



Storage Networking Industry Association



SNIA Emerald™ V4 Taxonomy

Technical White Paper
September 2019

Abstract: The SNIA Emerald™ Taxonomy used by the SNIA Emerald™ Power Efficiency Measurement Specification and the ENERGY STAR® for Data Center Storage Specification is being updated and expanded for use in future revisions of these specifications. This document describes the changes and their rationale and is a vehicle for soliciting review and comment on these changes in advance of their being incorporated into these specifications.

USAGE

The SNIA hereby grants permission for individuals to use this document for personal use only, and for corporations and other business entities to use this document for internal use only (including internal copying, distribution, and display) provided that:

1. Any text, diagram, chart, table or definition reproduced shall be reproduced in its entirety with no alteration, and,
2. Any document, printed or electronic, in which material from this document (or any portion hereof) is reproduced shall acknowledge the SNIA copyright on that material, and shall credit the SNIA for granting permission for its reuse.

Other than as explicitly provided above, you may not make any commercial use of this document, sell any or this entire document, or distribute this document to third parties. All rights not explicitly granted are expressly reserved to SNIA.

Permission to use this document for purposes other than those enumerated above may be requested by e-mailing tcmd@snia.org. Please include the identity of the requesting individual and/or company and a brief description of the purpose, nature, and scope of the requested use.

All code fragments, scripts, data tables, and sample code in this SNIA document are made available under the following license:

BSD 3-Clause Software License

Copyright (c) 2019, The Storage Networking Industry Association.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of The Storage Networking Industry Association (SNIA) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

DISCLAIMER

The information contained in this publication is subject to change without notice. The SNIA makes no warranty of any kind with regard to this specification, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The SNIA shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this specification.

Suggestions for revisions should be directed to <http://www.snia.org/feedback/>.

Copyright © 2019 SNIA. All rights reserved. All other trademarks or registered trademarks are the property of their respective owners.

Executive Summary

In order to remain current, the SNIA Emerald™ Taxonomy used by the SNIA Emerald™ Power Efficiency Measurement Specification and the ENERGY STAR® for Data Center Storage Specification is being updated and expanded to better describe today's storage systems by adding explicit support for storage systems based on persistent Solid State Storage.

In addition, an Online 1.5 is added to address simple JBODs, several taxonomy attributes are modified, and definitions are added for clarity.

Table of Contents

INTRODUCTION..... 6

EXISTING TAXONOMY 6

INTRODUCING STORAGE SETS 8

TAXONOMY CLASSIFICATION TABLES AND ATTRIBUTE DEFINITIONS 11

 Table Modifications and Additions 11

 Taxonomy Attribute Definitions 14

 Storage Controller..... 14

 Integrated Storage Controller..... 14

 Storage Protection 14

 System Capacity 14

SUMMARY..... 16

FOR MORE INFORMATION..... 17

INTRODUCTION

In order to remain current, the SNIA Emerald™ Taxonomy is being updated and expanded to better describe today’s storage systems. This will serve to stay abreast with expected new revisions of the SNIA Emerald™ Power Efficiency Measurement Specification and the ENERGY STAR® for Data Center Storage Specification.

The storage industry has introduced many products based on various types of non-volatile Solid State Storage. In the current SNIA Emerald Taxonomy, these are included within the same Categories that include hard disks drives, tape, etc. In order to better address these products, new Categories are added to the planned Taxonomy. To organize these new Categories, a new top level construct Set is added to the taxonomy. The planned new Taxonomy separately treats non-volatile Solid State Storage through the concept of the NVSS (Non-Volatile Solid State) Set and its Categories.

The planned new Taxonomy builds upon the current Taxonomy via the introduction of Storage Sets, which add NVSS (Non-Volatile Solid State) to the existing Disk, Removable and Virtual Media types of storage devices. In addition, an Online 1.5 is added to address simple JBODs, several taxonomy attributes are modified, and definitions are added for clarity. The new Taxonomy does not address combined storage/network/compute or software-defined/hyper-converged. These constructs are gray areas beyond the scope of pure storage systems.

