Draft ETSI EN 303 804 V0.0.9 (2025-01)



Environmental Engineering (EE); Energy efficiency metrics and measurement methods for data storage equipment

Document Preview

<u>CTSI EN 303 804 V0.0.9 (2025-01)</u>

https://standards.iteh.ai/catalog/standards/etsi/1c49df1e-3425-4d41-97bd-eae1906465d3/etsi-en-303-804-v0-0-9-2025-01

Reference

DEN/EE-EEPS44

Keywords

data storage, data storage equipment, energy efficiency, environment, measurement, metrics, storage

ETSI

650 Route des Lucioles F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B Association à but non lucratif enregistrée à la Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from the ETSI Search & Browse Standards application.

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format on ETSI deliver repository.

Users should be aware that the present document may be revised or have its status changed, this information is available in the <u>Milestones listing</u>.

If you find errors in the present document, please send your comments to the relevant service listed under <u>Committee Support Staff</u>.

If you find a security vulnerability in the present document, please report it through our <u>Coordinated Vulnerability Disclosure (CVD)</u> program.

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.

No recommendation as to products and services or vendors is made or should be implied.

In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI. The copyright and the foregoing restriction extend to reproduction in all media.

> © ETSI 2025. All rights reserved.

Contents

Intel	lectual Property Rights	5			
Fore	word	5			
Mod	al verbs terminology	5			
Intro	duction	6			
1	Scope				
2	References	7			
2.1	Normative references	7			
2.2	Informative references	8			
3	Definition of terms, symbols, abbreviations and conversions	8			
3.1	Terms	8			
3.2	Symbols	12			
3.3	Abbreviations	12			
3.4	Conversions	13			
4	Taxonomy				
4.1	Categories	13			
4.2	Classifications	14			
4.2.1	Attributes	14			
4.2.1	1 Access Paradigm	14			
1.2.1	2 Access Pattern	14			
1.2.1	3 Consumer/Component	14			
1.2.1	4 FBA/CKD Support	14			
1.2.1	5 MaxTTFD	14			
1.2.1	.6 Media Type	15			
4.2.1	.7 Multi-host Shareability	15			
1.2.1	8 No SPOF.	15			
4.2.1	9 Non-Disruptive Serviceability	15			
1.2.1	10 Robotics	15			
1.2.1	.11 Stable Storage Support				
4.2.1	12 Storage Controller	15			
4.2.1	13 Storage Protection	15			
+.2.1.	14 System Capacity	15			
+.2.2	Classifications overview	10			
12.3	Unine disk	18			
12.4	Neal-Olline Ulsk	19			
+.2.3 1 2 6	Virtual Media Library				
12.0	Disk access NVSS	20			
1.2.8	Memory access NVSS				
-	Materia	21			
) - 1	Derformen er metrie				
).1 こつ	Performance metric				
).∠ 5 3	Foregy afficiency matric				
,	Active state energy efficiency metric	21 21			
5.3.2	Idle state energy efficiency metric				
·	Tast definition and execution rules	22			
ן 1	Test setup				
).1 5 1 1	Configuration	22 22			
5.1.1 5.1.7	Finvironment	22 วา			
5.1.2 5.1.2	Dower	22 72			
614	RAS	·····∠3 ??			
6.1 5	Measurement equipment requirement				
< 1 C	EliT Consistency	24			

