

NVMe SSD Classification

Cameron Brett

Mark Carlson

SNIA CMSI Solid State Drive Special Interest Group

Revised February 2023

Table of Contents

Introduction..... 1

Enterprise Server..... 1

 Applications/Use Cases 1

 Form Factors..... 1

 Power..... 1

 Performance 1

 Latency 2

 Endurance..... 2

 Capacity 2

 Management..... 2

 Features..... 2

Enterprise Storage..... 2

 Applications/Use Cases 2

 Form Factors..... 2

 Power..... 3

 Performance 3

 Latency 3

 Endurance..... 3

 Capacity 3

 Management..... 3

 Features..... 3

Low Latency SSDs..... 4

 Applications/Use Cases 4

 Form Factors..... 4

 Power..... 4

 Performance 4

 Latency 5

 Endurance..... 5

 Capacity 5

Management 5

Features..... 5

Data Center/Hyperscale Server..... 5

 Applications/Use Cases 6

 Form Factors..... 6

 Power..... 6

 Performance 6

 Latency 6

 Endurance..... 6

 Capacity 6

 Management..... 7

 Features..... 7

Data Center/Hyperscale Storage..... 7

 Applications/Use Cases 7

 Form Factors..... 7

 Power..... 7

 Performance 7

 Latency 8

 Endurance..... 8

 Capacity 8

 Management..... 8

 Features..... 8

Introduction

NVM Express technology is driving the next generation of SSDs for data centers. NVMe™ SSDs are not limited by legacy form factors or protocols, and they can better address the needs of servers and storage in enterprise and hyperscale data centers. With the varied and dedicated use cases for NVMe SSDs, there are two classifications that generally define and segment these classes: Enterprise and Data Center.

This white paper describes use cases for NVMe SSDs used in enterprise servers, enterprise storage, low latency SSD, data center/hyperscale servers, and data center/hyperscale storage environments. Each of these use cases has distinct requirements and application environments, which are covered in requirements for form factors, power, performance, latency, endurance, capacity, management, and features. There will always be exceptions to these guidelines, but these are generally fit profiles for NVMe SSD classification

The paper is designed as a companion to the NVMe SSD Classification table on the SNIA.org website at <https://www.snia.org/technology-focus-areas/physical-storage/nvme-ssd-classification>.

Enterprise Server

An enterprise comprises all aspects of an organization's assets, systems, services, and functions. A server is an [intelligent device](#), usually a computer, that provides services to other intelligent devices, usually other computers or appliances. Some recent examples of enterprise servers are Dell EMC PowerEdge, HPE ProLiant, and Lenovo ThinkServers.

Applications/Use Cases

Standalone or rack servers and their applications are commonly used. Example applications include Online Transactional Processing (OLTP), which is I/O intensive; and large block, which includes images and video files as well as multi-server environments, such as cluster servers.

Form Factors

Form factors for NVMe SSDs in enterprise server environments include M.2 transitioning to E1.S for boot, 2.5-inch (U.2), E3.S, and E3.L for data and caching. More information on form factors can be found at <https://www.snia.org/forums/cmsi/knowledge/formfactors>

Power

Power requirements for NVMe SSDs in enterprise servers are 8.25-11W for M.2, 12-25W for 2.5-inch (U.2), and 25-40W for EDSFF.

Performance

Performance for NVMe SSDs in enterprise servers typically are among the highest available with the lowest latencies. It is common to need both the highest Input/Output operations per second (IOPS) and throughput since virtualized environments tend to have a mix of IOPS and throughput demands.

Latency

In enterprise servers, low latency is important, especially for I/O intensive applications where the lower latency the better. Quality of service (QoS) is not as important as lowest latency possible.

Endurance

In enterprise server environments, look for one Drive Write per Day (DWPD) for read-intensive environments, three DWPD for mixed-use environments, and five to ten DWPD for write-intensive environments.

Capacity

Typical capacity for NVMe SSDs in enterprise server environments are 1-15.36 TB.

Management

Enterprise servers are placed in environments where the customer controls what software is running on their server. They deploy a baseboard management controller (BMC) to manage the server devices over an internal management bus. Management of storage attached on a network needs to take place on the same network, as it is the only connection to the device. In an environment with NVMe SSDs a management product using SNIA Swordfish™ should be used. A SNIA Swordfish client would manage the storage through the BMC's external interface.

Features

Typical features found in NVMe SSDs in enterprise server environments are security/encryption, manageability, single-port, and RAID.

Enterprise Storage

Enterprise storage environments contain centralized repositories for business information that provide common data management, protection, and data sharing functions through connections to computer systems.