EXISTING TAXONOMY

The existing Storage Taxonomy is described in the SNIA Emerald Power Efficiency Measurement Specification Version 3.0.3; the high-level table from this specification is included below as Table 1.

Level	Category			
	Online (Sect 5.5)	Near-Online (Sect 5.5)	Removable Media Library (Sect 5.7)	Virtual Media Library (Sect 5.8)
Consumer/ Component ^a	Online 1	Near-Online 1	Removable 1	Virtual 1
Low-end	Online 2	Near-Online 2	Removable 2	Virtual 2
Mid-range	Online 3	Near-Online 3	Removable 3	Virtual 3
	Online 4			
High-end	Online 5	Near-Online 5	Removable 5	Virtual 5
Mainframe	Online 6	Near-Online 6	Removable 6	Virtual 6
^a Entries in this level of taxonomy include both consumer products and data-center components (e.g., stand-alone tape drives)				

Table 1. Existing high-level taxonomy table from the SNIA Emerald v3.0.3 Specification

This taxonomy was published prior to the publication of Energy Star for Data Center Storage version 1.0. It had been envisioned that Online 2 would be used for JBOD and very low-end storage systems. When Energy Star published their rules for Online 2, they included a requirement for storage protection (such

as a version of RAID) to be included in the Online 2 product. This excluded the low level JBOD from the requirement. To accommodate this, SNIA is proposing to add a level for low level JBOD devices that will be designated as Online 1.5. The JBOD will only be categorized as Online 1.5 as by its very nature, a JBOD will meet the Online requirements. An Online 1.5 JBOD can be used in a Near-Online system or a Virtual Media Library, but will become part of the system and not be a stand-alone product.

The updated taxonomy table below incorporates the JBOD Level.

Level	Category			
	Online (Sect 5.5)	Near-Online (Sect 5.5)	Removable Media Library (Sect 5.7)	Virtual Media Library (Sect 5.8)
Consumer/ Component ^a	Online 1	Near-Online 1	Removable 1	Virtual 1
JBOD	Online 1.5			
Low-end	Online 2	Near-Online 2	Removable 2	Virtual 2
Mid-range	Online 3	Near-Online 3	Removable 3	Virtual 3
	Online 4			
High-end	Online 5	Near-Online 5	Removable 5	Virtual 5
Mainframe	Online 6	Near-Online 6	Removable 6	Virtual 6
^a Entries in this level of taxonomy include both consumer products and data-center components (e.g., stand-alone tape drives)				

Table 2. Proposed JBOD-level addition to Taxonomy

JBOD (abbreviation of “just a bunch of disks/drives”) is a “simple” JBOD, an architecture that does not have storage protection capabilities built in. The disks in a JBOD array can function as their own individual volumes or can be connected or *spanned*, to form a single logical volume. A spanned volume provides no redundancy, so failure of a single hard drive amounts to failure of the whole logical volume. Redundancy for resilience and/or bandwidth improvement may be provided at a higher level.

The taxonomy developed when the SNIA Emerald™ specification was originally released was developed based on the existing technology at that time (2011) and was adopted by ENERGY STAR for datacenter storage in 2013. At that time, Solid State Storage (Flash based memory devices) for Enterprise Storage Systems mainly consisted of Solid State Disks (SSDs) and were too expensive to be used as main storage. They were primarily used as cache. Since that time, SSDs have grown in capacity and their prices have come down to where there are many SSD only systems. In addition, there are now many newer systems that are using Flash based memory devices as internal storage for persistent data storage. These devices are either directly connected to the memory busses or reside on PCI cards within the storage controller.

The above changes to the storage architecture are the motivation for an expanded taxonomy. SSDs have different performance characteristics than conventional magnetic and optical disk drives. Flash based memory devices using either PCI or direct connection to the processor memory buss represent a new connectivity and performance characteristic that was not envisioned in the original taxonomy. Due to these changes, SNIA has developed an expanded taxonomy to better deal with these new technologies.

