Storage Networking Industry Association
Technical White Paper

Storage Device-Level Power Efficiency Measurement for Cloud, Datacenter & Enterprise Storage

October 27, 2022

Abstract: Data storage power efficiency measurements are essential for evaluating power, performance and storage operating costs. Data center operating costs include compute, network, storage, and facility-side systems and resources. The ISO/IEC 24091:2019 IT Specification for Data Center Storage (ISO 24091) Standard is based on the SNIA Emerald™ Power Efficiency Measurement Specification V3.0 created by the SNIA Green Storage Technical Work Group (Green Storage TWG) for measuring Enterprise **storage system-level** power efficiency. The Green Storage TWG is now developing a complementary **storage device-level** specification for a standardized, easy-to-use, cost-effective method for power efficiency measurement.

This White Paper describes storage device-level (individual storage drive) power efficiency test methodologies, the device test harness, the use of workloads, calculating metrics and reporting device test results. The Green Storage TWG is seeking industry regulatory bodies as well as technical and supply chain stakeholders to participate in the further development of the device-level specification with the contribution of technical resources, short-term loan of in-production storage devices and IO capture of real world enterprise application workloads. Interested parties should contact the Green Storage TWG at emerald@snia.org.

USAGE

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/.

Table of Contents

Executive Summary

The SNIA Emerald™ Power Efficiency Measurement Specification is a building block specification for industry regulatory bodies, international standards and IT storage professionals created by the SNIA Green Storage Technical Work Group (Green Storage TWG). SNIA Emerald V3.0 is embodied in ISO/IEC 24091:2019 IT Power Efficiency Measurement Specification for Data Center Storage (ISO 24091). ISO 24091 is referenced by regional regulatory bodies and programs including the European Union Lot 9 program and the Japan TopRunner program. SNIA Emerald V4.0 is referenced by the United States Environmental Protection Agency's Energy Star Program for Data Center Storage. The SNIA Emerald™ Power Efficiency Measurement Specification defines enterprise storage system-level test methodologies, a storage taxonomy, the test and measurement harness, system workloads and power efficiency metrics.

While the SNIA Emerald specification defines **storage system-level** power efficiency, no direct measurement of individual storage devices (drives) is provided or defined. The Green Storage TWG is now developing a specification for measuring the power efficiency of **storage devices.** The working name for this new specification used through this document is Storage Device-Level Power Efficiency Measurement (SDLPEM).

Storage device-level power efficiency tests allow IT professionals to derive how much power is consumed by individual storage devices. SDLPEM will define the test, the test and measurement harness, the metrics, the workloads and the test methodology to measure individual storage device power efficiency.

Storage power efficiency measurement can be made in one of two ways: 1. At the **device-level** - by connecting an external power board server (XPBS) test harness to individual storage device power input ports; and 2. At the **system-level** - by connecting an external power meter (XPM) to the storage server(s) power input(s). The XPBS and XPM collect power measurements during host Idle and while processing host IO commands and enables the reporting of power efficiency in I/O Operations per second (IOPS) per watt and MiB/s per watt.

SDLPEM provides easy-to-use, low cost, standardized tests and specifies how to use the XPBS test harness to measure individual storage device power efficiency**.** This allows the IT professional to understand the actual power efficiency impact of storage drive scaling for specific use cases and architectures instead of relying on system-level power and limited workload types. SDLPEM measurement reports will further augment the coarse-level device manufacturers' device specification sheets for fine-grained power budget planning.

SDLPEM power measurements are taken while applying synthetic workloads (such as Random 8 KiB Read/Write and Sequential 128 KiB Read/Write) and/or Real World Workloads at the storage device-level. Real World Workloads are captured on deployed storage architectures during actual application use. Real World Workloads are a collection of IO Streams observed at the file system-level and block-level and are referred to as SPECstorage™ Solution 2020 workloads and Real World Workload (RWW) IO captures. Each SPECstorage™ Solution 2020 workload and RWW IO capture can be used to automatically create portable test scripts based on the sequence and occurrence of the observed IO streams. Portable test scripts can be run on storage systems other than those on which the workload was captured.

