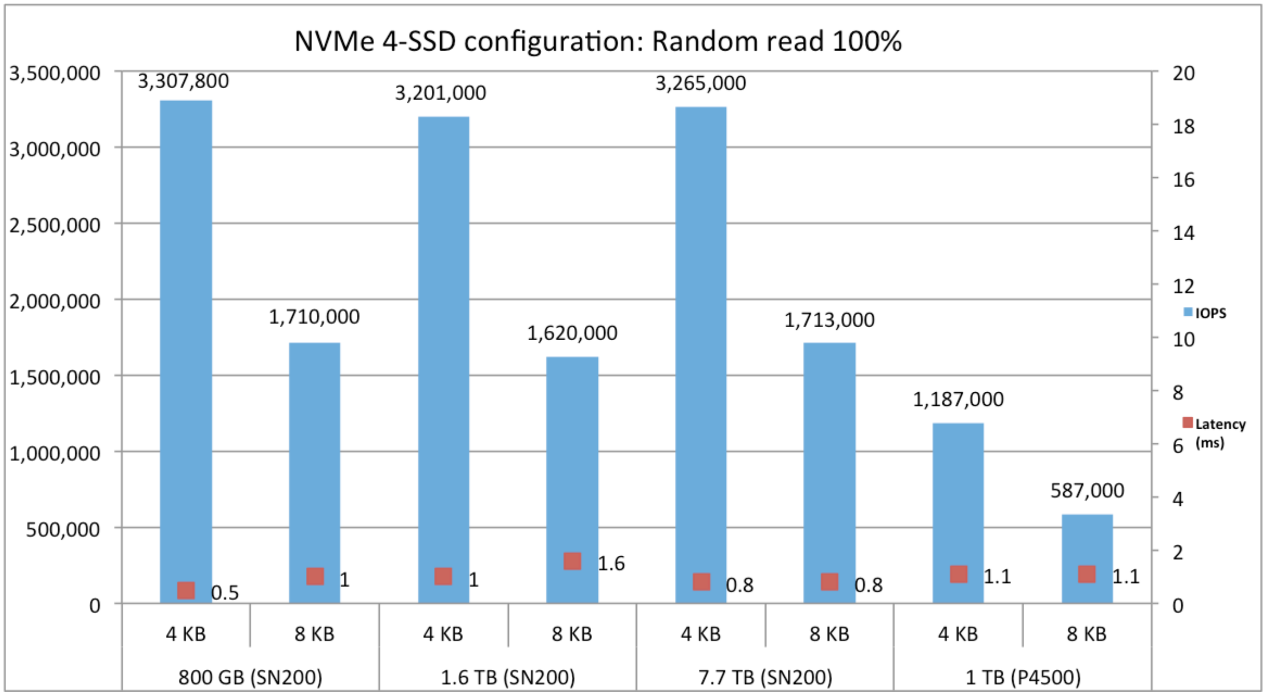

Figure 13 shows the performance of the NVMe SSDs under test with a four-disk configuration for a 100 percent random write access pattern. The graph shows that the performance values achieved for front-facing and rear-facing SFF NVMe PCIe SSDs for 1.6-TB SN200 NVMe SSDs is 700,000 IOPS for a 4-KB block size.

Figure 13. NVMe four-SSD configuration: Random write 100 percent

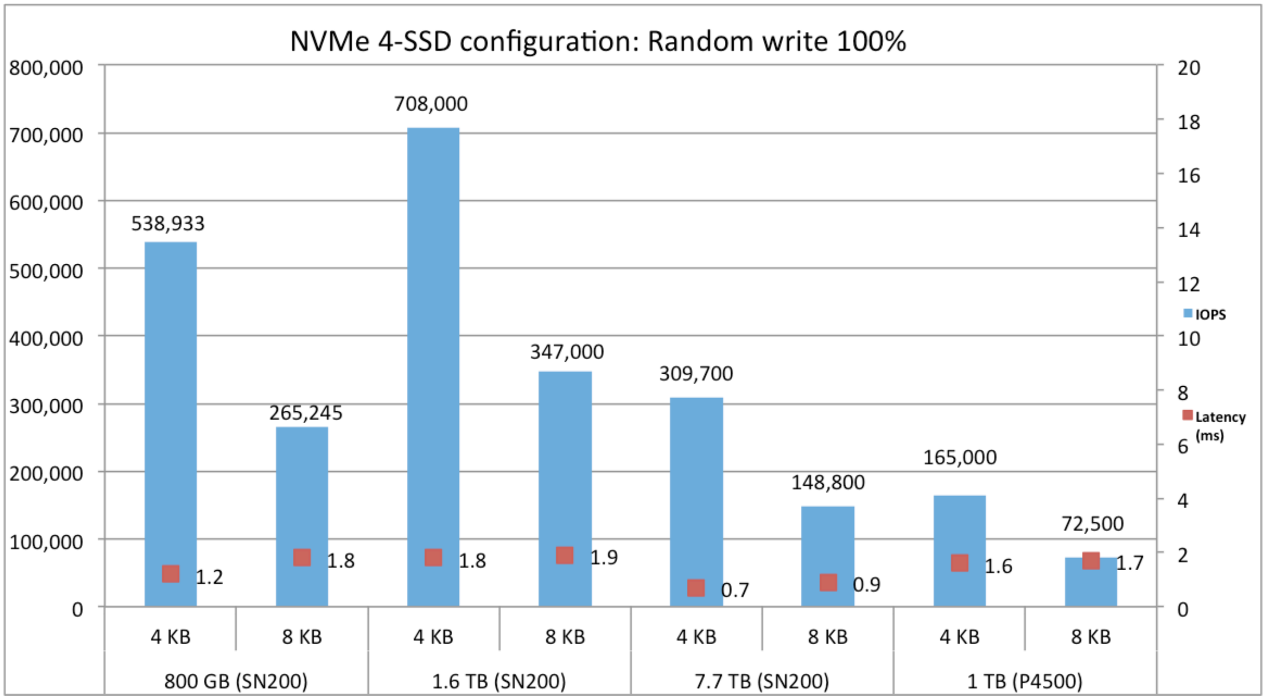

Figure 14 shows the performance of the NVMe SSDs under test with a four-disk configuration for a random read:write 70:30 percent access pattern. The performance of the 1.6-TB SN200 NVMe SSD is more than 1.3 million IOPS with a 4-KB block size and higher than that for other NVMe SSDs.

