INTERNATIONAL STANDARD

ISO/IEC 21836

First edition
2020-08

Information technology — Data centres — Server energy effectiveness metric

Technologies de l'information — Centres de données — Grandeurs de mesure de l'efficacité énergétique des serveurs

Reference number
ISO/IEC 21836:2020(E)

© ISO/IEC 2020



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2020

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

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

Published in Switzerland

Contents

Page

Foreword ..... **v**

Introduction ..... **vi**

1 Scope **1**

2 Normative references **1**

3 Terms, definitions and abbreviated terms **1**
3.1 Terms and definitions 1
3.2 Abbreviated terms 9

4 Applicability of Server Energy Effectiveness Metric (SEEM) **10**
4.1 General 10
4.2 Applicability of SERTv2 10
4.3 Determination of applicability for an "implementer-specified" metric 11

5 Determination of Server Energy Effectiveness Metric (SEEM) **11**
5.1 General 11
5.2 Power supply requirements 12
5.3 SERTv2 active state energy effectiveness metric 12
5.4 Determination of "implementer-specified" metrics 14
5.4.1 General 14
5.4.2 "Implementer-specified" active state metrics 14
5.4.3 "Implementer-specified" idle state metric 15

6 SEEM implementation **16**
6.1 General 16
6.2 Implementation steps for SEEM 16
6.3 Implementation steps for SERTv2 16
6.4 Implementation steps for "implementer-specified" metrics 16
6.5 Threshold selection 17

7 Server testing **18**
7.1 Configuration 18
7.2 Environment 18
7.3 Power analyser calibration requirement 18
7.4 Power requirements 18
7.5 SEEM idle state power test method 19
7.6 Testing Variance 19
7.6.1 General 19
7.6.2 Software Testing Variance Specification 20
7.6.3 SEEM Hardware Testing Variance Specification 20
7.7 Technical support 20
7.7.1 Technical support requirements 20
7.7.2 SPEC SERTv2 technical support 20
7.8 Server product family and special configuration testing 21
7.8.1 General 21
7.8.2 Single configuration testing 22
7.8.3 SERTv2 product family and special configurations 22
7.8.4 Product family and special configurations for "implementer-specified" metrics 24
7.9 "Implementer-specified" metric creation testing 24
7.9.1 General 24
7.9.2 Functionality and reproducibility testing 24
7.9.3 Real-world energy savings testing 26

8 Reporting SEEM **26**
8.1 Server categories 26
8.1.1 General 26
8.1.2 SERTv2 server category definitions 26

© ISO/IEC 2020 – All rights reserved

8.1.3 "Implementer-specified" metric server categories 29
8.2 Documentation 29
8.2.1 General 29
8.2.2 Implementation documentation 29
8.2.3 SEEM end user documentation 30
9 Minor update acceptance testing 33
9.1 General 33
9.2 SEEM minor update testing and requirements 33
9.2.1 SEEM component versions 33
9.2.2 Power analysers 33
9.2.3 Temperature sensors 35
9.3 SERTv2 minor update testing and requirements 35
9.3.1 General 35
9.3.2 SERTv2 Component versions 35
9.3.3 SERTv2 36
9.3.4 PTDaemon 36
9.3.5 SERTv2 Client Configuration XML 36
9.4 "Implementer-specified" metric minor update requirements 44
Annex A (normative) **SERTv2 pre-approved CPU architectures, operating systems, and Java Virtual Machines (JVMs) 46**
Annex B (normative) **SEEM pre-approved power analysers and temperature sensors 57**
Annex C (normative) **SERTv2 required documentation fields 59**
Annex D (normative) **Deployed power assessment 61**
Bibliography 69

 © ISO/IEC 2020 – All rights reserved

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. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (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 Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 39, *Sustainability, IT & Data Centres*.

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.

© ISO/IEC 2020 – All rights reserved

Introduction

The global economy is now totally reliant on information and communication technologies (ICT) and the associated generation, transmission, dissemination, computation and storage of digital data. While the internet backbone carries the traffic, it is data centres which find themselves at the nodes and hubs of a wide variety of both private enterprise and shared/collocation facilities. With the large and continually increasing data capacity demands placed on data centres worldwide, efficient use of data centre energy is an extremely important strategy for managing environmental, cost, electrical grid capacity and other impacts.

