
**Information technology — Data centres
— Server energy effectiveness metric**

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

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 21836:2020



STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 21836:2020



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

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

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.

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.

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.

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.]

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 Steamroller³⁾.

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 8180³⁾.

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.

[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)

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)

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.

[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)

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

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

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	server 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),

- 5) contain a CPU from a CPU architecture class which is not supported by the SERTv2 test method,
- 6) do not support balanced and evenly distributed memory population (same number and type of DIMMs per memory channel and per CPU),
- 7) are only sold with a DC power supply, or
- 8) are quantum computers.

SERTv2 does not properly assess the energy effectiveness of servers which:

- 1) ship with one or more integrated APA,
- 2) are high-performance computing systems,
- 3) are fully fault tolerant servers,
- 4) are large servers,
- 5) are server appliances,
- 6) are storage products,
- 7) are blade storage products, or
- 8) are large network equipment.

4.3 Determination of applicability for an “implementer-specified” metric

For all “implementer-specified” metrics (as described in 5.4), the implementer shall specify to which types of servers the metric applies, in accordance with all of the following requirements:

- 1) SEEM metrics are only applicable to servers.
- 2) SEEM implementations shall only have one active state metric which applies to a given server category. The addition of one idle state metric to each server category is also allowed.
- 3) No active state “implementer-specified” metrics shall apply to servers for which the SERTv2 metric applies (see 4.2).
- 4) The test method shall function on all applicable servers (see 7.9.2).
- 5) The implementer shall only apply a metric to servers where there is an expectation that a better energy effectiveness score indicates improved real-world energy effectiveness during typical usage (see 7.9.3).

NOTE See 4.2 as an example of the server applicability specified for SERTv2.

5 Determination of Server Energy Effectiveness Metric (SEEM)

5.1 General

SEEM is a complete methodology to assess and report the energy effectiveness of a server. The SEEM metric(s) is/are an active state and optional idle state numeric result(s) that quantifies a server’s energy effectiveness, based on the measurements of the corresponding SEEM test method. For active state metrics, the larger the value, the less energy per unit of data processing the server consumes when performing data processing similar to the applicable metric’s test method. For idle state metrics, smaller is better, indicating that the server consumes less energy while ready to perform, but not performing any data processing.

All SEEM implementation shall include an active state metric for each server type, and an optional idle state metric may be added to any server type for which an active state metric is included. Idle state metrics shall not be included for server types which do not also have an active state metric.

For servers where SERTv2 applies (see 4.2), the active state portion of SEEM shall be equal to the numeric overall result of SPEC SERT 2.x.x, referred to in this document as SERTv2.

For servers where SERTv2 does not apply, the active state portion of SEEM may be equal to the numeric result of an alternate active state test method, if specified by the implementer.

Implementers may choose to add an idle state measurement to a SEEM implementation, which shall be reported as a separate numeric result from the main active state result.

5.2 Power supply requirements

As part of a SEEM implementation, for all servers which include alternating current (AC) power supplies, the implementer may require either only reporting or reporting and a minimum required 80 PLUS® power supply efficiency certification level. If a minimum 80 PLUS efficiency level is required by the implementer, the implementer may also specify additional power supply energy efficiency requirements, such as an efficiency requirement at a percent loading level not included in 80 PLUS. If an implementation includes any power supply efficiency requirements, a minimum 80 PLUS certification level shall be specified. All power supplies in tested servers shall meet all implementer specified energy efficiency certification levels.

NOTE Information about power supply efficiency requirements, the test protocol used to certify power supplies, and the process to get new power supplies certified, is available at [17].

5.3 SERTv2 active state energy effectiveness metric

The SERTv2 metric is a bigger-is-better number which indicates the overall energy effectiveness of a server. SERTv2 measures the energy effectiveness of an array of types of data processing related to typical server usage and simulates periods when a server is only partially utilized. When the server is not fully utilized, most servers intermittently enter periods of idle state, resulting in the SERTv2 metric effectively including an assessment of a server's idle state energy consumption. The SERTv2 test method consists of four components which shall be used to accurately obtain a SERTv2 result. These are SERT, PTDaemon, the Client Configuration XML file, and the SPEC SERT Run and Reporting Rules.

The SERTv2 worklets are designed to model typical server usage, however real-world results may vary. The energy consumption measured by this method should not be assumed to represent the energy consumption of other applications on the same hardware.

The numerical value of the SERTv2 metric shall be determined and described by [Formula \(1\)](#) using the SPEC SERTv2 test method.