Energy Consumption for Cloud, Data Center & Enterprise Storage

Data center energy consumption growth has been well documented since the early 2000's. A key publication is the 2016 US Data Center Energy Usage Report[1] (the Report) by Lawrence Berkeley National Laboratory sponsored by the US Department of Energy.

The Report includes historical and forecast data center energy consumption including breakdowns for server, storage, network and infrastructure equipment. According to the Report, while the growth rate of data center energy has declined, data centers constituted an estimated 1.8% of total U.S. electrical energy consumption in 2014 and data center electrical power consumption increased by about 4% from 2010-2014. The continued forecasted decrease in data center energy usage growth rate is attributed primarily to ongoing improvements in data center energy efficiency. Figure 1, taken from the Report, illustrates these findings.

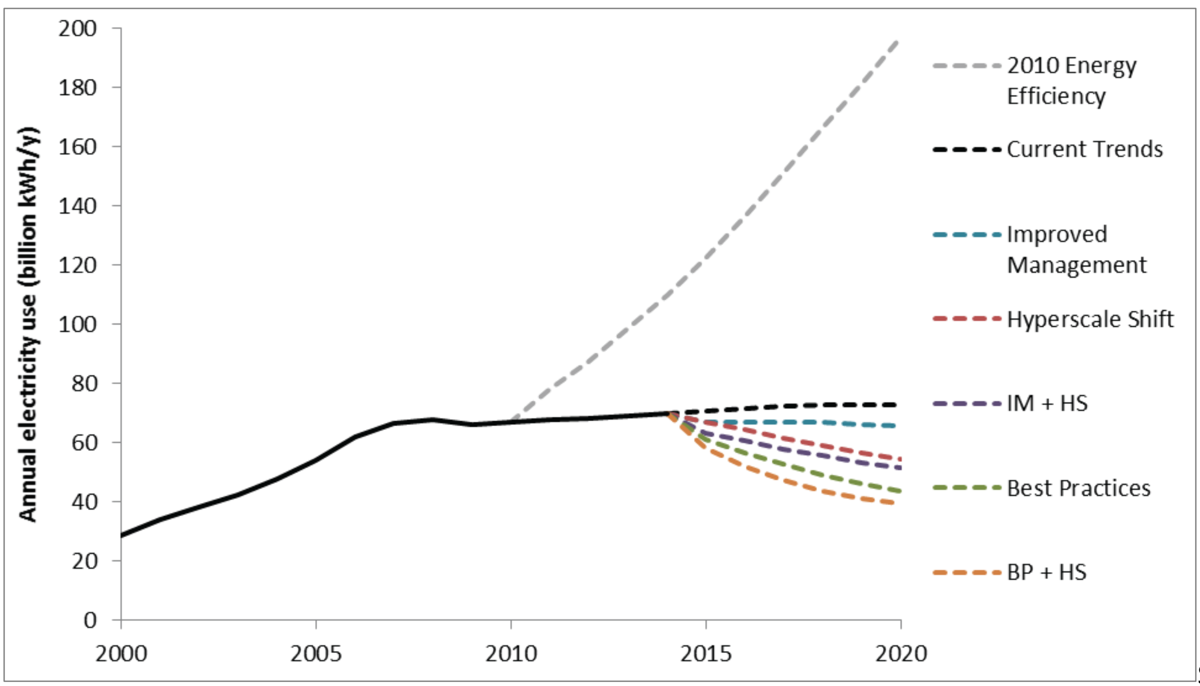

Se

Figure 1 – Projected Data Center Total Electrical Use – US Data Center Energy Usage Report, June 2016

The Report indicates that improvements in energy efficiency come from server, storage, network and infrastructure equipment improvements.

Efficiency improvements are expected to further optimize power consumption, e.g., the Report states "*Storage devices are becoming more efficient on a per-drive basis, with growth in drive storage capacity projected to outpace increases in data storage demand by 2020, ultimately reducing the number of physical drives needed throughout data centers.*"

[1] United States Data Center Energy Usage Report, June 2016 - Lawrence Berkeley National Laboratory LBNL-1005775, www.osti.gov/biblio/1372902

The Report bases projections of data center (DC) storage system power consumption on estimates of the average power consumed by Hard Disk Drives (HDDs) and Solid State Drives (SSDs) of an estimated storage capacity based on selected workload types. SSDs were estimated to comprise 47% of the total 2020 data center installed storage drive count.