4

]	History		44		
1	Annex G (informative):	Bibliography	43		
1	Annex F (normative):	File Access workload IO Profile	41		
1	Annex E (normative):	Block Access workload IO Profile			
1	Annex D (informative):	SNIA Emerald ^{1M} Power Efficiency Measurement			
1	A mex C (informative).				
	Annex C (informative).	CTS Lite Device	37		
1	Annex B (informative):	SPECstorage [®] Solution 2020			
/stan	Annex A (informative):	BenchDEE Benchmark	303-804-v0- 3 5 ⁹ -		
Ç	9.3.2 scoring rules	<u>etsi en 303 804 v0.0.9 (2023-01)</u>			
(9.3.1 Information obtaining				
Ģ	9.3 Evaluation method	Evaluation methodology.			
Ģ	9.2.2 Energy saving f	eatures in active state			
Ģ	9.2.1 Energy saving f	eatures at idle and low loadlevel state			
Ģ	9.2 Energy saving featu	irettms.//standards.itah.ai)			
Ģ	9.1 General				
	9 Energy saving level of	evaluation. Teh Standards			
2	8 Test report	-			
-	7.7 Total energy efficie	ency score			
-	7.6 Idle state energy ef	Idle state energy efficiency			
-	7.5 Active state Period	ic energy efficiency	30		
-	7.5 Measurement men 7.4 Sensitivity analysis	Sonsitivity analysis			
	7.2 Measurement not it	cs and intervals			
	7.1 Measurement for id	ulve state efficiency			
-	7 1 Measurement for a	ntive state efficiency	29 20		
,	7 Measurement		20		
6	6.4.3 Test procedure	for idle power at normal operating conditions			
(6.4.2 File Access wor	kload Test procedure	28		
6	6.4.1 Block Access w	orkload Test procedure	27		
(6.4 Test procedures		27		
(6.3.4 Workload Gene	rator			
(6.3.3.6 Genomics W	/orkload			
e	6.3.3.5 AI Image W	orkload			
6	6.3.3.4 Electronic D	Design Automation Workload			
é	6.3.3.3 Video Data	Acquisition Workload	20		
6	6.3.3.2 Software Bu	ilding Workload			
	6.3.3.1 Overview of	F Workloads	23		
	6.3.3 File Access IO	Drofiles	23		
(6.3.1 General	General			
(6.3 Workload	Workload			
6	6.2 Equipment Under 7	Equipment Under Test (EUT) Configuration			
			a :		

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: "Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards", which is available from the ETSI Secretariat. Latest updates are available on the ETSI IPR online database.

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECTTM, **PLUGTESTSTM**, **UMTSTM** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPPTM**, **LTETM** and **5GTM** logo are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2MTM** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM**[®] and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword (https://standards.iteh.ai)

This draft European Standard (EN) has been produced by ETSI Technical Committee Environmental Engineering (EE), and is now submitted for the combined Public Enquiry and Vote phase of the ETSI EN Approval Procedure (ENAP).

ETSI EN 303 804 V0.0.9 (2025-01)

Date of latest announcement of this EN (doa):	3 months after ETSI publication
Date of latest publication of new National Standard	
or endorsement of this EN (dop/e):	6 months after doa
Date of withdrawal of any conflicting National Standard (dow):	6 months after doa

Modal verbs terminology

In the present document "shall", "shall not", "should", "should not", "may", "need not", "will", "will not", "can" and "cannot" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"must" and "must not" are NOT allowed in ETSI deliverables except when used in direct citation.

Introduction

The present document specifies:

1) A metric for the assessment of energy efficiency of Data Storage Equipment (DSE).

The metric for energy efficiency of DSE is different from server's. For DSE it is important to evaluate the performance of data operation in the unit of Input/Output operations or data throughput per second, while for servers it is mainly to evaluate the performance of computing.

For comparison, evaluations should be conducted across similar types of categories of DSE. The present document categorizes DSE to address applicability, configuration groupings to represent a family of DSE to address the broad range of custom configurations possible within each family, and tool revision control to ensure comparability and consistency of the resulting metric value.

2) Test and evaluation methodologies for the assessment of energy efficiency of DSE.

The present document formalizes the tools, conditions and calculations used to generate a single figure of merit of DSE. The present document formalizes the methodology for evaluating energy saving level of DSE from the perspective of supported energy saving feature.

The present document considers some benchmark documents on energy efficiency of DSE.

iTeh Standards (https://standards.iteh.ai) Document Preview

ETSI EN 303 804 V0.0.9 (2025-01)

nttps://standards.iteh.ai/catalog/standards/etsi/1c49df1e-3425-4d41-97bd-eae1906465d3/etsi-en-303-804-v0-0-9-2025-01

1 Scope

The present document is based upon Energy Efficiency Benchmark for Data storage products.

The present document specifies:

- 1) an active state metric, test conditions and product family configuration for the assessment of energy efficiency of DSE using reliable, accurate and reproducible measurement methods;
- 2) an idle state metric and the calculation of the idle state power;
- 3) a measurement method of the active state power;
- 4) a measurement method of the idle state power;
- 5) the measurement and calculation of the maximum power;
- 6) the measurement and calculation of the operating condition class, the ASHRAE validation, using reliable, accurate and reproducible measurement methods, which take into account the recognized state of the art;
- 7) requirements for equipment to perform the measurements and analysis;
- 8) documentation and reporting requirements;
- 9) evaluation methodology for energy saving level from the perspective of supported energy saving features.

