



What has Changed?

SNIA Emerald™ Power Efficiency Measurement Specification V2.1.1 to V3.0.2



SNIA Emerald™

July 11, 2018

About the SNIA

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

About the SNIA Green Storage Initiative

SNIA's Green Storage Initiative (GSI) is focused on advancing energy efficiency and conservation for data center networked storage technologies in an effort to minimize the environmental impact of data storage operations. SNIA's Green Storage activities take place in two separate working bodies, the SNIA Green Storage Technical Working Group (TWG) and the Green Storage Initiative. The TWG is focused on developing repeatable and fair test methodologies and metrics for enterprise storage systems through which energy consumption and efficiency can be measured. The Green Storage Initiative is focused on creating and publicizing best practices to the industry for energy efficient storage networking, promoting storage-centric applications that reduce storage footprint and associated power requirements, and educating regulatory bodies and testing organizations to apply test methodologies and best practices.

About the SNIA Emerald™ Program

The SNIA Emerald™ Program is a vendor-neutral, public service to the storage industry, IT community, and regulatory body community that is sponsored and operated by the SNIA GSI. The program supports the use and evolution of the SNIA Emerald™ Power Efficiency Measurement Specification. The measurement procedure and test metrics are documented in the SNIA Emerald™ Power Efficiency Measurement Specification, which is developed, released, and maintained by the Green Storage TWG under the guidance of the GSI. GSI produces education programs and materials for testers to consistently and competently use the SNIA Emerald™ Power Efficiency Measurement Specification.

The EPA ENERGY STAR® Data Center Storage Program is based on the methodology defined in the Specification and offers another vehicle for publication of product test results created in accordance with the Specification. Some national regulatory bodies cross-reference the EPA ENERGY STAR Program for their needs, while other national regulatory bodies around the world are aware of the SNIA Emerald™ Specification and in the future, may base their programs on the methodology and metrics.



Copyright © 2018 Storage Networking Industry Association.

The information contained in this publication is subject to change without notice. This guide represents a "best effort" attempt by the SNIA Green Storage Technical Working Group to provide guidance to those implementing the *SNIA Emerald™ Power Efficiency Measurement Specification*, and the guide may be updated or replaced at any time. The SNIA shall not be liable for errors contained herein.

Suggestions for revisions to this guide and questions concerning implementation of the *SNIA Emerald™ Power Efficiency Measurement Specification* can be directed (via email) to emerald@snia.org.

What has Changed? SNIA Emerald™ Power Efficiency Measurement Specification V2.1.1 to V3.0.2

[This page deliberately blank.]



Contents

| | |
|--|----------|
| I Introduction | I |
| 1.1 Audience | I |
| 1.2 References..... | I |
| 2 Changes from V2.1.1 to V3.0.2..... | 2 |
| 2.1 Title page and Usage..... | 2 |
| 2.2 Scope | 2 |
| 2.3 Normative References..... | 3 |
| 2.4 Definitions, Symbols, Abbreviations, and Conventions (Clause 4)..... | 3 |
| 2.4.1 Definitions (Subclause 4.2)..... | 3 |
| 2.4.2 Acronyms and Abbreviations (Subclause 4.3) | 3 |
| 2.4.3 Expression of Provisions (Subclause 4.4) | 3 |
| 2.4.4 Conventions (Subclause 4.5) | 3 |
| 2.5 Taxonomy (Clause 5) | 3 |
| 2.5.1 Introduction (Subclause 5.1)..... | 3 |
| 2.5.2 Category Attributes (Subclause 5.2.2) | 4 |
| 2.5.3 Taxonomy Rules (Subclause 5.4)..... | 4 |
| 2.5.4 Online Category (Subclause 5.5)..... | 5 |
| 2.5.5 Near-Online Category (Subclause 5.6)..... | 5 |
| 2.5.6 Removable Media Library Category (Subclause 5.7) | 5 |
| 2.5.7 Virtual Media Library Category (Subclause 5.8) | 5 |
| 2.6 Capacity Optimization (Clause 6)..... | 5 |
| 2.6.1 COMs Characterized (Subclause 6.3) | 5 |
| 2.7 Test Definition and Execution Rules (Clause 7) | 6 |
| 2.7.1 File Access Execution Overview (Subclause 7.1.2) | 6 |
| 2.7.2 Configuration Guidelines (Subclause 7.2.1, Figures 1 and 2) | 6 |
| 2.7.3 Power (Subclause 7.2.7) | 8 |
| 2.7.4 Environmental (Subclause 7.2.8.1 to 7.2.8.3)..... | 8 |
| 2.7.5 Power and Temperature Instrumentation (Subclause 7.2.9) | 8 |
| 2.7.6 Workload Generator (Subclause 7.2.10) | 8 |
| 2.7.7 Instrumentation (Subclause 7.2.11)..... | 9 |
| 2.7.8 IO Profiles (Subclause 7.2.12) | 9 |
| 2.7.9 Block Access Online and Near-Online Tests (Subclause 7.3) | 9 |