Figure 14. NVMe four-SSD configuration: Random read:write 70:30 percent

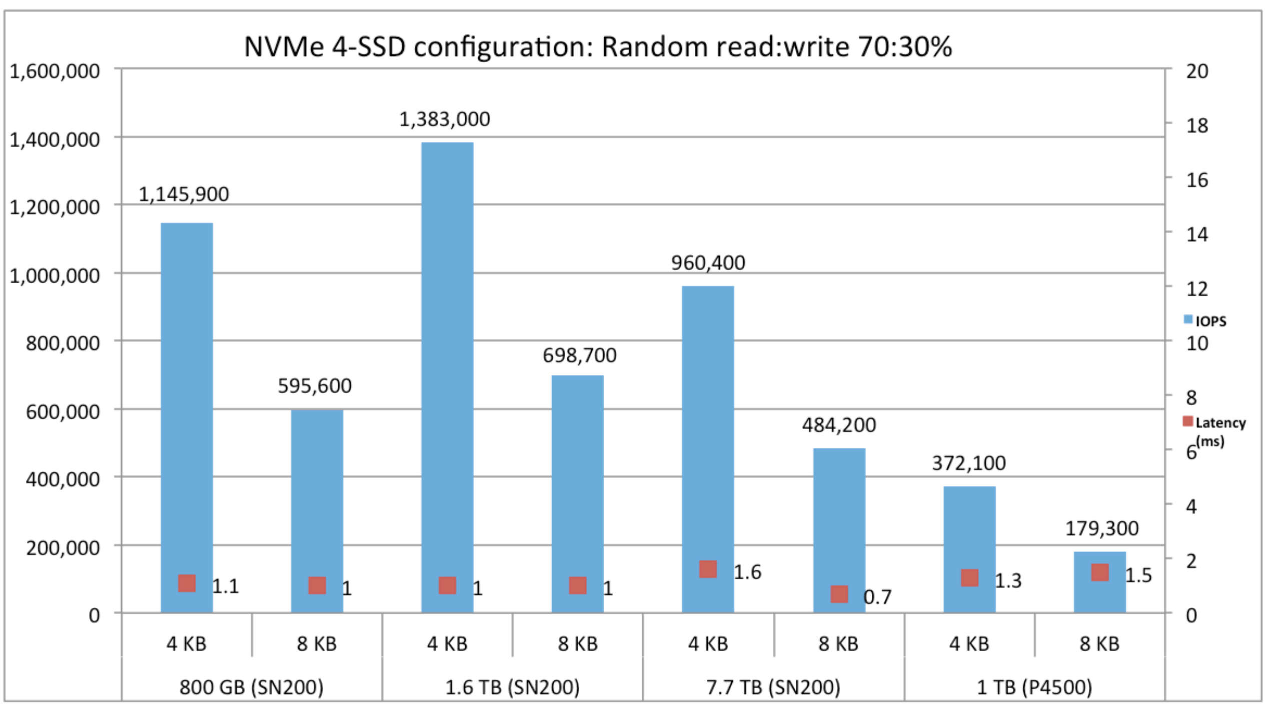

Figure 15 shows the performance of the NVMe SSDs under test with a four-disk configuration for a 100 percent sequential read access pattern. The graph shows that performance for SN200 and P4500 SFF NVMe SSDs is more than 13,000 MBps for the 256-KB and 1-MB block sizes.

