



Flash Memory Summit

How Much Power Does Your Storage System or Device Consume?

Wayne M. Adams
Chairman Emeritus
Co-chair Green Storage Initiative
SNIA
www.snia.org



SNIA Emerald™

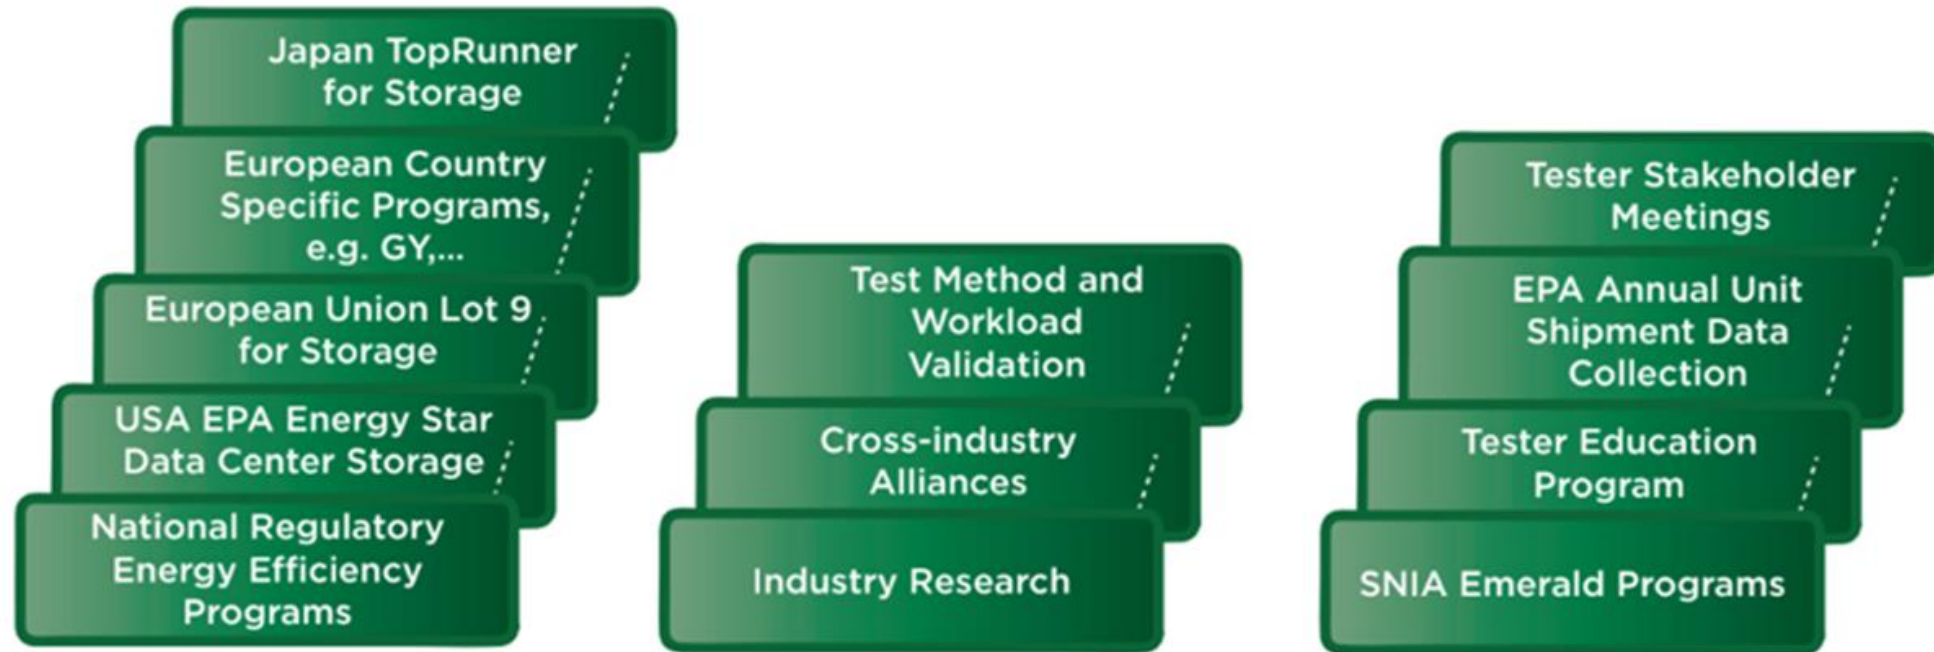
Brief Agenda of Topics

- Storage system power efficiency characterization with SNIA Emerald
 - Industry data with EPA Energy Star Data Center Storage Program
 - Workloads and workload generator evolution
