



SNIA Emerald™ Power Efficiency Measurement Specification for Data Center Storage Devices

Version 0.0.49

ABSTRACT: This document describes a standardized method to assess the power efficiency of storage devices used in data center applications in both active and idle states of operation. A taxonomy is defined that classifies storage devices in terms of operational profiles and supported features. Test definition and execution rules for measuring the power efficiency of each taxonomy category are defined; these include test sequence, test configuration, instrumentation, benchmark driver, IO profiles, measurement interval, and metric stability assessment. 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.

Publication of this Working Draft for review and comment has been approved by the Green Storage Technical Work Group. This draft represents a “best effort” attempt by the Green Storage Technical Work Group to reach preliminary consensus, and it may be updated, replaced, or made obsolete at any time. This document should not be used as reference material or cited as other than a “work in progress.” Suggestions for revisions should be directed to <http://www.snia.org/feedback/>.

Working Draft

June 23, 2025

USAGE

Copyright © 2025 Storage Networking Industry Association. All rights reserved. All other trademarks or registered trademarks are the property of their respective owners.

Storage Networking Industry Association (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 SNIA copyright on that material, and shall credit SNIA for granting permission for its reuse.

Other than as explicitly provided above, you may not make any commercial use of this document or any portion thereof, 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 © 2025, 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, or 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. 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. 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.

DRAFT

Revision History

Revision	Date	Changes
V0.0.36	August 19, 2024	Working Draft approved by Technical Council for release for public review and periodic call for IP.
V 00.49	June 23, 2025	Working Draft approved by Technical Council for release for public review, call for IP, and SNIA membership vote.

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

CONTACTING SNIA

SNIA Emerald™ Program Web Site

The SNIA Emerald™ Program web site is <https://www.snia.org/emerald>. SNIA Emerald™ Program-related downloads are available at <https://www.snia.org/emerald/download>, the *Download Materials for SNIA Emerald™ Testing* web page.

SNIA Address

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent via the SNIA Feedback Portal at <http://www.snia.org/feedback> or by mail to SNIA, 5201 Great America Parkway, Suite 320, Santa Clara, CA, 95054 U.S.A.

DRAFT

Acknowledgements

The SNIA Green Storage Technical Working Group, which developed and reviewed this document, recognizes the significant contributions made by the following individuals and companies:

Contributors

Individuals		Companies and Organizations
Wayne Adams	Santosh Kumar	Calypso Systems, Inc.
Ivan Alexeev	Suchang Kim	Dave Thiel, Consultant LLC
Bill Bailey	Tai Lee	Dell Inc.
Don Capps	Chris Lueth	Ellisys Group SA (SANBlaze / SerialTek / Quarch)
Jeffrey Doolittle	Andrew Norrie	Hewlett Packard Enterprise
John Geldman	Keith Orsak	Huawei Technologies Co Ltd.
Donald Goddard	Chuck Paridon	lozone
Michael Gregoire	Mike Peeler	Kioxia Corporation
Patrick Hery	Carlos Pratt	NetApp
Eden Kim	David Slik	Seagate Technology
	David Thiel	S.K. Hynix
		SNIA Japan
		Solidigm
		Toshiba America Electronic Components, Inc.
		Ultrastrattech

Contents

Introduction	12
1 Scope	13
1.1 Introduction	13
1.2 Purpose of a SNIA Emerald™ Device Power Efficiency Measurement Test Result	13
2 Normative References	14
3 Definitions, Symbols, Abbreviations, and Conventions	15
3.1 Overview	15
3.2 Definitions	15
3.3 Symbols and Abbreviated Terms	19
3.4 Expression of Provisions	19
3.5 Storage Capacity Conventions	20
4 Taxonomy	21
4.1 Introduction	21
4.2 Taxonomy Categories	21
4.3 Taxonomy Classifications	24
4.4 Client SSD Classifications	24
4.5 Taxonomy Details	25
4.6 Taxonomy Rules	28
5 Test Definition and Execution Rules	29
5.1 Overview	29
5.2 Basic Test Flow	29
5.3 General Requirements	29
6 Test Process Concepts	31
6.1 Introduction	31
6.2 Steady State	31
6.3 Purge	32
6.4 Pre-Fill and Conditioning	32
6.5 Active Range	33
6.6 Data Patterns	33
6.7 Ready Idle	33
6.8 Caching	33
7 Test Steps	34
7.1 Overview	34
7.2 Single Continuous Test	34
7.3 Self-Optimizing TOIO Sweep	35

7.4	Test Step Definitions	36
7.5	Test Software.....	38
8	Test Workloads	39
8.1	Overview.....	39
8.2	Random Read and Random Write	39
8.3	Sequential Read and Sequential Write	39
8.4	Complex Workload	39
8.5	Ready Idle Workload	41
9	Metrics.....	42
9.1	Primary Metrics.....	42
9.2	Power Efficiency Metrics	42
9.3	Average Response Time	43
9.4	Average Power	43
9.5	Operations Rate	43
9.6	Measurement Interval.....	44
9.7	Reporting Granularity	44
10	Workload Generator Requirements	45
10.1	Purpose	45
10.2	Procedure	45
11	Device Power Analyzer.....	46
11.1	Purpose	46
11.2	Procedure	46
12	General Environmental Measurement Recommendations	47
13	Reporting Requirements	48
13.1	General Reporting Requirements.....	48
13.2	Test Reporting	49
Annex A	(Informative) CTS Lite Report Example	50
A.1	Introduction	50
A.2	P1: IOPS and Power vs. Time.....	51
A.3	P2: Throughput and Power vs. Time	52
A.4	P3: Pre-Fill Test Step	53
A.5	P4: Steady State Verification.....	54
A.6	P5: IOPS and Power vs. Test Step	55
A.7	P6: Throughput and Power vs Segment	56
A.8	P7: IOPS And IOPS per Watt vs. Test Steps (Segments)	57
A.9	P8: Throughput and Throughput/Watt Vs. Test Steps (Segments).....	58
A.10	P9: Complex – DI Curves	59

A.11	P10: Random Write – DI Curves	60
A.12	P11: Random Read – DI Curves	61
A.13	P12: Sequential Write – DI Curves.....	62
A.14	P13: Sequential Read – DI Curves	63
A.15	P14: All Outside DI Curves	64
A.16	P15: Ready Idle Power vs. Time	65
A.17	P16: IOPS/Watt vs. Time.....	66
A.18	P17: MiB/s/Watt vs. Time	67
Annex B	(Informative) Test Methodologies	68
B.1	Introduction	68
B.2	Test Flow and Methodology	68
B.3	Substitutable Test Workloads for Future Revisions	68
B.4	Steady State Algorithm	68
B.5	Demand Intensity Curve	69
B.6	Demand Intensity Outside Curve.....	69
B.7	Report Generation	73
Annex C	(Normative) Workload Generator	74
C.1	Required Workload Generator.....	74
C.2	Workload Generator Requirements.....	74
Annex D	(Normative) Required Test Software	75
Annex E	(Normative) Device Power Analyzer.....	76
E.1	Purpose	76
E.2	General Description	76
E.3	General Specifications.....	76
E.4	Qualified Device Power Analyzers	76
Annex F	(Normative) Device Power Supply Requirements	78
F.1	Purpose	78
F.2	General Description	78
F.3	General Specifications.....	78
Bibliography	79

List of Tables, Figures and Equations

Table 1 – Normative References	14
Table 2 – Storage Capacity Conventions	20
Figure 1 – Taxonomy Classes - Relative Performance	22
Figure 2 – SSD Taxonomy	23
Figure 3 – HDD Taxonomy	23
Table 3 – SSD Category, Client Classifications	26
Table 4 – SSD Category, Data Center Classifications	27
Table 5 – HDD Classifications	28
Figure 4 – Test Configuration	30
Figure 5 – Conditioning Run and Steady State Check	32
Table 6 – Test Steps	35
Table 7 – IO Streams Within the Complex Workload	40
Table 8 – IO Transfer Size Within the Complex Workload For 512 Byte Native Devices	41
Table 9 – IO Transfer Size Within the Complex Workload For 4 KiB Native Devices	41
Equation 9-1: Power Efficiency, Ready Idle	42
Equation 9-2: Power Efficiency, Active	43
Equation 9-3: Average Power	43
Table 10 – Test Results	49
Figure A.1 – Report Plots Tab	50
Figure A.2 – IOPS and Power vs. Time	51
Figure A.3 – Throughput and Power vs. Time	52
Figure A.4 – Pre-Fill	53
Figure A.5 - Conditioning Run with Steady State	54
Figure A.6 – IOPS and Power vs. Test Steps (Segments)	55
Figure A.7 – Throughput and Power vs. Segments	56
Figure A.8 – IOPS and IOPS/Watt vs. Test Steps (Segments)	57
Figure A.9 – Throughput and Throughput/Watt vs. Test Steps (Segments)	58
Figure A.10 – Complex Workload – DI Curves	59
Figure A.11 – Random Write – DI Curves	60
Figure A.12 – Random Read – DI Curves	61
Figure A.13 – Sequential Write – DI Curves	62
Figure A.14 – Sequential Read – DI Curves	63
Figure A.15 – All Outside DI Curves	64
Figure A.16 – Ready Idle Power vs. Time	65
Figure A.17 – IOPS/Watt vs. Time	66
Figure A.18 – MiB/s/Watt vs. Time	67
Table B.1 – Workload Name Mapping	68
Figure B.1 – Demand Intensity Curve	69
Figure B.2 – DI Curve Complex	70
Figure B.3 – DI OS Curve Complex	71

Figure B.4 – All DI OS Curves	72
Figure B.5 – Audit File Playback	73
Table E.1 – Device Power Analyzer Requirements	76
Table E.2 – Qualified Device Power Analyzers	77
Table F.1 – Power Supply Requirements	78

DRAFT

Introduction

Introduction

The rapid growth in data storage and management has led to the development of more advanced storage devices, which are designed to meet the demands for increased capacity, performance, and reliability. However, as storage devices continue to evolve, the need for optimizing their power efficiency has become increasingly important.

To address this challenge, it is necessary to have a comprehensive and standardized method for measuring the power efficiency of individual storage devices. This document provides a detailed description of the procedures, protocols, and metrics that are used to measure the power efficiency of individual storage devices.

This document defines the requirements for conducting accurate and consistent power efficiency measurements, including the test environment, equipment, and data collection and reduction methods. This document serves as a reference for manufacturers, system integrators, regulatory bodies, test organizations and information technology professionals who are interested in evaluating the power efficiency of individual storage devices.

The methodologies of this document are based on those of

- International Standard ISO/IEC 24091:2019 *Information technology — Power efficiency measurement specification for data center storage*^[9] and
- SNIA Standard *SNIA Emerald™ Power Efficiency Measurement Specification* Version 4.0.0^[10]

with modifications that include a self-optimizing procedure and Principles of Demand Intensity and performance evolution based on

- SNIA Standard *Solid State Storage (SSS) Performance Test Specification (PTS)* Version 2.0.2^[7] and
- SNIA Standard *Real World Storage Workload (RWSW) Performance Test Specification for Datacenter Storage* Version 1.0^[8].

1 Scope

1.1 Introduction

This document defines a standardized method to assess the energy efficiency of individual storage devices in both active and idle states of operation. A taxonomy is defined that classifies storage devices 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 steady state assessment. Resulting power efficiency metrics are defined as ratios of idle capacity or active operations during a selected steady state measurement interval to the average measured power.

An individual storage device is a storage product that is contained in a standard device form factor, connects to a standard computer bus, supports a recognized storage device protocol, and presents storage to a host server as a single block device.

A storage device is said to be in an active state when the storage device is processing externally initiated requests for data transfer. A storage device is said to be in steady state when the workload-dependent performance is relatively time-invariant. A storage device is said to be in ready idle state when it is not receiving host IO commands.

This document specifies:

- A generalized taxonomy for storage devices (Clause 4) that addresses both solid state and magnetic disk storage devices;
- Measurement and data collection methodology for assessing the power efficiency of storage devices in both active and idle states (Clauses 5 6 7 , 8 10 and 11 ;
- Metrics describing storage device power efficiency (Clause 9); and
- Required disclosures for a test result published as a SNIA Emerald™ Device Power Efficiency Measurement¹ test result (Clause 13).

The test process defined by this document applies only to devices that support a 512 byte or 4096 byte logical blocks.

This document is intended for use by individuals and companies engaged in assessing the power utilization of storage devices.

1.2 Purpose of a SNIA Emerald™ Device Power Efficiency Measurement Test Result

The purpose of a SNIA Emerald™ Device Power Efficiency Measurement test result is to provide a reproducible and standardized assessment of the energy efficiency of storage devices used in data center applications in both active and ready idle states.