NOTE The analysis of component sensitivity of real-world workloads, energy effectiveness rankings of hundreds of servers using different weighting schemes, and deployed power analysis ([Annex D](#)) suggests that using weights of 65 % for CPU, 30 % for Memory, and 5 % for Storage most accurately models typical server usage. Four *SPEC SERT Metric Weighting Studies*^[1] were conducted to determine these values.

All software and testing components in the measured system under test (SUT) shall be SEEM compliant, as described in [9.2.1](#).

$$SERTv2 = \exp(0.65 * \ln[Eff_{CPU}] + 0.30 * \ln[Eff_{Mem}] + 0.05 * \ln[Eff_{Stor}]) \quad (1)$$

where

$\exp$ is the exponential function;

Eff_{CPU} is the energy effectiveness of the SERT CPU workload;

Eff_{Mem} is the energy effectiveness of the SERT Memory workload;

Eff_{Stor} is the energy effectiveness of the SERT Storage workload;

$\ln$ is the natural logarithm.

CPU, Memory, and Storage are three workloads which are calculated as described by [Formula \(2\)](#).

$$Eff_{wd} = \exp\left(\frac{1}{n} \times \sum_{i=1}^n \ln(Eff_{wt,i})\right) \quad (2)$$

where

Eff_{wd} is the effectiveness of workloads;

$\exp$ is the exponential function;

N is the number of worklets in a given workload;

$\ln$ is the natural logarithm;

$Eff_{wt,i}$ is the effectiveness of worklets for a given workload.

— the CPU workload has seven worklets (Compress, CryptoAES, LU, SHA256, SOR, Sort, and SS);

— the Memory workload has two worklets (Flood3 and Capacity3);

— the Storage workload has two worklets (Random and Sequential).

Worklet effectiveness is calculated as described by [Formula \(3\)](#).

$$Eff_{wt} = \exp\left(\frac{1}{m} \times \sum_{i=1}^m \ln(Eff_{li})\right) \times 1000 \quad (3)$$

where

$Eff_{wt,i}$ is the effectiveness of worklets;

$\exp$ is the exponential function;

M is the number of load levels for a given worklet;

Ln is the natural logarithm;

Eff_l is calculated by dividing the normalized performance of each load level by the power consumed at that load level.

The memory worklets do not use [Formula \(3\)](#) to calculate Eff_{wt} , but instead use a calculation provided in the *SPEC SERT 2.x.x Design Document*^[2].

Details on the implementation of each worklet are available in the *SPEC SERT 2.x.x Design Document*^[2].

5.4 Determination of “implementer-specified” metrics

5.4.1 General

SEEM allows implementers to select test methods for servers where SERTv2 is not applicable (see [4.2](#)), if they satisfy the requirements of [5.4.2](#) for active state metrics and [5.4.3](#) for idle state metrics. Since the general requirements of [5.4.2](#) and [5.4.3](#) are only a subset of analyses performed on SERTv2 and do not rigorously ensure a test method is as real-world relevant as SERTv2, and for the purposes of result comparability and standardization across SEEM regulations and programmes, implementers of SEEM shall only use SERTv2 for applicable servers.

Alternate active state test methods shall only be used for servers where SERTv2 is not applicable. The creation of idle state power metrics in addition to an active state metric is allowed but not recommended.

5.4.2 “Implementer-specified” active state metrics

The development of a metric, and the corresponding test method used to measure it, is a complex process. If due diligence is not taken during the design, implementation, and use of the test method, the metric may not indicate real-world energy effectiveness.

All “implementer-specified” active state metrics shall satisfy the following requirements.

- 1) Each active state metric shall be a single numerical result which combines performance and energy measurements to indicate the energy effectiveness of the SUT.
- 2) An active state metric shall be a bigger-is-better value, where an increasing score indicates increased server energy effectiveness.
- 3) The active state metric shall include reporting of a single performance result which indicates the overall performance of the server when running relevant workloads.
- 4) The active state metric shall include reporting of a single energy or average power result which indicates the energy consumption of the entire server when running the metric.
- 5) Each metric shall be analysed to verify it accurately indicates real-world energy effectiveness for applicable servers (see [7.9.3](#)).
- 6) Results shall be reproducible (see [7.9.2](#)).
- 7) Metrics shall not artificially limit or favour any certain type of server.

Further, the test method used to determine the “implementer-specified” metric shall satisfy the following requirements.