Figure 15. NVMe four-SSD configuration: Sequential read 100 percent

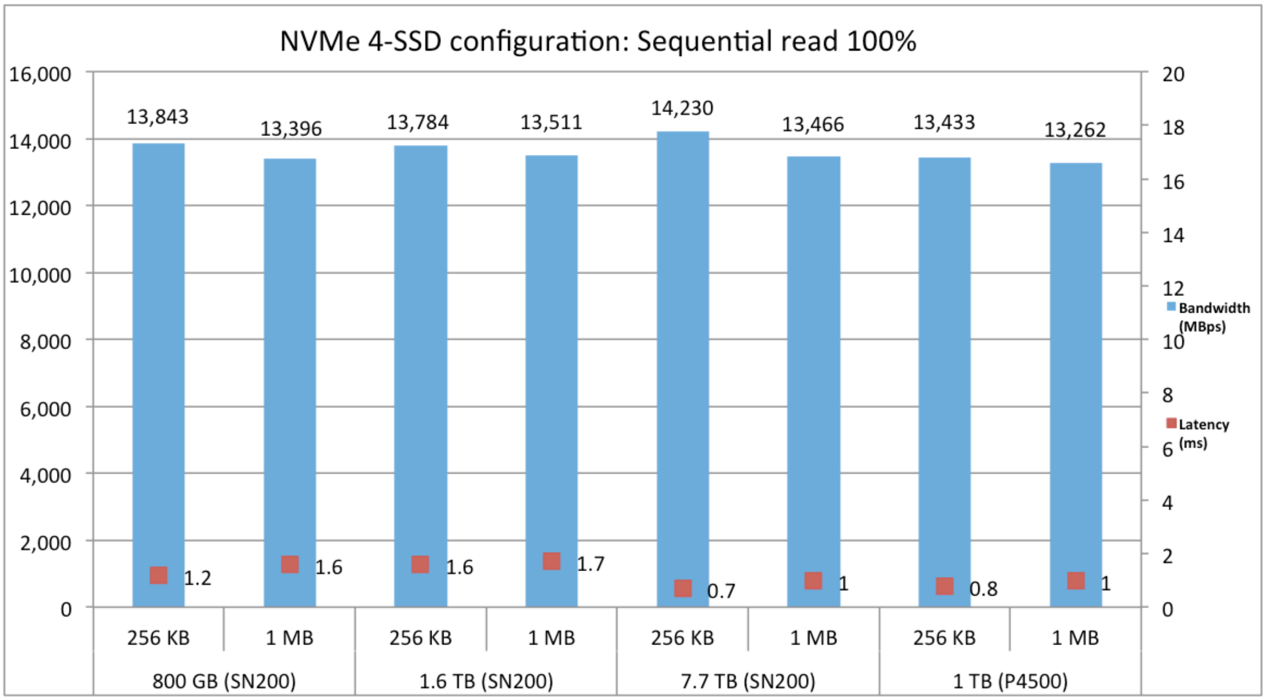

Figure 16 shows the performance of the NVMe SSDs under test with a four-disk configuration for a 100 percent sequential write access pattern. The graph shows that performance for 1.6-TB and 7.7-TB SN200 NVMe SSDs is more than 8500 MBps for the 256-KB and 1-MB block sizes.

Figure 16. NVMe four-SSD configuration: Sequential write 100 percent

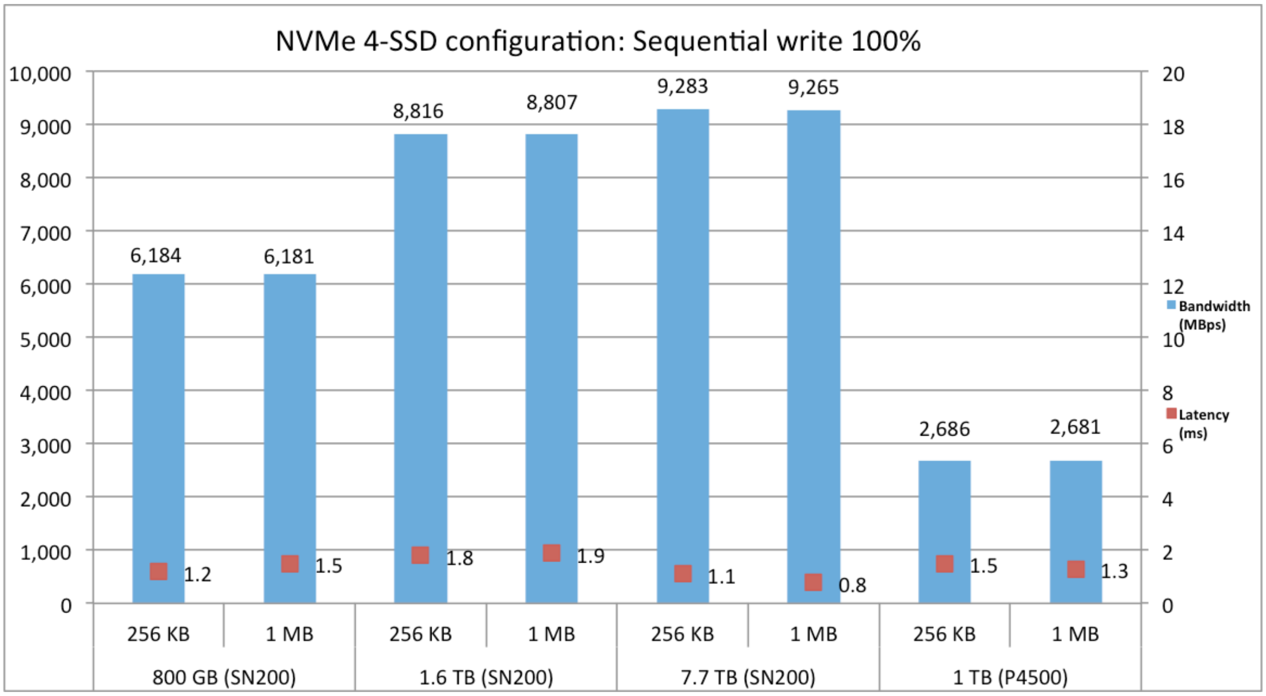

NVMe performance with four half-height half-length cards

The I/O performance characterization tests include NVMe HHHL PCIe cards, which provides better performance than individual NVMe SFF SSD disks. The cards used in this testing were 1.6-TB SN260 NVMe HHHL PCIe 3.0 x8 cards. These cards were used as top-tier storage to accelerate database and high-frequency workloads.

From the below tests it is observed that performance is linearly scaling for all access patterns when number of cards increase from two cards to four cards configuration.

Figure 17 shows the performance of the NVMe HHHL cards under test with a two-card and four-card configuration for a 100 percent random read access pattern. The graph shows the comparative performance values achieved to help customers understand the performance trade-off when choosing an NVMe HHHL card. For 1.6-TB SN260 NVMe HHHL cards, performance for two cards is more than 2.3 million IOPS, and for four cards it is more than 4.7 million IOPS, with a 4-KB block size.