| | | |
|--------|---|----|
| 2.7.10 | File Access Online and Near-Online Active Test (Subclause 7.4)..... | 9 |
| 2.7.11 | Block and File Access Ready Idle test (Subclause 7.5) | 9 |
| 2.7.12 | Block and File Access Capacity Optimization Test (Subclause 7.6)..... | 9 |
| 2.7.13 | Removable Media Library Testing (Subclause 7.7) and Virtual Media Library Testing (Subclause 7.8) | 9 |
| 2.8 | Metrics (Clause 8) | 9 |
| 2.8.1 | File Access Primary Metrics (Subclause 8.3)..... | 9 |
| 2.8.2 | File Access Active Test (Subclause 8.4.3)..... | 9 |
| 2.9 | Disclosure Requirements (Clause 9) | 10 |
| 2.10 | A Note on Annexes..... | 10 |
| 2.11 | Annex A (normative)..... | 10 |
| 2.12 | Annex B (normative) | 10 |
| 2.13 | Annex C (normative)..... | 11 |
| 2.14 | Annex D (normative)..... | 11 |
| 2.15 | Annex E (normative)..... | 11 |
| 2.16 | Annex F (normative)..... | 12 |
| 2.17 | Annex G (normative)..... | 12 |
| 2.18 | Annex H (informative)..... | 12 |
| 2.19 | Annex I (normative)..... | 12 |
| 2.20 | Annex J (informative)..... | 12 |
| 2.21 | Bibliography..... | 12 |



I Introduction

This document provides a summary overview of the changes made to the SNIA Emerald™ Power Efficiency Measurement Specification V2.1.1 (referred to within this document as simply *Specification V2.1.1*), developed as part of the SNIA Emerald™ Program, which result in the SNIA Emerald™ Power Efficiency Measurement Specification V3.0.2 (referred to within this document as simply *Specification V3.0.2*). *Specification V3.0.2* is a major technical content upgrade from *Specification V2.1.1*.

Any conflict between this document and the *Specifications* shall defer to the *Specifications*.

1.1 Audience

The target audience of this document includes organizations and individuals planning for and testing in accordance with *Specification V3.0.2*. An organization or individual performing such testing is referred to as a test sponsor.

Test sponsors are urged to carefully read the *Specification V3.0.2* in order to become aware of all changes.

Important note: If the tester or evaluators are not familiar at least with *Specification V2.1.1*, we highly encourage that prior to reading this document you become familiar with *Specification V3.0.2*. Then this paper can be used to gain an understanding of the evolution of the specification.

1.2 References

Specification V2.1.1 and *Specification V3.0.2* are available on the SNIA website:
http://www.snia.org/tech_activities/standards/curr_standards/emerald.

2 Changes from V2.1.1 to V3.0.2

In response to an increase in demand for storage products offering file based data access, the SNIA Green TWG, has developed a series of extensions in *Specification V3.0.2*. These extensions include file access workloads, test procedures, metric collection, and reporting.