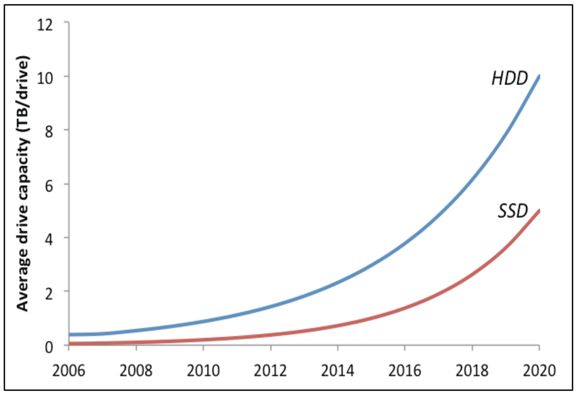

Figure 2 – Estimated Average Capacity DC Storage Drives

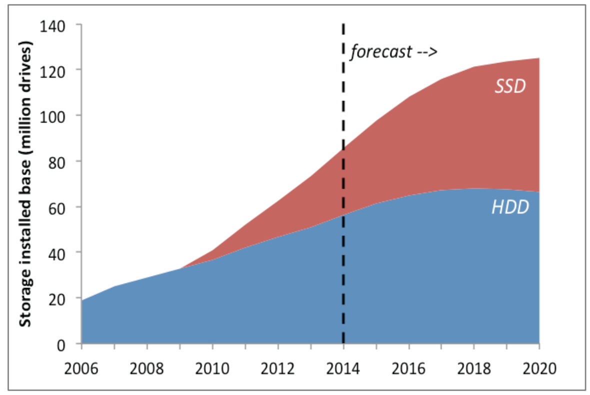

Figure 3 – Total US DC Installed Storage Drive Count

The Report estimates Seagate HDD power consumption to be 8.6 watts/drive in 2015 and 6.5 watts/drive in 2020. The Report assumes SSD power consumption to be 6 watts/drive.

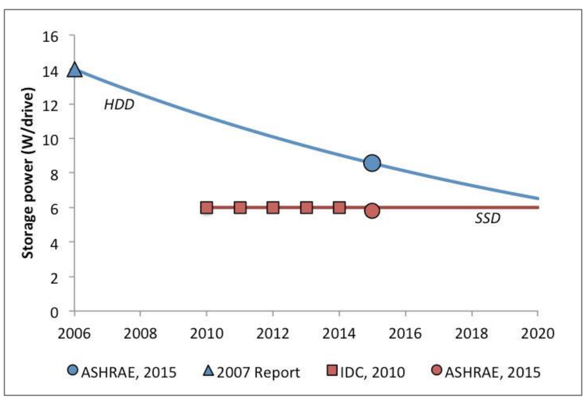

Figure 4 – Average Wattage US DC Storage Drives

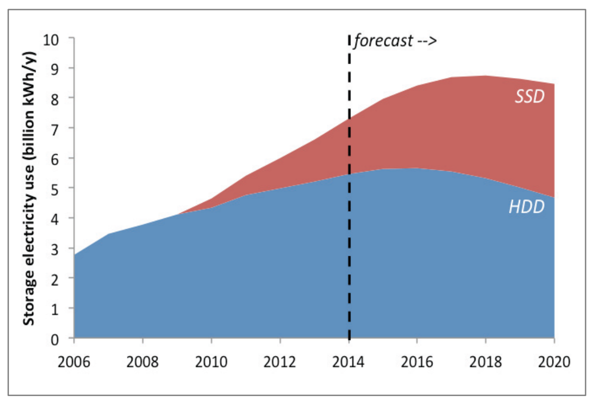

Figure 5 – Total US DC Storage Electricity Consumption

The Report states "*Electricity consumption for all data center storage is calculated as the product of the estimated installed base of drives and the assumed power consumption per drive. Additional operational consumption was then added to account for the controller and associated components required to operate external storage systems*".