The ISO/IEC 30134 series specifies data centre energy effectiveness key performance indicators (KPI) to help data centre operators measure and improve specific aspects of data centre energy effectiveness. ISO/IEC 30134-4 in particular defines a method to measure the peak capacity and utilization of servers operating in a data centre using operator selected benchmarks. However, it does not provide a method for comparing individual server energy effectiveness across data centres, and as stated in ISO/IEC 30134-4, "should not be used to set regulations for a data centre or individual server". There is stakeholder demand for an international standard to measure the energy effectiveness of servers before procurement and installation, particularly for use in worldwide server energy effectiveness regulations and programmes.

This document provides a server energy effectiveness metric (SEEM) to measure and report the energy effectiveness of specific server designs and configurations. This document will be useful to stakeholders, including vendors, users and governments, from the design verification testing phase all the way through conformance verification, procurement and operation. Organizations that wish to establish conformance or reporting programmes will find that the test methods and scoring specified in this document will save them significant time and effort in implementing such programmes. Standardization across such programmes will allow vendors to comply to stakeholder requirements more quickly and efficiently.

For applicable servers, this document builds upon the widely adopted Server Efficiency Rating Tool (SERT™)[1)] suite developed by the Standard Performance Evaluation Corporation (SPEC®)[2)] benchmark consortium, as the energy effectiveness metric and test method. For servers where SERT is not applicable, this document provides requirements for the creation of alternate server energy effectiveness metrics, referred to as "implementer-specified" metrics.

1) SERT is a trademark of the Standard Performance Evaluation Corporation. This information is given for the convenience of users of this document. References to SERT do not constitute an endorsement by ISO/IEC.

2) SPEC is a trademark of the Standard Performance Evaluation Corporation. This information is given for the convenience of users of this document. References to SPEC do not constitute an endorsement by ISO/IEC.

© ISO/IEC 2020 – All rights reserved

Information technology — Data centres — Server energy effectiveness metric

1 Scope

This document specifies a measurement method to assess and report the energy effectiveness of a computer server. This document does not set any pass/fail criteria for servers.

2 Normative references

There are no normative references in this document.

3 Terms, definitions and abbreviated terms

For the purposes of this document, the following terms and definitions 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/

3.1 Terms and definitions

3.1.1
64-bit CPU
CPU (3.1.12) which has data path widths, *memory* (3.1.34) addressing, registers, and other architectural features which are 64-bits wide

3.1.2
active state
operational state in which the *server* (3.1.49) is carrying out data processing

Note 1 to entry: an example is data retrieval from *memory* (3.1.34), cache, or storage while awaiting further input over the network.

3.1.3
auxiliary processing accelerator
additional compute device installed in the computer *server* (3.1.49) that handles parallelized *workloads* (3.1.60) in conjunction with the *CPU* (3.1.12)

3.1.4
blade chassis
enclosure that contains shared resources for the operation of *blade servers* (3.1.5), *blade storage* (3.1.6), and other blade form-factor devices

Note 1 to entry: Shared resources provided by a chassis may include power supplies, data storage and hardware for DC power distribution, thermal management, system management and network services.

© ISO/IEC 2020 – All rights reserved

3.1.5
blade server
server (3.1.49) that is designed for use in a *blade chassis* (3.1.4)

Note 1 to entry: A blade server is a high-density device that functions as an independent *server* (3.1.49) and includes at least one *processor* (3.1.40) and *system memory* (3.1.54), but is dependent upon shared *blade chassis* (3.1.4) resources (e.g. power supplies, cooling) for operation.

3.1.6
blade storage
storage device that is designed for use in a *blade chassis* (3.1.4) that is dependent upon shared *blade chassis* (3.1.4) resources, like power supplies or cooling, for operation

3.1.7
buffered memory
circuitry between the server's *memory* (3.1.34) and memory controller to either increase memory capacity, increase bandwidth, and/or reduce the electrical load on the memory controller

3.1.8
coefficient of determination
statistic used to determine the strength of a fit between a mathematical model and a set of observed data values

[SOURCE: ISO 15551-1:2015, 3.26, modified — Note 1 to entry has been removed.]

3.1.9
coefficient of variation
CV
standard deviation divided by the mean

[SOURCE: ISO 3534-1:2006, 2.38, modified — Note 1 to entry has been removed.]

3.1.10
configuration
interrelated functional and physical characteristics of a product defined in product configuration information

Note 1 to entry: This document employs the following configurations: *low-end* (3.1.33), *high-end* (3.1.25) and typical.