Throughout the text of *Specification V3.0.2*, more precise language has been introduced, for example “Work Load Generator” and “Solution Under Test”. When you read any of these terms, please refer to the definitions section of *Specification V3.0.2*.

NOTE 1: All references (to Clauses, figures, and tables) in this Clause pertain to the *Specification V3.0.2* unless otherwise noted. As a part of the evolution of this document some of the tables and references are updated and will be found in a more appropriate Clause of the document.

NOTE 2: Test sponsors should not expect any difference in SUT metric results between the Clauses of *Specification V2.1.1* and the block access Clauses of *Specification V3.0.2*, given all other test conditions are the same (e.g., SUT configuration, initial conditions, test environment, benchmark driver type, revision, script parameters, etc.). The major difference is the addition of the File Access workloads, test procedures, and metric collection.

2.1 Title page and Usage

Updated the Abstract and publication explanation for *Specification V3.0.2*.

Introduced SPEC materials used.

Other non-content significant updates were done, for example, copyright updates, acknowledgements, etc.

2.2 Scope

Next is a list of subclauses that were updated so they would reflect the main additions and modifications that *Specification V3.0.2* has compared with *Specification V2.1.1*. The updates are short enough that is not necessary to include them here, however in the interest of avoiding confusion it is highly recommended that they be read at least once.

- ✓ Abstract
- ✓ Introduction
- ✓ Current revision
- ✓ Purpose
- ✓ Measurement guidelines



2.3 Normative References

Updates were done to better indicate the reader where these can be found. Some old references were eliminated. The current reference table is shown here:

| Author/Owner | Title | Revision | URL |
|--------------|----------------------------|------------------------|---|
| ISO/IEC | ISO/IEC Directives Part II | Edition 7.0 2016-05 | https://www.iso.org/directives-and-policies.html |

2.4 Definitions, Symbols, Abbreviations, and Conventions (Clause 4)

2.4.1 Definitions (Subclause 4.2)

Due to the short updates, it is suggested that the reader go directly to this subclause.

2.4.2 Acronyms and Abbreviations (Subclause 4.3)

SUT was removed.

2.4.3 Expression of Provisions (Subclause 4.4)

This was added based on ISO/IEC Directives Part 2.

2.4.4 Conventions (Subclause 4.5)

Updates were made to the section. Due to its importance, it is highly recommended for the reader to visit them.

2.5 Taxonomy (Clause 5)

2.5.1 Introduction (Subclause 5.1)

It is important to note the omission of the Adjunct Product and the Interconnect Element as shown in the table below:

| Category Level | Online (see 5.5) | Near-Online (see 5.6) | Removable Media Library (see 5.7) | Virtual Media Library (see 5.8) |
|----------------------------------|------------------|-----------------------|-----------------------------------|---------------------------------|
| Consumer/ Component ^a | Online 1 | Near- Online 1 | Removable 1 | Virtual 1 |
| Low-end | Online 2 | Near- Online 2 | Removable 2 | Virtual 2 |
| Mid-range | Online 3 | Near- Online 3 | Removable 3 | Virtual 3 |
| | Online 4 | | | |
| High-end | Online 5 | Near- Online 5 | Removable 5 | Virtual 5 |
| Mainframe | Online 6 | Near- Online 6 | Removable 6 | Virtual 6 |

^a Entries in this level of taxonomy include both consumer products and data-center components (e.g., stand-alone tape drives)

2.5.2 Category Attributes (Subclause 5.2.2)

Due to the removals mentioned in the previous section, the attributes table is simpler as shown below:

| Attribute | Category | | | |
|--------------------------|--------------------|--------------------|-------------------------|-----------------------|
| | Online | Near-Online | Removable Media Library | Virtual Media Library |
| Access Pattern | Random/ Sequential | Random/ Sequential | Sequential | Sequential |
| MaxTTFD (t) ^a | t < 80 ms | t > 80 ms | t > 80 ms t < 5 min | t < 80 ms |

2.5.3 Taxonomy Rules (Subclause 5.4)

Important updates about how the product shall be classified were introduced.