Note: A SNIA Emerald™ Device Power Efficiency Measurement test result provides a high-level assessment of the power efficiency of the tested storage device in specific idle and active states. Actual performance and power consumption in storage deployments are highly dependent upon workload, environmental and usage parameters.

¹ SNIA Emerald™ is a trademark of SNIA.

2 Normative References

The following documents are referred to in the text in a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

Table 1 lists these documents.

Table 1 – Normative References

Author/Owner	Title	Revision	URL
ISO/IEC	<i>ISO/IEC Directives Part II</i>	Ninth edition, 2021	https://www.iso.org/directives-and-policies.html

3 Definitions, Symbols, Abbreviations, and Conventions

3.1 Overview

For the purposes of this document, the terms and definitions given in *The SNIA Dictionary*^[11] and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <https://www.iso.org/obp>
- IEC Electropedia: available at <http://www.electropedia.org/>

The terms and definitions defined in this document are based on those found in *The SNIA Dictionary*^[11]. They have been extended, as needed, for use in this document. In cases where the current definitions in the SNIA dictionary conflict with those presented in this document, the definitions in this document shall apply.

3.2 Definitions

3.2.1

Active Range

range of LBA's that are accessed by the test code

3.2.2

ART ceiling

maximum allowable ART

3.2.3

average response time

the summation of a set of response times divided by the number of IO requests in the set

3.2.4

cache

temporary data storage, not directly addressable by end-user applications, used to store data for expedited access to or from slower media

3.2.5

Complex Workload

workload consisting of a mix of different IO sizes, read/write mixes, and access patterns with a skewed access across a range of LBAs

3.2.6

Conditioning

process of writing data to the device to prepare it for Steady State measurement

3.2.7

device under test

storage device that is undergoing test

3.2.8

DI Curve

Demand Intensity Curve in total OIO versus ART & IOP

3.2.9

Enterprise and Data Center Standard Form Factor (EDSFF)

industry standard form factors defined by SNIA Specification SFF-TA-1006^[3], SNIA Specification SFF-TA-1008^[4], and SNIA Specification SFF-TA-1009^[5].

3.2.10

Fresh Out-of-the-Box (FOB)

state of device prior to being put into service for the first time

3.2.11

large form factor (LFF)

industry standard 3.5 inch large form factor device defined by SNIA Specification SFF-8301^[1].

3.2.12

Logical Block Address (LBA)

address of a logical block

3.2.13

non-volatile storage

storage that retains data in the event of the unexpected loss of power

3.2.14

OIO/Thread

number of OIO per Thread

3.2.15

outstanding IO (OIO)

IO operations issued by a workload generator and awaiting completion

3.2.16

power efficiency

ratio of useful work to the power required to do the work

3.2.17

Pre-Fill

test step that writes data to a DUT to prepare that DUT for subsequent Conditioning and test steps

3.2.18

Purge

process of returning a storage device to a state in which subsequent writes execute, as closely as possible, as if the device had never been used and does not contain any valid data

3.2.19

ready idle

operational state of a device where the device is able to execute an arbitrary IO request and no IO requests are in progress

3.2.20

response time

measured time from the start of an IO request to the completion of the IO request

3.2.21

Round

one complete pass through all the prescribed test intervals for any given steady state calculation

3.2.22

sequential read

IO load consisting of consecutively issued read requests to logically adjacent data

3.2.23

sequential write

IO load consisting of consecutively issued write requests to logically adjacent data

3.2.24

small form factor (SFF)

industry standard 2.5 inch small form factor device defined by SNIA Specification SFF-8201^[2]

3.2.25

SNIA Emerald™ Device Power Efficiency Measurement

test performed according to this document

3.2.26

SNIA Emerald™ Device Power Efficiency Measurement test result

result of a test performed according to this document

3.2.27

Solid State Drive (SSD)

storage capability built from non-volatile solid state storage

3.2.28

stable storage

storage that maintains data across power failures or other transient failures

3.2.29

Steady State

state where the workload-dependent performance becomes relatively time invariant

3.2.30

Steady State Measurement Window

measurement interval, measured in Rounds, during which steady state test data is collected

3.2.31

storage device

solid state drive or hard disk drive

3.2.32

storage server

physical computing configuration on which the DUT is installed and the test software is run

3.2.33

Test Software

software used to execute the workloads, test steps, power aggregation, data reduction, and reporting

3.2.34

test sponsor

operator, company, or agent that performs a test according to this document

3.2.35

Thread Count (TC)

number of Threads specified by a workload generator

3.2.36

Total OIO (TOIO)

number of outstanding IO Operations

Note: Total OIO is calculated (Thread Count) * (Queue Depth)

3.2.37

user capacity

logical block size times the number of logical blocks accessible by a workload generator

3.2.38

workload generator

component of the Test Software that creates the command sequences applied to the DUT

3.3 Symbols and Abbreviated Terms

ART	average response time
CTS Lite	Test Software to be used with this document
DI	Demand Intensity
DUT	Device Under Test
EDSFF	Enterprise and Data Center Standard Form Factor
FOB	Fresh-Out-of-the-Box
HDD	Hard Disk Drive
IOPS	I/O operations per second
IOPS/W	IOPS per watt
LBA	logical block address
LFF	large form factor
MiB/s	mebibytes per second
MiB/s/W	mebibytes per second per watt
ms	millisecond
OIO	outstanding IO
QD	queue depth
R/W	read/write
SCSI	Small Computer System Interface
SFF	small form factor
SSD	solid state drive
TC	Thread Count
TOIO	total outstanding IO
TP	throughput
USB	Universal Serial Bus

3.4 Expression of Provisions

This document uses the verbal forms for expressions of provisions as defined by ISO/IEC Directives Part 2 (Ninth edition, 2021). These verbal forms include shall, shall not, should, should not, may, can, cannot, and must.

3.5 Storage Capacity Conventions

The associated units and abbreviations used in this document are shown in Table 2.

Table 2 – Storage Capacity Conventions

Storage Capacity Conventions					
Decimal			Binary		
Name	Abbreviation	Value	Name	Abbreviation	Value
kilobyte	kB	10^3 B	kibibyte	KiB	2^{10} B
megabyte	MB	10^6 B	mebibyte	MiB	2^{20} B
gigabyte	GB	10^9 B	gibibyte	GiB	2^{30} B
terabyte	TB	10^{12} B	tebibyte	TiB	2^{40} B

In this document

- Storage capacities are represented in decimal (base-10).
- IO transfer sizes and offsets are represented in binary (base-2).
- IO throughput is reported in binary (base-2).

4 Taxonomy

4.1 Introduction

This Clause defines a market taxonomy that classifies block access individual storage devices in terms of operational profile and supported features. The purpose of the taxonomy is to define specific groups into which storage devices can be placed for test comparison.

Individual storage devices taxonomy classes are defined broadly for hard disk drive (HDD) and solid state storage device (SSD) products. Within each class, additional categories are provided for client and enterprise class devices across SATA, SAS, or NVMe protocols, that are housed in SFF (2.5 inch), LFF (3.5 inch), EDSFF, or M.2 standard storage device form factors.

While this taxonomy is broad and defines a framework for devices that range from consumer solutions to enterprise installations, it is not intended to address all storage device types. For example, SSD storage devices in this specification are limited to NAND Flash based media and specifically exclude Persistent Memory devices.

The taxonomy is structured as a 2-level hierarchy of Category and Classification. A Category is a broad grouping of devices such as HDD and SSD. A Classification is a level of device sophistication, size, and complexity that corresponds to market delineations. The Classifications distinguish performance, power utilization, and lifespan.

The taxonomy addresses a broader range of storage devices than is covered by the power efficiency tests defined in this document.

4.2 Taxonomy Categories

4.2.1 General

Taxonomy categories define broad market segments that can be used to group devices that share common functionality or performance requirements, and within which meaningful device comparisons can be undertaken. This document defines two broad taxonomy categories (SSD and HDD) summarized in Figure 1.

Each of the categories contains additional classifications. It is intended that storage devices tested to this specification shall only be compared to storage devices that share the same category and classification.

While HDD and SSD categories have unique differentiators, devices satisfying either category can be used in similar applications. The relative performance of HDD and SSD storage devices is shown in Figure 1.

Each class covers a range of sustained IOPS and average response time performance. These are general estimates of relative class performance.

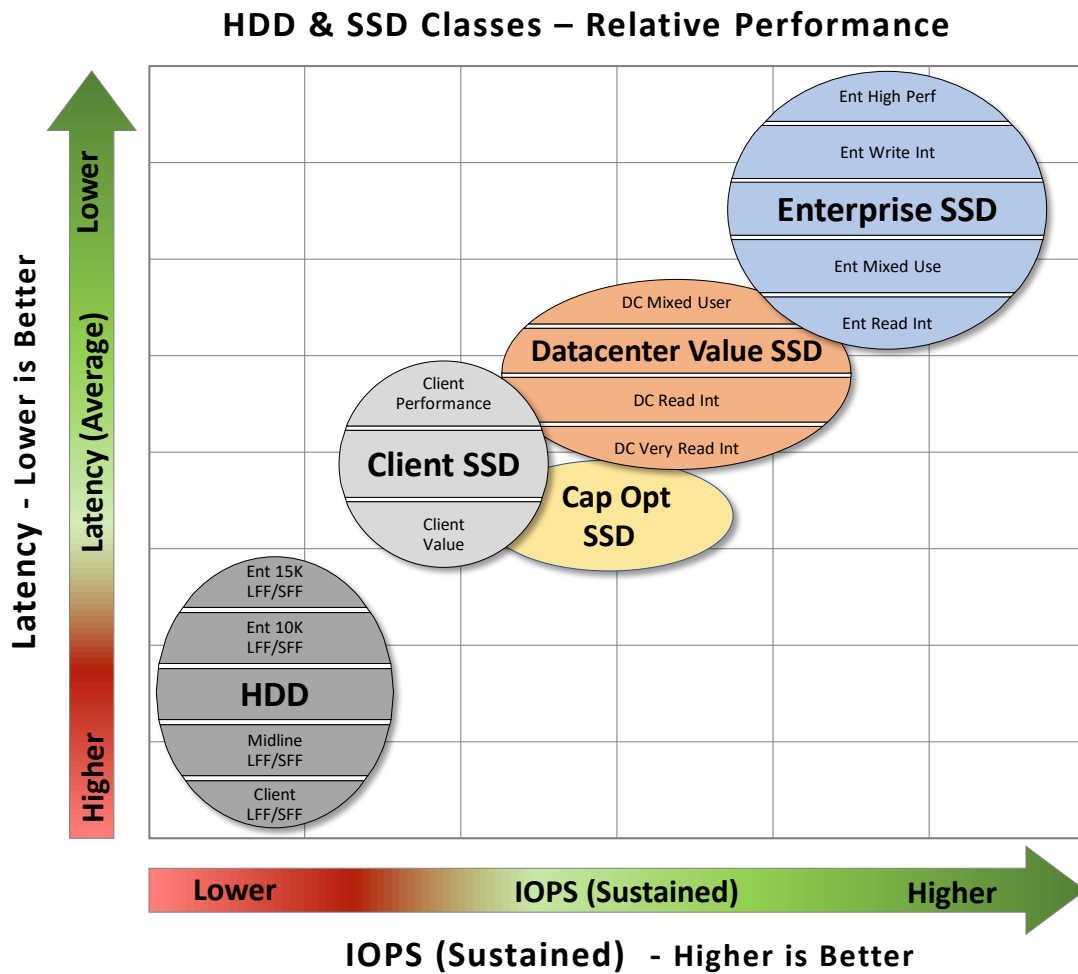


Figure 1 – Taxonomy Classes - Relative Performance

Taxonomy charts specific to HDD and SSD classes are listed in Figure 2 and Figure 3 and show relative power, life cycle, and performance.