INTRODUCING STORAGE SETS

The existing Emerald taxonomy defines four broad storage categories: Online, Near-Online, Removable Media Library, and Virtual Media Library. In order to accommodate the newer storage systems of an expanded taxonomy, the general term “Set” is adopted to reference the type of storage device used in these systems. Thus, Storage Sets now encompass both existing Disk and Removable Media Sets, and the new NVSS Set. The Disk Set consists of spinning devices found in all Online and Near-Online systems, and the Removable & Virtual Media Set consists of tape devices found in all Removable and Virtual Media Library systems. The category levels are unchanged from the existing taxonomy, see Figure 1. The new NVSS Set is introduced to include all storage systems based on non-volatile Solid State Storage. These storage systems may support either fabric or memory access.

Category levels for the NVSS Set remain similar, see Figure 2.

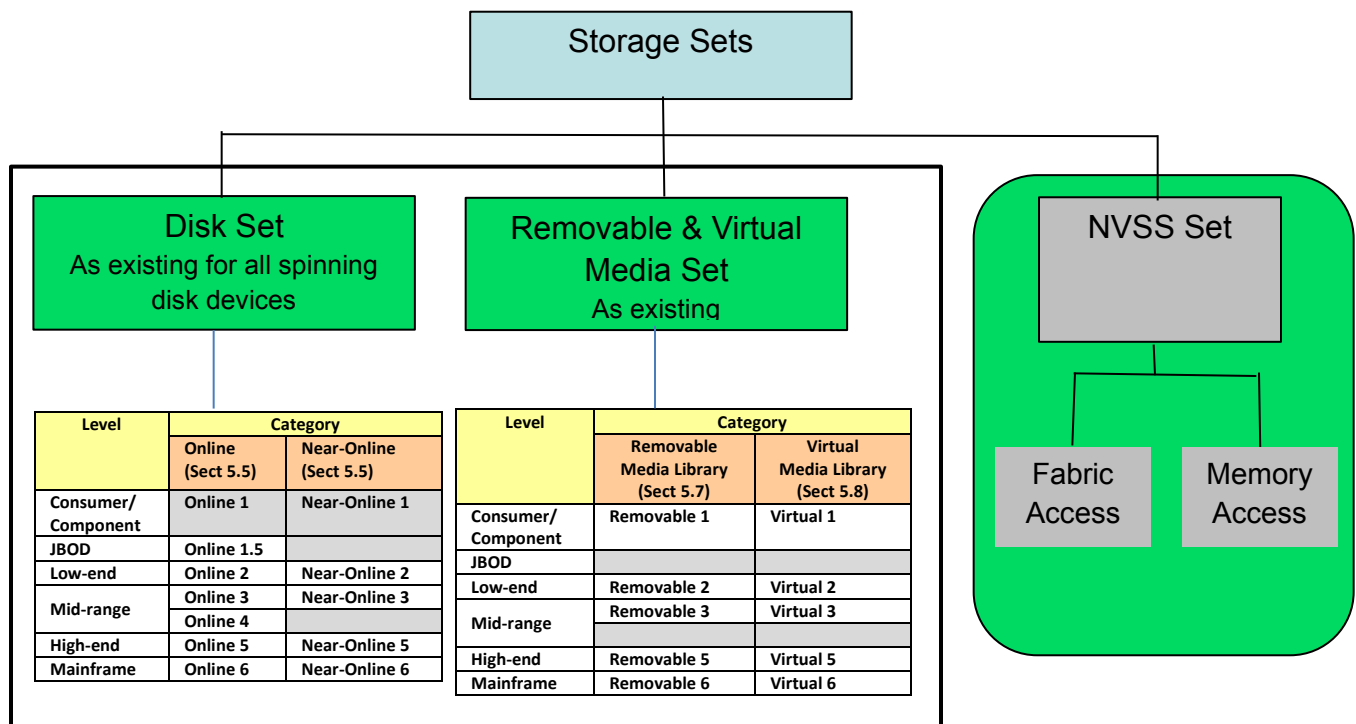


Figure 1. Introducing Storage Sets

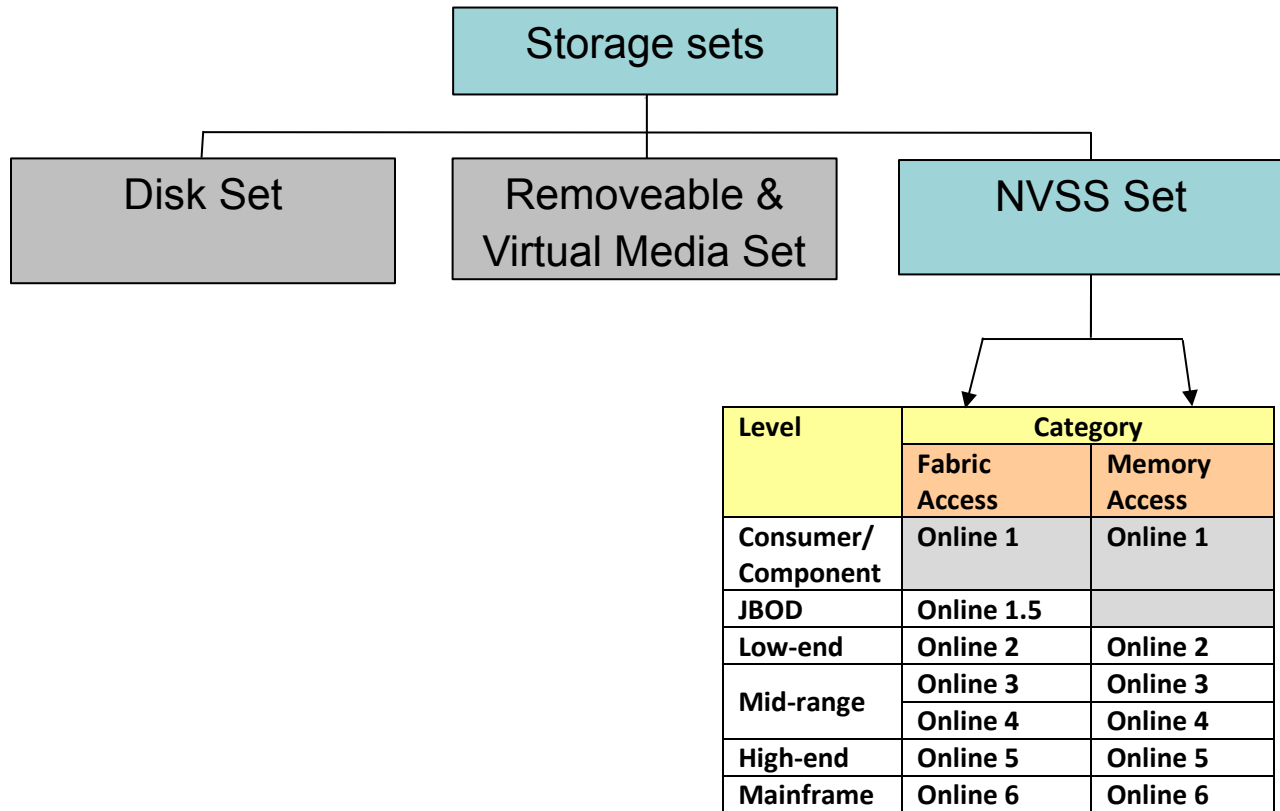


Figure 2. NVSS Set

Since the physical devices associated with non-volatile Solid State Storage-based storage system products can be very similar, the new version of the taxonomy has the location, connectivity and data transfer primitives of the devices as the differentiator. Depicted below (Figure 3) are two alternate implementations of products. The vertical column on the left illustrates products implemented within a network connected array. The workload for this configuration is applied using the traditional IO (fabric access) primitives such as open, read and write. The horizontal row to the right illustrates PM products installed on the memory bus of a server. In this case, memory access primitives, such as load and store, are used to transfer data to and from the devices.