Figure 17. NVMe HHHL configuration: Random read 100 percent

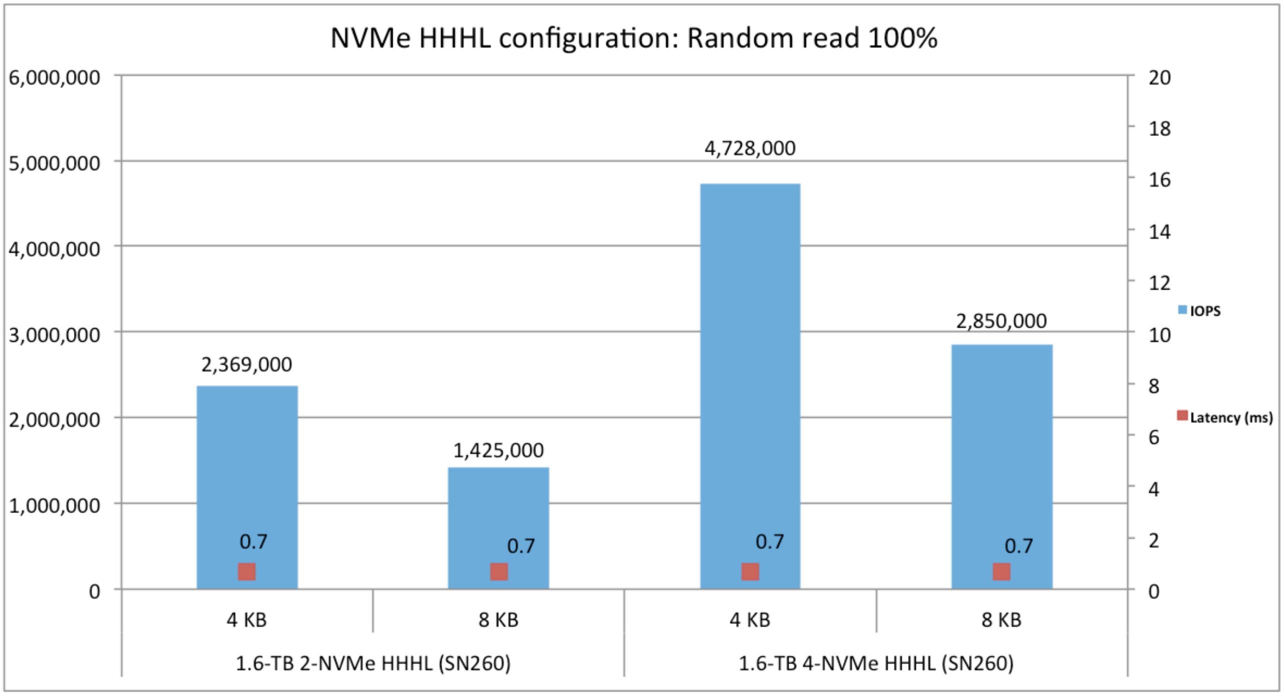

Figure 18 shows the performance of the NVMe HHHL cards under test with a two-card and four-card configuration for a 100 percent random write access pattern. The graph shows that performance for a four-card configuration is more than 750,000 IOPS with a 4-KB block size.

Figure 18. NVMe HHHL configuration: Random write 100 percent

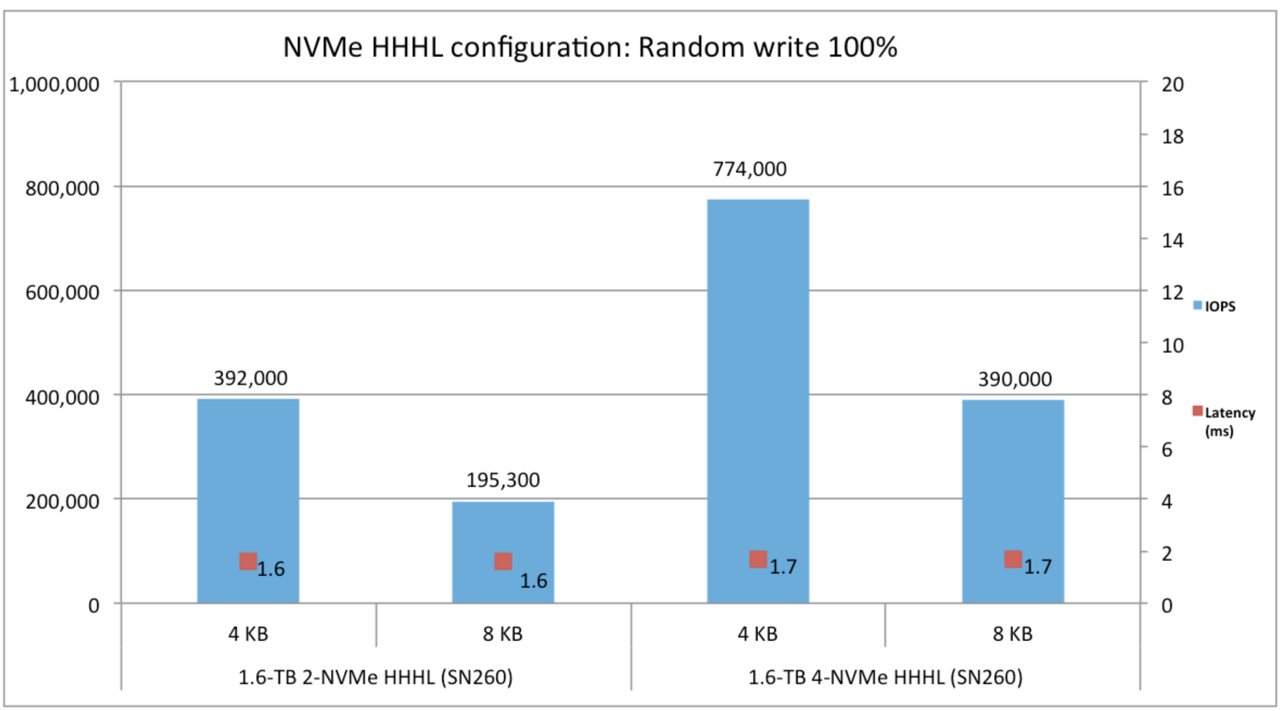

Figure 19 shows the performance of the NVMe HHHL cards under test with a two-card and four-card configuration for a 100 percent random read:write 70:30 percent access pattern. The graph shows that performance for the 1.6-TB SN260 NVMe HHHL for two cards is 1 million IOPS, and for four cards it is 2 million IOPS, with 4-KB block size.