Continued improvements in storage system power consumption depend upon meaningful measurements of power consumption to provide data to storage equipment vendors, IT professionals, and regulatory bodies. Actual storage device performance and energy consumption behavior is highly dependent upon precise workload, environmental and usage parameters that require actual measurement of storage devices to relevant use case workloads (see section *Preliminary SDLPEM Power Efficiency Data & Reports*).

Storage Device-Level Power Efficiency Measurement Testing

The proposed SDLPEM specification for power efficiency testing of storage devices envisions the use of synthetic tests used in the SNIA Emerald system-level power efficiency specification as well as new Real World Workload tests.

SNIA Emerald Synthetic Workload

The synthetic SNIA Emerald test applies a combination of access patterns across several LBA hot bands followed by the application of Random 8 KiB block size Writes and Reads and then Sequential 256 KiB block size Writes and Reads.

SNIA Real World Workloads (RWWs)

RWWs are emulations of the storage device IO operations that occur during real world application use (for example, data base queries, video capture and playback and virtual storefronts). An IO capture process is used to record this IO activity. RWWs may be captured at application-level, file system-level, storage system-level, or storage device-level. For the purposes of this white paper, the storage device-level is the capture point. The behavior recorded can later be replayed, thus creating a stimulus to correlate with that of the original running application. The concept streamlines the creation of defendable stimuli to expand the scope of workloads used for power efficiency measurements. For more information on Real World Workloads, see the SNIA CMSI White Paper "Introduction to Real World Workloads – A Primer."[2]

Currently, there are several reference RWWs that are available on the SNIA Compute, Memory, and Storage Initiative (CMSI) website[3]. This is the beginning of a collection of RWWs that may be incorporated as future workloads for power efficiency evaluation. SNIA Reference Real World Workloads[4] can be viewed, downloaded and replayed.

---

[2] SNIA CMSI, Introduction to Real Workloads – A Primer, www.snia.org/technology-focus-areas/physical-storage/nvme-ssd-classification/real-world-storage-workloads
[3] SNIA CMSI website: www.snia.org/forums/cmsi
[4] SNIA Reference Real World Workloads website: www.snia.org/technology-focus-areas/physical-storage/real-world-workloads/reference-real-world-workloads

ISO 24091 & SNIA Emerald Storage System-Level Test Specification

The SNIA Emerald Power Efficiency Measurement Specification[5] defines a standard method to assess the power efficiency of enterprise storage systems in both active and idle states of operation. A taxonomy is defined that classifies storage systems in terms of operational profiles and supported features. Test definition and execution rules for measuring the power efficiency are defined; these include test sequence, test configuration, instrumentation, benchmark driver, IO profiles, measurement interval and metric stability assessment. Qualitative heuristic tests are defined to verify the existence of several Capacity Optimization Methods. Resulting power efficiency metrics are defined as ratios of idle capacity or active operations during a selected stable measurement interval to the average measured power.

In greater detail, the specification defines methodologies and metrics for the evaluation of the related performance and energy consumption of storage systems in specific active and idle states. Storage systems and components are said to be in an “active” state when they are processing externally initiated, application-level requests for data transfer between host(s) and the storage system(s). For the purposes of the specification, idle is defined as “ready idle”, in which storage systems and components are configured, powered up, connected to one or more hosts and capable of satisfying externally initiated, application-level initiated IO requests within normal response time constraints, but with no such IO requests being submitted.

The SNIA Emerald system-level test methodology measures power at the AC power input to a storage system; the proposed storage device-level test methodology will measure DC power at the device power connector.

[5] SNIA Emerald™ Power Efficiency Measurement Specification, Version 4.0.0, July 3, 2020 www.snia.org/tech_activities/standards/curr_standards/emerald

Proposed SNIA Emerald Storage Device-Level Test Specification

The Green Storage TWG is developing a Storage Device-Level Power Efficiency Measurement (SDPLEM) specification. The SDLPEM is intended to closely follow the SNIA Emerald system-level specification. However, the SDLPEM will focus on the measurement of individual storage devices using both synthetic and Real World Workloads rather than system-level power efficiency measurement.