Note 2 to entry: For *server* (3.1.49) products, a configuration is one of many possible combinations of components including *CPU* (3.1.12), storage devices, *memory* (3.1.34) size, and capacity and input/output devices for a single *server* (3.1.49) product within a larger product family. There are a large number of possible configurations within a product family.

[SOURCE: ISO 17599:2015, 3.15, modified — Notes 1 and Note 2 to entry have been added.]

3.1.11
core
component of a *processor* (3.1.40) which can independently receive instructions and takes actions or performs calculations in response

3.1.12
CPU
central processing element with functions for interpreting and executing instructions

Note 1 to entry: In this document, cache *memory* (3.1.34) is included with the CPU.

Note 2 to entry: This document uses the terms CPU and *processor* (3.1.40) interchangeably.

[SOURCE: ISO/IEC 14576:1999, 2.1.9, modified — Note 1 and Note 2 to entry have been added.]

© ISO/IEC 2020 – All rights reserved

3.1.13
CPU architecture
CPU (3.1.12) design with significant similarities to other CPU architectures within the same *CPU architecture class* (3.1.14)

Note 1 to entry: CPU architectures are used to create *CPU models* (3.1.15) which are often released in a similar timeframe.

Note 2 to entry: Examples of different CPU architectures in the same *CPU architecture class* (3.1.14) are Intel® Haswell, Intel Broadwell, and Intel Skylake, or separately, AMD® Bulldozer, AMD Piledriver, and AMD Steamroller[3)].

3.1.14
CPU architecture class
group of one or more *CPU architectures* (3.1.13) which share the same instruction set architecture and in which newer architecture designs are derived from previous architecture designs

Note 1 to entry: Within a CPU architecture class, the initial *CPU architecture* (3.1.13) is, for the most part, a new design, and subsequent *CPU architectures* (3.1.13) are derived from preceding *CPU architectures* (3.1.13).

Note 2 to entry: Examples of different CPU architecture classes are ARM® v8-A and AMD EPYC®[3].

Note 3 to entry: In certain cases, software programmes need to be recompiled for use with different CPU architecture classes.

3.1.15
CPU model
specific *CPU* (3.1.12) that is sold in the marketplace

Note 1 to entry: All *CPU* (3.1.12) of the same model share the same technical characteristics, such as *core* (3.1.11) frequencies and core counts, and can be used interchangeably.

Note 2 to entry: Examples of different CPU models are AMD EYPC 7601, AMD EYPC 7251 and Intel Xeon Platinum 8180[3].

3.1.16
CPU nominal frequency
CPU core clock frequency, which is the main frequency used in naming, marketing and selling the *CPU* (3.1.12)

3.1.17
data averaging interval
for a *power analyser* (3.1.38), the time period over which all samples captured by the high-speed sampling electronics of the analyser are averaged to provide a set of measured data

3.1.18
double data rate
DDR
computer bus characteristic of transferring data on the rising and falling edges of the clock signal, resulting in twice the data bandwidth at a specific clock frequency, versus a single data rate bus

3.1.19
end user
person or persons who will ultimately be using the system for its intended purpose

Note 1 to entry: For the purposes of this document, the end user refers to a SEEM end user, which is the entity applying for certification of a *server* (3.1.49) model to a SEEM conformant regulation or programme. For example, if server manufacturer A was submitting a server model to ENERGY STAR for certification, server manufacturer A would be the end user.

---

3) AMD and EPYC are trademarks of Advanced Micro Devices, Intel is a trademark of the Intel Corporation and ARM is a trademark of Arm Limited. This information is given for the convenience of users of this document. References to AMD, EPYC, Intel and ARM do not constitute an endorsement by ISO/IEC.

© ISO/IEC 2020 – All rights reserved

[SOURCE: ISO/IEC 19770-5:2015, 3.13, modified — Note 1 to entry removed, new Note 1 to entry added.]

3.1.20
energy effectiveness
measure of the amount of data processing performed for a given amount of energy consumed

Note 1 to entry: For the purposes of this document, energy effectiveness is equivalent to the term energy efficiency as used in *server* (3.1.49) compliance regulations and programmes.

3.1.21
expansion auxiliary processing accelerator
expansion APA
auxiliary processing accelerator that is an add-in card installed in a general-purpose add-in expansion slot.

Note 1 to entry: An expansion APA add-in card may include one or more APAs and/or separate, dedicated removable switches.

EXAMPLE A GPGPU installed in a PCI-e slot.

