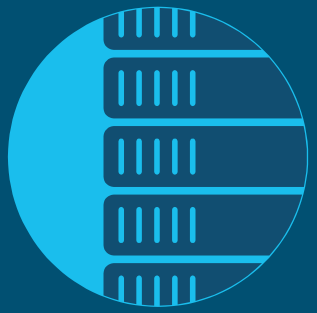

NVMe and NVMe over Fibre Channel

Unleash the power of flash storage using Cisco UCS, Nexus, and MDS

What is NVMe?

NVM Express (NVMe) is an optimized, high-performance, scalable interface designed to work with current and the next-generation NVM technologies. The NVMe interface is defined to enable host software to communicate with nonvolatile memory over PCI Express (PCIe). NVMe has been designed from the ground up to offer:

- Large number of queues (up to 64,000) for command submission and completion. Every CPU core can have its own independent I/O queue pair without any need for locking and contention.
- Each queue is much deeper (up to 64,000 commands) compared to SAS or SATA. The high parallelism of the NVMe interface is designed to handle ultra-fast CPU and storage.
- Streamlined and simple command set, resulting in fewer CPU instructions needed to process I/O requests, providing more Input/Output instructions Per Second (IOPS) per CPU instruction cycle, and lower I/O latency in the host software stack.
- Advanced features for future software and hardware development.

Due to these and multiple other enhancements of the storage I/O stack, NVMe has helped to improve the performance of flash storage many times over, in ways that were not possible with the traditional technologies for many operating systems. (Refer to http://www.nvmexpress.org/ for performance data, a complete list of enhancements, and specifications.)

Why NVMe?

Recent years have seen tremendous growth in the adoption of Solid-State Drives (SSDs). These drives are based on Nonvolatile Memory (NVM). Unlike traditional spinning drives, with SSDs there are no rotating motors and moving heads. As a result, the read and write operations are extremely fast. Earlier technologies (such as SAS and SATA) that were architected for hard disk drives, are unable to take full advantage of SSDs' potential. This warranted the need for a new architecture: Non-Volatile Memory Express.

NVMe over fabrics

NVMe has also been standardized to extend its advantages over various fabrics to external and centralized storage. This is known as NVMe over Fabrics (NVMe-oF). NVMe-oF is of huge interest to enterprises with a large installed base of centralized storage accessible over a network. The NVMe-oF specification defines common frameworks for NVMe to work over a variety of transports, including Ethernet-, Infiniband-, and Fibre Channel–based networks. While NVMe and NVMe-oF specifications are defined by the NVM Express, Inc., the exact transport details are defined by the respective responsible bodies. For example, the NVMe over Fibre Channel (FC-NVMe) specification is defined by T11 committee of the International Committee on Information Technology Standards (INCITS).

Cisco's leadership and offerings

Cisco is a member of the board of directors of the promoter group of NVM Express, Inc. as well as a member of the board of directors for the Fibre Channel Industry Association, and has been involved with NVMe and FC-NVMe since their inception.

We announced support for NVMe storage on the Cisco Unified Computing System™ (Cisco UCS®) C-Series Rack Servers and B-Series Blade Servers in October 2016. Multicore CPUs on the Cisco UCS servers can utilize the full potential of NVMe-capable solid-state drives connected via PCIe. In April 2017, the support of NVMe over Fibre Channel was extended to Cisco UCS C-Series Rack Servers, MDS 9000 series switches, and Cisco Nexus® 5000 Series Switches. In November 2017, the support was extended to Nexus 93180YC-FX on Fibre Channel ports.

NVMe Storage for Cisco UCS C-Series Rack and B-Series Blade Servers

NVMe-capable flash storage inside the Cisco UCS C-Series Rack and B-Series Blade Servers is best suited to applications requiring ultra-high performance. Multicore CPUs are connected to NVMe flash storage using non-oversubscribed PCIe lanes, bringing storage as close as possible to the CPU (Figure 1). This results in reduced latency and increased IOPS, resulting in overall increased application performance.