Under the SNIA Emerald system-level specification, power efficiency is measured for the complete storage system (defined as the storage server with storage) whereas the proposed SDLPEM measures power on the individual device power port.

The SNIA Emerald system-level specification requires a system-level power analyzer while the SDLPEM requires an External Power Board Server (XPBS). The SNIA Emerald power analyzer sits between house power and the storage server power plug (See Figure 7 below).

The SDLPEM XPBS has an independent power supply that provides individual 12 V / 5 V power lines to the storage device power input (See Figure 8 below).

SDLPEM also defines various synthetic and Real World Workloads to be run while individual storage device power is measured. Each test provides power efficiency measurements in IOPS/W and MiB/s/W.

A more detailed discussion of synthetic and Real World Workload test methodologies, platform set up, test flow, test settings, test measurements and reporting requirements are outside the scope of this paper and is expected to be described in subsequent SNIA white paper(s).

Test Harness Comparison Between System-Level and Device-Level Tests

The SNIA Emerald system-level and the storage device-level specifications require the use of specific test harnesses.

Figure 7 illustrates the test configuration for SNIA Emerald system-level power efficiency tests. The Power Analyzer connects the AC power input of a storage system to an AC power source.

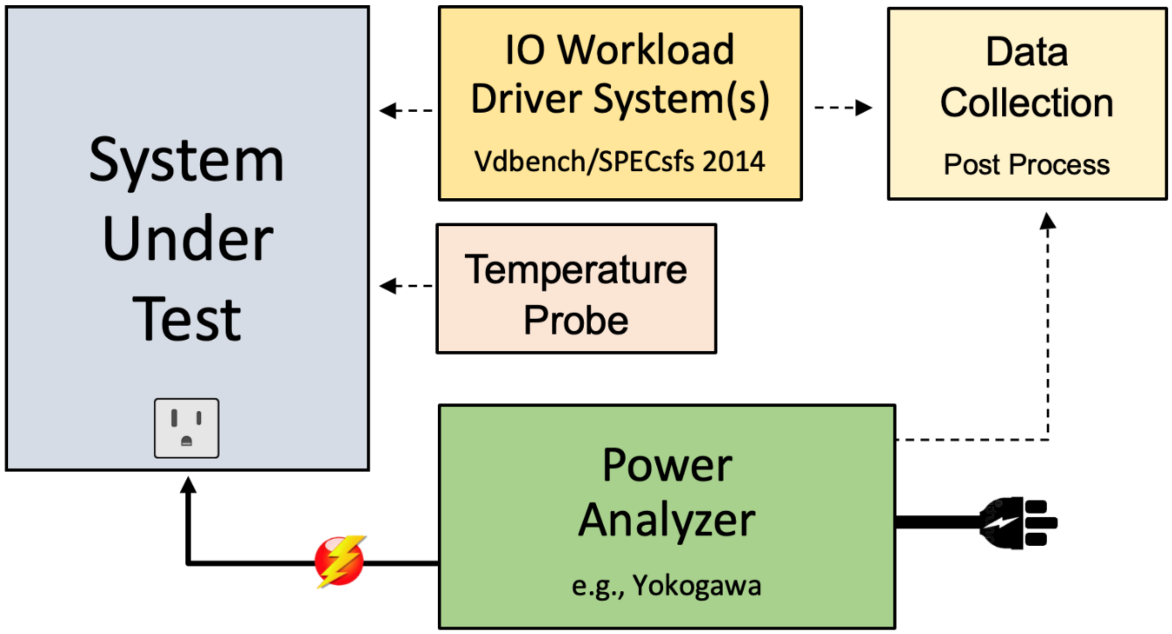

Figure 6 – System-Level Test Configuration

Figure 7 illustrates the test configuration for device-level power efficiency tests. The External Power Board Server (XPBS) provides power to the individual storage device power input port. This can be accomplished using a variety of connector types to match the storage device configuration (such a direct to SATA power port or direct to a storage device interposer board that is inserted in an appropriate motherboard bus slot).