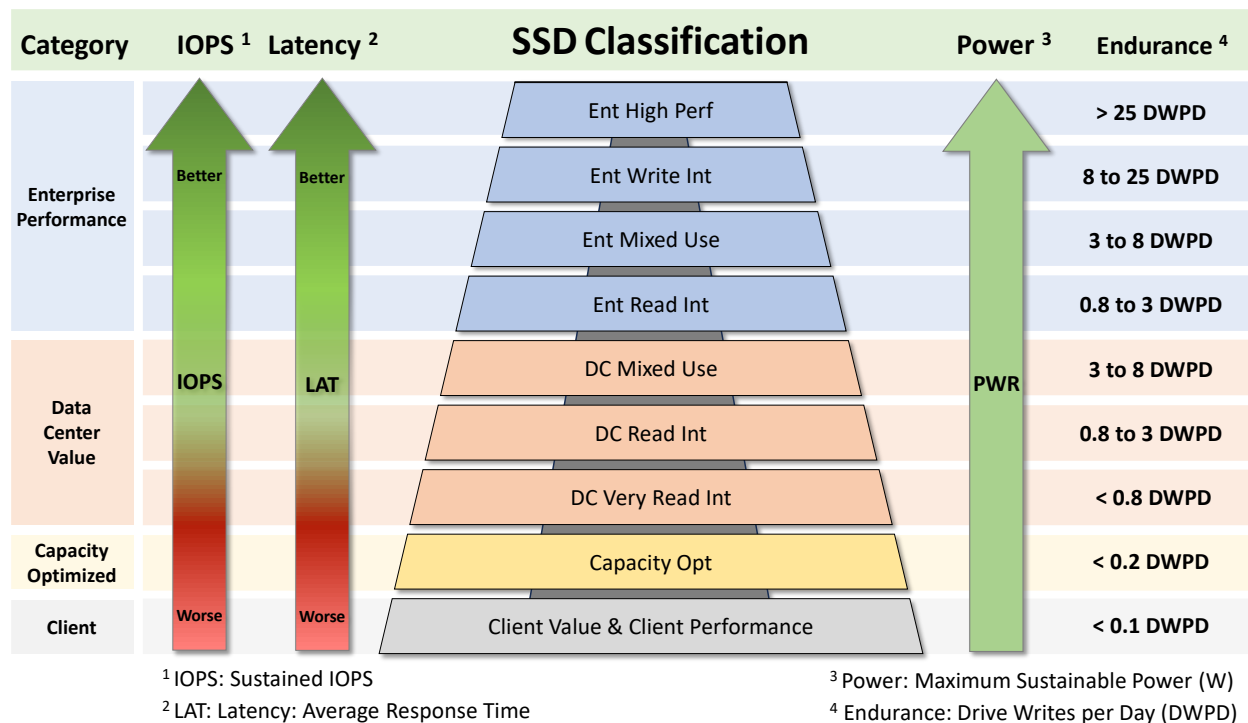


Figure 2 – SSD Taxonomy

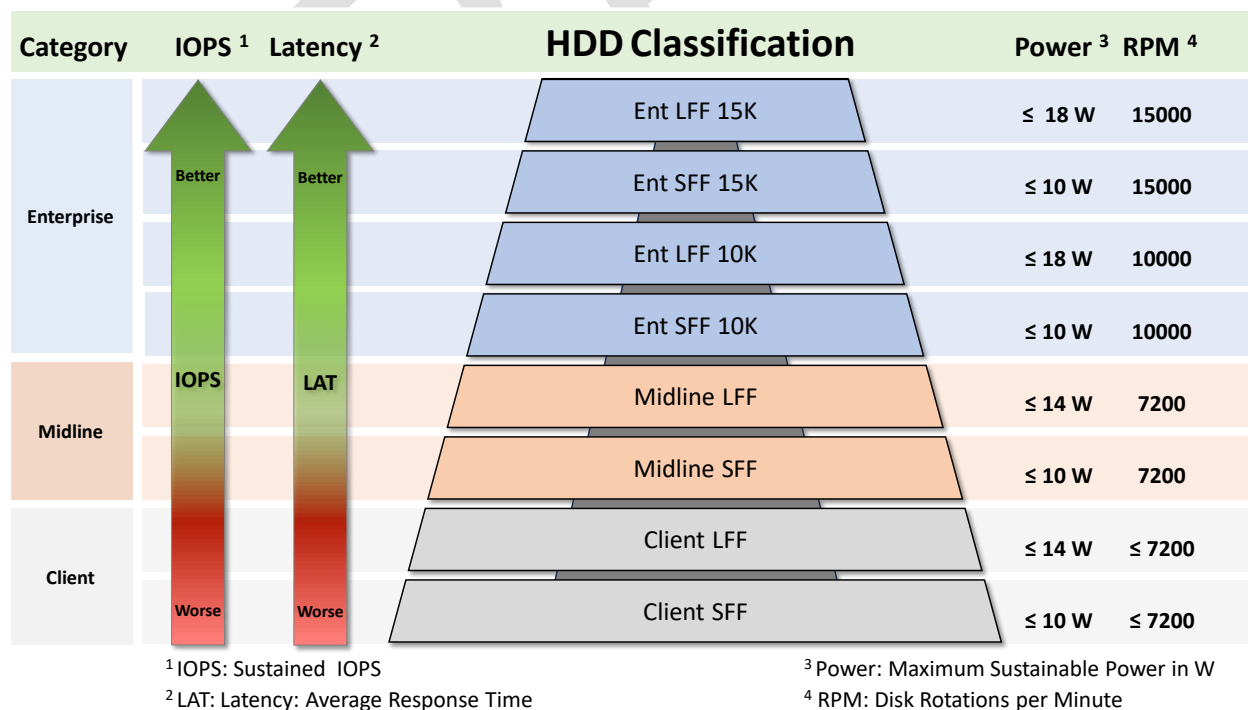


Figure 3 – HDD Taxonomy

4.3 Taxonomy Classifications

Classifications define combinations of settings or values for the attributes within a Category.

Taxonomy Classifications discriminate between different devices. The goal of the Classifications is to differentiate between devices of differing form factor, functionality, and target market.

4.3.1 SSD Category

4.3.1.1 Client SSD Classifications

Client devices are any of a wide variety of SSDs which are manufactured primarily for personal, family, household, or small-business purposes ranging from personal storage, personal computing, and boot to some Artificial Intelligence (AI) and Machine Learning (ML) applications. These devices typically have lower performance and endurance capabilities than those offered in Capacity Optimized, Data Center Value, or Enterprise Performance Classification SSDs.

- a) Client Value Classification devices consist of lower capacity and lower performance client SSDs.
- b) Client Performance Classification devices consist of higher capacity and higher performing client SSDs.

4.3.1.2 Capacity Optimized SSD Classification

Capacity Opt Classification devices are SSDs with lower performance and price than Enterprise Performance Classification devices. Capacity OPT devices are optimized for read intensive and large transfer applications where cost of capacity is critical. Common applications include AI/ML, Deep Learning, data streaming, and data analytics.

4.3.1.3 Data Center Value SSD Classifications

Data Center Value devices are any of a wide variety of cost optimized SSDs that are optimized for read intensive workloads. Common applications include Hyperscale, IoT, big data analytics, OLTP, and streaming media. These devices typically have lower performance and endurance capabilities than those offered by Enterprise Performance Classification devices.

- a) DC Very Read Int Classification devices are Data Center Value SSDs that are optimized for capacity and very read intensive workloads at the expense of performance and endurance.
- b) DC Read Int Classification devices are Data Center Value SSDs that are optimized for capacity and read intensive purposes with higher performance and endurance capabilities than DC Very Read Int Classification devices.
- c) DC Mixed Use Classification devices are Data Center Value SSDs that are optimized for higher performance and endurance purposes than other Data Center Value SSD Classification devices.

4.3.1.4 Enterprise Performance SSD Classifications

Enterprise Performance devices are any of a wide variety of performance and endurance optimized SSDs which are manufactured primarily for enterprise applications such as data streaming, web servers, analytics, High Performance Computing, caching and logging. These devices typically have higher throughput, IOPS, and lower response time performance and endurance capabilities than those offered in the Data Center Value Classifications.

- a) Ent Read Int Classification devices are Enterprise Performance SSDs that are optimized for capacity and read intensive workloads at the expense of performance and endurance.
- b) Ent Mixed Use Classification devices are Enterprise Performance SSDs that are performance and endurance optimized over Read Intensive Classification devices.

- c) Ent Write Int Classification devices are Enterprise Performance SSDs that are optimized for high performance and endurance optimized for applications such as Analytics, High Performance Computing, and server storage.
- d) Ent High Perf Classification devices are Enterprise Performance SSDs that are manufactured for Enterprise purposes requiring the highest endurance and performance capabilities.

4.3.2 HDD Category

4.3.2.1 Client HDD Classifications

Client devices are any of a variety of HDDs which are manufactured primarily for personal, family, household, or small-business purposes such as PC, video, audio, and gaming.

- a) Client SFF Classification devices are SFF client HDDs.
- b) Client LFF Classification devices are LFF client HDDs.

4.3.2.2 Midline HDD Classifications

Midline devices are any of a variety of HDDs which are manufactured primarily for business-critical applications such as backup, archive, and video surveillance.

- a) Midline SFF Classification devices are SFF midline HDDs.
- b) Midline LFF Classification devices are LFF midline HDDs.

4.3.2.3 Enterprise HDD Classifications

Enterprise devices are any of a variety of HDDs which are manufactured primarily for enterprise applications such as transaction processing and high-performance computing.

- a) Ent SFF 10K Classification devices are SFF 10000 RPM enterprise HDDs.
- b) Ent LFF 10K Classification devices are LFF 10000 RPM enterprise HDDs.
- c) Ent SFF 15K Classification devices are SFF 15000 RPM enterprise HDDs.
- d) Ent LFF 15K Classification devices are LFF 15000 RPM enterprise HDDs.

4.4 Taxonomy Details

4.4.1 SSD Classifications

Table 3 and Table 4 define the SSD taxonomy classifications and their attributes.

Table 3 – SSD Category, Client Classifications

Attribute	Category: SSD	
	Classification	
	Client	
	Client Value ^a	Client Performance ^a
Form Factor	M.2, EDSFF, SFF	M.2, EDSFF, SFF
Endurance (DWPD ^b)	< 0.1	< 0.1
Max Sustainable Power (W)	≤ 25	≤ 25
Read Response Time ^c (μs)	< 130	< 130
Write Response Time ^d (μs)	< 50	< 50
Multiple Power States	Low power	Low power
Media Type	NAND Flash	NAND Flash
Access Paradigm	Block	Block
Protocol	USB, SATA, NVMe	SATA, SAS, NVMe
Power Loss Protection ^e Required	No	No
Persistence ^f Required	No	No
Hot Plug Required	No	No
^a This document specifies a test procedure for Client Classification devices used in data center applications. ^b DWPD (drive writes per day) – amount of data written per day as a fraction of user capacity. ^c Read Response Time is the average response time for a random 4 KiB read, with a queue depth of 1 and thread count of 1. ^d Write Response Time is the average response time for a random 4 KiB write, with a queue depth of 1 and thread count of 1. ^e Power Loss Protection is a device technology consisting of hardware and/or software to prevent data loss on data residing in a volatile location in the event of an unexpected power loss scenario. ^f Persistence is the device property that writes are reported as complete only after data will be retained by the device in the event of an unexpected power loss.		

Table 4 – SSD Category, Data Center Classifications

Attribute	Category: SSD							
	Classification							
	Capacity Optimized	Data Center Value			Enterprise Performance			
	Capacity Opt	DC Very Read Int	DC Read Int	DC Mixed Use	Ent Read Int	Ent Mixed Use	Ent Write Int	Ent High Perf
Form Factor	EDSFF, SFF	EDSFF, SFF	EDSFF, SFF	EDSFF, SFF	EDSFF, SFF	EDSFF, SFF	EDSFF, SFF	EDSFF, SFF
Endurance (DWPD ^a)	< 0.2	< 0.8	0.8 to 3	3 to 8	0.8 to 3	3 to 8	8 to 25	> 25
Max Sustainable Power (W)	≤ 25	≤ 25	≤ 25	≤ 25	25 to 40	25 to 40	25 to 40	25 to 40
Read Response Time ^b (μs)	< 130	< 110	< 110	< 110	< 90	< 90	< 90	< 90
Write Response Time ^c (μs)	< 50	< 30	< 30	< 30	< 20	< 20	< 20	< 20
Multiple Power States	Low power	Low power	Low power	Low power	Low to high power	Low to high power	Low to high power	Low to high power
Media Type	NAND Flash	NAND Flash	NAND Flash	NAND Flash	NAND Flash	NAND Flash	NAND Flash	NAND Flash
Access Paradigm	Block	Block	Block	Block	Block	Block	Block	Block
Protocol	SATA, SAS, NVMe	SATA, SAS, NVMe	SATA, SAS, NVMe	SATA, SAS, NVMe	SAS, NVMe	SAS, NVMe	SAS, NVMe	SAS, NVMe
Power Loss Protection ^d Required	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Persistence ^e Required	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hot Plug Required	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

^a DWPD (drive writes per day) – amount of data written per day as a fraction of user capacity.

^b Read Response Time is the average response time for a random 4 KiB read, with a queue depth of 1 and thread count of 1.

^c Write Response Time is the average response time for a random 4 KiB write, with a queue depth of 1 and thread count of 1.