- 1) The workloads of the test method shall be relevant examples or representations of typical real-world workloads performed by applicable servers for the metric.
- 2) For metrics applying to servers which perform many types of data processing, an adequate and appropriate number of workloads to represent a cross-section of performed data processing shall be included.

- 3) It is strongly recommended, but not required, that different load levels are included in the test, similar to SERTv2.
- 4) Real-time system level power data shall be collected during the execution of the measured portions of the test method.
- 5) The test method shall include methods to verify that the test was successfully completed, and results are accurate.
- 6) The test method shall include automation and user interfaces to reduce the complexity of result collection.

5.4.3 “Implementer-specified” idle state metric

5.4.3.1 General

Idle state metrics indicate the average power consumption of a server when it is ready to, but not executing data processing. For a given server type, idle state metrics shall only be created for types of servers which have an active state metric. SEEM implementations shall not include standalone idle state metrics. This is because standalone idle state metrics tend to inadvertently incentivize more energy consumption than active state metrics. The *Green Grid Idle state vs. Active Metric Study*^[3], showed up to 35 % higher data centre energy consumption from the implementation of an idle state metric instead of an active state metric.

For a given server type, adding an idle state metric to an implementation which includes an active state metric is allowed but not recommended for several reasons. First, pages 15-21 of the *Green Grid Active vs. Idle state and Active Study*^[4], show that selecting servers based only on an active state metric resulted in more energy savings in the data centre than using both an active state and separate idle state metric. Second, idle state metrics are nearly impossible to successfully design and maintain. Creating an idle state metric is complex because by default, most added functionality increases a server's idle state power, but the extra capability is not included in the metric calculation. To avoid favouring the smallest, least capable servers, a manual system is needed to adjust the idle state power limit to compensate for the additional capability of optional or more capable components. This results in a complex system of idle state adders which needs to be often adjusted and updated as technology advances. Unfortunately, with the rate at which component technology is changing and the number of types of components, virtually all adder systems incorrectly compensate for some components or capabilities. This error can become the dominant differentiator between server results, instead of the server's actual idle state power consumption, causing a programme to incorrectly rank servers in energy consumption. To avoid changing requirements and non-comparable results, government regulations and programmes understandably tend to be slow to update adder systems, which worsens this problem.

For servers where SERTv2 is applicable, the addition of an idle state metric is even more strongly discouraged because idle state power is already measured and included as part of the SERTv2 active state energy effectiveness metric, as shown in the *Green Grid Active Efficiency Measurements Include Idle State Study*^[5].

5.4.3.2 Requirements for creation of an idle state metric

If an implementer of this document wishes to add an idle state metric to an implementation, the following steps and requirements shall be followed.

- 1) In a SEEM implementation, idle state metrics shall only be created for server types which also have an active state metric. Standalone idle state metrics shall not be created.
- 2) For all servers where SERTv2 applies (see 4.2), the idle state measurement result within SERTv2 shall be used, and for all other servers, the test method defined in 7.5 shall be used.
- 3) A component adder system shall be created to add idle state power credits for all server capabilities that use noticeable amounts of idle state power.

- 4) A process shall be established to review component adders at least every two years and make adjustments and additions to the adders to align with modern component technology.
- 5) A per component analysis shall be performed to identify if one idle state power credit can properly account for all variants of a component type, or if the component type needs to be subdivided and each division assigned a separate adder.
- 6) A compute adder shall be developed which calculates increasing idle state power credits for all servers with more CPU data processing capability. A performance metric relevant to the server type shall be identified by the implementer for use in calculating the compute adder.

6 SEEM implementation

6.1 General

The purpose of SEEM is to assess and report server energy effectiveness, and an important element of a successful SEEM implementation is the completion of the specified steps and decisions defined in this document. To ease implementation, this clause lists the requirements from the rest of this document and outlines the steps to create each type of SEEM implementation.

6.2 Implementation steps for SEEM

The following steps shall be completed for all SEEM implementations.

- 1) The implementer shall decide how many and which metrics to include in their regulation or programme. If the implementation applies to any servers for which the SERTv2 metric applies (see [4.2](#)), then the implementation steps in [6.3](#) shall be followed. If any active state metrics other than SERTv2 are included, then the implementation steps in [6.4](#) shall be followed.
- 2) For all server types with AC power supplies, decide if a power supply efficiency reporting or minimum efficiency requirement is included, as described in [5.2](#).
- 3) The implementer shall decide whether or not to add an idle state power metric to each active state metric in an implementation (see [5.4.3](#)).