Figure 1. NVMe Storage for Cisco UCS C-Series Rack and B-Series Blade Servers

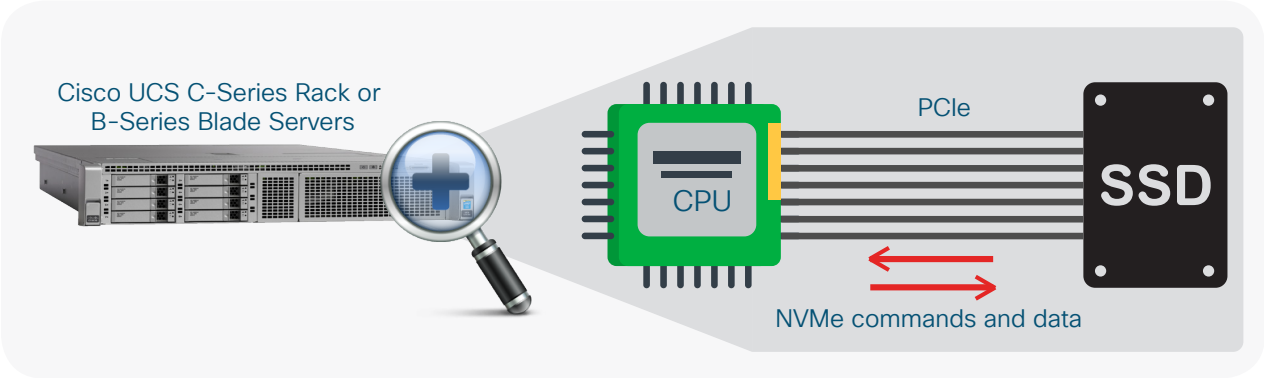

Cisco UCS C-Series Rack Servers and B-Series Blade Servers with internal NVMe flash storage offer following unique advantages:

- **Reduced total cost of ownership:** Increased application performance compared to spinning disk or non-NVMe flash storage, resulting in a reduced cost per IOPS and hence a further reduction in power, cooling, and overall footprint
- **Simplicity:** Integrated computing, networking, storage, and virtualization
- **Manageability:** Cisco UCS Manager capabilities for managing hundreds of Cisco UCS servers with NVMe flash storage
- **Flexibility:** Best-in-class capacity with the flexibility to choose from as little as 400 GB to as much as 30 TB of capacity in a 2-rack-unit server

NVMe over Fibre Channel using Cisco UCS, Cisco MDS, and Cisco Nexus solutions

Benefits of NVMe can be extended to centralized storage infrastructure over a SAN built using Cisco Nexus and MDS switches. NVMe initiators consisting of Cisco UCS C-Series Rack Servers installed with Broadcom (Emulex) or Cavium (QLogic) Host Bus Adapters (HBAs) can access NVMe targets over Fibre Channel or Fibre Channel over Ethernet (FCoE) fabrics (Figure 2). Fibre Channel is the preferred protocol for connecting all-flash arrays in today's data centers due to its performance, availability, scalability, and plug-and-play architecture. FC-NVMe offers the best of Fibre Channel and NVMe. You get the improved performance of NVMe along with the flexibility and the scalability of the shared storage architecture.

Figure 2. Seamless migration to FC-NVMe using Cisco Nexus and MDS

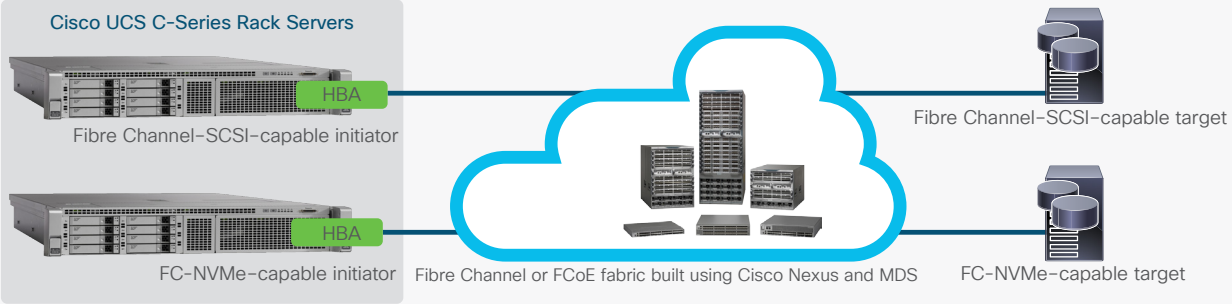

FC-NVMe on Cisco Nexus and MDS switches offers the following unique advantages:

- **Multiprotocol flexibility:** Cisco MDS and Nexus switches support NVMe and SCSI (Small Computer System Interface) simultaneously over Fibre Channel or FCoE fabrics. Enterprises can continue to use their existing infrastructure and roll out new FC-NVMe–capable end devices in phases, sharing the same Fibre Channel or FCoE SAN.
- **Seamless insertion:** NVMe support can be enabled on Cisco Nexus and MDS switches through a nondisruptive upgrade of Cisco NX-OS Software. No hardware changes are required.
- **Investment protection:** FC-NVMe is supported on current and future generations of switches. Your current investment goes a long way to support FC-NVMe without any extra SAN cost.
- **Superior architecture:** Large enterprises across the globe from various industry verticals trust Cisco MDS 9700 Series solutions for their mission-critical data centers due to their superior architecture. Fully redundant components, non-oversubscribed and nonblocking architecture, automatic isolation of failure domains, and exceptional capability to detect and automatically recover from SAN congestion are few of the top attributes that make these switches the best choice for high-demand storage infrastructures that support NVMe-capable workloads.
- **Integrated storage traffic visibility and analytics:** The 32-Gbps products in the Cisco MDS 9000 switch family offer Cisco SAN Telemetry Streaming. This solution provides line-rate, native, on-switch storage traffic visibility to help you to make data-driven decisions to increase the utilization of your infrastructure and remove bottlenecks. Cisco SAN Telemetry Streaming is based on a programmable network processing unit, which makes the solution ready for FC-NVMe traffic in future with just a non-disruptive software-only upgrade.
- **Strong ecosystem:** Cisco UCS C-Series Rack Servers are certified with Broadcom (Emulex) and Cavium (QLogic) HBAs. Additionally, Cisco UCS servers and Cisco Nexus, and MDS switches offer unique merits leading to stable and efficient storage infrastructure.

For more information

NVMe Storage for Cisco UCS C-Series Rack and B-Series Blade Servers Data Sheet

Cisco UCS data sheets and literature

Cisco MDS 9000 data sheets and literature

Cisco Nexus 9300-EX and 9300-FX platform switches data sheet

Cisco Nexus 5600 data sheets and literature

NVMe and NVMe over fabrics specifications

Emulex/Broadcom HBA

QLogic/Cavium HBA

C22-740196-00 03/18

Supported products

NVMe

Cisco UCS C-Series Rack Servers (M4 and M5 models)

Cisco UCS B-Series Blade Servers (M4 and M5 models)

FC-NVMe

Table 1. Support matrix for FC-NVMe on Cisco UCS C-Series Rack Servers, HBAs, Cisco Nexus and MDS switches.

Servers as initiators	HBA	Switches
Cisco UCS C-Series Rack Servers (M4 and M5 models)	Emulex/Broadcom LPe 32000 Emulex/Broadcom LPe 31000 QLogic/Cavium QLE 2700 QLogic/Cavium QLE2690	▪ Cisco MDS 9000 switch family (Fibre Channel and FCoE ports), NX-OS 8.1(1) onwards. ▪ Cisco Nexus 93180YC-FX (Fibre Channel and FCoE ports), NX-OS 7.0(3)I7(2) onwards. ▪ Cisco Nexus 5600 and Cisco Nexus 2348UPQ (Fibre Channel and FCoE ports), NX-OS 7.3(0)N1(1) onwards.

Notes:

- NVMe over FCoE is supported only on Inter-Switch Link (ISL) ports between Cisco Nexus or MDS switches, or both.
- As of February 2018, Fibre Channel and FCoE support on Nexus 93180YC-FX are available in NPV mode only.
- FC-NVMe is independent of the operational speed of the switch port. Higher speeds are recommended.