^d Power Loss Protection is a device technology consisting of hardware and/or software to prevent data loss on data residing in a volatile location in the event of an unexpected power loss scenario.

^e Persistence is the device property that writes are reported as complete only after data will be retained by the device in the event of an unexpected power loss.

4.4.2 HDD Classifications

Table 5 defines the HDD taxonomy classifications.

Table 5 – HDD Classifications

Attribute	Category: HDD							
	Classification							
	Client ^a		Midline		Enterprise			
	Client SFF	Client LFF	Midline SFF	Midline LFF	Ent SFF 10K	Ent LFF 10K	Ent SFF 15K	Ent LFF 15K
Form Factor	SFF	LFF	SFF	LFF	SFF	LFF	SFF	LFF
Workload Rate ^b (TB / year)	< 500	< 500	≥ 500	≥ 500	> 500	> 500	> 500	> 500
Rotational Rate (RPM)	≤ 7200	≤ 7200	< 10000	< 10000	≥ 10000	≥ 10000	15000	15000
Max Sustainable Power (W)	≤ 10	≤ 14	≤ 10	≤ 14	≤ 10	≤ 18	≤ 10	≤ 18
Power Loss Protection ^b Required	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Hot Plug Required	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Persistence ^c Required	No	No	Yes	Yes	Yes	Yes	Yes	Yes
^a This document specifies a test procedure for Client Classification devices used in data center applications. ^b Workload rate (TB / year) is maximum workload that the device is designed to accommodate. One TB = 10 ¹² bytes. ^c Power Loss Protection is a device technology consisting of hardware and/or software to prevent data loss on data residing in a volatile location in the event of an unexpected power loss scenario. ^d Persistence is the device property that writes are reported as complete only after data will be retained by the device in the event of an unexpected power loss.								

4.5 Taxonomy Rules

To be considered to be in a particular category and classification, a device shall satisfy all the attributes for the designated category and designated classification, i.e., satisfy all of the attributes of a column in one of the taxonomy tables, i.e., Table 3, Table 4, or Table 5.

If a device satisfies the attributes of multiple classifications, the device may be considered to be in any or all of those classifications.

If a device does not satisfy all of the row attribute parameters listed in the relevant classification column listed in Table 3, Table 4, or Table 5, then the device shall be considered to be in the taxonomy category Other and taxonomy classification Other.

5 Test Definition and Execution Rules

5.1 Overview

A SNIA Emerald™ Device Power Efficiency Measurement shall be a good faith effort to accurately characterize the power requirements of the tested device. The precise configuration used in a SNIA Emerald™ Device Power Efficiency Measurement is left to the test sponsor. Any commercially released or soon to be released components may be used, and a focus on new or emerging components or technologies is encouraged.

This Clause defines the test definition and execution rules that shall be used. Test specific parameter settings, reports, and other requirements are documented in other sections that follow.

5.2 Basic Test Flow

The Basic Test Flow defines test specific parameter settings and reporting requirements.

- 1) Establish physical and software configuration
 - a) Connect Device Power Analyzer to interposer or DUT power port.
 - b) Identify and prepare the DUT.
 - c) Record the configuration of the storage server's hardware and test software, and DUT.
- 2) Execute test
 - a) Start the test software.
 - b) Select the Device Power Analyzer.
 - c) Select the target DUT.
 - d) Select the test.
 - e) Enter user-selectable parameters.
 - f) Start test.
 - g) Wait for test completion.
- 3) Report data as specified in Clause 13 .

5.3 General Requirements

5.3.1 Configuration Guidelines

Figure 4 shows a DUT connection via a direct connection or an interposer board to the device power port.

The sample configuration shown in Figure 4 is provided as a guideline. Test sponsors may modify the configuration to meet their particular needs and equipment, provided no other requirement of this document is violated.

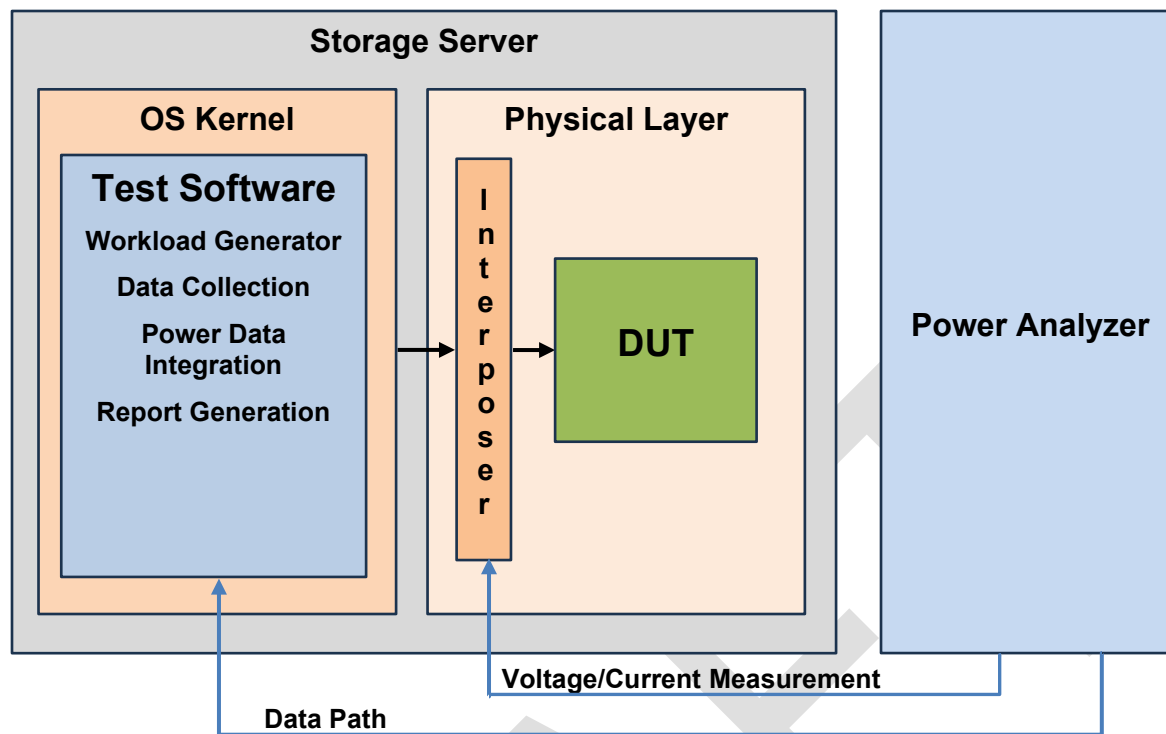


Figure 4 – Test Configuration

6 Test Process Concepts

6.1 Introduction

The performance of a storage device can be highly dependent on its prior usage, write history, caches, the pre-test state of the device and test parameters. This section defines key test process concepts. The next sections define Test Flow and Test Methodologies.

The test process defined by this document applies to storage devices used in data center applications.

6.2 Steady State

Storage devices that are Fresh Out-of-the-Box (FOB), or in an equivalent state, can exhibit an initial transient period of elevated performance, or peak performance.

For example, SSDs typically exhibit an initial peak performance, relative to the workload being applied, followed by a transition that evolves to Steady State.

The Steady State is a state where the workload-dependent performance is relatively time invariant.

It is important that the test data be gathered during a Steady State Measurement Window when the device is in Steady State, for two primary reasons:

- 1) To ensure that a device's initial performance (FOB or other peak) will not be reported as "typical", since this is transient behavior and not a meaningful indicator of the device's performance during the bulk of its operating life.
- 2) To enable test sponsors and reviewers to observe and understand trends. For example, oscillations around an average are "steady" in a sense, but might be a cause for concern.

Accordingly, the DUT shall achieve steady state wherein the linear least squares best-fit curve for five consecutive 30 minute Rounds does not exceed a $\pm 10\%$ deviation from average IOPS of the five Rounds and the slope of linear least squares best-fit curve over the 5 Rounds does not exceed 10 %. This steady state determination procedure is referred to as the 5 Rounds Steady State Determination.

Note: Steady State Measurement Window is bounded by the Round in which the device has been observed to have maintained Steady State for the specified number of Rounds (Round x), and five Rounds previous (Round $x-4$).

The steady state procedure is contained within the Conditioning Run (CR) test step. The Conditioning Run Steady State report is automatically generated in the Test Audit File and reported as Plot P4.

Figure 5 shows an example Plot P4 with seven 1800 s (30 minute) Conditioning Run Rounds for an SSD DUT. The 5 Round Steady State Window occurs within the Conditioning Run and is shown occurring during Rounds 3 through 7. Round 1 shows IOPS at FOB peak performance. Rounds 1 and 2 show IOPS settling to a stable value. Conditioning Run Rounds 3 through 7 show the Steady State (SS) Window. Here, the SS Window occurs during SS Window Rounds 1 through 5.




Figure 5 – Conditioning Run and Steady State Check

6.3 Purge

The purpose of Purge is to put the device in a consistent state in which subsequent writes execute, as closely as possible, as if the device had never been used and does not contain any valid data. The Purge step, if used, shall be executed prior to the Pre-Fill step to facilitate a clear demonstration of Steady State convergence behavior.

If the DUT does not support any kind of Purge method, then the fact that Purge was not supported/run shall be documented in the test report.

The Test Software allows the test sponsor to choose to purge or not to purge the DUT. If the test sponsor chooses to purge the DUT, the Test Software detects and applies a device purge (e.g., ATA: Security Erase, SCSI: Format Unit, NVMe: Sanitize) command.

The test sponsor shall disclose the PURGE method employed.

6.4 Pre-Fill and Conditioning

6.4.1 Pre-Fill

The Pre-Fill step is intended to prepare the DUT for subsequent steady state and test measurement steps by writing data to the DUT. The goal of Pre-Fill is to facilitate convergence to Steady State prior to running the Steady State measurement rounds of the test.

6.4.2 Conditioning Run

The Conditioning Run applies the Complex Workload followed by a Steady State procedure where the DUT is run for a sequence of consecutive 1800 s (30 minute) Rounds during which five (5) consecutive Rounds show relatively time invariant performance.

6.4.3 Steady State

Steady State criteria are intended to put the DUT in a condition where performance is relatively time invariant so that subsequent measurements are consistent and repeatable.

6.5 Active Range

Active Range refers to the span of logical block addresses (LBA) that are accessed by the test code.

It is important to test the performance characteristics of devices using workloads that issue IO across 100% of the LBA space of the device.

6.6 Data Patterns

The workload generator uses data patterns to generate test code workloads. It is important to use data with known and specified data compression ratio because the compressibility of the data impacts performance.

6.7 Ready Idle

Changes in ready idle power can indicate the impact of garbage collection, write amplification, and power throttling activities on power consumption.

Ready idle measurements shall be taken during a period of specified length when no host commands are issued to the DUT.

6.8 Caching

Memory caches in a data path significantly impact DUT performance. Therefore, DUT cache settings shall be reported.

7 Test Steps

7.1 Overview

7.1.1 DUT Configuration

The test sponsor shall identify one taxonomy set, category and classification for the DUT.

The DUT shall be configured to satisfy the requirements of the selected taxonomy set, category and classification.

The DUT shall represent a customer orderable configuration whose use within the selected one taxonomy set, category and classification is supported by the device vendor.

7.1.2 DUT Consistency

The physical and logical configuration of the DUT shall not be changed during a test unless explicitly allowed in the definition of a test step.

7.1.3 No Non-Test Activity

Other than booting/starting the device under test and any test equipment employed during the test, no substantive work shall be performed on the DUT between the test steps defined in this document, unless explicitly allowed in the definition of a test step.

7.1.4 Test Step Sequence

All test steps shall be executed as an uninterrupted sequence.

7.2 Single Continuous Test

The test steps shall be run as a single, continuous test with no time gaps or delays between the test steps. See Table 6.