2.5.4 Online Category (Subclause 5.5)

The best way to observe the main differences is by looking at the updated table:

| Attribute | Classification | | | | | |
|--|---------------------------------------|---------------------------------------|--------------------|--------------------|--------------------|--------------------|
| | Online 1 | Online 2 | Online 3 | Online 4 | Online 5 | Online 6 |
| Access Pattern | Random/ Sequential | Random/ Sequential | Random/ Sequential | Random/ Sequential | Random/ Sequential | Random/ Sequential |
| MaxTTFD (t) | t < 80 ms | t < 80 ms | t < 80 ms | t < 80 ms | t < 80 ms | t < 80 ms |
| Connectivity | Not Specified | Connected to single or multiple hosts | Network-connected | Network-connected | Network-connected | Network-connected |
| Consumer/ Component | Yes | No | No | No | No | No |
| Integrated Storage Controller | Optional | Optional | Required | Required | Required | Required |
| Storage Protection | Optional | Optional | Required | Required | Required | Required |
| No SPOF | Optional | Optional | Optional | Required | Required | Required |
| Stable storage support | Optional, unless Required by protocol | Optional, unless Required by protocol | Required | Required | Required | Required |
| Non-Disruptive Serviceability | Optional | Optional | Optional | Optional | Required | Required |
| FBA/CKD Support | Optional | Optional | Optional | Optional | Optional | Required |
| Maximum Supported Configuration ¹ | ≥ 1 | ≥ 4 | ≥ 12 | > 100 | >400 | >400 |

The Stable Storage Support condition has been added and the User Accessible Data condition removed.

2.5.5 Near-Online Category (Subclause 5.6)

Similarly to the online category, the User Accessible Data condition was removed.

2.5.6 Removable Media Library Category (Subclause 5.7)

Similarly to the online category, the User Accessible Data condition was removed.

2.5.7 Virtual Media Library Category (Subclause 5.8)

Similarly to the online category, the User Accessible Data condition was removed.

2.6 Capacity Optimization (Clause 6)

2.6.1 COMs Characterized (Subclause 6.3)

Was updated to include clarity on how it works. Please refer to *Specification V3.0.2* for more details.

2.7 Test Definition and Execution Rules (Clause 7)

This Clause has the most changes across the specification. The previous versions of the specification did not address file access storage.

Unless otherwise indicated all the changes indicated here are relevant to file access storage. Since there are four different scripts to run and, maybe, adjusted (transition activities between workloads) for each of their runs, this can add significant time to the overall test for file access systems. **Please refer to the User Guide for the SNIA Emerald™ Power Efficiency Measurement Specification V3.0.2 for detailed information.** The User Guide is available at http://sniaemerald.com/download/Spec_V3.0.

2.7.1 File Access Execution Overview (Subclause 7.1.2)

Added a file access execution overview that indicates the three basic premises for a file access Emerald test evaluation. If, in the past, you have evaluated a file system capable of block access, it had to be tested as an online block storage system. Now, the sponsor has the option to evaluate it as a file system.

2.7.2 Configuration Guidelines (Subclause 7.2.1, Figures 1 and 2)

Figure 1 is an improved sample configuration for **block**.