3.1.22
fully fault tolerant server
computer *server* (3.1.49) that is designed with complete hardware redundancy, in which every computing component is replicated between two nodes running identical and concurrent *workloads* (3.1.60)

Note 1 to entry: A fully fault tolerant server uses two systems to simultaneously and repetitively run a single workload for continuous availability in a mission critical application.

Note 2 to entry: An example of a fault tolerant server ; if one node fails or needs repair, the second node can run the workload alone to avoid downtime.

3.1.23
hardware threads
in a CPU *core* (3.1.11), the number of fully independent instruction streams which can be executed through SMT

3.1.24
high-performance computing system
HPC system
HPC server
computing system which is designed, marketed, sold, and optimized to execute highly parallel applications for high performance, deep learning, or artificial intelligence applications

Note 1 to entry: HPC systems consist of multiple clustered *servers* (3.1.49, primarily for increased computational capability, high speed inter-processing interconnects, large and high bandwidth *memory* (3.1.34) capability and often accelerators such as GPGPUs or FPGAs.

Note 2 to entry: HPC systems may be purposely built or assembled from more commonly available computer servers.

3.1.25
high-end configuration
server (3.1.49) equipped with a specific selection of high-performance components, which is required to be tested as part of measuring a *server product family* (3.1.51)

© ISO/IEC 2020 – All rights reserved

3.1.26
idle state
operational mode in which the OS and other software have completed loading, the *server* (3.1.49) is capable of completing *workload* (3.1.60) transactions, but no active state (3.1.2) workload transactions are requested or pending by the system

Note 1 to entry: In the idle state, the *server* (3.1.49) is operational, but not performing any useful data processing.

Note 2 to entry: For systems where Advanced Configuration and Power Interface (ACPI) has been implemented, idle state is the ACPI G0 global state and S0 sleep state.

3.1.27
idle state power
average *server* (3.1.49) power, in watts, when in *idle state* (3.1.26)

Note 1 to entry: *SERT* (3.1.46) provides a standard way to measure the idle state power of a server, which is included with the result output, and is in addition to power measurement while the *server* (3.1.44) is actively performing data processing.

3.1.28
implementer
entity that transforms specified designs into their physical realization

Note 1 to entry: For the purposes of this document, implementer is the entity which creates a selection or procurement program based on SEEM.

[SOURCE: IEC 62279:2015, 3.1.15, modified — Note 1 to entry added.]

3.1.29
integrated auxiliary processing accelerator
auxiliary processing accelerator that is integrated into the motherboard or *CPU* (3.1.12) package

3.1.30
large network equipment
network product which contains more than 11 network ports with a total line rate throughput of 12 Gb/s or more

3.1.31
large server
resilient/scalable *server* (3.1.49) which ships as a pre-integrated/pre-tested system housed in one or more full frames or racks and that includes a high connectivity I/O subsystem with a minimum of 32 dedicated I/O slots

3.1.32
load level
percentage of data processing relative to the maximum a *server* (3.1.49) can execute

Note 1 to entry: Load levels are typically used by benchmark designers to simulate situations where a system is receiving fewer data processing requests than it can execute.

Note 2 to entry: Load level is not necessarily the same as *CPU* (3.1.12) utilization.

3.1.33
low-end configuration
server (3.1.49) configuration which includes a specific selection of entry level components, which is required to be tested as part of measuring a *server product family* (3.1.51)

© ISO/IEC 2020 – All rights reserved

3.1.34
memory
any device associated with a computer that is used to store information such as programmes or data, in a digital form

Note 1 to entry: For the purposes of this document, the terms *memory* and *system memory* (3.1.54) are used to refer to memory DIMMs in *servers* (3.1.49) which provide temporary, fast data storage.

Note 2 to entry: At the time of drafting this document, memory DIMMs are the predominate type of memory modules used in servers. In this document, the term *memory DIMM* is used to refer to a server's memory modules and is not intended to exclude future types of memory modules or imply that statements referring to memory DIMMs would not apply to other types of memory modules.

[SOURCE: ISO 1213-1:1993, 11.3.32, modified, — Notes 1 and 2 to entry added.]

3.1.35
memory channel
independent interface in a computer which facilitates the communication of data between a *core's* (3.1.11) memory controller and installed memory DIMMs

Note 1 to entry: Modern computer *servers* (3.1.43) usually have a number of memory channels connected to different *CPUs* (3.1.12) or *cores* (3.1.11).

3.1.36
multi-node server
server (3.1.49) that is designed with two or more independent server nodes that share a single enclosure and one or more power supplies

Note 1 to entry: Power is distributed to all nodes through shared power supplies and nodes in a multi-node server are not designed to be hot-swappable.