Table 6 – Test Steps

Test Steps	
Test Step	Description
1. Setup Test Environment	Connect the DUT. Setup and start the Device Power Analyzer.
2. Purge	Apply the Purge, if supported, to the DUT.
3. Pre-Fill	Apply the SEQ 256 KiB Write Workload to the DUT for twice the user capacity of the DUT. For HDDs, this test step is optional.
4. Conditioning Run with Steady State	Apply the Complex Workload for a minimum of 5 Rounds of 1800 s (30 minutes) each. Execute Steady State procedure at Round 1 during the Steady State Conditioning Run. If the DUT is not in steady state, perform additional Rounds (up to a total of 25) until Steady State is achieved.
5. Complex	Apply the Complex Workload for a warm-up interval of 600 s (10 minutes) that includes a self-optimizing sweep of Total OIO followed by a measurement interval of 1800 s (30 minutes).
6. RND 8 KiB Write	Apply the RND 8 KiB Write Workload for a warm-up interval of 600 s (10 minutes) that includes a self-optimizing sweep of Total OIO followed by a measurement interval of 1800 s (30 minutes).
7. RND 8 KiB Read	Apply the RND 8 KiB Read Workload for a warm-up interval of 600 s (10 minutes) that includes a self-optimizing sweep of Total OIO followed by a measurement interval of 1800 s (30 minutes).
8. SEQ 256 KiB Write	Apply the SEQ 256 KiB Write Workload for a warm-up interval of 600 s (10 minutes) that includes a self-optimizing sweep of Total OIO followed by a measurement interval of 1800 s (30 minutes).
9. SEQ 256 KiB Read	Apply the SEQ 256 KiB Read Workload for a warm-up interval of 600 s (10 minutes) that includes a self-optimizing sweep of Total OIO followed by a measurement interval of 1800 s (30 minutes).
10. Ready Idle	Apply Ready Idle for an interval of 1800 s (30 minutes) and report test measurements.
11. Post Process and Report	Import the power data, post process the power data, and generate the report.

7.3 Self-Optimizing TOIO Sweep

For the Complex, RND 8 KiB Write, RND 8 KiB Read, SEQ 256 KiB Write, and SEQ 256 KiB Read test steps, a self-optimizing TOIO sweep is executed during each test step warm-up interval. The self-optimizing TOIO sweep determines the TOIO setting to be used during the corresponding test step measurement interval. The TOIO setting shall be the maximum IOPS point where the average response time is less than 20 ms.

The Test Software selects the TOIO setting to be used in the immediately following measurement interval.

Note: the self-optimizing TOIO sweep determines the optimal TOIO setting during a single test run instead of having to run multiple tests at different TOIO settings to achieve the same result.

7.4 Test Step Definitions

7.4.1 Setup Test Environment

Connect the DUT to the storage server. Connect the Device Power Analyzer to the DUT. Start the Device Power Analyzer. Start the Test Software.

7.4.2 Purge

If supported, the Purge command or equivalent may be applied to the DUT.

This step places the DUT in a Fresh-Out-of-Box (FOB) state as if no writes have occurred. If no Purge command is applied, "No Purge" shall be reported in the test reporting.

7.4.3 Pre-Fill

This test step prepares the DUT for convergence to steady state.

Apply the SEQ 256 KiB Write Workload to the DUT for twice the user capacity of the DUT.

For HDDs, this test step may be omitted. Note: If an HDD test has not performed the pre-fill test step and subsequently fails the steady state conditional run test step, then it might be beneficial to include the pre-fill test step when the test is rerun.

7.4.4 Conditioning Run With Steady State

The 5 Rounds Steady State Determination shall show relatively time invariant (stable) performance for five consecutive rounds. Apply the Complex Workload with a Thread Count of 4 and a Queue Depth of 32 to the DUT for a minimum of 5 Rounds of 1800 s (30 minutes) each. If the DUT has not achieved steady state for 5 consecutive Rounds, perform additional Rounds (up to a total of 25) until steady state is achieved for 5 consecutive Rounds. If steady state is not achieved after 25 Rounds, the data for this test step shall not be reported.

7.4.5 Complex

The IO steam composition of the Complex Workload is intended to emphasize the impact of caches by skewing IO Block Size distributions and adjusting the LBA Range accesses of the IO load.

The Complex test step starts with a 600 s (10 minute) warm-up interval that includes a self-optimizing TOIO sweep to determine the TOIO setting to be used during the measurement interval. The Complex Workload defined in Clause 8 shall be applied to the DUT using various Thread Count and Queue Depth combinations. The combinations shall be a Thread Count of each of 64, 48, 40, 32, 16, 8, 4, 2, and 1 and, for each Thread Count, a Queue Depth each of 96, 80, 64, 48, 40, 32, 16, 8, 4, 2, and 1. Each combination shall be of 5 s duration. Record IOPS and ART measurements for each combination.

Upon completion, plot a demand intensity outside curve and select the Thread Count and Queue Depth combination where IOPS are greatest while the response time is less than 20 ms.

As the measurement interval, apply the Complex Workload to the DUT using the selected Thread Count and Queue Depth for 1800 s (30 minutes) and record IOPS and ART measurements.

7.4.6 RND 8 KiB Write

The RND 8 KiB Write Workload is defined in Clause 8 .

The RND 8 KiB Write test step starts with a 600 s (10 minute) warm-up interval that includes a self-optimizing TOIO sweep to determine the TOIO setting to be used during the measurement interval. The RND 8 KiB Write Workload defined in Clause 8 shall be applied to the DUT using various Thread Count and Queue Depth combinations. The combinations shall be a Thread Count of each of 64, 48, 40, 32, 16, 8, 4, 2, and 1 and, for each Thread Count, a Queue Depth each of 96, 80, 64, 48, 40, 32, 16, 8, 4, 2, and

1. Each combination shall be of 5 s duration. Record IOPS and ART measurements for each combination.

Upon completion, plot a demand intensity outside curve and select the Thread Count and Queue Depth combination where IOPS are greatest while the response time is less than 20 ms.

As the measurement interval, apply the RND 8 KiB Write Workload to the DUT using the selected Thread Count and Queue Depth for 1800 s (30 minutes) and record IOPS and ART measurements.

7.4.7 RND 8 KiB Read

The RND 8 KiB Read Workload is defined in Clause 8 .

The RND 8 KiB Read test step starts with a 600 s (10 minute) warm-up interval that includes a self-optimizing TOIO sweep to determine the TOIO setting to be used during the measurement interval. The RND 8 KiB Read Workload defined in Clause 8 shall be applied to the DUT using various Thread Count and Queue Depth combinations. The combinations shall be a Thread Count of each of 64, 48, 40, 32, 16, 8, 4, 2, and 1 and, for each Thread Count, a Queue Depth each of 96, 80, 64, 48, 40, 32, 16, 8, 4, 2, and 1. Each combination shall be of 5 s duration. Record IOPS and ART measurements for each combination.

Upon completion, plot a demand intensity outside curve and select the Thread Count and Queue Depth combination where IOPS are greatest while the response time is less than 20 ms.

As the measurement interval, apply the RND 8 KiB Read Workload to the DUT using the selected Thread Count and Queue Depth for 1800 s (30 minutes) and record IOPS and ART measurements.

7.4.8 SEQ 256 KiB Write

The SEQ 256 KiB Write Workload is defined in Clause 8 .

The SEQ 256 KiB Write test step starts with a 600 s (10 minute) warm-up interval that includes a self-optimizing TOIO sweep to determine the TOIO setting to be used during the measurement interval. The SEQ 256 KiB Write Workload defined in Clause 8 shall be applied to the DUT using various Thread Count and Queue Depth combinations. The combinations shall be a Thread Count of each of 64, 48, 40, 32, 16, 8, 4, 2, and 1 and, for each Thread Count, a Queue Depth each of 96, 80, 64, 48, 40, 32, 16, 8, 4, 2, and 1. Each combination shall be of 5 s duration. Record IOPS and ART measurements for each combination.

Upon completion, plot a demand intensity outside curve and select the Thread Count and Queue Depth combination where IOPS are greatest while the response time is less than 20 ms.

As the measurement interval, apply the SEQ 256 KiB Write Workload to the DUT using the selected Thread Count and Queue Depth for 1800 s (30 minutes) and record IOPS and ART measurements.

7.4.9 SEQ 256 KiB Read

The SEQ 256 KiB Read Workload is defined in Clause 8 .

The SEQ 256 KiB Read test step starts with a 600 s (10 minute) warm-up interval that includes a self-optimizing TOIO sweep to determine the TOIO setting to be used during the measurement interval. The SEQ 256 KiB Read Workload defined in Clause 8 shall be applied to the DUT using various Thread Count and Queue Depth combinations. The combinations shall be a Thread Count of each of 64, 48, 40, 32, 16, 8, 4, 2, and 1 and, for each Thread Count, a Queue Depth each of 96, 80, 64, 48, 40, 32, 16, 8, 4, 2, and 1. Each combination shall be of 5 s duration. Record IOPS and ART measurements for each combination.

Upon completion, plot a demand intensity outside curve and select the Thread Count and Queue Depth combination where IOPS are greatest while the response time is less than 20 ms.

As the measurement interval, apply the SEQ 256 KiB Read Workload to the DUT using the selected Thread Count and Queue Depth for 1800 s (30 minutes) and record IOPS and ART measurements.

7.4.10 Ready Idle

The Ready Idle Workload is defined in Clause 8 .

Apply the Ready Idle Workload to the DUT for at least 1800 s (30 minutes). Then, as the measurement interval, apply the Ready Idle Workload to the DUT for at least 1800 s (30 minutes) and record the average power consumption over the measurement interval.

7.4.11 Post Process and Report

Import the power data and perform post-processing. Plot the reports and steady state rounds data as required by Clause 13 .

7.5 Test Software

The Test Software required by this document creates and applies the required workloads to the DUT. The Test Software is used to accurately capture and analyze the dynamic power characteristics of a DUT tested pursuant to this document. See Annex A, Annex B, Annex C, and Annex D.

8 Test Workloads

8.1 Overview

The particular IO stimuli used to drive the DUT during a test step are specified in terms of an IO profile (a.k.a. workload) made up of multiple attributes:

- Name: the name of the IO pattern for this stimulus. The identifier for the associated test step is included parenthetically, when appropriate;
- IO Size: the number of bytes requested by a given read or write operation;
- Read/Write Percentage: the mixture of read/write IO requests within an IO profile;
- Transfer Alignment: Minimum granularity of IO transfer addresses. All transfer addresses within an IO stream shall be a multiple of this value;
- Access Pattern: either one or the other of the following two alternatives:
 - Random: Randomly distributed throughout the address space of the device under test;
 - Sequential, as defined in Subclause 3.2.22 sequential read and Subclause 3.2.23 sequential write;
- Data Pattern: the data patterns written shall be uniformly distributed random data bits.

8.2 Random Read and Random Write

Random Reads and Random Writes refer to two different types of operations that can be performed on data stored in a storage medium.

Random Reads involve retrieving or reading data from non-contiguous or randomly chosen logical block addresses (LBAs) within a storage device.

Random Writes refer to the process of writing data to non-contiguous or randomly chosen logical block addresses (LBAs) within a storage device.

The RND 8 KiB Read Workload consists of Random Reads of 8 KiB data blocks from randomly chosen LBAs.

The RND 8 KiB Random Write Workload consists of Random Writes of 8 KiB data blocks from randomly chosen LBAs.

8.3 Sequential Read and Sequential Write

Sequential Reads and Sequential Writes refer to two different types of operations that can be performed on data stored in a storage device.

Sequential Reads refers to the process of reading data from a storage device in consecutive logical block address (LBA) order.

Sequential Writes refers to the process of writing data to a storage device in consecutive logical block address (LBA) order.

The SEQ 256 KiB Read Workload consists of Sequential Reads of 256 KiB data blocks.

The SEQ 256 KiB Write Workload consists of Sequential Writes of 256 KiB data blocks.

8.4 Complex Workload

The Complex Workload consists of a mix of different IO sizes, read/write mixes, and access patterns with a skewed access across a range of blocks.

The Complex Workload is adopted from the Hot Band IO Profile of the *ISO/IEC 24091:2019*^[9] and *SNIA Emerald™ Power Efficiency Measurement Specification Version 4.0.0*^[10] standards with terminology changes to match this document. Table 15, Table 16, and Table 17 of those documents correspond to Table 7, Table 8, and Table 9, respectively, of this document.

The goal of the Complex Workload is to provide continuity of workloads as defined in the *ISO/IEC 24091:2019*^[9] and *SNIA Emerald™ Power Efficiency Measurement Specification Version 4.0.0*^[10] standards and those defined in this document. Accordingly, the Complex Workload considers, among other things, the contribution of device mechanisms that exploit access locality, e.g., read caching.

Table 7 shows information concerning IO streams of the Complex Workload.