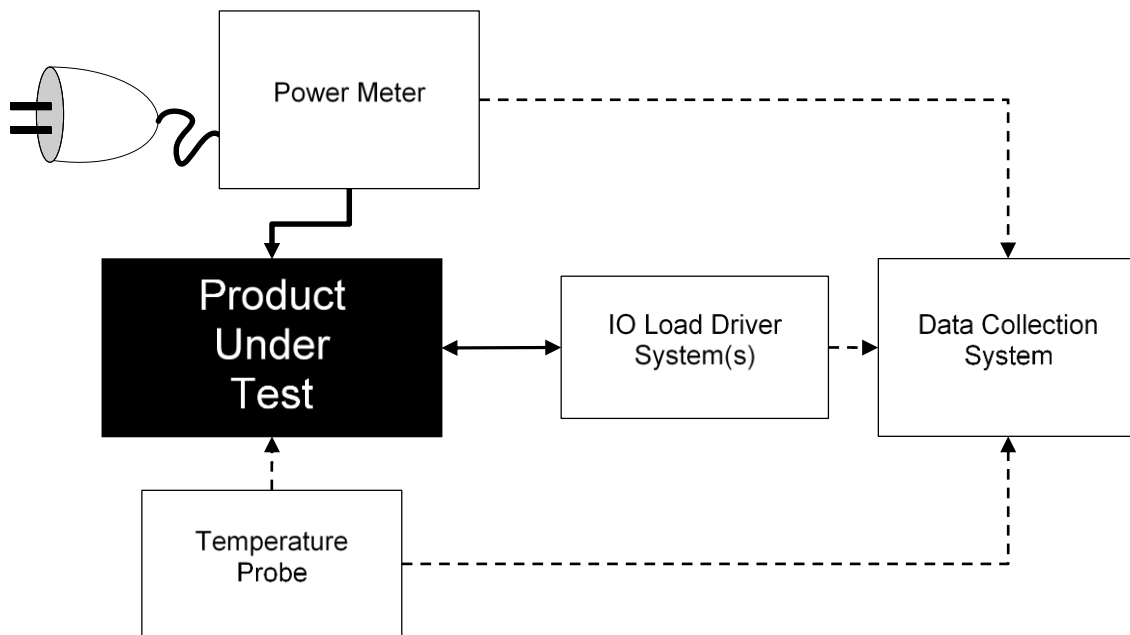
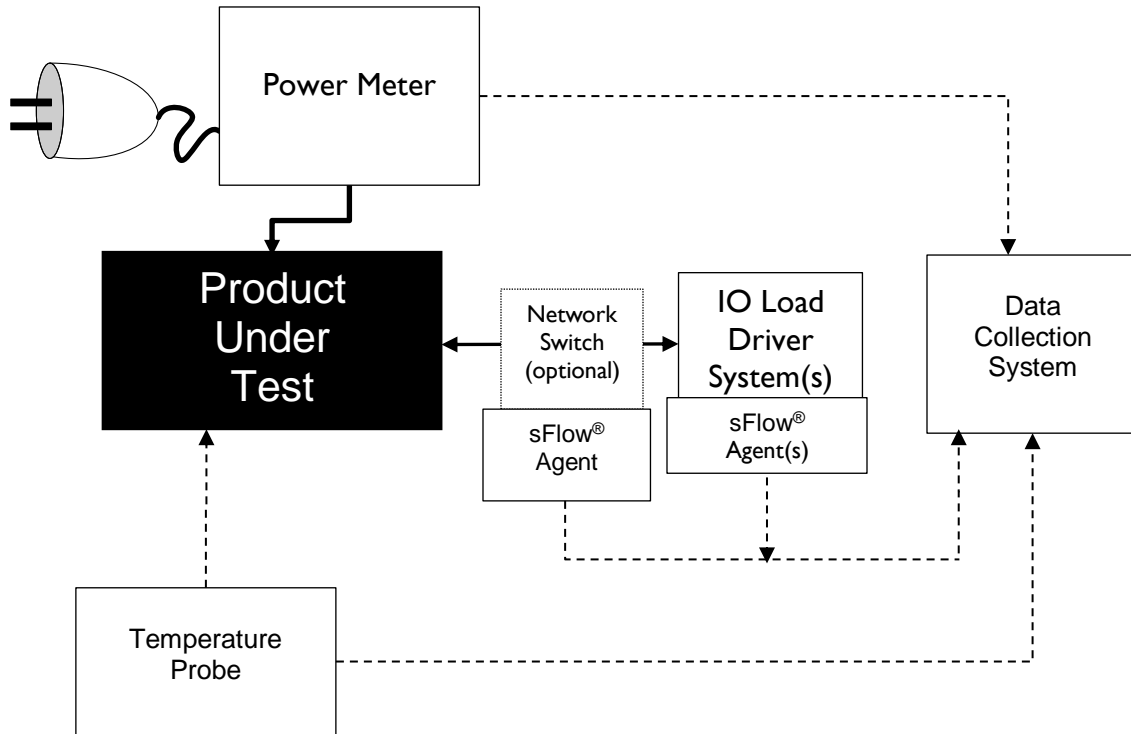




Figure 2 adds a diagram that depicts the connectivity for a file access storage system. Please observe that the key difference with block is the sFlow® connection.