Figure 19. NVMe HHHL configuration: Random read:write 70:30 percent

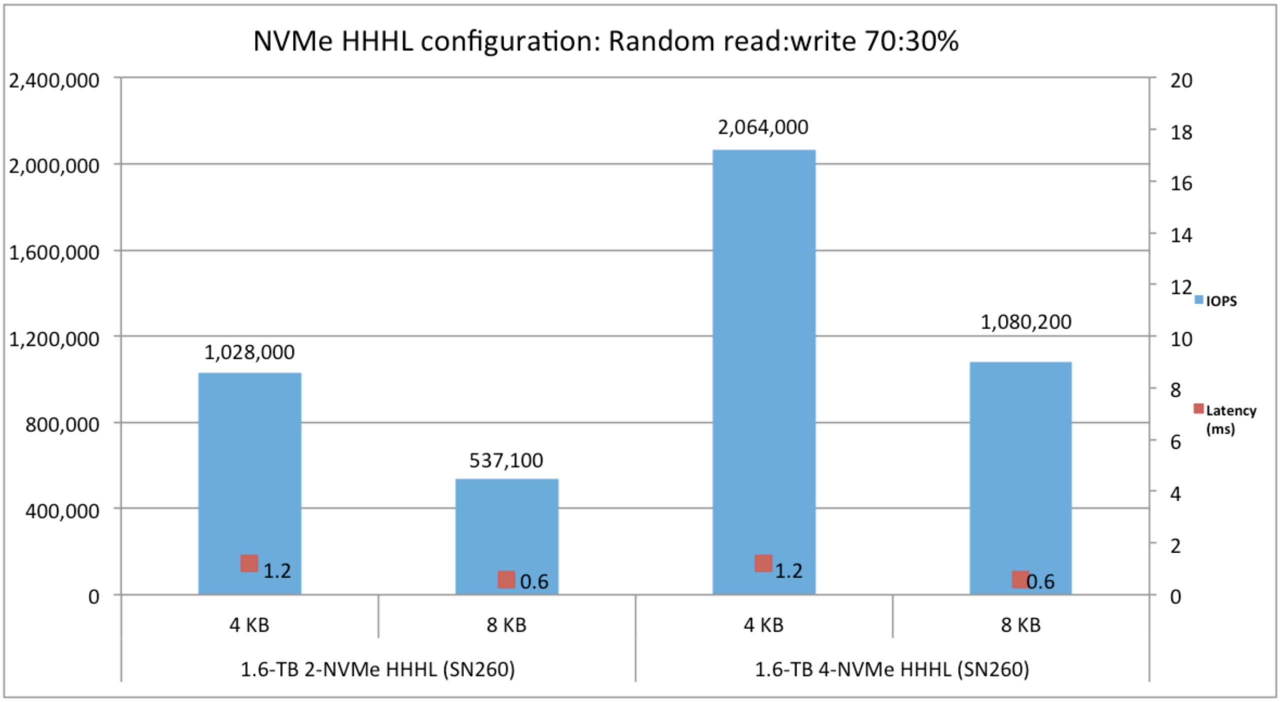

Figure 20 shows the performance of the NVMe HHHL cards under test with a two-card and four-card configuration for a 100 percent sequential read 100 percent access pattern. The graph shows that performance for the 1.6-TB SN260 NVMe HHHL for two cards is 13,000 MBps, and for four cards it is 26,000 MBps, with a 256-KB block size.

Figure 20. NVMe HHHL configuration: Sequential read 100 percent

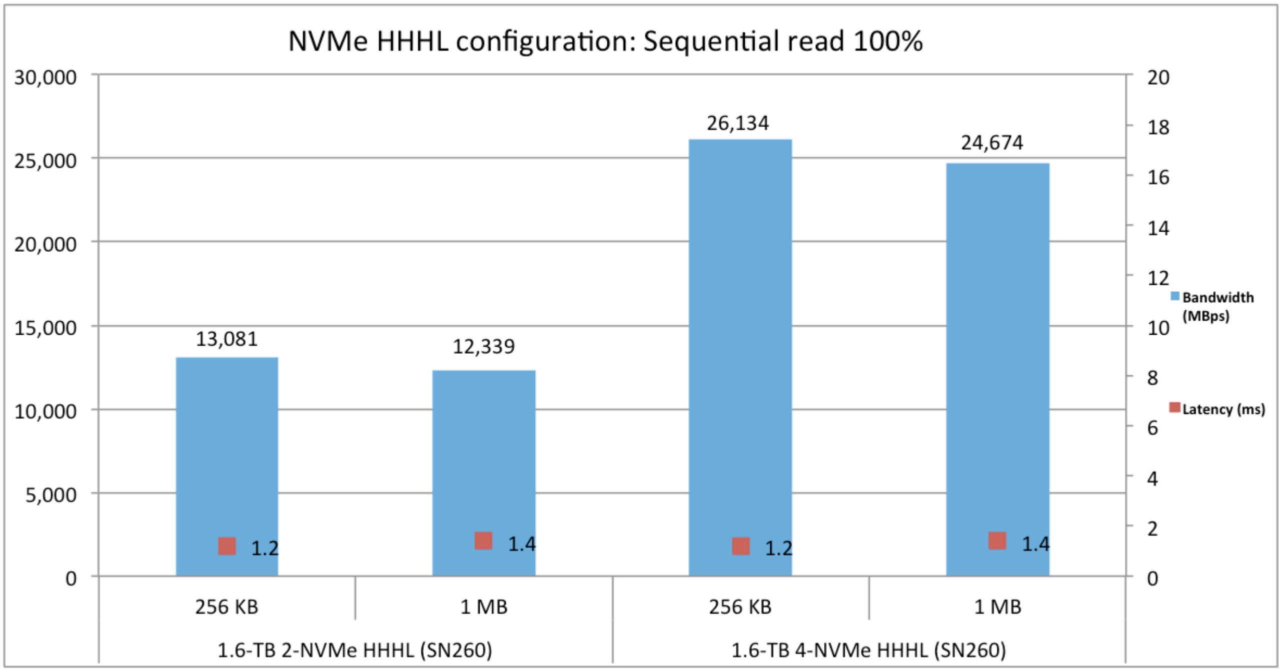

Figure 21 shows the performance of the NVMe HHHL cards under test with a two-card and four-card configuration for a 100 percent sequential write 100 percent access pattern. The graph shows that performance for two cards is more than 4500 MBps, and for four cards it is more than 9000 MBps, with 256-KB and 1-MB block sizes.