Applications/Use Cases

Traditional storage boxes are commonly used in enterprise storage. Applications include all-flash arrays where NVMe SSDs are used for both cache as well as data storage connected by a traditional Storage Area Network (SAN) or NVMe-over-Fabrics (NVMe-oF).

Form Factors

Example Form factors for NVMe SSDs in enterprise storage environments include 2.5-inch (U.2), E3, and E1. More information on form factors can be found at <https://www.snia.org/forums/cmsi/knowledge/formfactors>

Power

Power requirements for NVMe SSDs in enterprise storage are 8.25-11W for M.2, 12-25W for 2.5-inch (U.2), and 25-40W for E3 and E1.

Performance

Performance for NVMe SSDs in enterprise storage typically are among the highest available with the lowest latencies. It is common to need both the highest Input/Output operations per second (IOPS) and throughput since virtualized environments tend to have a mix of IOPS and throughput demands.

Latency

For enterprise storage, Quality of Service (QoS) or consistent latency becomes more important as more applications use the storage. Lower latency is better, but reliable and consistent latency is critical.

Endurance

In enterprise storage environments, look for one Drive Write per Day (DWPD) for read-intensive environments and three DWPD for mixed-use environments.

Capacity

Typical capacity for NVMe SSDs in enterprise storage environments are 1.92-30.72 TB.

Management

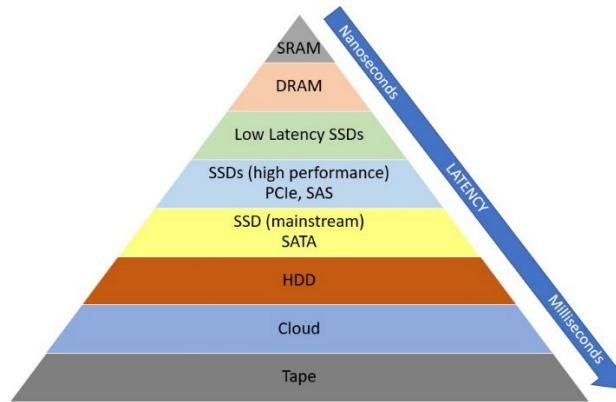
Like the enterprise server environment, enterprise storage may be managed through a BMC attached to a controller. This BMC may be the point of management for all the NVMe SSD drives. The network connection also allows a host to manage the storage through Swordfish, and the BMC through the NVMe Management Interface (NVMe-MI). SNIA Swordfish may also be used to manage multiple storage systems in a data center.

Features

Typical features found in NVMe SSDs in enterprise storage environments are dual-port x8/16 PCIe AIC, higher prices (greater than SANs) and capacities up to 30.72 TB.

Low Latency SSDs

Low latency SSDs can fulfill the niche where persistent memory is targeted, sitting between DRAM and Flash in the cost/performance “pyramid”.



They are usually accessed by infrastructure such as caches, filesystems, hypervisors and operating systems.

Applications/Use Cases

Applications in low latency SSD environments include read/write cache; tier 0 between DRAM and TLC SSD; write intensive workloads; write logging (metadata); low latency SSD plus QLC – lower TCO than TLC.

Form Factors

Example Form factors for NVMe SSDs in low latency SSD environments include 2.5-inch (U.2) and add in cards (AICs). More information on form factors can be found at <https://www.snia.org/forums/cmsi/knowledge/formfactors>

Power

Power requirements for NVMe SSDs in low latency SSD environments are 20-25W for 2.5-inch (U.2), and 9W for AIC.

Performance

NVMe SSDs in low latency SSD environments are used in high performance, write heavy workloads for best in class performance.

NVMe SSD Classification

Latency

For low latency SSDs, latency is 1/10th to 1/3 the latency of TLC SSDs.

Endurance

In low latency SSD environments, look for 30, 60, or 100 Drive Writes per Day (DWPD).

Capacity

Typical capacity for NVMe SSDs in low latency SSD environments are 375 -3200TB.

Management

Like the enterprise server environment, low latency SSDs may be managed through a BMC attached to a controller. This BMC may be the point of management for all the NVMe SSD drives. The network connection also allows a host to manage the storage through Swordfish, and the BMC through the NVMe Management Interface (NVMe-MI). SNIA Swordfish may also be used to manage multiple storage systems in a data center.

Features

Typical features found in NVMe SSDs in enterprise storage environments are dual-port (2.5-inch form factor), predictable low latency, and very high endurance.



