DRAFT INTERNATIONAL STANDARD
ISO/IEC DIS 21836

ISO/IEC JTC **1**/SC **39** Secretariat: **ANSI**

Voting begins on:
2019-03-07

Voting terminates on:
2019-05-30

Information Technology — Data Centres — Server Energy Effectiveness Metric

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

ICS: 35.020

This document is circulated as received from the committee secretariat.

Reference number
ISO/IEC DIS 21836:2019(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

CONTENTS

List of tables

List of figures

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

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 ISO documents 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 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 https://www.iso.org/iso-standards-and-patents.html).

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

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

The committee responsible for this document is ISO/IEC JTC 1, Information technology, SC 39, Sustainability for and by Information Technology.

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 to manage environmental, cost, electrical grid capacity and other impacts.

The ISO/IEC 30134 series of International Standards specify data centre energy effectiveness KPIs 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. It, however, does not provide a method of comparing individual server energy effectiveness across data centres and as stated in ISO/IEC 30134-4, “should not be used to set regulation 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 programs.

This standard 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 compliance verification, procurement and operation. Organizations that wish to establish compliance or reporting programs will find that the test methods and scoring specified in ISO/IEC 21836 will save them significant time and effort in implementing such programs. Standardization across such programs will allow vendors to comply to stakeholder requirements more quickly and efficiently.

For applicable servers, ISO/IEC 21836 builds upon the widely-adopted Server Efficiency Rating Tool™ (SERT), developed by the Standard Performance Evaluation Corporation (SPEC®), as the energy effectiveness metric and test method. For servers where SERT is not applicable, this standard provides requirements for the creation of alternate server energy effectiveness metrics, referred to as “implementer-specified” metrics.

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 standard does not set any pass/fail criteria for servers.

2 Normative references

The following documents are referred to in the text in such 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.

SPEC Server Efficiency Rating Tool (SERT) version 2.x.x

SPEC PTDaemon™ version 1.8.1 and later

SPEC SERT Run and Reporting Rules version 2.0.1 (20170929)

SPEC Client Configuration XML (https://www.spec.org/sert2/SERT-JVM_Options-2.0.html)

3 Definitions and abbreviations

3.1 Definitions

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:

- IEC Electropedia: http://www.electropedia.org/
- ISO Online browsing platform: http://www.iso.org/obp

3.1.1

64-bit server
server which includes a CPU which has data path widths, memory addressing, registers, and other architectural features which are 64-bits wide

3.1.2

active state
operational state in which the server is carrying out data processing in response to external requests

Note 1 to entry: Example: active processing, data seeking/retrieval from memory, cache, internal/external storage while awaiting further input over the network

3.1.3

active state power

average server power, in watts, when server is performing data processing

3.1.4

auxiliary processing accelerator

additional compute device installed in the computer server that handles parallelized workloads in conjunction with the CPU

3.1.5

blade chassis

enclosure that contains shared resources for the operation of blade servers, blade storage, 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

3.1.6

blade server

server that is designed for use in a blade chassis

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

3.1.7

blade storage

storage device that is designed for use in a blade chassis that is dependent upon shared blade chassis resources, like power supplies or cooling, for operation.

3.1.8

buffered memory

circuitry between the server's memory and memory controller to either increase memory capacity, increase bandwidth, and/or reduce the electrical load on the memory controller

3.1.9

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, Note 1 has been removed]

3.1.10

coefficient of variation

standard deviation divided by the mean

[SOURCE: ISO 3534-1:2006]

3.1.11

configuration

interrelated functional and physical characteristics of a product specified in product configuration information

Note 1 to entry: This standard employs the following configurations: low-end, high-end and typical

Note 2 to entry: For server products, a configuration is one of many possible combinations of components including CPU, storage devices, memory size and capacity and input/output devices for a single server 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, Note 1 and Note 2 have been added]

3.1.12

core

component of a processor which can independently receive instructions and takes actions or performs calculations in response

3.1.13

CPU

a central processing element with functions for interpreting and executing instructions.