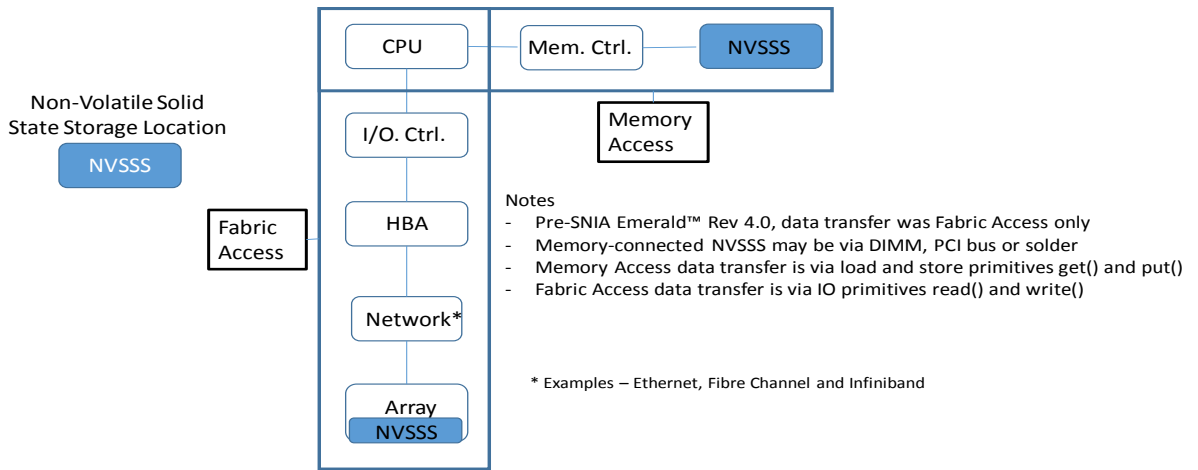


Figure 3. Alternate CPU – IO Path Illustrations

TAXONOMY CLASSIFICATION TABLES AND ATTRIBUTE DEFINITIONS

Table Modifications and Additions

Attribute	Classification					
	Online 1	Online 2	Online 3	Online 4	Online 5	Online 6
Access Pattern	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms
Connectivity	Not Specified	Direct-connected to single or multiple hosts	Network-connected	Network-connected	Network-connected	Network-connected
Consumer/ Component	Yes	No	No	No	No	No
Integrated Storage Controller	Optional	Optional	Required	Required	Required	Required
Storage Protection	Optional	Optional	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Required	Required	Required
Stable storage support	Optional, unless Required by protocol	Optional, unless Required by protocol	Required	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Required	Required
FBA/CKD Support	Optional	Optional	Optional	Optional	Optional	Required
Maximum Supported Configuration ^a	≥ 1	≥ 4	≥ 12	> 100	> 400	> 400
a. Maximum Supported Configuration does not apply to an all solid-state system that is not based on replaceable storage devices.						

Table 3. Existing Online Category

Based on the discussion above, the above classification table (Table 3) is being modified to become the Disk Set Online Classification Table (Table 4) by adding the Online 1.5 Classification and redefining the attribute requirements to reflect the discussed changes. Online 2 will add network connected to the connectivity options and require an integrated storage controller. This will bring Online 2 into alignment with the requirements in Energy Star for Data Center Storage. The Online 1.5 Classification will have fewer required attributes to reflect the definition of the simple JBOD above.

Attribute	Classification						
	Online 1	Online 1.5	Online 2	Online 3	Online 4	Online 5	Online 6
Access Pattern	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms
Multi-host Shareability	Not Specified	Ability to share with 2 or less hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts
Consumer/ Component	Yes	No	No	No	No	No	No
Integrated Storage Controller	Optional	NA	Required	Required	Required	Required	Required
Storage Protection	Optional	Not integrated	Required	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Optional	Required	Required	Required
Stable storage support	Optional, unless Required by protocol	Not integrated	Optional, unless Required by protocol	Required	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Optional	Required	Required
FBA/CKD Support	Optional	Not integrated	Optional	Optional	Optional	Optional	Required
System Capacity	≥ 1	≥ 4	≥ 4	≥ 12	> 100	> 400	> 400

Table 4. Modified table for the Disk Set Online Category

The classification table for NVSS Fabric Access (Table 5) is the same as the above Disk Set Online Classification table with the exception of the System Capacity attribute. The system capacity for NVSS Set Fabric Access is a much more complicated issue that it was for the Disk Set. Drives used in NVSS Set Fabric Access storage systems have a large range of capacities – as much as two orders of magnitude. A system with 12 high capacity SSDs can have a higher capacity than a system with over 3000 low capacity SSDs. This topic will be discussed later in this paper when we define the System Capacity attribute for the NVSS Set.