Table 7 – IO Streams Within the Complex Workload

IO Streams	% of workload	Read/Write Percentage	IO Size (KiB) ^a	Access Pattern	Usable Address Range
Write Stream 1	5	0/100	See Table 8 & 9	Sequential	0 % to 100 %
Write Stream 2	5	0/100	See Table 8 & 9	Sequential	0 % to 100 %
Write Stream 3	5	0/100	See Table 8 & 9	Sequential	0 % to 100 %
Read Stream 1	5	100/0	See Table 8 & 9	Sequential	0 % to 100 %
Read Stream 2	5	100/0	See Table 8 & 9	Sequential	0 % to 100 %
Read Stream 3	5	100/0	See Table 8 & 9	Sequential	0 % to 100 %
Read Stream 4	5	100/0	See Table 8 & 9	Sequential	0 % to 100 %
Read Stream 5	5	100/0	See Table 8 & 9	Sequential	0 % to 100 %
Uniform Random	6	50/50	See Table 8 & 9	Random	0 % to 100 %
Hot Band 1	28	70/30	See Table 8 & 9	Random	10 % to 18 %
Hot Band 2	14	70/30	See Table 8 & 9	Random	32 % to 40 %
Hot Band 3	7	70/30	See Table 8 & 9	Random	55 % to 63 %
Hot Band 4	5	70/30	See Table 8 & 9	Random	80 % to 88 %
^a For a DUT using native or emulated 512 B sectors, see Table 8. For a DUT using native 4 KiB sectors, see Table 9.					

The IO transfer size used within the Complex Workload is listed in Table 8 and Table 9.

Table 8 – IO Transfer Size Within the Complex Workload For 512 Byte Native Devices

Transfer Size	Streaming Write	Streaming Read	Uniform	Hot Band
512 B			2 %	2 %
1 KiB			2 %	2 %
4 KiB	29 %	29 %	27 %	27 %
8 KiB	33 %	33 %	31 %	31 %
16 KiB	6 %	6 %	5 %	5 %
32 KiB	5 %	5 %	5 %	5 %
48 KiB			1 %	1 %
56 KiB			1 %	1 %
60 KiB			2 %	2 %
64 KiB	22 %	22 %	20 %	20 %
128 KiB	3 %	3 %	2 %	2 %
256 KiB	2 %	2 %	2 %	2 %

Table 9 – IO Transfer Size Within the Complex Workload For 4 KiB Native Devices

Transfer Size	Streaming Write	Streaming Read	Uniform	Hot Band
4 KiB	29 %	29 %	31 %	31 %
8 KiB	33 %	33 %	31 %	31 %
16 KiB	6 %	6 %	5 %	5 %
32 KiB	5 %	5 %	5 %	5 %
48 KiB			1 %	1 %
56 KiB			1 %	1 %
60 KiB			2 %	2 %
64 KiB	22 %	22 %	20 %	20 %
128 KiB	3 %	3 %	2 %	2 %
256 KiB	2 %	2 %	2 %	2 %

8.5 Ready Idle Workload

The Ready Idle Workload is the empty workload, i.e., there are no I/O issued to the DUT.

9 Metrics

9.1 Primary Metrics

The primary metrics are the principal power efficiency results of the active test steps (see Subclause 9.2). They are:

- 1) EP_{CPR} : Complex Workload IO Rate Power Efficiency (IOPS/W)
- 2) EP_{CMT} : Complex Workload IO Throughput Power Efficiency (MiB/s/W)
- 3) EP_{RRR} : Random 8K Write IO Rate Power Efficiency (IOPS/W)
- 4) EP_{RRT} : Random 8K Write IO Throughput Power Efficiency (MiB/s/W)
- 5) EP_{RWR} : Random 8K Read IO Rate Power Efficiency (IOPS/W)
- 6) EP_{RMT} : Random 8K Read IO Throughput Power Efficiency (MiB/s/W)
- 7) EP_{SRR} : Sequential 256K Write IO Rate Power Efficiency (IOPS/W)
- 8) EP_{SRT} : Sequential 256K Write IO Throughput Power Efficiency (MiB/s/W)
- 9) EP_{SWR} : Sequential 256K Read IO Rate Power Efficiency (IOPS/W)
- 10) EP_{SWT} : Sequential 256K Read IO Throughput Power Efficiency (MiB/s/W)
- 11) EP_{RI} : Ready Idle Average Power (GB/W)

9.2 Power Efficiency Metrics

9.2.1 Ready Idle Metric

The Ready Idle metric represents the amount of user capacity supported per watt of power required by the DUT. It is calculated as shown in Equation 9-1, as the ratio of:

- The user capacity of the DUT, expressed in GB;
- The average power, from the Ready Idle test step, expressed in watts.

Equation 9-1: Power Efficiency, Ready Idle

$$EP_{RI} = \frac{C_R}{PA_{RI}(1800)}$$

Where:

- EP_{RI} is the power efficiency metric for the Ready Idle test step;
- C_R is the user capacity of the DUT (see Subclause 3.2.37); and
- $PA_{RI}(1800)$ is the average power over the 1800 s (30 minute) measurement interval for the Ready Idle test step.

9.2.2 Active Test Steps

The active tests steps are

- 1) Complex,
- 2) RND 8 KiB Write,
- 3) RND 8 KiB Read,
- 4) SEQ 256 KiB Write, and
- 5) SEQ 256 KiB Read.

For each active test step, the power efficiency metric represents the rate of data transfer supported per watt required by the DUT during a selected measurement interval. It is calculated, as shown in Equation 9-2, as the ratio of:

- The operations rate or throughput rate, during the measurement interval of the active test step, expressed in IOPS or MiB/s;
- The average power, during the measurement interval of the active test step, expressed in watts.

Equation 9-2: Power Efficiency, Active

$$EP_i = \frac{O_i(1800)}{PA_i(1800)}$$

Where:

- EP_i is the power metric for an active test step;
- $PA_i(1800)$ is the average power over the 1800 s (30 minute) measurement interval for the active test step i ;
- $O_i(1800)$ is the average operations rate over the 1800 s (30 minute) measurement interval for the active test step i .

9.3 Average Response Time

The average response time (ART) for an active test step is calculated over a specified time interval T in seconds.

9.4 Average Power

The average power for a test step i , $PA_i(T)$, is the arithmetic average of sampled power measurements taken over a specified time interval T , as illustrated in Equation 9-3.

Equation 9-3: Average Power

$$PA_i(T) = \frac{\sum W_s}{n}$$

Where:

- $PA_i(T)$ is the average power during test step i , taken over a time interval of T seconds;
- W_s is power in watts measured at each sampling interval s taken during the time interval T ;
- n is the number of samples gathered by the power meter during the time interval T .

9.5 Operations Rate

The operations rate for a test step i , $O_i(T)$, is a measure of the average rate of completed work over a specified time interval T .

For random workloads, i.e., Complex, RND 8 KiB Write, and RND 8 KiB Read, the operations rate is the average rate of IO operation completions during time interval T .

For sequential workloads, i.e., SEQ 256 KiB Write and SEQ 256 KiB Read, the operations rate is the average rate of data transfer in mebibytes per second (MiB/s) during time interval T .

To provide a uniform basis for the power efficiency metrics, these two different measures of operations rate are both represented by $O_i(T)$.

9.6 Measurement Interval

Each test step specifies a duration for the measurement interval. A measurement interval is a subset of a test step during which the data underlying a specific metric or calculation is gathered.

9.7 Reporting Granularity

The power efficiency metrics shall be reported to three significant digits.

DRAFT

10 Workload Generator Requirements

10.1 Purpose

The Workload Generator shall be capable of generating the IO test data necessary to carry out the measurement procedures specified by this document.

Annex C specifies the Workload Generator that shall be used.

10.2 Procedure

The workload generator shall operate within the Test Software. The workload generator shall produce IO data patterns and fill memory buffers that the Test Software uses to define and create test IO streams that are applied to a DUT.

11 Device Power Analyzer

11.1 Purpose

The Device Power Analyzer is used to accurately capture and analyze the dynamic power characteristics of DUTs tested pursuant to this specification.

The Device Power Analyzer incorporates voltage and current measurement apparatus.

Annex E contains Device Power Analyzer requirements and a list of qualified Device Power Analyzers.

11.2 Procedure

Setup the Device Power Analyzer to measure input voltage and current to the DUT.

The Device Power Analyzer shall be started prior to a test and continue to be active throughout the test.

The power data timestamp shall be synchronized with the test performance data timestamp to a resolution of at least 1 s.

The Device Power Analyzer shall be configured to take voltage and current measurement samples at a rate of at least 1000 samples per second.

The Device Power Analyzer shall report data at least once every second. The data shall include a timestamp and average voltage, current, and power consumption for each power rail.

12 General Environmental Measurement Recommendations

The DUT shall be in an environment consistent with the specified operating environment for the DUT.

DRAFT

13 Reporting Requirements

13.1 General Reporting Requirements

This Clause lists the information that shall be disclosed in every published test result identified as an SNIA Emerald™ Device Power Efficiency Measurement test result. Test results identified as an SNIA Emerald™ Device Power Efficiency Measurement test result shall be generated according to this document.

When units are specified after an item, that item shall be reported in terms of the specified unit.

The items listed in Subclauses 13.1.1, 13.1.2, 13.1.3, 13.1.4, 13.1.5 and 13.1.6 shall be included in the test report. The test report shall also include all components in the data path that affect the performance of the DUT and/or test repeatability.

13.1.1 Administrative

- 1) Test date
- 2) Report date
- 3) Test ID number
- 4) Test sponsor name
- 5) Auditor name, if applicable
- 6) Test specification version

13.1.2 Device Under Test

- 1) Manufacturer
- 2) Model number
- 3) Serial number
- 4) Firmware revision
- 5) User capacity in GB or TB
- 6) Interface, e.g. U2
- 7) Protocol, e.g., SAS or NVMe
- 8) Form factor, e.g., SFF
- 9) Media type, e.g., QLC, NAND, or Flash
- 10) Native block size, i.e., 512 B or 4096 B
- 11) Block size emulation, e.g., 512e, in use during the test (Yes or No)

13.1.3 Primary Test Metrics

- 1) EP_{CPR} : Complex Workload IO Rate Power Efficiency (IOPS/W)
- 2) EP_{CMT} : Complex Workload IO Throughput Power Efficiency (MiB/s/W)
- 3) EP_{RRR} : Random 8K Write IO Rate Power Efficiency (IOPS/W)
- 4) EP_{RRT} : Random 8K Write IO Throughput Power Efficiency (MiB/s/W)
- 5) EP_{RWR} : Random 8K Read IO Rate Power Efficiency (IOPS/W)
- 6) EP_{RMT} : Random 8K Read IO Throughput Power Efficiency (MiB/s/W)
- 7) EP_{SRR} : Sequential 256K Write IO Rate Power Efficiency (IOPS/W)
- 8) EP_{SRT} : Sequential 256K Write IO Throughput Power Efficiency (MiB/s/W)
- 9) EP_{SWR} : Sequential 256K Read IO Rate Power Efficiency (IOPS/W)
- 10) EP_{SWT} : Sequential 256K Read IO Throughput Power Efficiency (MiB/s/W)
- 11) EP_{RI} : Ready Idle Average Power (GB/W)

13.1.4 Storage Server Hardware

- 1) Manufacturer and model number
- 2) System board manufacturer and model number
- 3) For each CPU, CPU model, number of cores, clock rate, cache size, and number of sockets
- 4) System board I/O bus type and version (e.g. PCIe Gen 4)
- 5) DRAM quantity in GiB and DRAM speed
- 6) Physical Interface to DUT
- 7) HBA card connecting DUT manufacturer and model number

13.1.5 Storage Server Software

- 1) Operating system name and kernel version
- 2) File system name and version
- 3) CTS Lite version

13.1.6 Device Power Analyzer

- 1) Manufacturer and model Number
- 2) Device Power Analyzer serial number
- 3) DUT connection interposer board
- 4) DUT connection cable

13.2 Test Reporting

A SNIA Emerald™ Device Power Efficiency Measurement test result shall include the information defined by Table 10.

Table 10 – Test Results

SNIA Emerald™ Device Power Efficiency Measurement Required Test Results							
Primary Test Results							
		Workload					
Metric	Units	Complex	RND 8 KiB Write	RND 8 KiB Read	SEQ 256 KiB Write	SEQ 256 KiB Read	Ready Idle
Power Efficiency	IOPS/W	<value>	<value>	<value>	<value>	<value>	N/A
	MiB/s/W	<value>	<value>	<value>	<value>	<value>	N/A
	GB/W	N/A	N/A	N/A	N/A	N/A	<value>
Supporting Data							
Average Power	mW	<value>	<value>	<value>	<value>	<value>	<value>
Performance	IOPS	<value>	<value>	<value>	<value>	<value>	N/A
	MiB/s	<value>	<value>	<value>	<value>	<value>	N/A
Thread Count	Number	<value>	<value>	<value>	<value>	<value>	
Queue Depth	Number	<value>	<value>	<value>	<value>	<value>	
Test Parameters							
Parameter	Units	Data		Parameter	Units	Data	
Steady State	Yes/No	<value>		Purge	Yes/No	<value>	
Pre-Fill	Yes/No	<value>		Active LBA Range	Percent	<value>	

Annex A (Informative) CTS Lite Report Example

A.1 Introduction

CTS Lite generates a series of reports, examples of which are shown in this Clause.