Figure 21. NVMe HHHL configuration: Sequential write 100 percent

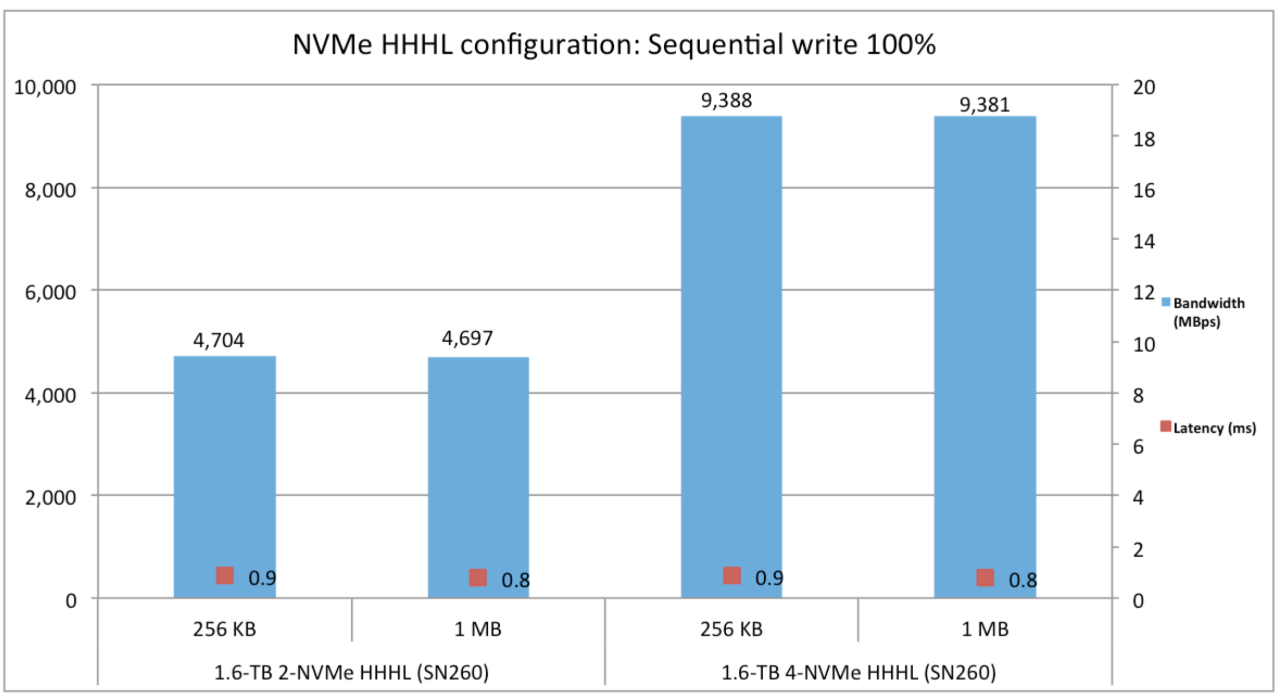

For more information

For additional information, see:

- https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m5-sff-specsheet.pdf
- https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m5-sff-specsheet.pdf

Appendix: Test environment

Table 6 lists the details of the server under test.

Table 6. Server properties

Name	Value
Product name	Cisco UCS C220 M5SX and Cisco UCS C240 M5SX
CPU	CPU: 2 x 2.60-GHz Intel Xeon Gold 6126 processor
Number of cores	12
Number of threads	24
Total memory	384 GB
Memory DIMMs (12)	16 GB x 24 DIMM channels
Memory speed	2666 MHz
Network controller	Cisco LOM X550-T2; 2 x 10-Gbps interfaces
VIC adapter	Cisco UCS VIC 1387 mLOM 40-Gbps SFP+
RAID controller	Cisco 12-Gbps modular RAID controller with 2-GB flash-backed write cache (FBWC; UCSC-RAID-M5)
NVMe SFF SSDs	• 800-GB SN200 NVMe High-Performance High-Endurance SSD (UCSC-NVMEHW-H800) • 1.6-TB SN200 NVMe High-Performance High-Endurance SSD (UCSC-NVMEHW-H1600) • 7.7-TB SN200 NVMe High-Performance High-Endurance SSD (UCSC-NVMEHW-H7680) • 1-TB P4500 NVMe High-Performance Value-Endurance SSD (UCSC-NVMEHW-I1000)
NVMe HHHL card	• 1.6-TB SN260 NVMe Extreme-Performance High-Endurance HHHL card (UCSC-F-H16003)

Table 7 lists the recommended server BIOS settings for a standalone Cisco UCS C-Series Rack Server for NVMe performance tests.

Table 7. BIOS settings for standalone rack server

Name	Value
BIOS version	Release 4.0.1f.0.0906180119
Cisco Integrated Management Controller (IMC) version	Release 4.0(1b)
Cores enabled	All
Intel Hyper-Threading Technology (All)	Enable
Execute disable bit	Enable
Intel Virtualization Technology (VT)	Enable
Hardware prefetcher	Enable
Adjacent-cache-line prefetcher	Enable
Data cache unit (DCU) streamer	Enable
DCU IP prefetcher	Enable

Name	Value
Extended Cisco Application Policy Infrastructure Controller (APIC)	Disable
Non-uniform memory access (NUMA)	Enable
Sub-NUMA Clustering (SNC)	Enable
IMC interleaving	1-way interleave
Mirror mode	Disable
Patrol scrub	Disable
Intel VT for Directed I/O (VT-d)	Enable
Interrupt remapping	Enable
Pass-through direct memory access (DMA)	Disable
Address Translation Services (ATS)	Enable
Posted interrupt	Enable
Coherency support	Disable
Intel SpeedStep Technology (P-states)	Enable
Enhanced Intel SpeedStep Technology (**EIST) PSD function**	HW_ALL
Boot performance mode	Maximum performance
Energy-efficient Intel Turbo	Disable
Intel Turbo mode	Enable
Hardware P-states	Native mode
Enhanced Performance Profile (EPP) profile	Balanced performance
Autonomous core	Disable
CPU C6 report	Disable
Enhanced halt state (C1E)	Disable
OS Advanced Configuration and Power Interface (**ACPI) Cx-state**	ACPI C2
Package C-state	C0/C1 state
Power performance tuning	OS controls EPB
PECI PCB EPB	OS controls EPB
Workload configuration	Balanced

Figure 22 shows the recommended server BIOS settings for a Cisco UCS C-Series Rack Server for NVMe performance tests managed by Cisco UCS.