Product Under Test Consistency, No Non-Test Consistency, and Test Sequence (Subclauses 7.2.4 – 7.2.6)

These Clauses add clarity to some of the questions a practitioner may have on what is and what is not allowed during an evaluation run.

2.7.3 Power (Subclause 7.2.7)

IMPORTANT: This is for all storage, BLOCK ACCESS and FILE ACCESS.

The table power requirement is updated to consider a division based upon power rating.

Table 9 - Input Power Requirements (Products with Nameplate Rated Power ≤ 1500W)

| Supply Voltage | Phases | Voltage Tolerance | Maximum Total Harmonic Distortion |
|-----------------------------|--------|-------------------|-----------------------------------|
| 100V ac, 115 V ac, 230 V ac | 1 | ±1 % | 2.0 % |
| 200V ac, 208V ac, 400V ac | 3 | ±1 % | 2.0 % |

Table 10 - Input Power Requirements (Products with Nameplate Rated Power > 1500W)

| Supply Voltage | Phases | Voltage Tolerance | Maximum Total Harmonic Distortion |
|-----------------------------|--------|-------------------|-----------------------------------|
| 100V ac, 115 V ac, 230 V ac | 1 | ±5 % | 5.0 % |
| 200V ac, 208V ac, 400V ac | 3 | ±5 % | 5.0 % |

If you have done testing in your lab, this information is crucial to setting up the test environment for the evaluation.

2.7.4 Environmental (Subclause 7.2.8.1 to 7.2.8.3)

Updated the environmental ranges, important for the formal test run.

2.7.5 Power and Temperature Instrumentation (Subclause 7.2.9)

This clause replaces 7.3.5 from *Specification V2.1.1*.

It is important to note the new range allowed and levels of confidence indicated in this clause.

2.7.6 Workload Generator (Subclause 7.2.10)

This clause eliminates the use of the term “Benchmark Driver” and replaces clause 7.3.7 from *Specification V2.1.1*.

This clause now includes two Workload Generators, one for Block Access and one for File Access.

In the case of Block Access, not much has changed from *Specification V2.1.1*, please read subclause 7.2.10.1 of *Specification V3.0.2*.

7.2.10.2 is dedicated to File Access. It is highly recommended that you read in detail this new subclause.



2.7.7 Instrumentation (Subclause 7.2.11)

This clause indicates how to do the data collection according to the type of test done. Subclause 7.2.11.1 covers block access; file access is covered by 7.2.11.2

2.7.8 IO Profiles (Subclause 7.2.12)

Like the previous clause, it is divided on block access and file access. The file access portion (7.2.12.2) is completely new and it is suggested that you read and understand this part in order to be able to successfully evaluate a file access system.

2.7.9 Block Access Online and Near-Online Tests (Subclause 7.3)

There are many changes that make this part of the *Specification V2.1.1* clearer. Some of the data previously required to be collected is not required anymore (see subclause 7.3.3.3).

2.7.10 File Access Online and Near-Online Active Test (Subclause 7.4)

This entire clause is new.

If you are an Emerald practitioner and this is the first time you will do file access testing, it is highly recommended that you learn the details of the whole clause before attempting the testing of a file access storage product.

2.7.11 Block and File Access Ready Idle test (Subclause 7.5)

This clause expands the Ready Idle testing to include File Access.

2.7.12 Block and File Access Capacity Optimization Test (Subclause 7.6)

It combines the usage of COMs previously dedicated to Block.

It clarifies how they should be used for the testing of their existence.