6.3 Implementation steps for SERTv2

To include SERTv2 in a SEEM implementation, the following steps shall be applied:

- 1) Decide whether the implementation requires energy effectiveness thresholds, and if so, per category, select a threshold which satisfies the requirements of [6.5](#).
- 2) Select one of the category system options in [8.1.2](#), and specify if any categories will be combined.
- 3) Specify any additional required documentation fields (see [8.2.2.4](#)).
- 4) Publish a report showing how the implementation complies with all SEEM requirements (see [8.2.2.2](#)).

6.4 Implementation steps for “implementer-specified” metrics

In order to include an “implementer-specified” metric in a SEEM implementation, the steps of this subclause shall be followed. If multiple “implementer-specified” metrics are included in a SEEM implementation, the following steps shall be applied for each metric.

- 1) Develop a metric and test method which satisfies the design requirements of [5.4](#).
- 2) Specify the applicable server types (see [4.3](#)).

- 3) Specify reporting categories (see [8.1.3](#)).
- 4) Define family configurations (see [7.8.4](#)).
- 5) Define any special testing configurations (see [7.8.4](#)).
- 6) Decide whether the implementation will require energy effectiveness thresholds, and if so, per category, select a threshold which satisfies the requirements of [6.5](#).
- 7) Specify the type of power that is compatible with the test method (see [7.4](#)).
- 8) Identify the technical support provider and provide a copy of the technical support commitment (see [7.7.1](#)).
- 9) Test and analyse the metric and test method, as specified in [7.9](#).
- 10) For all active state metrics, specify the Software Testing Variance (see [7.6.2.1](#)).
- 11) Publish a report showing how the implementation complies with all SEEM requirements (see [8.2.2.3](#)).
- 12) Specify the metric's reporting requirements (see [8.2.2.5](#)).
- 13) Specify the required test method components (see [5.3](#) as an example).
- 14) Write test plans with required outcomes for any components that will be allowed to be updated (see [9.4](#)).

6.5 Threshold selection

The implementer may choose to specify required active state and idle state (if included) threshold(s) per server category, or combination of categories, for each SEEM metric. If selected, only servers which score higher than the selected active state threshold for their category and lower than the required idle state threshold (if included), shall be certified under the SEEM regulation or programme. If any thresholds are included in an implementation, they shall satisfy the following requirements, with the following exception. If SEEM is implemented as an update or direct replacement to an existing regulation or programme, the threshold determination method used in the previous version of the regulation or programme may also be used.

- 1) For each metric in a SEEM implementation, per category, measurements shall be collected from at least 12 server product families for SERTv2, and at least 4 server product families for "implementer-specified" active state metrics. Selected servers shall include CPUs from a CPU architecture class released in the last 3 years, and, per category, shall represent a cross-section of the majority of applicable server types. SERTv2 measurements should include the majority of CPU Architecture Classes, CPU Architectures, server manufacturers, operating systems (OS), server component configurations (such as storage devices, dual in-line memory (DIMM) types and sizes, network cards), and software components (such as Java Virtual Machines [JVM]).

NOTE For SERTv2, a link to a database with more than 700 server results is available^[6].

- 2) Select a percentage of servers which are targeted to pass the regulation or programme.
- 3) For each category (or group of categories if any categories are combined), create a list of servers that meet the requirements of that category. Include only the lowest scoring server from the three servers in each server product family.
- 4) Per category, rank the selected servers by metric score, and multiply the number of servers by the target percentage (rounding down). Select the server whose rank matches this value.
- 5) One threshold shall be selected for each category. The selected threshold shall be $\pm 10\%$ of the score of the server selected in Step 4). Adjustments from the selected server's score shall only be based on factors such as server database diversity, server age, or number of servers tested.

7 Server testing

7.1 Configuration

All basic input/output system (BIOS), software, and hardware configurations shall be configured as they would be received by a server end customer.

Special Cases:

- 1) If a server model is always received by customers without an OS, a compatible OS may be installed, and all default OS installation and configuration settings shall be used.
- 2) If using default OS installation and configuration settings causes the server to not be able to be tested using the applicable SEEM test method, non-default settings required for functionality may be used.
- 3) If different server end customers receive the same server configuration with different settings, then the setting received by the most server end customers shall be used.
- 4) If a server OS, BIOS, or other component requires users to select settings on first boot and a selection menu includes a default value (i.e., would be selected by hitting enter, selecting next, or similar), then this default selection shall be used. If the menu has no default (i.e., an active state selection has to be made before proceeding), then the tester shall use the setting that is expected to be most commonly selected by server end customers.