3.1.37
normalized
dividing a set of numeric values by one or more numeric value(s)

Note 1 to entry: In this standard normalization is performed for two purposes. One, to adjust benchmark measurement results to a common numeric scale to ease comparability and combination. Two, dividing benchmark results by an arbitrary constant to enable sharing of results while obfuscating the actual value of benchmark results.

3.1.38
power analyser
device used to measure energy consumption of a system under test

3.1.39
PTDaemon
software tool to interface with and control a *power analyser* (3.1.38) or temperature sensor during measurement intervals, providing an interface between the supported power analysers and *SERT* (3.1.46)

Note 1 to entry: The PTDaemon software allows for automatic power and thermal data collection throughout a test run and is the source of the detailed power data in the SERT result summary.

3.1.40
processor
in a computer, functional unit that interprets and executes instructions

Note 1 to entry: the processor is the *CPU* (3.1.12) of the computer *server* (3.1.49). A typical CPU is a physical package to be installed on the server motherboard via a *socket* (3.1.52) or direct solder attachment. The CPU package may include one or more processor *cores* (3.1.11).

Note 2 to entry: This document uses the term processor and CPU interchangeably.

© ISO/IEC 2020 – All rights reserved

[SOURCE: ISO/IEC 2382:2015, 2122866, Note 1 and Note 2 to entry added.]

3.1.41
programmable load
test equipment or instrument which emulates DC or AC resistance loads normally required to perform functional tests of a system under test

3.1.42
quantum computer
use of quantum phenomena for computational purposes

3.1.43
rack server
computer *server* (3.1.49) that is designed for deployment in a standard 19-inch data centre rack as defined by EIA-310, IEC 60297, or DIN 41494

3.1.44
redundant power supply
additional power supply added to a *server* (3.1.49) which can fully power the server if the main power supply fails

Note 1 to entry: During normal operation, redundant power supplies can be configured to either provide power to the server concurrently with the main power supply, or remain on standby and only provide power in the event the main power supply fails.

3.1.45
resilient server
computer *server* (3.1.49) designed with extensive reliability, availability, serviceability, and scalability features integrated in the micro architecture of the system, *CPU* (3.1.12) and chipset

3.1.46
Server Efficiency Rating Tool
SERT
performance and power software measurement tool created by the SPEC benchmark standards consortium

Note 1 to entry: SERT was specifically designed for use in government sponsored *server* (3.1.49) energy efficiency programmes.

Note 2 to entry: SERT has components that run on the system under test and a controller system, and interfaces with a *power analyser* (3.1.38) connected between the electrical *socket* (3.1.52) and server power supply.

Note 3 to entry: Detailed performance and power data is collected while running server *worklets* (3.1.59) at different *load levels* (3.1.32), and these measurements are combined into an overall weighted server energy efficiency score.

3.1.47
SERT Client Configuration XML
file which contains the required *worklet* (3.1.59) settings for each supported OS, JVM, and *CPU architecture* (3.1.13) in *SERT* (3.1.46)

Note 1 to entry: Ensuring proper JVM tuning parameters specific to each supported OS, JVM, and CPU architecture is an important element for accurate comparisons across disparate *server* (3.1.49) platforms.

3.1.48
SERT Run and Reporting Rules
set of requirements established to ensure accurate, reproducible, and comparable measurements across *servers* (3.1.49) when running *SERT* (3.1.46)

© ISO/IEC 2020 – All rights reserved

3.1.49
server
physical system unit (Host) composed of *CPUs* (3.1.12), *Memory* (3.1.34), Storage, PSUs, Fans, and I/O that provides computational services to workstations, to personal computers or to other functional units in a network

Note 1 to entry: Servers provide services and manage networked resources for client devices and are sold through enterprise channels for use in data centres or office/corporate environments. They are primarily accessed via network connections, instead of directly connected user input devices such as a keyboard or mouse. Servers are designed for and listed as supporting one or more computer enterprise operating systems (OS) and/or hypervisors. They are targeted to run user-installed applications, typically enterprise in nature, and provide support for error-correcting code (ECC) and/or *buffered memory* (3.1.7) (including both buffered dual in-line memory modules [DIMMs] and buffered on board [BOB] *configurations* [3.1.10]). Servers are designed such that all *processors* (3.1.36) have access to shared *system memory* (3.1.54) and are visible to a single OS or hypervisor.