2.7.13 Removable Media Library Testing (Subclause 7.7) and Virtual Media Library Testing (Subclause 7.8)

While for the most part these sections remain very much the same as the *Specification V2.1.1*, some **key** changes that will make the tester's life much easier were introduced. Some of these are like block and include data that is not necessary to collected any more.

2.8 Metrics (Clause 8)

2.8.1 File Access Primary Metrics (Subclause 8.3)

This is an entirely new section introducing the metrics to be used by File Access products.

2.8.2 File Access Active Test (Subclause 8.4.3)

This is a new section that gives details on what is necessary to test file access product power efficiency. It is highly encouraged that you read this subclause in detail.

2.9 Disclosure Requirements (Clause 9)

This is a new clause for *Specification V3.0.2*. It lists all the information required to be disclosed for an Emerald result (or results) to be published. It is strongly suggested that the practitioner, regardless of experience, read this section in detail.

2.10 A Note on Annexes

All the **normative** annexes for the relevant category are a must read for anyone involved in the testing of a storage product.

2.11 Annex A (normative)

Some important normative references were updated. It is of interest for the tester to review the Alternate Meter Usage (A 2) portion of the Annex.

2.12 Annex B (normative)

Important changes to the Measurements Requirements were done, as shown in the following tables:

Table B-1 Online and Near-Online Block Access Summary

| Test | Power and Temperature | | Workload Generator Data Collection | | Minimum Test Duration (minutes) |
|--------------|-------------------------------------|--|--|-------------------------------|---------------------------------|
| | Power PA _i (·) (seconds) | Temperature Recording Interval (seconds) | Metric | Collection interval (seconds) | |
| Conditioning | 60 | 10 | Average Response Time RTA _{sc} (milli-seconds) | 60 | 720 |
| Active | 60 | 10 | 1) Operations Rate O _i (IO/s or MiB/s) 2) Average Response Time RTA _i (milli-seconds) | 60 | 40 |
| Ready Idle | 60 | 10 | N/A | N/A | 120 |



Table B-2 Removable and Virtual Media Library Summary

| Test | Power and Temperature | | Workload Generator Data Collection | | Minimum Test Duration (minutes) |
|--------------|-------------------------------------|--|--|-------------------------------|---------------------------------|
| | Power PA _i (·) (seconds) | Temperature Recording Interval (seconds) | Metric | Collection interval (seconds) | |
| Conditioning | 60 | 10 | 1) Average throughput for each drive (MiB/s) 2) Operations Rate O _{sc} (MiB/s) | 60 | 14 |
| Active | 60 | 10 | 1) Average throughput for each drive (MiB/s) 2) Operations Rate O _i (MiB/s) | 60 | 30 |
| Ready Idle | 60 | 10 | N/A | N/A | 120 |

Table B-3 Online and Near-Online File Access Summary

| Test | Power and Temperature | | sFlow Data Collection | | Minimum Test Duration (minutes) |
|--------------------------|-------------------------------------|--|--|-------------------------------|---------------------------------|
| | Power PA _i (·) (seconds) | Temperature Recording Interval (seconds) | Metric | Collection interval (seconds) | |
| INIT | 10 | 10 | Operations Rate O _i (MiB/s) | 10 | N/A |
| Warm-up – per load point | 10 | 10 | Operations Rate O _i (MiB/s) | 10 | 5 |
| Active – per load point | 10 | 10 | Operations Rate O _i (MiB/s) | 10 | 5 |
| Ready Idle | 10 | 10 | N/A | N/A | 120 |

2.13 Annex C (normative)

New annex on Stable Storage.

2.14 Annex D (normative)

New annex on Vdbench.

2.15 Annex E (normative)

New annex on Vdbench Test Scripts.

2.16 Annex F (normative)

New annex on the SPEC SFS® 2014 Benchmark.

2.17 Annex G (normative)

New annex on the File Access IO Load Driver Configuration File.

2.18 Annex H (informative)

New annex on sFlow®.

2.19 Annex I (normative)

New annex on the COM Test Data Set Generator.

2.20 Annex J (informative)

New annex on gzip.

2.21 Bibliography

Added this for reference and further reading.