- Storage device characterization with Calypso Test Platform
 - Workload capture programs and test methods
- Opportunity, challenges, and planning points for IT
- Additional Resources

Overview of SNIA Emerald™



Flash Memory Summit

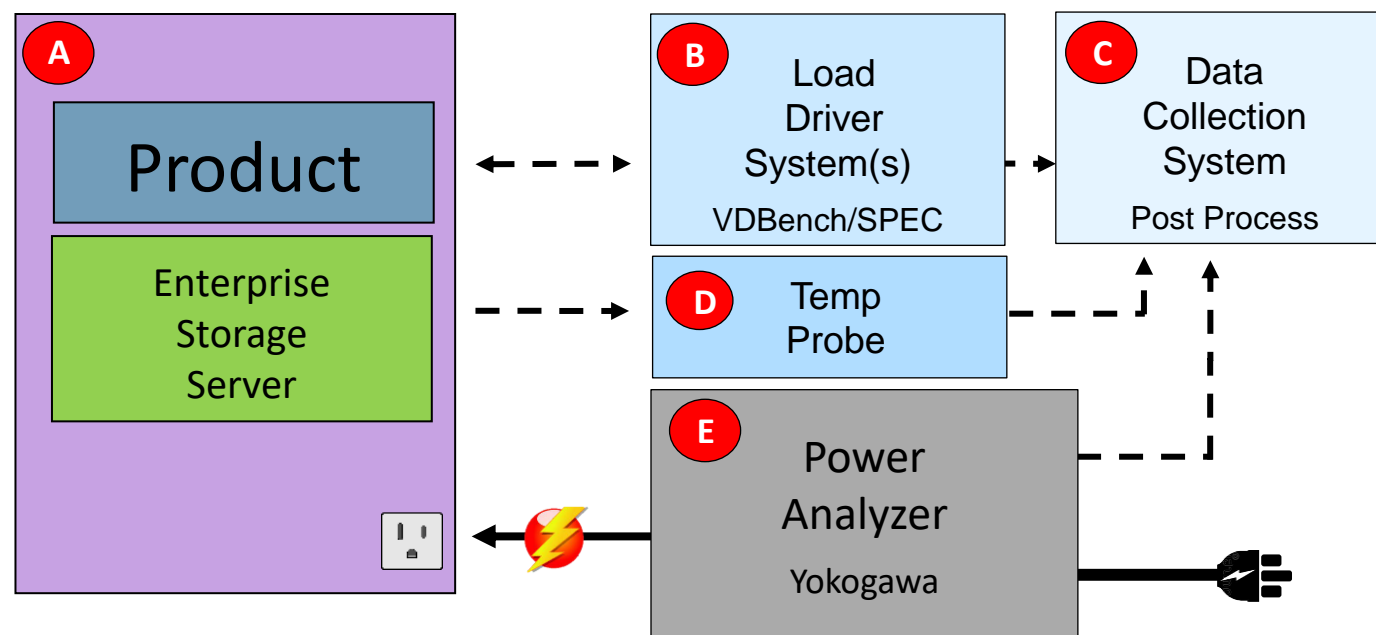


SNIA Emerald™ Power Efficiency Measurement Specification
for Enterprise Storage Systems
ISO/IEC 24091 – Information Technology – Power Efficiency
Measurement Specification for Data Center Storage