Illustration: Low Latency SSD Drive - FL Series SSD, with XL-FLASH™ memory, in a 2.5" 15mm Z-height form factor. (courtesy of KIOXIA)

Data Center/Hyperscale Server

Large data center customers are known as hyperscalers. Nearly all of these customers build their own storage systems using Software Defined Storage (SDS) and commodity components, assembled in racks. Hyperscale servers are included in these systems.

Applications/Use Cases

NVMe SSDs are used to boot the servers securely (i.e., M.2) and support the master boot block feature for this purpose. The drives in a cloud data storage system are also direct attached but exported via software defined storage.

Form Factors

Different form factors have different ranges of power requirements. Hyperscalers are working with the new E1.S form factor(s) because of the densities allowed. Hyperscalers can and have also been using M.2 and U.2. More information on form factors can be found at <https://www.snia.org/forums/cmsi/knowledge/formfactors>

Power

Power requirements for NVMe SSDs in data center/hyperscale server are 8.25-11W for M.2, 12-25W for E1.S, and 12-25W for U.2.

Performance

Hyperscaler environments want to see consistent latency performance and make sure drives scale their IOPS performance with increasing capacity. Because they control the software implementing the storage system, they can tune the performance to the drives' capabilities.

Latency

Tail latencies are those operations which take much longer than the average. This issue causes problems for upper layers of the stack waiting for that last completion. This shows through to the hyperscalers' customers as long page loads. New features of NVMe provide means by which to minimize or eliminate tail latency. One of the critical criteria for data center SSDs is consistent and reliable latency, to ensure predictable performance.

Endurance

Endurance is the key feature for hyperscaler environments. They create data centers and want the components of that data center to last until they build a data center to replace it. The amount of data typically written during a single day is equal to its capacity. One Drive write per day (DWPD) is the most common endurance.

Capacity

Typical capacity for NVMe SSDs in data center/hyperscale server environments are 1- 4 TB in 2021.

Management

Hyperscalers run their own software on their servers and can therefore use host agents via NVMe Command Line Interface (NVMe-CLI CLI) and scripting. While hyperscaler servers may have a BMC, the BMC only handles minimal power and cooling management. Minimal support for NVMe-MI basic is all that is needed.

Features

Hyperscalers handle redundancy and availability at the upper layers of their own software. Therefore, they do not need components with “no single point of failure”. This makes the components cheaper when they have multiple copies of the data elsewhere. The quality of service for reads is key in reducing tail latency. NVMe new features of NVM Sets and I/O determinism (IOD) also support QoS. Data center class SSDs for servers are typically single-port.

Data Center/Hyperscale Storage

Large data center customers are known as Hyperscalers. Nearly all of these customers build their own storage systems using Software Defined Storage (SDS) and commodity components, assembled in racks.

Applications/Use Cases

Hyperscalers are now moving to deploying JBOF/EBOF systems with drives network attached to the host, creating a disaggregated compute, storage and network environment. This allows the storage and server capacities to scale separately. IT also improves density. The system software that provides storage is evolving in this direction as well.

Form Factors

With the evolution of new enterprise storage features, new form factors such as the E1.L and E1.S are appearing in these external chassis. Both density and a matching between CPU power and SSD capability is needed and will need to accommodate new technologies.

Power

The added power allowed in these form factors is increasing both performance and density.

Performance

NVMe SSDs offer performance for IOPS/TB scaling and mixed workloads. Performance tends not to be as high as enterprise but require consistent and predictable performance.

Latency

For data center/hyperscale storage, read prioritization is most important without long tails. As with performance, consistency and predictability is critical for hyperscale applications.

Endurance

Endurance is a key feature for hyperscaler environments. They create data centers and want the components of that data center to last until they build a data center to replace it. The amount of data typically written during a single day is equal to its capacity. One Drive writes per day (DWPD)

Capacity

Data center/hyperscale storage capacities range from 4-8 TB.

Management

Hyperscale storage systems are built from servers typically with direct attached drives. These drives are managed by host agents and again with minimal support from the BMC.

Features

Many features of hyperscale storage are similar to those of hyperscale servers. With hyperscale servers deploying JBOF/EBOF systems with drives network attached to the host, this allows the storage and server capacities to scale separately. It also improves density. The system software that provides storage is evolving in this direction as well.

About SNIA

SNIA is a not-for-profit global organization, made up of member companies spanning the global storage market. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For more information, visit <http://www.snia.org>.



SNIA

5201 Great America Pkwy, Ste 320 • Santa Clara, CA 95054 • Phone: 719-694-1380 • Fax: 719-694-1389 • www.snia.org

© SNIA. All Rights Reserved.