The present document addresses DSE.

The present document is applicable at the energy efficiency of:

- online storage;
- nearline storage.

The present document defines metric for the assessment of energy efficiency of DSE and related testing methodology considering data storage equipment HW and system.

TSI EN 303 804 V0.0.9 (2025-01)

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found in the ETSI docbox.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

- [1] <u>ANSI INCITS 400-2004</u>: "Information Technology SCSI Object-Based Storage Device Commands (OSD)".
- [2] <u>ANSI INCITS 458-2011</u>: "Information Technology SCSI Object Based Storage Device Commands- 2 (OSD 2)".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

- [i.1] "SNIA EmeraldTM Power Efficiency Measurement Specification", Version 4.0.0.
- [i.2] <u>BenchDEE</u>.
- [i.3] "<u>SPECstorage[®] Solution 2020</u>".
- [i.4] "CTS Lite Device Test Power Efficiency Measurement".
- [i.5] <u>ASHRAE TC 9.9 Reference Card</u>: "Equipment Thermal Guidelines for Data Processing Environments".

3 Definition of terms, symbols, abbreviations and conversions

3.1 Terms (https://standards.iteh.ai)

For the purposes of the present document, the following terms apply:

a.c.-d.c. power supply unit: power supply unit that converts line-voltage alternating current (a.c.) input power into one or more direct current (d.c.) power outputs

active state: operational state of a DSE (as opposed to the idle state) in which the DSE is carrying out work in response to prior or concurrent external requests (e.g. instruction over the network), mainly data operation work

auto-tiering: policy-based system that automatically places and moves data across tiers to optimize performance service levels, cost targets, and overall energy consumption

NOTE: Each storage tier may comprise different storage technologies, offering varying performance, cost, and energy consumption characteristics.

blade storage: storage device that is designed for use in a blade chassis and that is dependent upon shared blade chassis resources (e.g. power supplies, cooling) for operation

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

committed data: data that has been written to stable storage

compression: process of encoding data to reduce its size

controller system: computer or computer server that manages a benchmark evaluation process

Count-Key-Data (**CKD**): disk data organization model in which the disk is assumed to consist of a fixed number of tracks, each having a maximum data capacity

NOTE: The CKD architecture derives its name from the record format, which consists of a field containing the number of bytes of data and a record address, an optional key field by which particular records can be easily recognized, and the data itself.

data deduplication: replacement of multiple copies of data at variable levels of granularity with references to a shared copy in order to save storage space and/or bandwidth

Data Storage Equipment (DSE): collective term for disk drives, solid state drives and modules, tape cartridges, and any other mechanisms providing non-volatile data storage

data storage server: enterprise storage device which contains the same components as a computer server together with ≥ 10 storage devices and software (vendor or 3rd party) that supports storage system connectivity, capacity optimization management, virtualized storage environment and software defined storage

NOTE: Supporting features are described by the product's datasheet description and are either accompanied with vendor specific utilities and/or commercially available software supporting these functions.

data storage system: fully-functional system that supplies data storage services to clients and devices attached directly or through a network

- NOTE 1: A storage system may be composed of integrated storage controllers, storage devices, embedded network elements, software, and other devices.
- NOTE 2: While storage system may contain one or more embedded processors, these processors do not execute user-supplied software applications but may execute data-specific applications (e.g. data replication, backup utilities, data compression, install agents).
- NOTE 3: Components and subsystems that are an integral part of the storage product architecture (e.g. to provide internal communications between controllers and disks) are considered to be part of the storage system.
- NOTE 4: Components that are normally associated with a storage environment (e.g. devices required for operation of an external Storage Area Network) are not considered to be part of the storage system.

d.c.-d.c. power supply unit: power supply unit that converts line-voltage direct current (d.c.) input power to one or more d.c. outputs

NOTE: For purposes of the present document, a d.c-d.c. converter that is internal to a data storage system and is used to convert a low voltage d.c. (e.g. 12 VDC) into other d.c. power outputs for use by components inside is not considered a d.c-d.c. power supply unit.