SNIA Emerald™ & IEC/ISO 24091:2019

US EPA Energy Star, EU Lot 9, and Japan TopRunner Referenced

- Standardized Methodology to Assess Power Efficiency
- Commercial Storage Products (Systems/Servers)
- Idle Power States: Idle, Active
- IO Power States: VdBench Block Synthetic Workloads (RND/SEQ 8K, 128K RW)
- IO Power States: File System SPEC SFS 2014 Real World Workloads
- Power Measured at the Product Server AC Input
- System considerations for Capacity Optimization and RAS



E Yokogawa WT 210 Power Analyzer

Detailed Taxonomy



Flash Memory Summit

Attribute	Set					
	Disk		RVML		NVSS	
	Category					
	Online	Near-Online	Removable Media Library	Virtual Media Library	Disk Access	Memory Access
Access Pattern	Random/ Sequential	Random/ Sequential	Sequential	Sequential	Random/ Sequential	Random
MaxTTFD	≤ 80 ms	> 80 ms	≤ 5 min	≤ 80 ms	≤ 80 ms	≤ 80 ms
Media Type	Magnetic disk	Magnetic disk	Magnetic tape, optical disk	Magnetic disk, Solid State Storage	Solid State Storage + optional magnetic disk ^a	Solid State Storage
Access Paradigm	Block, File, Object	Block, File, Object	Block	Block	Block, File, Object	Memory
^a Allows a purely Solid State Storage system or a hybrid Solid State Storage and magnetic disk system.						

Attribute	Category: Disk Set Online						
	Classification: Online						
	1 ^a	1.5	2	3	4	5	6
Multi-host Shareability	Not Specified	Ability to share with 1 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	Ability to share with 2 or more hosts
Consumer/Component	Yes	No	No	No	No	No	No
Storage Controller	Optional	Optional	Required	Required	Required	Required	Required
Storage Protection	Optional	No	Required	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Optional	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Optional	Required	Required
FBA/CKD Support	Optional	No	Optional	Optional	Optional	Optional	Required
System Capacity (number of drives) ^b	≥ 1	≥ 4	≥ 4	≥ 12	> 100	> 400	> 400

EPA Energy Star Testing – since July 2021

DCS V2.0 – Threshold for Block; not File yet

energystar.gov/productfinder/product/certified-data-center-storage/results

Fox News - Breakin... DW The Daily Wire BOS Boston.com Worcester - massliv... Speedtest by Ookla... Shared Streets

☐ Disk Set Online 2 (20)
☐ Disk Set Online 3 (6)
☐ Disk Set Online 4 (5)
☐ NVSS Disk Set Online 2 (8)
☐ NVSS Disk Set Online 3 (1)
☐ NVSS Disk Set Online 4 (16)
☒ Do not filter

Product Type: NVSS Disk Set Online 4
Storage Controller Configuration: Scale-Out Storage Workload Optimization Type
Storage Model Connectivity: File I/O
Capacity Optimized Method Available (COMs):
Delta Snapshots, Thin Provisioning, Compression, Data Deduplication

CLICK FOR PRODUCT DETAILS

Brand Name

☐ DELL (14)
☐ DELL EMC (5)
☐ Dell EMC Unity
☐ Dell EMC Unity
☐ Dell EMC Unity
☐ Dell EMC Unity
☐ Dell EMC Unity
☐ Hewlett Packard
[Show more](#)

Storage Controller

☐ Scale-Out Storage
☐ Scale-Up Storage (43)

Storage Controller Configuration

☐ Scale-Out Storage (13)
☐ Scale-Up Storage (43)

Workload Optimization Type

☐ Composite (11)
☐ Streaming (18)
☐ Transaction (27)

Storage Model Connectivity

☐ Block I/O (43)
☐ File I/O (13)

Capacity Optimized Method Available (COMs)

☐ Thin Provisioning (48)
☐ Data Deduplication (29)
☐ Compression (29)
☐ Delta Snapshots (23)