Note 1 to entry: In this standard, cache memory is included with the CPU

Note 2 to entry: This standard uses the terms CPU and processor interchangeably

[SOURCE: ISO 14576:1999, 2.1.9 – modified, Definition is shortened, Note 1 and Note 2 have been added]

3.1.14

CPU architecture

CPU design with significant similarities to other CPU architectures within the same CPU architecture class

Note 1 to entry: CPU architectures are used to create CPU models which are often released in a similar timeframe

Note 2 to entry: Examples of CPU architectures in the same CPU architecture class are Intel® Haswell, Intel Broadwell, and Intel Skylake, or separately, AMD® Bulldozer, AMD Piledriver, and AMD Steamroller

3.1.15

CPU architecture class

group of one or more CPU architectures which share the same instruction set architecture and in which newer architectures designs are derived from previous architecture designs

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

Note 2 to entry: Examples of different CPU architecture classes are ARMv8-A® and AMD EPYC®

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

3.1.16

CPU model

specific CPU that is sold in the marketplace

Note 1 to entry: All CPUs of the same model share the same technical characteristics, such as core 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

3.1.17

CPU nominal frequency

the CPU core clock frequency which is the main frequency used in naming, marketing, and selling the CPU

3.1.18

data averaging internal

for a power analyser, 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

3.1.19

double data rate

transfers data on the rising and falling edges of the memory clock signal, resulting in twice the memory bandwidth capacity at a specific clock frequency

3.1.20

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 standard, the end user is the entity applying for certification of a server model to a SEEM compliant regulation or program. For example, if server manufacturer A was submitting a server model to ENERGY STAR for certification, server manufacturer A would be the end user.

[SOURCE: ISO/IEC/IEEE 24765:2010, 3.990, Note 1 has been added]

3.1.21

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 standard energy effectiveness is equivalent to the term energy efficiency as used in server compliance regulations and programs such as ENERGY STAR®.

3.1.22

expansion auxiliary processing accelerator

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

Note 2 to entry: example is a GPGPU installed in a PCI-e slot

3.1.23

fair use

uses of content that are considered valid defences to copyright infringement, such as for criticism or educational purposes

Note 1 to entry: For the purposes of this standard, fair use refers to the guidelines the developer of a metric has specified related to how results can be used

[SOURCE: ISO 19153:2014, 4.14 – modified, original Note 1 remove, Note 1 added]

3.1.24

fully fault tolerant server

computer server that is designed with complete hardware redundancy, in which every computing component is replicated between two nodes running identical and concurrent workloads

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

3.1.25

hardware threads

in a CPU core, the number of fully independent instruction streams which can be executed through SMT

3.1.26

high performance computing (HPC) system

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, primarily for increased computational capability, high speed inter-processing interconnects, large and high bandwidth memory 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

3.1.27

high-end configuration

server equipped with a specific selection of high performance components, which is required to be tested as part of measuring a server product family

3.1.28

idle state

operational mode in which the OS and other software have completed loading, the server is capable of completing workload transactions, but no active workload transactions are requested or pending by the system

Note 1 to entry: In the idle state, the server 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.

3.1.29

idle state power

average server power, in watts, when in idle state

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

3.1.30

implementer

entity that transforms specified designs into their physical realization

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

[SOURCE: IEC 62279:2015, 3.1.15, Note 1 added]

3.1.31

integrated auxiliary processing accelerator

auxiliary processing accelerator that is integrated into the motherboard or CPU package

3.1.32

large network equipment

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

3.1.33

large server

resilient/scalable server 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

3.1.34

load level

percentage of data processing relative to the maximum a server 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 utilization

3.1.35

low-end configuration

server configuration which includes a specific selection of entry level components, which is required to be tested as part of measuring a server product family

3.1.36

memory channel

independent interface in a computer which facilitates the communication of data between a core's memory controller and installed memory DIMMs

Note 1 to entry: Modern computer servers usually have a number of memory channels connected to different CPUs or cores