Figure 22. BIOS settings for rack server managed by Cisco UCS

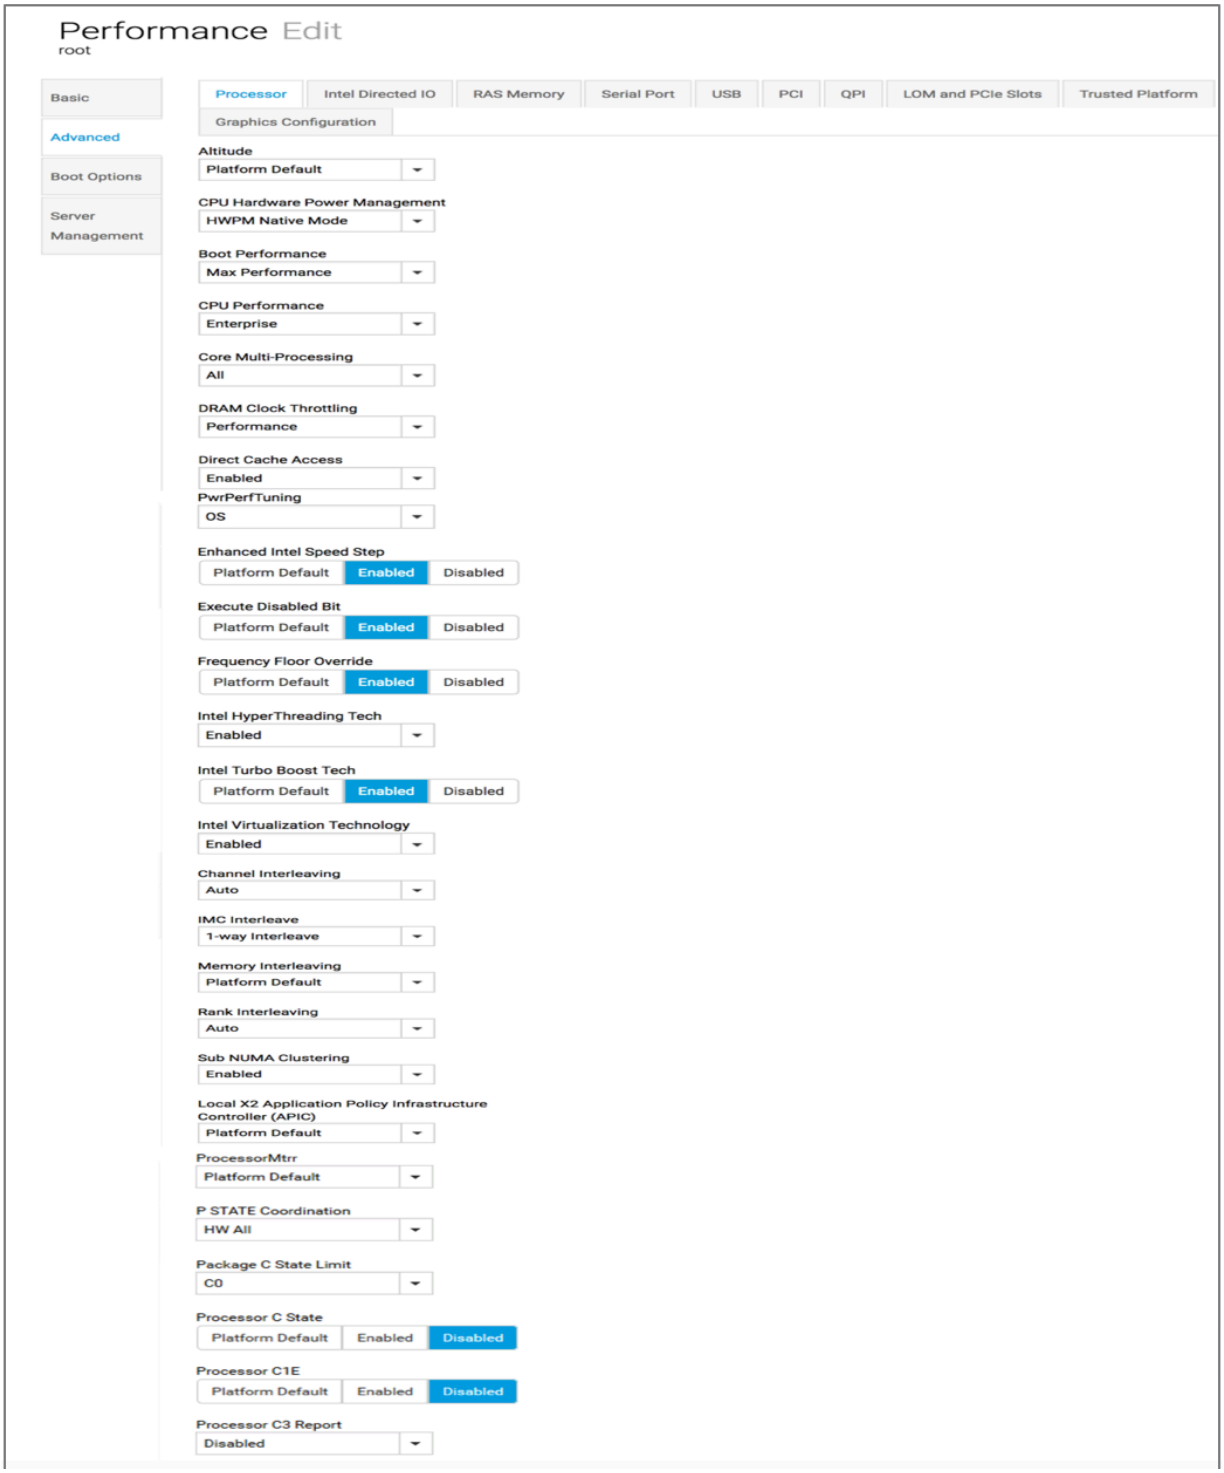

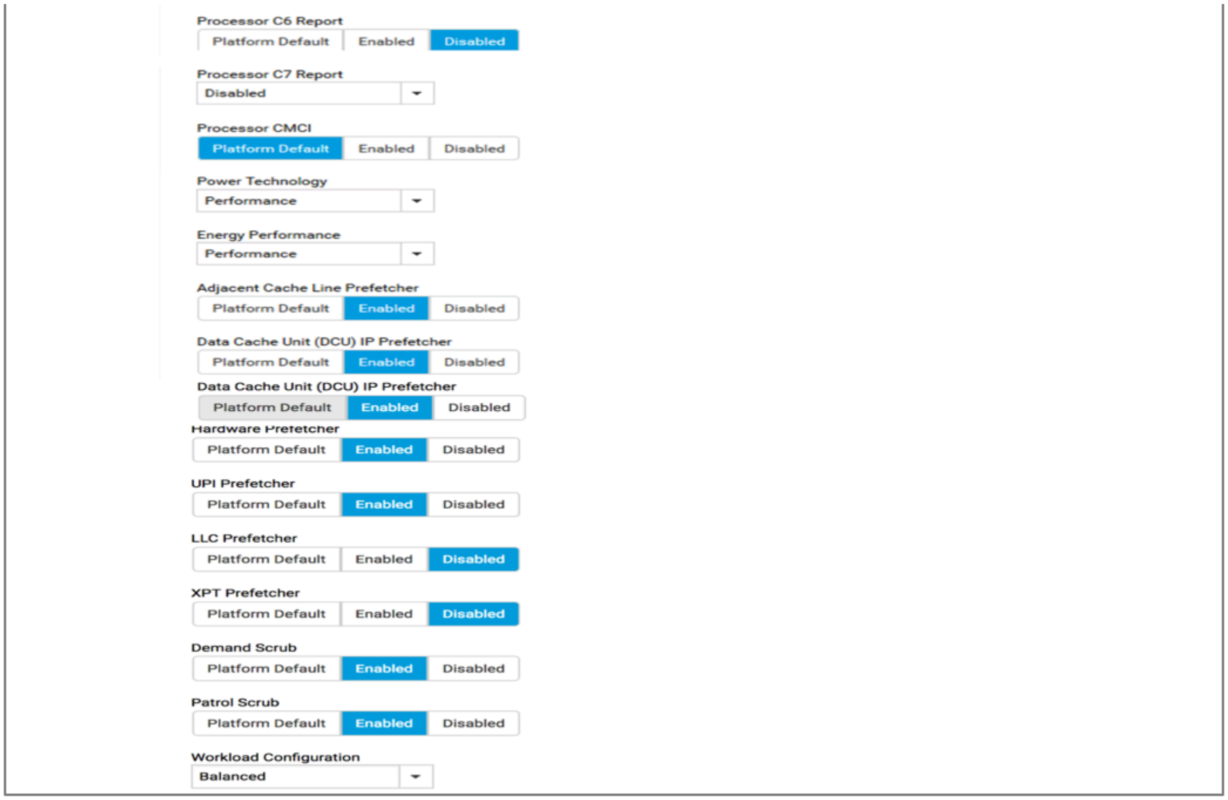