Markets

☒ United States (56)
☐ Canada (44)

Brand Name

D

☐ DELL (14)
☐ DELL EMC (5)
☐ Dell EMC Unity XT380 (1)
☐ Dell EMC Unity XT380F (1)
☐ Dell EMC Unity XT480 (1)
☐ Dell EMC Unity XT480F (1)

I

☐ IBM (6)

L

☐ Lenovo (2)

S

☐ Seagate (5)

V

☐ Veritas Technologies LLC (1)
☐ Viking Enterprise Solutions (1)

H

☐ Hewlett Packard Enterprise (3)
☐ HPE (4)
☐ HPE Primera (3)

N

☐ NetApp (5)
☐ NetApp, Inc. (3)

Capacity Optimized Method Available (COMs):

7 2022 56 test reports (incl 2021 reports)
7 2021 33 test reports (DCS V2.0)

7 2020 224 test reports (DCS V1.0/V1.1)
9 2017 163 test reports

V1.0 had no threshold, spanned 4+ years of data test reports to establish fair Block IO thresholds in V2.0)

Vendors test in-house or contract with independent test lab. MET

Comprehensive Data Analysis in Joint Whitepaper from SNIA/TGG:
ENERGY EFFICIENT DATA CENTER STORAGE: AN ASSESSMENT OF STORAGE PRODUCT POWER EFFICIENCY



Directions for SNIA Emerald 5.0

Evolve the Test Tools


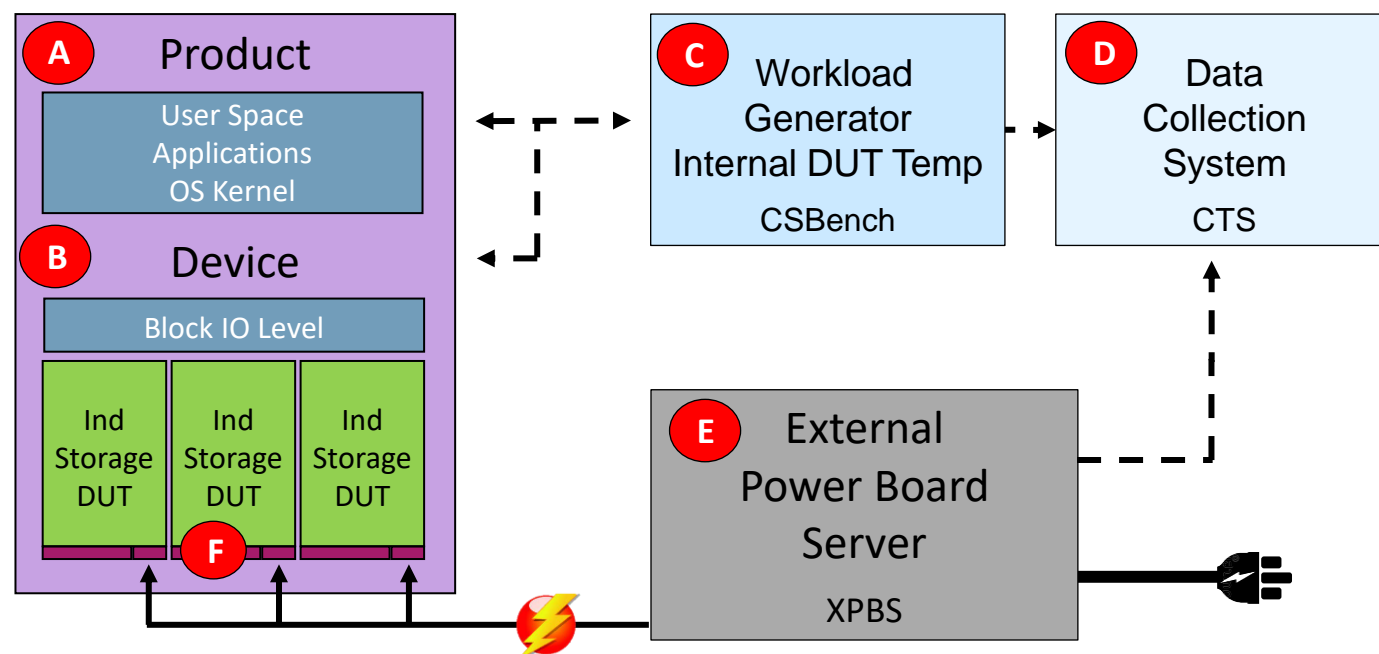
- SPEC SFS 2014 → SPECstorage Solution 2020
 - New workloads
 - Retain 2 workloads
- VdBench → Calypso Test Suite
 - Run the same VdBench scripts for the 4 corner tests
 - Collect and review enterprise data center real time work IO streams
- Work with EPA (for DCS v2.x /3.0)
 - assess whether Block IO thresholds need adjusted to correlate the old workload generators and the new generators
 - assess if enough industry data for File IO to establish thresholds