7.2 Environment

The following environmental conditions shall be met for all SEEM testing. Air flow shall not be overtly directed in the vicinity of the measured equipment in a way that would be inconsistent with normal data centre practices.

- 1) Ambient temperature equal to or greater than 20 °C.
- 2) Ambient temperature below the documented maximum operating temperature of the SUT.
- 3) Elevation within the documented operating specification of the SUT.
- 4) Humidity within the documented operating specification of the SUT.

NOTE The intent is to discourage extreme environments that can artificially impact energy consumption or performance of the server, before and during the measurement.

7.3 Power analyser calibration requirement

The power analyser used for testing shall have been calibrated within the 12 months preceding the test date.

7.4 Power requirements

SERTv2 testing shall use power within one of the listed frequency and voltage ranges.

AC Frequency: ± 1 % of 50 Hz or 60 Hz

Voltage: ± 5 % of 100 V, 110 V, 120 V, 200 V, 208 V, 220 V, 230 V, 240 V or 400 V

For “implementer-specified” metrics, the implementer shall specify the allowed power types and ranges.

7.5 SEEM idle state power test method

While idle state power metrics are not recommended (see 5.4.3.1), an implementer may choose to add an idle state power requirement or idle state power reporting requirement to a SEEM implementation. For all servers where SERTv2 applies (see 4.2), the idle state measurement within SERTv2 shall be used, and for all other servers, the test method in this subclause shall be used. It is expected that the result from the idle state measurement built into SERTv2 is comparable to the result of this test method.

- 1) Obtain a power analyser which is listed in Annex B or satisfies the requirements of 9.2.2.
- 2) Correctly configure the power analyser and measurement logging software (often distributed by the power analyser manufacturer) to log the server's total energy consumption (volts and amps) at least one time per second. Measurement logging software shall be run on a separate system from the SUT to avoid affecting the measurement.
- 3) With one possible exception, configure the SUT as described in 7.1. If the default power configuration causes the server to enter a limited functionality low power state, such as sleep or hibernate, before the end of the measurement period, the setting to delay entering such states until after the measurement period, shall be set.
- 4) Boot the server into the operating system and do not manually make any modifications to the system, such as opening or closing application or services. The SPEC SERTv2 application can be opened if it is used to measure idle state power.
- 5) Let the system sit idle state for between 5 and 6 mins.
- 6) Measure the server's energy consumption for 1 min, collecting at least one measurement per second.
- 7) $SEEM_{idle\ state}$ for the SUT shall equal the arithmetic mean of power, in watts, during the 1-min measurement period.

7.6 Testing Variance

7.6.1 General

Server performance and power measurements inherently have result differences when repeated, called variance. If the same benchmark is run many times on the same server, it will not obtain the same performance or power result each iteration, and this difference in score is referred to as software variance in this document. Further, if two servers are built with the same types of components, small differences between the components can cause changes in the server's performance and power consumption, causing hardware variance. Some components are even designed to benefit from variations between components of the same type. For example, some CPU models include the ability for a given CPU to identify its temperature and energy consumption levels in real time and use this information to increase its frequency to the individual CPU's maximum ability. This allows a CPU which has better power consumption characteristics to outperform a CPU of the same model with inferior power consumption characteristics.

When implementing SEEM, it is important to understand and accommodate these testing variances to avoid unintended consequences. A tester may obtain a valid passing score on a server model by a small margin, but later, test the same server model for verification and obtain a failing score, solely because of variance.

All SEEM implementations shall include a testing variance maximum percentage, which shall be the sum of the software and hardware testing variance maximum percentages. For the purposes of verifying compliance with a SEEM regulation that includes an active state threshold, a server model shall be considered to pass verification if its score is greater than the threshold minus the sum of the hardware and software testing variance maximum percentages. For purposes of verifying compliance with a SEEM regulation that includes an idle state power threshold, a server model shall be considered to pass verification if its score is greater than the threshold minus the hardware testing maximum.

It is important to understand that testing variance is only applicable to verification that a server passes the threshold after it has previously been shown to pass the threshold. As specified in 7.8, to pass, server configurations and server product families are required to pass the specified threshold, and there shall be no reduction in requirements for testing variances.

