INTERNATIONAL STANDARD

ISO/IEC 24091

First edition
2019-11

Information technology — Power efficiency measurement specification for data center storage

Reference number
ISO/IEC 24091:2019(E)

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Fax: +41 22 749 09 47
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents) or the IEC list of patent declarations received (see http://patents.iec.ch).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html.

This document was prepared by SNIA (as SNIA Emerald™ Power Efficiency Measurement Specification V3.0.3) and drafted in accordance with its editorial rules. It was adopted, under the JTC 1 PAS procedure, by Joint Technical Committee ISO/IEC JTC 1, *Information technology*.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

Contents

List of Tables, Figures and Equations

1 Introduction

There is a growing awareness of the environmental impact of IT equipment use. This impact takes several forms: the energy expended in equipment manufacture and distribution, the impact of materials reclamation, and the energy consumed in operation and cooling of the equipment. IT equipment users of all kinds now wish to make their IT operations as energy efficient as possible. This new priority can be driven by one or more of several requirements:

- Rising energy costs have made power and cooling expenses a more significant percentage of total cost of ownership of server and storage equipment;
- Some data centers are physically unable to add more power and cooling load, which means that new applications and data can only be brought on if old ones are retired or consolidated onto new, more efficient configurations;
- Increased regulatory and societal pressures provide incentives for companies to lower their total energy footprints. For many companies, IT is a significant portion of overall energy consumption, and corporate Green goals can only be achieved by reducing IT's energy needs or by making operations more efficient.

IT equipment users will seek advice on the most energy efficient approach to getting their work done. It is not practical for customers to test a wide range of storage products and architectures for themselves. A more effective approach is to create a collection of standard metrics that allow IT architects to objectively compare a range of possible solutions. This objective, metric-based approach has a dual impact:

- Users can select the mode of storage usage that accomplishes their work objectives with the lowest overall energy consumption;
- Companies will be driven to innovate and compete in the development of energy efficient products as measured by the standard yardsticks.

2 Scope

2.1 Abstract

This document describes a standardized method to assess the energy efficiency of commercial storage products in both active and idle states of operation. A taxonomy is defined that classifies storage products in terms of operational profiles and supported features. Test definition and execution rules for measuring the power efficiency of each taxonomy category are described; these include test sequence, test configuration, instrumentation, benchmark driver, IO profiles, measurement interval, and metric stability assessment. Qualitative heuristic tests are defined to verify the existence of several capacity optimization methods. Resulting power efficiency metrics are defined as ratios of idle capacity or active operations during a selected stable measurement interval to the average measured power.

2.2 Introduction

This document defines methodologies and metrics for the evaluation of the related performance and energy consumption of storage products in specific active and idle states.

Storage products and components are said to be in an “active” state when they are processing externally initiated, application-level requests for data transfer between host(s) and the storage product(s). For purposes of this document, idle is defined as “ready idle”, in which storage systems and components are configured, powered up, connected to one or more hosts and capable of satisfying externally initiated, application-level initiated IO requests within normal response time constraints, but no such IO requests are being submitted.

2.3 Current Revision

This document addresses storage products supporting block or file data access. Block access and file access refer to the type of service provided typically by Storage Area Network (SAN) and Network Attached Storage (NAS) systems, respectively. It is not appropriate to use this document to ascertain power efficiency for anything other than these two access modes. This document includes:

- A generalized taxonomy for storage products (clause 5);
- An assessment mechanism for software-based Capacity Optimization Methods (clause 6);
- Measurement and data collection guidelines for assessing the power efficiency of block- and file-based storage products in both active and ready idle states (clause 7);
- Metrics describing storage product power efficiency (clause 8);
- Required disclosures for a test result published as a SNIA Emerald™[1] Power Efficiency Measurement test result (clause 9).

2.4 Purpose

The purpose of a SNIA Emerald™ Power Efficiency Measurement is to provide a reproducible and standardized assessment of the energy efficiency of commercial storage products in both active and ready idle states.

1. Tested systems shall be comprised of commercially released products and components;
2. Tested systems shall employ settings, parameters, and configurations that would allow end-users to achieve power efficiency levels equivalent to the published result;
3. All data published as an SNIA Emerald™ Power Efficiency Measurement test result shall be gathered from test execution conducted according to this document;

[1] SNIA Emerald™ is a trademark of the Storage Networking Industry Association. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO of the product named.

4. Test execution shall complete in its entirety and without test failure or test error messages;
5. Software features which invoke, generate, or use software designed specifically for the test shall not be used. Configuration options chosen for test execution shall be options that are generally recommended for the customer;
6. Energy for powering the equipment shall be provided by the power mains that are being monitored (not by an internal UPS).

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

2.5 Disclaimer

A SNIA Emerald™ Power Efficiency Measurement test result provides a high-level assessment of the energy efficiency of the tested system in specific ready idle and active states. It is not an attempt to precisely model or reproduce any specific installation.

Actual performance and energy consumption behavior is highly dependent upon precise workload, environmental and usage parameters. While a SNIA Emerald™ Power Efficiency Measurement test result is intended to provide a realistic and reproducible assessment of the relative power efficiency of a system across a broad range of configurations and usage patterns, it cannot completely match the precise needs of any one specific installation.

3 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	*ISO/IEC Directives Part II*	Eighth edition, 2018	https://www.iso.org/directives-and-policies.html