SNIA Device Level Power Efficiency Measurement

Under Development by SNIA Solid State & Green Storage TWGs



- Standardized Device Level Power Efficiency Measurement
- Commercial Storage Devices (SSD/HDD)
- Idle Power States: Cold, Warm, Active and Cool Down
- IO Power States: Synthetic & Real-World Workloads
- File System and Block level
- Power Measurement at the Individual Storage Device



E Calypso XPBS – External Power Board Server




Storage Device Level Power Efficiency is Highly Dependent on:

1

SSD vs HDD

Temp, Power
IO Stream Dependent

Perf, Temp, Power
Idle v Active Dependent



1. Type of Storage: SSD and HDD
2. Type of Workload: Real World and Synthetic
3. Power Efficiency Measurement: IOPS/W
4. RAS: Features & Settings
(Reliability, Availability, Serviceability), 5 9s Quality of Service

2

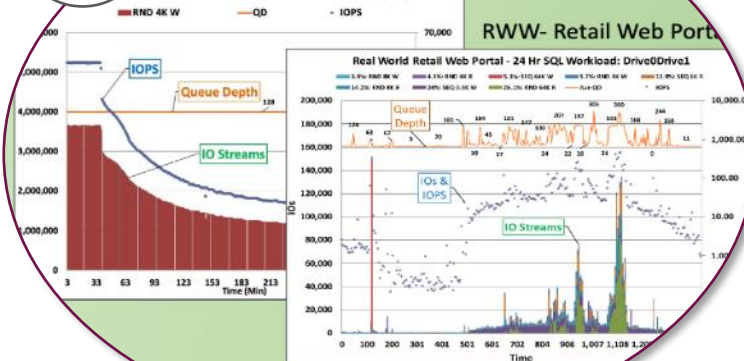
Synthetic - RND 4K W

RND 4K W Benchmark: IOPS IO's v Time - Fixed QD (128)

■ RND 4K W ■ QD ■ IOPS


RWW- Retail Web Portal

Real World Retail Web Portal - 24 Hr SQL Workload: Drive0Drive1



3

Composite vs Individual IO Stream




4

RAS - Response Time QoS Histogram

P15 Max Point T6/Q1 CLP, IOPS=305906, 1195 MB/s, MRT=5.77 ms


■ Latency ■ % Confidence ■ ART = 0.02 ■ 99.9% = 0.02 ■ 99.99% = 0.2 ■ 99.999% = 0.64



Real World Workload Replay-Ind. Streams Test: Test Flow



Summit



Test Flow:

1. Beg (Cold) Idle (i¹) - 60 sec Dur
2. Ind. Stream Segments - (9) Ind. IO Streams x 60 sec
3. Warm Idle (i²) - After IO Streams x 60 sec
4. Replay Segment - 9 IO Stream Sequence x 290 5 sec steps
5. Active (Hot) Idle (i³) - 60 sec idle between Steps 60 & 61
6. Begin Cool Down Idle (i⁴) - 60 sec idle after step 290
7. Cool Down - 10 min following step 290
8. End Idle (i⁵) - 60 Sec idle after 10 min cool down