Attribute	Classification						
	Online 1	Online 1.5	Online 2	Online 3	Online 4	Online 5	Online 6
Access Pattern	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms
Multi-host Shareability	Not Specified	Ability to share with 2 or less hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts
Consumer/ Component	Yes	No	No	No	No	No	No
Integrated Storage Controller	Optional	NA	Required	Required	Required	Required	Required
Storage Protection	Optional	Not integrated	Required	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Optional	Required	Required	Required
Stable storage support	Optional, unless Required by protocol	Not integrated	Optional, unless Required by protocol	Required	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Optional	Required	Required
FBA/CKD Support	Optional	Not integrated	Optional	Optional	Optional	Optional	Required
System Capacity (make this a qualitative description to help guide)	"home"	"JBOD"	"very small"	"small"	"medium"	"large"	"mainframe"

Table 5. Draft table for NVSS Set Fabric Access

The Classification table for NVSS Set Memory Access (Table 6) is different from the NVSS Set Fabric Access in several ways. There is no Online 1.5 classification for Memory Access as Memory Access requires direct connection with the memory bus of the controller. There is no need for an integrated controller attribute as that is not optional for memory access as stated above. System Capacity will be discussed later when the System Capacity attribute is defined for the NVSS Set.

Attribute	Classification					
	Online 1	Online 2	Online 3	Online 4	Online 5	Online 6
Access Pattern	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms
Multi-host Shareability	Not Specified	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts	Ability to share with 2 or more hosts
Consumer/ Component	Yes	No	No	No	No	No
Storage Protection	Optional	Optional	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Required	Required	Required
Stable storage support	Optional, unless Required by protocol	Optional, unless Required by protocol	Required	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Required	Required
System Capacity (make this a qualitative description to help guide)	"home"	"very small"	"small"	"medium"	"large"	"mainframe"

Table 6. Draft table for NVSS Set Memory Access

Taxonomy Attribute Definitions

Storage Controller

A Storage Controller is a device for handling storage requests that includes a processor or sequencer programmed to autonomously process a substantial portion of IO requests directed to storage devices (Ref. SNIA Emerald™ Power Efficiency Measurement Specification Version 3.0.3).

Integrated Storage Controller

An Integrated Storage Controller is a Storage controller and storage media integrated into the same physical enclosure(s). May allow for expansion to additional enclosures.

Storage Protection

Storage Protection is the assurance that all completed IO operations will be preserved in the event of power loss or storage device failure. This assurance can be provided by a combination of hardware and/or software, e.g., RAID, NVRAM, disk sparing, background disk scrubbing, and/or background media scan.

System Capacity

System Capacity is a metric used to help discriminate between different systems. The goal of the metric is to differentiate between system sizes. The system sizes that we are attempting to identify are in

general the following; Consumer (Online 1), JBOD (Online 1.5), Entry level (Online 2), Mid-Range (Online 3 and Online 4), High End (Online 5), and Mainframe (Online 6). We further recognized that the Entry Level should be broken down into very small and small systems. In general, the very small systems were thought to have very little expandability whereas the small systems could have limited expandability.

For Hard Drive systems, we decided to use the number of disk drives a system could support as the means of identifying these separations. Consumer systems are systems that are purchased for home or very small office applications and usually contain a small number of disk drives. They can be direct connected or connected via a network, but are not used in Data Center applications.

JBOD devices are a single shelf of disk drives. A single shelf can be as small as a 1U shelf with 4 drives up to a multi-U shelf with a large number of drives.

Very small entry level systems (Online 2) tend to be a single shelf with one or two integrated controllers and limited expandability. They may have as few as 4 drives and may not support more than 4 drives, Small entry level systems generally have the ability to support more drives than the very small entry level systems and can generally support some level of expansion. They are generally able to support 12 or more drives in their maximum configuration.