EXAMPLE If a SERTv2 implementation requires a SERTv2 active state score of 10, a 5 % software testing variance (as per 7.6.2.1) and a 5 % hardware testing variance (as per 7.6.3), then a server model being tested for certification would be required to meet a score of 10. If the implementer later obtained a server of the same model and configuration to verify compliance, the second server would be required to score 9 ($10 - 10 \cdot 0,05 - 10 \cdot 0,05$) or higher to pass verification.

7.6.2 Software Testing Variance Specification

7.6.2.1 SERTv2 Software Testing Variance

All SEEM implementations that include SERTv2 and a threshold shall include a 5 % software testing variance.

7.6.2.2 “Implementer-specified” Software Testing Variance

All “implementer-specified” SEEM implementations which include an active state energy effectiveness threshold shall specify a software testing variance per metric, and implementers shall include this variance in calculating if a server passes verification.

The software testing variance shall be a percentage up to 10 % which is equal to the measured CV percentage result of the variance testing step of 7.9.2, rounded up to the nearest whole percentage.

7.6.3 SEEM Hardware Testing Variance Specification

All SEEM implementations that include thresholds shall include a 5 % hardware testing variance.

7.7 Technical support

7.7.1 Technical support requirements

The test methods for all SEEM metrics shall have an identified technical support provider and statement of technical support commitment. For “implementer-specified” metrics, the implementer shall identify the provider and provide their technical support commitment (see 7.7.2 as an example). The scope and impact of different SEEM implementations will likely vary significantly, so the appropriate level and type of technical support versus cost, for a given programme, should be considered by the implementer.

7.7.2 SPEC SERTv2 technical support

For SERTv2, SPEC is the technical support provider, and details about SPEC support is included in the *SPEC SERT 2.x.x Sample License Agreement*^[7]. Details are directly quoted here (in italics) for informational purposes only.

a) Support type

SPEC has no obligation to provide updates, modifications or new releases of the Materials to USER under this Agreement. SERT suite Support Subscribers are provided technical support in English via email or the SPEC Support forum; <https://www.spec.org/forums/>. In some cases, SPEC may be able to provide support in languages other than English.

b) Response time

SPEC will provide an initial response to technical support questions within three (3) business days, excluding US federal holidays.

c) *Length of support*

- i) *Support will be offered on the major version of the SERT suite referenced in the latest revision of ISO/IEC 21836 until a different major version of the suite is referenced by a newer revision of ISO/IEC 21836, or SPEC terminates the SERT suite's inclusion in ISO/IEC 21836 (see e)), whichever occurs first.*
- ii) *When a different major version of the SERT suite is referenced in a new revision of ISO/IEC 21836, support for the previously referenced major version will be offered for four (4) years from the date of publication of the new revision of ISO/IEC 21836.*
- iii) *When a new minor version of the SERT suite has been shown by SPEC to be compliant with ISO/IEC 21836, support for the previous compliant minor version will be offered for nine (9) months.*

d) *New platform support*

Unless the inclusion of the SERT suite in ISO/IEC 21836 is terminated, SPEC will accept proposals for the SERT suite acceptance of additional CPU architecture classes, CPU architectures, JVMs, or OS versions during a version's entire support offering period. These updates may require an updated release of the SERT suite. Update acceptance proposals require the contribution of engineering resources to SPEC, either by the CPU silicon vendor or their nominated alternate, as described in the SERT suite's Client Configuration Acceptance Process (<https://www.spec.org/sert>).

e) *Termination processes*

SPEC may choose at any time and for any reason, in its sole discretion, to terminate the inclusion of the SERT suite in ISO/IEC 21836. In this case, SPEC will:

- i) *Provide notice both publicly and by email to ISO/IEC JTC 1, and*
- ii) *At SPEC's sole option, either:*
 - *Offer support for the final version of the SERT suite included in ISO/IEC 21836 for four (4) years from the effective date of the notice, or*
 - *Provide the SERT suite source code and necessary dependencies, such as the PTDaemon too, to a designated third-party which commits to make the SERT suite publicly available under commercial terms which are fair, reasonable, and non-discriminatory*
 - *Open source the SERT suite source code and necessary dependencies.*

f) *Support Pricing*

SPEC provides support under commercial terms that are fair, reasonable, and non-discriminatory. Each SERT suite license purchase includes support for one year from the date of purchase. Information about obtaining additional support is available at <https://www.spec.org/order.html>.