Directions for SNIA Storage Device Level Work

- **Test Specification can be of benefit for (apples to apples comparisons)**
 - System vendors supply chain in device selection and custom workloads
 - Cloud service vendors who roll their own systems with business/application specific workloads
 - Specification target completion is before year end 2022
- **Potential Regulator interest – longer term**
 - Japan TopRunner could evolve from idle only HDD; include SSD and active workload
 - Re-validate EPA Data Center Server Spec assumptions for storage device “adders” rules of thumb (currently based on HDD from several years ago)
- **Increase SNIA IO library with more real-world workloads for public use**
- **No plan for a device-level repository of test results, mostly due to the device-level program is focused on custom workloads**

Challenges and Opportunity for Operations Measurements



Flash Memory Summit

- Methods described characterize a system and or device as though it is in a production environment, however pre-deployment
- Current power measurement methods in operations tend to be focused on the rack, regardless of what is in the rack – e.g. server, network, storage, not tied to any workload (CPU, network, or storage)
- Though servers and storage are instrumented with DMTF Redfish and SNIA Swordfish management interfaces (including power and IO), no universal/open DCIM platform has embraced system level energy usage reporting.
- Hyper-converge systems considered to difficult to assess with server/network/storage, and SDS mix-n-match

A few planning points, some you may have tackled for those directly involved in IT and Business Ops

- Manage the size of your dataset(s);
- Leverage as many capacity optimization products and tools as possible;
- Matched to the requirements for data protection and response time; copies of data moved onto nearline or offline storage.
- **What is your data deletion strategy/policy to curb ballooning storage requirements...while AI/ML, IoT, and other growth drivers mature ?
- Future climate and emissions reporting for public corporations will need disclosure of your outsourced services, so you will need to work with your cloud vendor(s) of choice to receive energy usage reporting along with the resource usage if they are storing data for you.

Signup for Test Tool Demo Day and Stop by the SNIA Exhibit Area at FMS



Flash Memory Summit



SNIA Emerald™


**SNIA Emerald V5.0
Next Version
Test and Measurement
Tool Demonstrations**

SNIA: GSI | GREEN STORAGE

**SPECstorage Solution 2020
for System File IO**

**Calypso Test Platform and
Suite for System Block IO
and Device Level IO**

TestMyWorkload IO Capture



spec®

**CALYPSO
Systems**

Attention test and measurement engineers for

- SNIA Emerald System testing for EPA Energy Star
- Supply chain storage device procurement
- Cloud architects' custom hardware platforms

**Sign up for the SNIA Hosted end-of-September
Test and Measurement Tool Virtual Demo Day**

**Scan the QR Code
or visit www.sniaemerald.com**

**With your signup, you'll be provided more
details for the date and time, along with
participation logistics**



Additional Resources

- SNIA Emerald™ Specification, Test Kit, and Training
 - www.sniaemerald.com , emerald@snia.org
- IEC/ISO 29041:2019
 - <https://www.iso.org/standard/77801.html>
- Joint TGG(TheGreenGrid®) / SNIA Whitepaper #86:
 - ENERGY EFFICIENT DATA CENTER STORAGE:
AN ASSESSMENT OF STORAGE PRODUCT POWER EFFICIENCY
 - <https://www.thegreengrid.org/en/resources/library-and-tools>
- SNIA Whitepaper, estimated release - October 2022
 - Storage Device- Level Power Efficiency Measurement for
Cloud, Datacenter, and Enterprise Storage
 - https://www.snia.org/sites/default/files/SSSI/Introduction%20to%20RWW%20-%20A%20Primer%20v2.1.2.Final_.pdf
- SNIA CMSI White Paper: Real World Workloads – A Primer
 - <https://www.snia.org/educational-library>
- SNIA IO Traces and Workload Library
 - <http://iota.snia.org/>



Thank You

Time for Q&A

www.sniamerald.com

emerald@snia.org

Wayne M. Adams