Mid-Range (Online 3 and Online 4) systems are larger systems that can support a higher level of expansion with the ability to support 100 drives or more in their maximum configuration.

High End (Online 5) and Mainframe (Online 6) systems are very large systems that can support large numbers of drives. They can support 400 drives or more in their maximum configuration.

We did not specify the maximum number of drives for any of these system differentiators, as it is possible for them to support large numbers of drives. The other attributes of the classifications may force a system that would be considered as large (Disk Set Online 5) to be placed into a lower classification due to lack of support for one or more of the attributes required for Disk Set Online 5. Similar for Disk Set Online 4 and Disk Set Online 3. If the system does not support all of the attributes of a given classification, it must be placed in a classification that its features support. In general, a system should be placed in the highest classification that its features support.

System capacity is a much harder attribute to define for NVSS Set systems. Persistent Solid State Storage devices have a wide range of capacities. This is especially true for the devices (Solid State Drives, SSDs) used in NVSS Set Fabric Attach storage systems. SSDs currently range in capacity from a few hundred megabytes (200 MB) to many terabytes (63 TB or higher). A system with 12 63 TB devices has a higher capacity than a system with 3000 200 MB devices. An NVSS Set Memory Attach storage system is even harder to classify as the capacity is harder to differentiate by any unit other than actual capacity. The problem with using actual capacity is that it becomes a moving target. Capacity increases per device every few months, so setting a capacity limit today would be obsolete in a couple of years and setting the capacity limit based on projection would not apply to current systems. Therefore, a different way of specifying System Capacity for NVSS Set storage systems is needed.

Based on the above considerations, SNIA plans to replace the quantitative measurement used for Hard Drive systems with a qualitative description for NVSS Set storage systems. For the NVSS Set Fabric Attach Classifications, we provide the following qualitative descriptions:

- JBOD: A single SSD drive enclosure with no storage protection integrated into the enclosure and limited to one or two hosts.
- Very small: An entry level system with relatively small capacity and limited expandability.
- Small: An entry level system with relatively higher capacity and more expandability than the very small system.
- Medium: A higher capacity system that may have higher expandability than the small capacity system and with significantly more capabilities than the small systems.
- Large: A system with very high capacity and greater expandability compared to the medium and small systems intended for large Enterprises.
- Mainframe: A large system intended for Mainframe applications.

For the NVSS Set Memory Attach Classifications, we provide the following qualitative descriptions:

- Very small: An entry level system with relatively small capacity.
- Small: An entry level system with relatively higher capacity than the very small system
- Medium: A higher capacity system than the small capacity system and with significantly more capabilities than the small systems.
- Large: A system with very high capacity compared to the medium and small systems intended for large Enterprises.
- Mainframe: A large system intended for Mainframe applications.

SUMMARY

The Emerald taxonomy has been updated and expanded to better encompass newer storage systems. An Online 1.5 has been added to designate a simple JBOD, and to differentiate from Online 2 which remains a low-end storage system having storage protection. Solid State Storage (e.g., flash-based memory devices) have become mainstream as persistent data storage. The taxonomy has been expanded to include the NVSS Set in addition to the Disk Set and Removeable & Virtual Media Set. The NVSS Set can have connectivity as either Fabric Access or Memory Access. Due to the very large span of capacities, and frequent drive capacity increases, the System Capacity attribute for non-volatile Solid State Storage systems is now being defined by qualitative guidelines such as “small” or “large.” Also, the taxonomy has provided definitions for terms that needed clarification.

FOR MORE INFORMATION

For more information on the SNIA Emerald™ program, visit www.sniaemerald.com and www.snia.org/forums/green.

The SNIA Emerald™ Power Efficiency Measurement Specification Version 3.0.3 is available at www.sniaemerald.com/download.

For information on the USA EPA ENERGY STAR® Data Center Storage program, visit www.energystar.gov/products/data_center_equipment/data_center_storage/partners.

For information on the SNIA (Storage Networking Industry Association), visit www.snia.org.