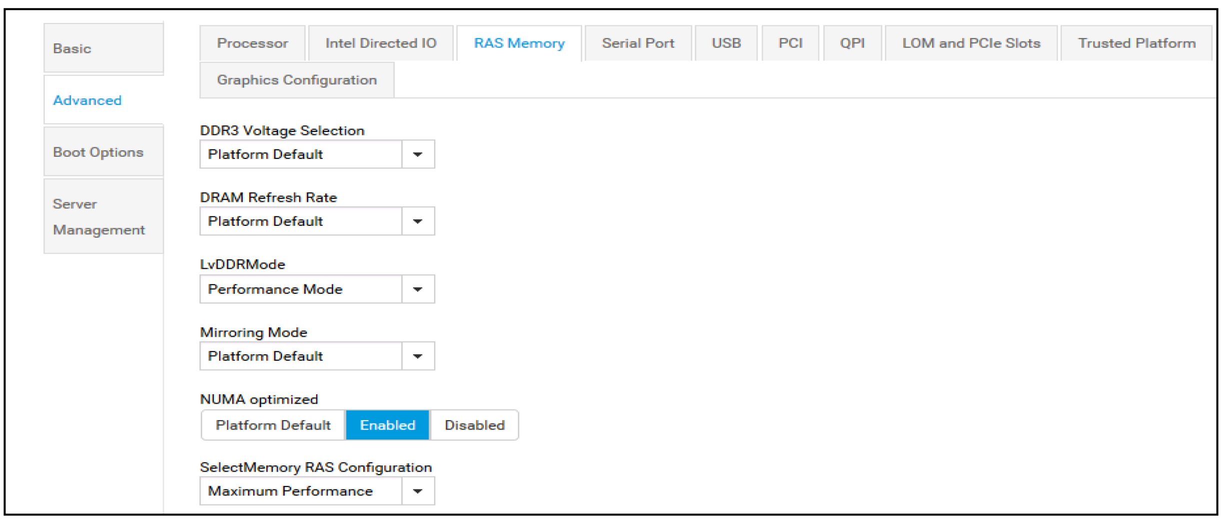

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

C11-741525-00 10/18