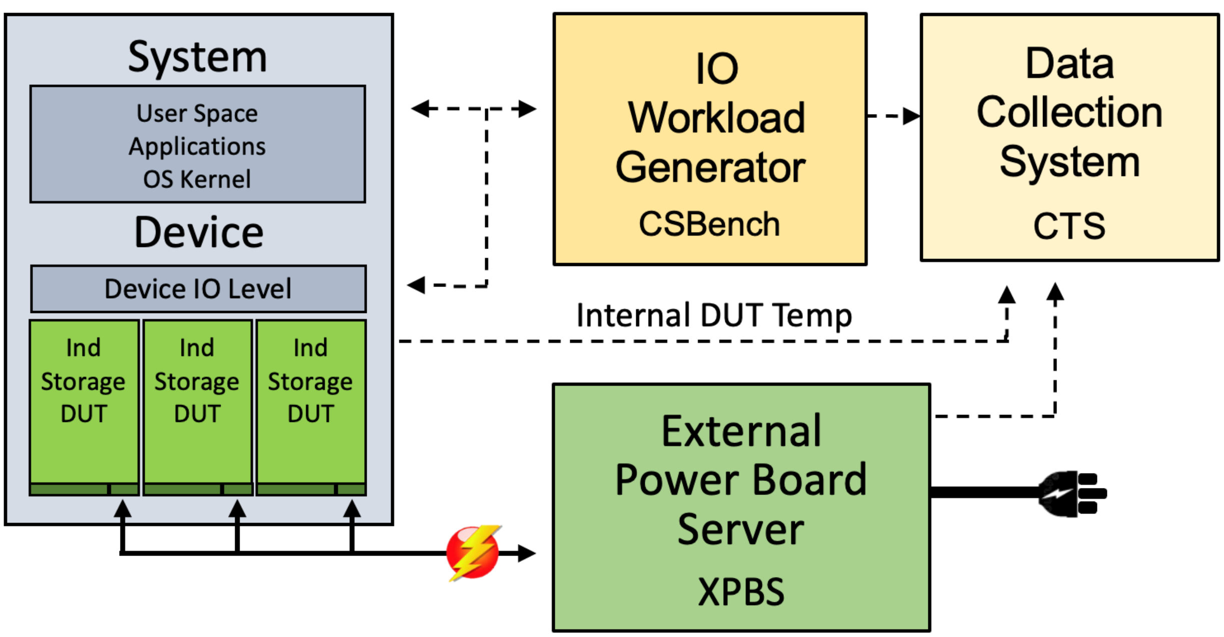

Figure 7 – Device-Level Test Configuration

Preliminary SDLPEM Power Efficiency Data & Reports

The following power efficiency data is based on Green Storage TWG and SNIA Solid State Storage TWG preliminary testing for both SAS HDD and NVMe SSD devices in 3 drive RAID 0 configurations.

The storage server was a SNIA Reference Test Platform 5.0[6] which is a PCIe Gen 3 motherboard, dual 2.1 GHz 8 core CPUs, 32 GB of 2400 MHz DDR4 ECC RAM, CentOS 7.0 OS and using Calypso Test Systems, Inc. Calypso Test Suite (CTS) test software with the CTS workload generator.

The target storage is set up as three drive RAID 0 LUNS using software RAID 0 (striped, N=3 (whole drives), block size = 64 KiB). The HDD RAID 0 LUN is comprised of 3 SAS 600 GB HDDs connected via a 12 Gb/s HBA. The SSD RAID 0 LUN is comprised of 3 NVMe 1920 GB SSDs via single drive interposer cards inserted in the motherboard PCIe Gen 3 bus slots.

The tests run are Green Storage TWG power efficiency tests that apply a range of Random/Sequential block sizes over a range of Read/Write mixes. Results are recorded for individual IO stream IOPS, MiB/s, average power in W and power efficiency in IOPS/W and MiB/s/W.

The following Figure 8 and Figure 9 present SAS 15 K RPM HDD data. Plots are shown for Random 8 KiB Read/Write and Sequential 128 KiB Read/Write for both 3 drive RAID 0 and Single HDD storage.