7.8 Server product family and special configuration testing

7.8.1 General

Servers are configured and sold in many different configurations. In order to reduce the amount of required testing, it is allowed to test three configurations of a server product family as a proxy for testing all the configurations of that family. Only if all three configurations obtain the required threshold (minimum score), shall the entire family of configurations pass. For verification testing, all other configurations in the family shall also meet the required threshold. If a server model is sometimes received by the end customer without all of the CPU sockets populated (e.g., 2 socket server with only 1 CPU), this configuration is considered a different server family, and does not pass as part of the product family of servers with all CPU sockets populated. Server models with partial CPU sockets shall be tested as a separate product family, or single configuration, and shall meet the threshold designated for the number of CPUs populated.

If a single configuration or an entire server product family are not received by an end customer with a storage device, a compatible storage device may be added to the configuration for testing purposes.

Testing of single configurations is also allowed (see 7.8.2). Additionally, some form factors of servers, such as blade, multi-node, or servers which include APAs, require special test procedures as described in 3) and 7.8.3.6.

7.8.2 Single configuration testing

The server shall be configured as follows:

- 1) The server hardware shall be configured in the exact configuration which the tester wishes to certify.
- 2) All memory channels shall be equally populated with at least one identical DIMM and a total system memory capacity that is greater than or equal to the minimum capacity required by the SERT test (see Formula (4)). If this server configuration does not meet these requirements, then SERTv2 was not designed to function on this server (see 4.2) and this server configuration cannot be tested in this configuration.

$$\text{SERTv2 Minimum System Memory Size} = 2\text{GB} + 496\text{MB} \times \text{System Threads} \quad (4)$$

where *System Threads* is calculated in Formula (5).

7.8.3 SERTv2 product family and special configurations

7.8.3.1 General

For SERTv2, the following family configurations and test methods for blade, multi-node, and servers with APAs shall be used. All configurations requirements of 7.1 shall be followed.

All memory channels shall be equally populated with at least one identical DIMM and a total system memory capacity that is greater than or equal to the minimum capacity required by the SERT test (see Formula 4). If this server configuration does not meet these requirements, then SERTv2 was not designed to function on this server (see 4.2) and this server configuration cannot be tested in this configuration.

7.8.3.2 High-end configuration

The server shall be configured as follows:

- 1) CPU model with the largest product of core count, hardware threads per core and CPU nominal frequency.
- 2) System memory capacity (GB) equal to or greater than 3,0 times the number of *System Threads*, as calculated in Formula (5). If no servers of this product family could be received by an end customer with a memory configuration that meets this requirement, a configuration that is not received by an end customer may be used.

$$\text{System Threads} = \# \text{ of CPUs} \times \# \text{ of cores per CPU} \times \# \text{ of hardware threads per core} \quad (5)$$

7.8.3.3 Low-end configuration

The server shall be configured as follows:

- 1) CPU model with the smallest product of core count, hardware threads per core, and CPU nominal frequency.

- 2) System memory capacity (GB) equal to or greater than 1,0 times the number of *System Threads*, as calculated in [Formula \(5\)](#). If no servers of this product family could be received by an end customer with a memory configuration that meets this requirement, a configuration that is not received by an end customer may be used.

7.8.3.4 Typical configuration

The server shall be configured as follows:

- 1) CPU model with a product of core count, hardware threads per core, and CPU nominal frequency larger than the CPU used in the low-end configuration and smaller than used in the high-end configuration. If no servers of this product family use CPUs that were not used in the high-end and low-end configurations, the CPU from the low-end configuration shall be used.
- 2) A server configuration which is a deployed product with high volume sales.
- 3) System memory capacity (GB) equal to or greater than 2,0 times the number of *System Threads*, as calculated in [Formula \(5\)](#). If no servers of this product family could be received by an end customer with a memory configuration that meets this requirement, a configuration that is not received by an end customer may be used.

7.8.3.5 Blade and multi-node testing

For multi-node systems, the unit under test shall be tested in a fully populated multi-node chassis, and the energy consumption value shall be the total energy consumed by the multi-node chassis divided by the number of nodes. All individual node installed in the chassis shall be identical, sharing the same configuration, and all memory channels shall be equally populated with at least one identical DIMM. If no servers of this product family ship with a memory configuration that meets these requirements, a memory configuration that does not ship with this product family may be used.

For blade systems, the unit under test shall be tested for blade server energy consumption in the half-populated chassis configuration, and the chassis shall be populated as follows.