Note 2 to entry: For the purpose of determining which types of computing systems are considered servers in this document, both the definition of servers and Note 1 to entry apply.

Note 3 to entry: Services may be dedicated services or shared services.

3.1.50
server appliance
computer *server* (3.1.49) that is bundled with a pre-installed OS and application software that is used to perform a dedicated function or set of tightly coupled functions

3.1.51
server product family
group of computer servers sharing one chassis and motherboard which often contains a large number of possible hardware and software *configurations* (3.1.10)

3.1.52
socket
server interface designed for the installation of a *processor* (3.1.36)

3.1.53
storage product
fully functional storage system that supplies data retention services to other devices attached directly or through a network

Note 1 to entry: A storage server can run on more than one non-vendor specific software which is designed to support storage system connectivity, COM deployments and virtualized storage environments arrayed in a software defined storage network.

3.1.54
system memory
volatile data storage which is accessible and shared by all of the *cores* (3.1.11) of a computer

3.1.55
threshold
level or numeric value at which a change occurs

Note 1 to entry: In this document, this term refers to a numeric result of a test metric which an *implementer* (3.1.28) determines is required to pass a programme or regulation.

3.1.56
tower server
a computer *server* (3.1.49) designed for use in a standalone, non-rack mountable chassis

© ISO/IEC 2020 – All rights reserved

3.1.57
variance
measure of the spread of a statistical distribution

Note 1 to entry: For the purposes of this standard, variance is used to refer to the change in benchmark results between two different runs of the benchmark.

[SOURCE: ISO/IEC 19795-1:2006, 4.8.1, modified — original notes and symbol removed, Note 1 to entry has been added.]

3.1.58
version
particular form or variation of a resource that differs from other instantiations of the resource in at least one aspect or item of information

Note 1 to entry: The version number(s) or letter(s) before the first decimal point or space is called the *major version* and any portion of the version which is not the *major version* is called the *minor version.*

EXAMPLE In the version designated as "v1.2.3," the major version identification is "1", and the minor version identification is "2.3".

[SOURCE: ISO 24619:2011, 3.1.9, modified — original notes removed, Note 1 and Note 2 to entry have been added.]

3.1.59
worklet
parts of a *workload* (3.1.60) consisting of specific code sequences which are executed during testing

3.1.60
workload
group of *worklets* (3.1.59) which share common attributes and are combined into an overall result

Note 1 to entry: *SERT* (3.1.46) includes *CPU* (3.1.12), *Memory* (3.1.34), and Storage workloads.

Note 2 to entry: A product family has common family attributes, which are a set of features common to all models/*configurations* (3.1.10) that are in the family.

3.2 Abbreviated terms

AC	alternating current
ACPI	advanced configuration and power interface
BIOS	basic input/output system
COM	capacity optimization management
DC	direct current
DIMM	dual inline memory module
ECC	error-correcting code
EXP	exponential function
FPGA	field programmable gate array
GB	Gigabytes
GPGPU	general purpose graphics processing unit

© ISO/IEC 2020 – All rights reserved

GUI	graphical user interface
HDD	hard disk drive
I/O	input/output
JVM	Java virtual machine
LN	natural logarithm
MB	Megabytes
OS	operating system
PCI-e	peripheral component interconnect express
PSU	power supply unit
RAS	reliability, availability, and serviceability
RMS	root mean square
SSD	solid state drive
SEEM	sever energy effectiveness metric
SMT	simultaneous multi-threading
SPEC	server performance evaluation corporation
SUT	system under test
XML	extensible mark-up language

4 Applicability of Server Energy Effectiveness Metric (SEEM)

4.1 General

SEEM applies to all computer servers, as defined in 3.1.49. However, the SERTv2 metric (defined in Clause 5) only applies to the types of servers where it properly functions and accurately assesses server energy effectiveness.

For servers to which SERTv2 does not apply, an implementer of this document may develop an alternate test methodology (see 5.4).

4.2 Applicability of SERTv2

SERTv2 shall apply to all servers except those listed in this subclause, where it either was not designed to function or does not properly assess the server's energy effectiveness.

SERTv2 was not designed to function on servers which:

1) only support installation of less than 8 GB of system memory,

2) only support installation of system memory less than SERTv2 Minimum System Memory Size [see Formula (4)],

3) do not include a 64-bit CPU,

4) contain more than 4 CPU sockets (per blade for blade servers and per node for multi-node servers),

© ISO/IEC 2020 – All rights reserved