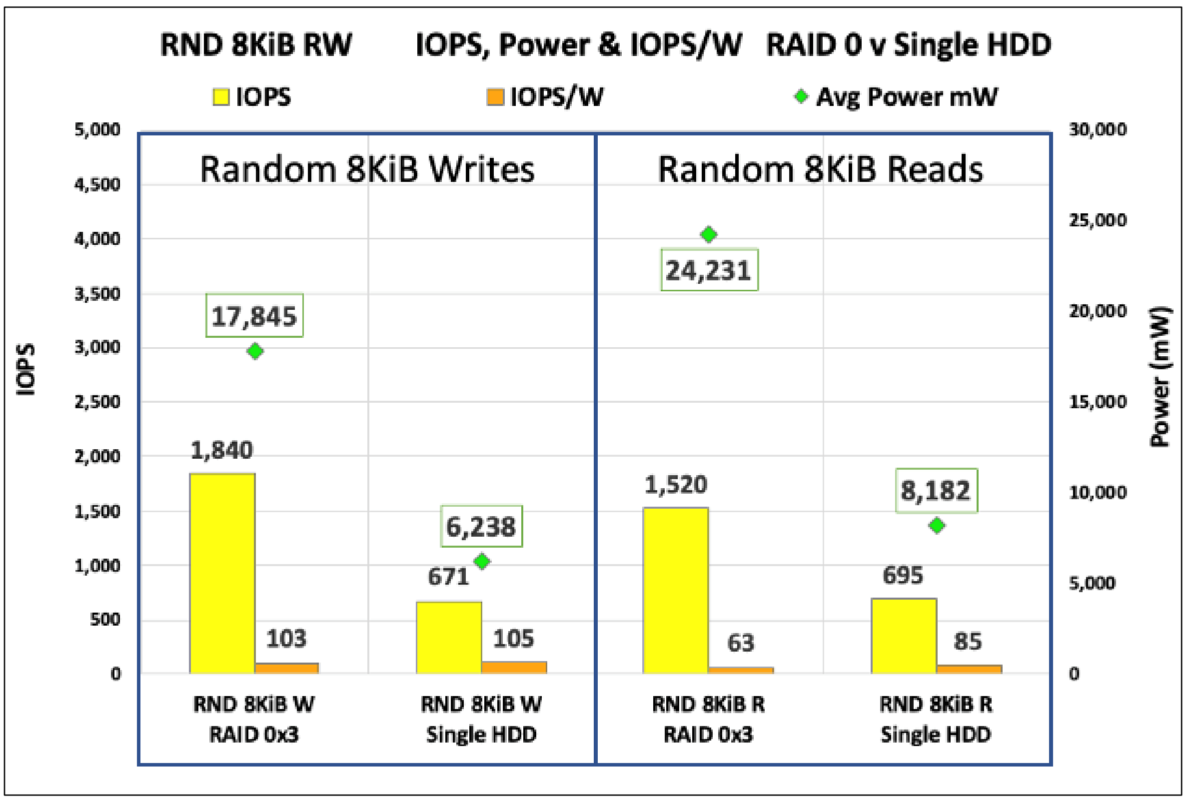

Figure 8 – Random 8 KiB Read/Write: IOPS, Power & IOPS/W

---

[6] Reference Test Platform, www.snia.org/forums/sssi/rtp

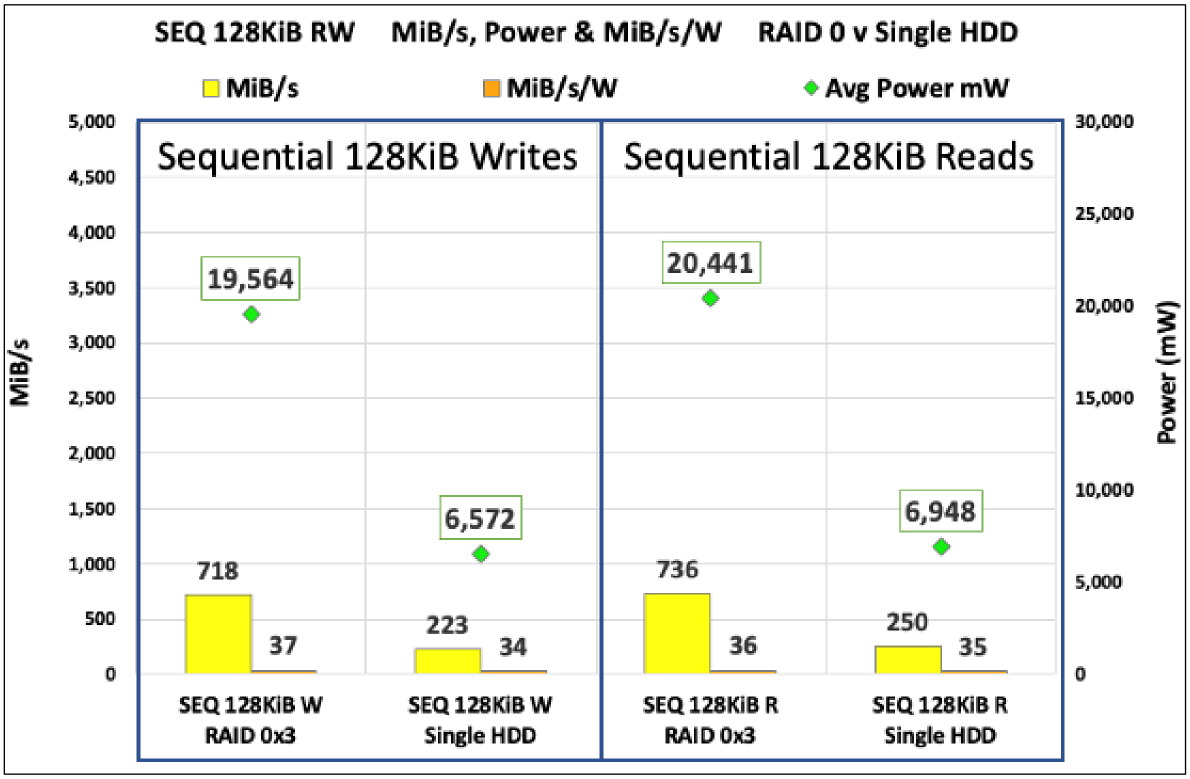

Figure 9 – Sequential 128 KiB Read/Write: MiB/s, Power & MiB/s/W

Conclusion

The SNIA Emerald™ Power Efficiency Measurement Specification is a building block specification for industry regulatory bodies, international standards, and IT storage professionals created by the SNIA Green Storage Technical Work Group (Green Storage TWG). SNIA Emerald V3.0 is embodied in ISO/IEC 24091:2019 IT Power Efficiency Measurement Specification for Data Center Storage (ISO 24091). While the SNIA Emerald specification defines **storage system-level** power efficiency, no direct measurement of individual **storage devices** (drives) is provided or defined.

The SNIA Green Storage Technical Work Group (Green Storage TWG), that developed the SNIA Emerald™ Power Efficiency Measurement Specification for measuring the power efficiency of storage systems, is developing a specification for measuring the power efficiency of **storage devices.** The working name for this new specification used through this document is **SDLPEM**.

Storage device-level power efficiency tests allow IT professionals to derive how much power is consumed by individual storage devices. SDLPEM defines the test, the test harness, workloads and test methodology to measure individual storage device power efficiency. Informed decisions can be made to optimize the power efficiency of storage systems based on the characteristics of the devices contained therein.

The main difference between the system-level and device-level test environments is the measurement boundary and equipment associated with obtaining power consumption metrics. It is intended that new workloads and workload generators will be introduced and be usable in future revisions of the storage system-level test specification while providing consistent system-level test results.

Lastly, the results of several IT reports provide the motivation behind pursuing the development of the specification, followed with the metrics and reports from the actual environment.

For More Information

For more information on the SNIA Emerald™ program, visit www.sniaemerald.com.

The SNIA Emerald™ Power Efficiency Measurement Specification Version 4.0.0 is available at www.sniaemerald.com/download.

The SNIA CMSI White Paper "Real World Workloads – A Primer" is available at www.snia.org/technology-focus-areas/physical-storage/nvme-ssd-classification/real-world-storage-workloads.

For information on the SNIA Green Storage Initiative (GSI), visit www.snia.org/forums/green.

For information on the SNIA Compute, Memory, and Storage Initiative (CMSI), visit www.snia.org/forums/cmsi.

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