CTS Lite automatically generates a Read-only Audit file that contains report plots labelled P1 – P17. Figure A.1 shows the list of report plots generated by the Test Software. Corresponding reports are also contained in editable Microsoft Excel format report files.

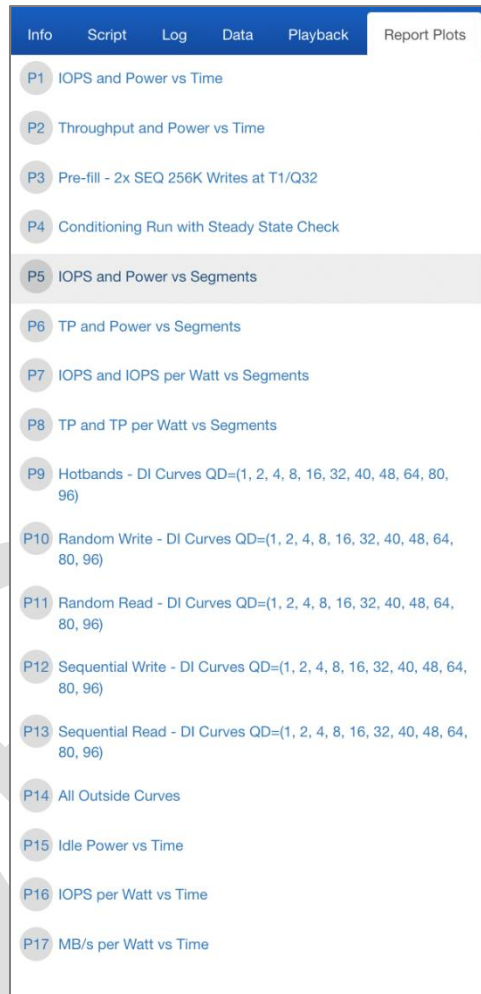


Figure A.1 – Report Plots Tab

CTS Lite generates 17 reports plots and are grouped as follows:

- | | | |
|--------------------|---------------|-----------------------------|
| 1. Test Validation | Test Flow: | P1, P2, P3, P4, P5, P6 |
| 2. Data Reporting | Test Results: | P7, P8, P15, P16, P17 |
| 3. Data Analysis | DI Curve: | P9, P10, P11, P12, P13, P14 |

The following Subclauses present an example of each of the 17 report plots.

A.2 P1: IOPS and Power vs. Time

An example of P1 IOPS and Power vs. Time is shown in Figure A.2.

P1 shows the continuous and uninterrupted test flow in IOPS from the start of the test to the end of the test. The Ready Idle segment shows only power, indicated by a dotted line.




Figure A.2 – IOPS and Power vs. Time

A.3 P2: Throughput and Power vs. Time

An example of P2 Throughput and Power vs Time is shown in Figure A.3.

P2 shows the continuous and uninterrupted test flow in MiB/s from the start of the test to the end of the test. The Ready Idle segment shows only power, indicated by a dotted line.




Figure A.3 – Throughput and Power vs. Time

A.4 P3: Pre-Fill Test Step

An example of P3 Pre-fill – 2x SEQ 256K Writes at T1/Q32 is shown in Figure A.4.

P3 shows the throughput of the DUT while twice the user capacity is filled using SEQ 256 KiB Writes at a TOIO of 32 (Thread Count 1 / Queue Depth 32).




Figure A.4 – Pre-Fill

A.5 P4: Steady State Verification

An example of P4 Condition Run with Steady State Check is shown in Figure A.5.

P4 documents the convergence of the Complex Workload to steady state.




Figure A.5 - Conditioning Run with Steady State

A.6 P5: IOPS and Power vs. Test Step

An example of P5 IOPS and Power vs Test Steps (Segments) is shown in Figure A.6.

P5 shows the average IOPS, average power, and TOIO for each active test step workload measurement interval and average power only for the Ready Idle measurement interval.




Figure A.6 – IOPS and Power vs. Test Steps (Segments)

A.7 P6: Throughput and Power vs Segment

An example of P6 TP and Power vs Segments is shown Figure A.7.

P6 shows the average throughput, average power, and TOIO (Total OIO = Thread Count times Queue Depth) for each active test step measurement interval and average power only for the Ready Idle measurement interval.




Figure A.7 – Throughput and Power vs. Segments

A.8 P7: IOPS And IOPS per Watt vs. Test Steps (Segments)

An example of P7 IOPS and IOPS per Watt vs Segments is shown in Figure A.8.

P7 shows the average IOPS and average IOPS/watt for each workload.




Figure A.8 – IOPS and IOPS/Watt vs. Test Steps (Segments)

A.9 P8: Throughput and Throughput/Watt Vs. Test Steps (Segments)

An example of P8 TP and TP per Watt vs Test Steps (Segments) is shown in Figure A.9.

P8 shows average throughput and average throughput/watt for each workload.




Figure A.9 – Throughput and Throughput/Watt vs. Test Steps (Segments)

A.10 P9: Complex – DI Curves

An example of P9 Multi-stream – DI Curves is shown in Figure A.10.

P9 shows a DI Curve for each Thread Count / Queue Depth value of the Complex Workload and the DI OS Curve.




Figure A.10 – Complex Workload – DI Curves

A.11 P10: Random Write – DI Curves

An example of P10 Random Write – DI Curves is shown in Figure A.11.

P10 shows a DI Curve for each TOIO value of the RND 8 KiB Write workload and the DI OS Curve.




Figure A.11 – Random Write – DI Curves

A.12 P11: Random Read – DI Curves

An example of P11 Random Read – DI Curves is shown in Figure A.12.

P11 shows DI Curves for each TOIO value of the RND 8 KiB Read workload and the DI OS Curve.




Figure A.12 – Random Read – DI Curves

A.13 P12: Sequential Write – DI Curves

An example of P12 Sequential Write – DI Curves is shown in Figure A.13.

P12 shows DI Curves for each TOIO value of the SEQ 256 KiB Write workload and the DI OS Curve




Figure A.13 – Sequential Write – DI Curves

A.14 P13: Sequential Read – DI Curves

An example of P13 Sequential Read – DI Curves is shown in Figure A.14.

P13 shows DI Curves for each TOIO value of the SEQ 256 KiB Read workload and the DI OS Curve.




Figure A.14 – Sequential Read – DI Curves

A.15 P14: All Outside DI Curves

An example of P14 All Outside Curves is shown in Figure A.15.

P14 shows an Outside DI Curve for each workload.




Figure A.15 – All Outside DI Curves

A.16 P15: Ready Idle Power vs. Time

An example of P15 Ready Idle Ave Power v Time is shown in Figure A.16.

P15 shows Ready Idle Power vs. Time.




Figure A.16 – Ready Idle Power vs. Time

A.17 P16: IOPS/Watt vs. Time

An example of P16 IOPS per Watt vs Time is shown in Figure A.17.

P16 shows continuous IOPS/watt vs. time for each test step.




Figure A.17 – IOPS/Watt vs. Time

A.18 P17: MiB/s/Watt vs. Time

An example of P17 MiB/s per Watt vs Time is shown in Figure A.18.

P17 shows continuous MiB/s/watt vs. time for each test step.




Figure A.18 – MiB/s/Watt vs. Time

Annex B (Informative) Test Methodologies

B.1 Introduction

Principles guiding test processes derived from existing standards^{[9][10][7][8]} include Steady State algorithms, block size sequencing, NAND flash-based FOB and write history, impact of thread count and queue depth and the impact of workloads on HDD and SSD device performance.

B.2 Test Flow and Methodology

The basic test flow and workloads are based on the ISO/IEC 24091:2019^[9] and *SNIA Emerald™ Power Efficiency Measurement Specification Version 4.0.0*^[10] for storage systems. Table B.1 shows the correspondence between the names used for the workloads in this document and those used in ISO/IEC 24091:2019 and in *SNIA Emerald™ Power Efficiency Measurement Specification*.

Table B.1 – Workload Name Mapping

ISO/IEC 24091:2019 and SNIA Emerald Specification	This Document
Hot Band IO Profile	Complex Workload
Random Read	RND 8 KiB Read Workload
Random Write	RND 8 KiB Write Workload
Sequential Read	SEQ 256 KiB Read Workload
Sequential Write	SEQ 256 KiB Write Workload
ready idle	Ready Idle Workload

The 12 hour Conditioning Run of standards^{[9][10]} is modified to a minimum of five 1800 s (30 minutes) rounds totalling 9000 s (2.5 hours) up to a maximum duration of twenty five 1800 s (30 minutes) rounds totalling 45000 s (12.5 hours), depending on the Steady State determination criteria.

The 20 ms average response time ceiling of standards^{[9][10]} is adopted for the Conditioning Run IOPS performance.

A self-optimizing TOIO sweep eliminates the need (implicit in [9] and [10]) to run multiple tests to determine the TOIO setting with the highest IOPS at less than 20 ms average response time.

B.3 Substitutable Test Workloads for Future Revisions

This basic test flow applies five specific workloads (Complex, Random 8 KiB Write, Random 8 KiB Read, Sequential 256KiB Write and Sequential 256 KiB Read) and is intended to provide easy comparison to power efficiency measurements defined by existing standards^{[9][10]} for measuring storage system power efficiency.

The test flow, methodologies and reporting requirements of this document are structured to enable future revisions where one or more of the test workloads are substituted with new workloads deemed relevant and significant by industry and governing bodies. The Test Software allows any given workload to be substituted with a different test workload and provides the ability to set test parameters and settings for future tests released under this test methodology.

B.4 Steady State Algorithm

The 5 Rounds Steady State Determination algorithm in the Conditioning run ensures the DUT achieves relatively time invariant performance of the Complex Workload.

B.5 Demand Intensity Curve

The Self-optimizing TOIO Sweep is based on a Demand Intensity Curve for IOPS and Average Response Times plotted as a function of varying TOIO.

Demand Intensity (DI), a.k.a., TOIO, is the product of Thread Count and Queue Depth for a given applied workload. The plotting of IOPS at different TOIO values against IOPS and average response time generates a Demand Intensity Curve. As DI is increased, the IO rate in IOPS and the average response times (ART) increases until the storage becomes “saturated”. At this saturation point, IOPS no longer increases (and may regress or “fold back”) while ART begins to dramatically rise.

By setting a given ART ceiling, the Test Software determines the Thread Count and Queue Depth settings that generate the highest IOPS for a given workload while staying below the designated ART ceiling.

Figure B.1 is an example Demand Intensity Curve for an 1800 GB HDD. IOPS are on the x-axis with ART in ms is on the y-axis. The DI Thread Count points are 2, 4, 8, 16, 32, 48, 64, 80 and 96 with the ART ceiling shown at 20 ms. Queue Depth is set to 1 for each Thread Count in this example.

In this DI Curve, the optimal TOIO setting is Thread Count of 48 which yields 2559 IOPS at an ART of 18.75 ms. Accordingly, the measurement interval for this Complex Workload as determined by the self-optimizing TOIO sweep would be set at a TOIO of 48.