dedupable: property that a collection of data is said to possess if the needed storage capacity for the data is reduced significantly by data deduplication

delta snapshot: type of point in time copy that preserves the state of data at an instant in time, by storing only those blocks that are different from an already existing full copy of the data

direct-connected: storage designed to be under the control of a single host, or multiple hosts in a non-shared environment

efficiency: ratio of workload output to the resource input to the system

energy efficiency: ratio of workload output to the power input to the system

file: abstract data object made up of:

- a) an ordered sequence of data bytes stored on a disk or tape;
- b) a symbolic name by which the object can be uniquely identified; and
- c) a set of properties, such as ownership and access permissions;

that allow the object to be managed by a file system or backup manager

file system: software component that imposes structure on the address space of one or more physical or virtual disks so that applications may deal more conveniently with abstract named data objects of variable size (files)

Fixed Block Architecture (FBA): model of disks in which storage space is organized as linear, dense address spaces of blocks of a fixed size

NOTE: Fixed block architecture is the disk model on which SCSI is predicated.

formatted capacity: total number of bytes available to be written after a system or device has been formatted for use, e.g. by an object store, file system or block services manager

NOTE: Formatted capacity is less than or equal to raw capacity. It does not include areas set aside for system use, spares, RAID parity areas, checksum space, host- or file system-level remapping, "right sizing" of disks, disk labelling, caches, file system metadata, and so on. However, it may include areas that are normally reserved such as snapshot set-asides if they can alternatively be configured for ordinary data storage by the storage administrator.

free space: amount of additional irreducible data that can be written to the product under test as configured

hot band: simulation of naturally occurring areas of storage address space accessed more frequently than is typical

hybrid Solid State Storage and magnetic disk system: storage system whose formatted capacity is provided by a combination of Solid State Storage and magnetic disk storage devices

hyper converged system: software-defined set of resources with at least two of the following: compute, storage, networking and/or virtualization

ICT equipment: equipment providing data storage, processing and transport services

NOTE: A combination of Information Technology Equipment and Network Telecommunications Equipment.

idle state: operational state of a DSE (as opposed to the active state) in which the DSE is not performing any useful work

NOTE 1: The DSE is capable of completing workload transactions, but no active workload transactions are requested or pending by the system.

NOTE 2: For systems where ACPI standards are applicable, idle state correlates only to ACPI System Level S0.

I/O device: device which provides data input and output capability between a DSE and other devices

NOTE: An I/O device may be integral to the DSE motherboard or may be connected to the motherboard via expansion slots.

I/O intensity: measure of the number of IOPS requested by a load generator

NOTE: IO intensity is phrased as a percentage of selected maximum IOPS level that satisfies the timing requirement(s) for a taxonomy category.

I/O port: physical circuitry within an I/O device where an independent I/O session can be established

NOTE: A port is not the same as a connector receptacle; it is possible that a single connector receptacle can service multiple ports of the same interface.

irreducible data: data that is neither compressible nor dedupable

Just a Bunch Of Disks (JBOD): shelf that does not provide storage protection

NOTE: A JBOD is a "simple" "just a bunch of disks," an architecture that does not provide storage protection capabilities. The storage devices of a JBOD can function as individual volumes or can be combined to form a single logical volume by a simple controller. A JBOD provides no redundancy or resilience, so failure of a single storage device amounts to failure of a whole logical volume. A JBOD is a single shelf.

load generator: hardware and software environment executing the workload generator to drive the product under test during measurements

Logical Unit (LU): entity within a SCSI target that executes IO commands

Logical Unit Number (LUN): synonym for logical unit

maximum power: peak sustained or root means square power consumption value while operating the worst case functions

Maximum Time to First Data (MaxTTFD): maximum time required to start receiving data from a storage system to satisfy a read request for arbitrary data

memory: server component external to the processor in which information is stored for immediate use by the processor

11

motherboard: main circuit board of the DSE typically accommodating the processor, memory, expansion slots and enabling the attachment of additional circuit boards

network-connected: storage designed to be connected to a host via a network protocol (e.g. TCP/IP, IB, and FC)

non-disruptive serviceability: support for continued availability of data during all service operations

- NOTE 1: Examples of non-disruptive serviceability are FRU replacement, code patches, software/firmware upgrades, configuration changes, data migrations, and system expansion done during production time.
- NOTE 2: Service operations may result in performance impacts to data availability, but will not result in a loss of access.