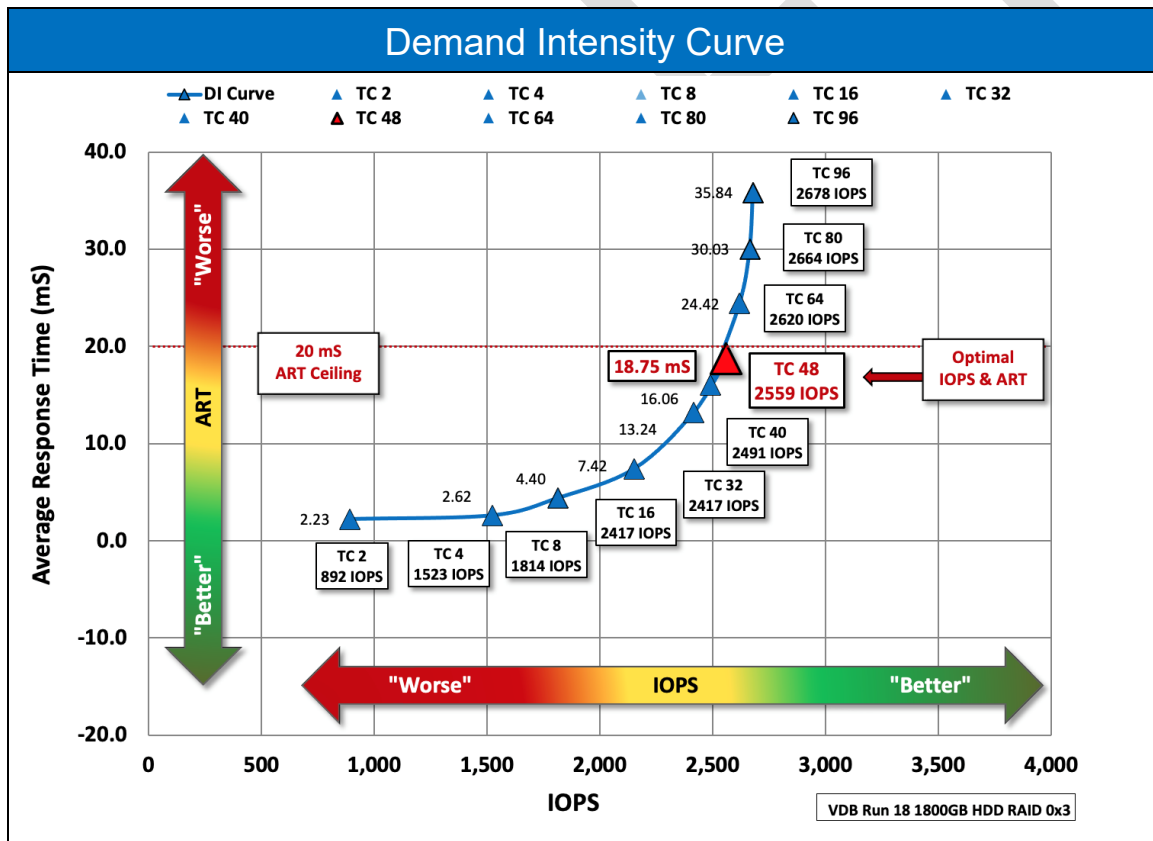


Figure B.1 – Demand Intensity Curve

B.6 Demand Intensity Outside Curve

A Demand Intensity Curve is generated for each Thread Count and Queue Depth combination. Therefore, multiple DI Curves are generated for a multi-Thread Count, multi-Queue Depth self-optimizing TOIO sweep.

Annex B (Informative) Test Methodologies

A Demand Intensity Outside Curve (DI OS Curve) is the best fit curve for all of the individual Thread Count/Queue Depth DI Curves. A DI OS Curve shows the increasing IOPS and ART relative to a given ART ceiling and is useful for observing the overall performance saturation for a given workload. Figure B.2 is an example plot of an HDD device with a Complex DI Curve for each TOIO value.




Figure B.2 – DI Curve Complex

Figure B.3 is an example DI OS Curve for the Complex Workload. The individual OS Curves are de-selected to only show the Complex DI OS Curve. The optimal TOIO value is 32 at Thread Count = 8 and Queue Depth = 4.




Figure B.3 – DI OS Curve Complex

Figure B.4 is an example All DI OS Curves (Complex, Random 8 KiB Write, Random 8 KiB Read, Sequential 256 KiB Write and Sequential 256 KiB Read).




Figure B.4 – All DI OS Curves

B.7 Report Generation

The Test Software automatically generates a read-only html audit file report and an editable Microsoft Excel format report file. Both report formats provide required information for test review and process validation. The audit file report also provides a dynamic playback file. Figure B.5 is an example screen shot of a dynamic playback file.

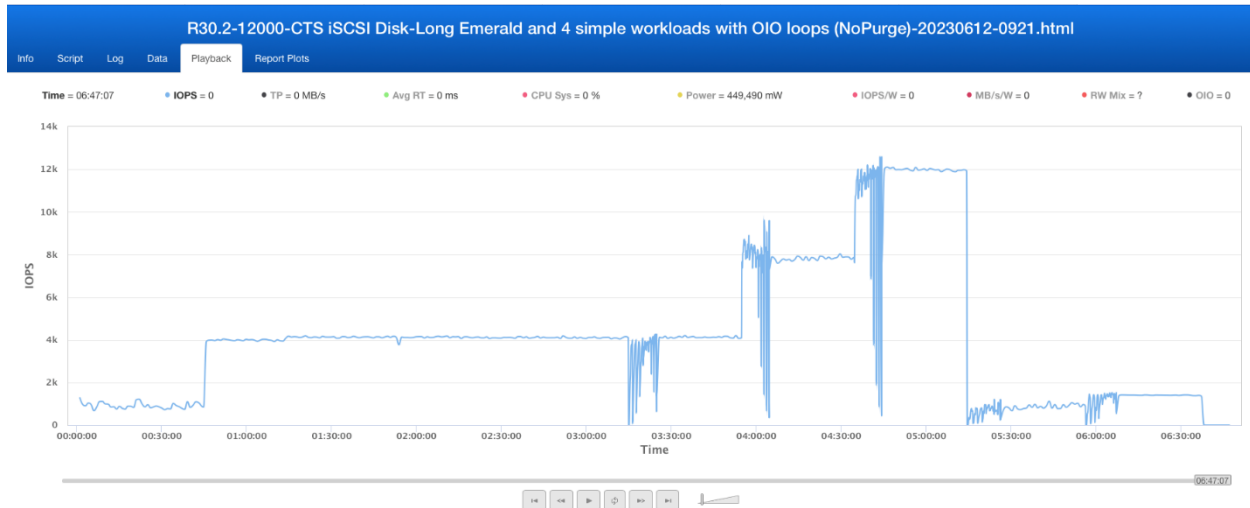


Figure B.5 – Audit File Playback

Annex C (Normative) Workload Generator

C.1 Required Workload Generator

The Test Software (see Annex D) contains the Workload Generator that shall be used.

C.2 Workload Generator Requirements

The Workload Generator shall be able to:

- 1) test storage devices pursuant to the requirements of this document;
- 2) act as workload stimulus generator as well as data recorder;
- 3) issue Random and Sequential block level I/O;
- 4) restrict LBA accesses to a particular range of available user LBA space;
- 5) limit “total unique LBAs used” to a specific value;
- 6) randomly distribute a number of equally sized LBA segments across the test active range;
- 7) set R/W percentage mix;
- 8) set Random/Sequential IO percentage mix
- 9) set IO Transfer Size;
- 10) set data pattern compression ratio;
- 11) generate and maintain multiple outstanding IO requests and ensure that all steps in the test sequence can be executed immediately one after the other, to ensure that devices are not recovering between processing steps, unless recovery is the explicit goal of the test;
- 12) provide output that can be used to derive IOPS, MiB/s, maximum response time, and average response time (if TOIO is 1) within some measurement period;
- 13) run a TOIO sweep (in Thread Count and Queue Depth), save results in an array, generate a Demand Intensity Curve (DI Curve) based on a 20 ms ART Ceiling, and set the workload segment TOIO to the values ascertained in the aforementioned DI Curve;
- 14) integrate and aggregate power measurements (time stamp, voltage (V), and current (mA); and
- 15) All IOs issued shall complete successfully.

The random function for generating random LBA numbers during random IO tests shall be:

- 1) seedable;
- 2) have an output ≥ 48 bits; and
- 3) deliver a uniform random distribution independent of capacity.

Annex D (Normative) Required Test Software

CTS Lite is the Test Software that shall be used for a test performed in accordance with this document.

CTS Lite is an IO performance test software product developed by Calypso Systems, Inc. This document describes the use of CTS Lite as the Workload Generator. CTS Lite consists of device block access test software and an associated workload generator. CTS Lite executes the test specified by this document.

The CTS Lite software specified on the *Download Material for SNIA Emerald™ Testing*^[12] web page for use with this document shall be used as the Workload Generator. The web page also provides access to the CTS Lite device user guide and other related testing materials.

CTS Lite single seat and enterprise licenses are available for purchase from Calypso Systems, Inc.^[13] The CTS Lite web page^[13] shows license terms and fees and provides access to the CTS Lite software and related materials.

Annex E (Normative) Device Power Analyzer

E.1 Purpose

The purpose of this Clause is to outline the requirements for the Device Power Analyzer for measuring power consumption in the DUT. The Device Power Analyzer is used to accurately capture and analyze the dynamic power characteristics of storage devices tested pursuant to this document.

E.2 General Description

The Device Power Analyzer measures the voltage and current used by the DUT.

The Device Power Analyzer shall report the average voltage, average current, and average power consumption of each rail with an associated timestamp each reporting period.

The Device Power Analyzer shall aggregate the power consumption of all of the rails of the DUT and report the aggregated average power consumption with an associated timestamp each reporting period.

E.3 General Specifications

Table E.1 defines the specifications that the Device Power Analyzer shall meet.

Table E.1 – Device Power Analyzer Requirements

Device Power Analyzer Specifications		
Item	Requirement	Description
Meter Current Range		Shall cover the current levels of the DUT during read, write, and idle states.
Voltage Accuracy	$2 \text{ mV} \pm 1\% \text{ of reading}$	
Current Resolution	$\leq 1 \text{ mV}$	
Current Accuracy	$2 \text{ mA} \pm 1\% \text{ above } 1 \text{ mA}$ $25 \text{ uA} \pm 1\% \text{ below } 1 \text{ mA}$	
Current Resolution	$\leq 1 \text{ mA}$	
Meter Sampling Rate	$\geq 1 \text{ kHz}$	
Reporting Period	$\leq 1 \text{ s}$	
Integration with Data Acquisition Systems	Integration with the Test Software or provide a data file accepted by the Test Software.	

E.4 Qualified Device Power Analyzers

A Device Power Analyzer from the list in Table E.2 shall be used.

Table E.2 – Qualified Device Power Analyzers

Qualified Power Analyzers		
Manufacturer	Model	Description
Quarch Technology ^[14]	QTL1999	Single port Programmable power module
Quarch Technology ^[14]	QTL1995	Multi-power Programmable Power Module
Quarch Technology ^[14]	QTL2312	Single port Power Analysis Module

Annex F (Normative) Device Power Supply Requirements

F.1 Purpose

The purpose of this Clause is to outline the requirements for the power supplied the DUT.

F.2 General Description

The power supply shall provide the voltage and current required by each of the power rails of the DUT during a test.

F.3 General Specifications

Table F.1 defines the specifications that the power supply shall meet.

Table F.1 – Power Supply Requirements

Item	Requirement
Power Supply Voltage Range	Provide the voltage levels of each of the power rails of the DUT remain within the DUT specified voltage range throughout a test.
Power Supply Current	Provide requirements of each power rail of the DUT.
Rail Voltage Accuracy	$\pm 5\%$ of rail nominal voltage.

Bibliography

- [1] SNIA (2018). SFF-8301 *Specification for 3.5" Form Factor Drive Dimensions*. Available from <https://www.snia.org/technology-communities/sff/specifications>.
- [2] SNIA (2018). SFF-8201 *Specification for 2.5" Form Factor Drive Dimensions*. Available from <https://www.snia.org/technology-communities/sff/specifications>.
- [3] SNIA (2021). SFF-TA-1006 *Enterprise and Datacenter 1U Short Device Form Factor Pin (E1.S)*. Available from <https://www.snia.org/technology-communities/sff/specifications>.
- [4] SNIA (2023). SFF-TA-1008 *Enterprise and Datacenter Standard Form Factor*. Available from <https://www.snia.org/technology-communities/sff/specifications>.
- [5] SNIA (2025). SFF-TA-1009 *Enterprise and Datacenter Standard Form Factor Pin and Signal Specification*. Available from <https://www.snia.org/technology-communities/sff/specifications>.
- [6] SNIA (2017). SFF-8639 *Specification for Multifunction 6X Unshielded Connector*. Available from <https://www.snia.org/technology-communities/sff/specifications>.
- [7] SNIA (2020). SNIA Standard *Solid State Storage (SSS) Performance Test Specification (PTS) Version 2.0.2*. Available from https://www.snia.org/tech_activities/standards/curr_standards/pts.
- [8] SNIA (2018). SNIA Standard *Real World Storage Workload (RWSW) Performance Test Specification for Datacenter Storage Version 1.0.7*. Available from <https://www.snia.org/rws>.
- [9] ISO (2019). International Standard ISO/IEC 24091:2019 *Information technology — Power efficiency measurement specification for data center storage*. Available from <https://www.iso.org/>.
- [10] SNIA (2020). SNIA Standard *SNIA Emerald™ Power Efficiency Measurement Specification Version 4.0.0*. Available from https://www.snia.org/tech_activities/standards/curr_standards/emerald.
- [11] SNIA. *The SNIA Dictionary*. Available from <https://www.snia.org/dictionary>.
- [12] SNIA. *Download Material for SNIA Emerald™ Testing*. Available from <https://www.snia.org/emerald/download>.
- [13] Calypso Systems, Inc. *CTS Lite*. Available from <https://www.ctlite.com>.
- [14] Quarch Technology Ltd. *Quarch Technology*. Available from <https://www.quarch.com>.