non-volatile: property of retaining data in the event of the loss of power

normalized performance: relative performance values calibrated to a baseline common to the set of equipment being evaluated

parity RAID: collective term used to refer to Berkeley RAID Levels 3, 4, 5 and 6

permanent storage: data storage media which can retain data indefinitely without a power source

power supply unit: self-contained device, physically separable from the motherboard of the computer server, that converts a.c. or d.c. input power to one or more d.c. power outputs for powering the computer server via a removable or hard-wired electrical connection

raw capacity: sum of the raw, unformatted, uncompressed capacity of each of the storage device in the product under test

ready idle: operational state in which a system is capable of satisfying an arbitrary IO request within the response time and MaxTTFD constraints of its selected taxonomy category, but no user-initiated IO requests are being submitted to the system

NOTE: In the ready idle state, background I/O activity, autonomously initiated by the solution under test, may take place.

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

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

shelf: modular enclosure suitable for installation in a rack

NOTE: A shelf typically houses storage devices, storage controllers, power supplies, and cooling devices. A shelf typically has a pre-wired backplane that carries power and I/O interconnect signals to the housed components.

Single Point of Failure (SPOF): one component or path in a system, the failure of which would make the system inoperable or data inaccessible

solid state drive: storage device that uses memory chips instead of rotating magnetic platters for data storage

solid state storage: storage capability built from non-volatile solid state electronic devices

storage controller: device for handling storage requests that includes a processor or sequencer programmed to autonomously process a substantial portion of IO requests directed to storage devices

NOTE: This definition is specifically intended to exclude aggregating storage elements such as RAID array subsystems, robotic tape libraries, filers, and file servers. Also excluded are storage devices which are not directly accessible by end-user application programs, and are instead employed as a form of internal cache.

storage protection: any combination of hardware and software (e.g. RAID, NVRAM, disk sparing and background disk scrubbing or media scan) that assures that all completed IO operations will be preserved in the event of power loss or storage device failure

system crash: hardware or software failure which causes data to be temporarily unavailable and which requires a reboot of one or more hardware components and/or re-initialization of one or more software components in order for data access to be restored

Uninterruptible Power Supply (UPS): combination of convertors, switches, and energy storage devices (such as batteries) constituting a power system for maintaining continuity of load power in case of input power failure

virtual drive: removable media storage device, e.g. tape drive, that is emulated using other storage devices

weighted geometric mean: geometric mean calculated using a predetermined factor for each of the elements prior to aggregation

worklet: synthetic software routine, using real application functions focused on a particular type of data operation activity, which stresses a particular characteristic of the system

workload generator: software used in the load generator to drive the product under test during measurement

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

a.c., AC	Alternating Current
ACPI	Advanced Configuration and Power Interface
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
BIOS	Basic Input/Output System for and and a if the ail
CPU	Central Processor Unit
d.c., DC	Direct Current
DSE	Data Storage Equipment
ESC	Energy Saving Cooling
ETSI	European Telecommunications Standards Institute
EUT	Equipment Under Test EN 303 804 V0.0.9 (2025-01)
FBAteh.ai/cata	Fixed Block Architecture 9df1e-3425-4d41-97bd-eae1906465d3/etsi-en-303-804-v0-0-9-2025-01
FRU	Field-Replaceable Unit
GB	GigaByte
HDD	Hard Disk Drive
I/O	Input/Output
ICT	Information and Communication Technology
ID	IDentification
INCITS	InterNational Committee for Information Technology Standards
IOPS	Input/Output operation Per Second
JBOD	Just a Bunch Of Disks
MaxTTFD	Maximum Time To First Data
NVRAM	Non-Volatile Random Access Memory
NVSS	Non-Volatile Solid State
OS	Operating System
OSD	Object-based Storage Device
PDU	Power Distribution Unit
PSU	Power Supply Unit
RAID	Redundant Array of Independent Disks
RAS	Reliability, Availability, and Serviceability
RMS	Root Mean Square
RVML	Removable and Virtual Media Library
SCSI	Small Computer System Interface
SPOF	Single Point of Failure
SSD	Solid State Drive
UPS	Uninterruptible Power Supply
VDA	Video Data Acquisition