- 1) All memory channels shall be equally populated with at least one identical DIMM.
- 2) All individual blade servers installed in the chassis shall be identical, sharing the same configuration.
- 3) Half chassis population shall be determined as follows.
 - a) The number of blade servers required to populate half the number of single-wide blade server slots available in the blade chassis shall be calculated.
 - b) For blade chassis having multiple power domains, the number of power domains that is closest to filling half of the chassis shall be chosen. If there are two choices that are equally close to filling half of the chassis, the test shall be performed with the domain or combination of domains which use a higher number of blade servers.
 - c) All user manual or manufacturer recommendations for partially populating the chassis, which may include disconnecting some of the power supplies and cooling fans for the unpopulated power domains, should be followed.
 - d) If user manual recommendations are not available or are incomplete, then the following guidance shall be used.
 - i) Completely populate the power domains.
 - ii) If possible, disconnect the power supplies and cooling fans for unpopulated power domains.
 - iii) Fill all empty bays with blanking panels or an equivalent airflow restriction for the duration of testing.

7.8.3.6 Expansion APA testing

For all configurations, SERTv2 testing shall be conducted with any offered expansion APAs removed from the server. Where an APA relies on a separate Peripheral Component Interface Express (PCI-e) switch for communication between the APA and CPU, the separate PCI-e card(s) or riser(s) shall be removed for SERTv2 testing of all configurations.

7.8.4 Product family and special configurations for “implementer-specified” metrics

For “implementer-specified” metrics, the implementer shall define a set of family and special configurations.

Selected family configurations shall:

- 1) significantly reduce the number of tested server configuration permutations,
- 2) maintain a high probability that untested configurations in the same family also pass the metric threshold,
- 3) include separate families for groups of configurations which include unique components which significantly affect the metric result.

Not including any families for special configurations is allowed if no applicable servers include unique characteristics which significantly affect the metric result.

For each special configuration an implementer specifies, changes to the default test method shall be specified to allow the results of testing the special configurations to be comparable to the rest of the servers in the same category. If an appropriate special test method cannot be devised, a separate category shall be created for each special configuration (see [8.1.3](#)).

NOTE Subclause [7.8.3](#) provides an example set of family ([7.8.3.2](#), [7.8.3.3](#) and [7.8.3.4](#)) and special configurations ([7.8.3.5](#) and [7.8.3.6](#)) which satisfy the above requirements and can be used as an example.

To further explain how this example, [7.8.3](#), satisfies the requirement of this subclause ([7.8.4](#)):

Subclause [7.8.3](#) satisfies [7.8.4 2](#)) in the following way. SERTv2 metric is sensitive to differences in server compute and memory capability. The requirement of using high-end and low-end CPU models, as well as different memory capacities, ensures a broad range of configurations with different performance scores is tested, reducing the chances that an intermediate configuration would not pass the requirement.

Subclause [7.8.3](#) satisfies [7.8.4 3](#)) in the following way. Blade servers share power consuming hardware components across multiple servers and comparison with typical servers would have significantly altered comparability, so test method modifications were developed ([7.8.3.5](#)) and blade and multi-node servers were put into their own category ([Table 1](#), Categories C1, G1, and J1). Similarly, servers with APAs, which SERTv2 does not test, would be significantly disadvantaged as APAs consume considerable power, so [7.8.3.6](#) allows special testing procedures.

7.9 “Implementer-specified” metric creation testing

7.9.1 General

When developing a new metric, as per the requirements of [5.4.2](#), testing and analysis is needed to determine the type of servers where the test method functions, and where it properly assesses real-world energy savings. This subclause describes these testing requirements. If some requirements cannot be met by the test method, adjustment to the test method or server applicability may be alternatives.

7.9.2 Functionality and reproducibility testing

The test method for all “implementer-specified” SEEM metrics shall be thoroughly validated in the following ways. This subclause specifies a minimum test plan, but significantly more testing is recommended before finalizing a new metric. If one test point satisfies the requirements of multiple

test requirements, the test does not need to be repeated. If required outcomes are not met by an “implementer-specified” metric, modification of the metric is required.

EXAMPLE If the measured software variance CV is greater than the allowed 10 %, then the metric would require modification to reduce variance, such as increasing the number of iterations or removing the most variant portions of the benchmark.

The test plan is as follows:

1) Hardware functionality testing

Run the test method, including power and temperature measurement, on a representative cross-section of applicable servers and server configurations. Include, as a minimum, the following:

- a) one CPU from each applicable CPU architecture class,
- b) one server from each category (8.1.3),
- c) servers using two types of power (7.4), unless only one type is supported,
- d) one server from each server family (7.8.4),
- e) if any special configurations are specified, one server of each special configuration type (7.8.4).

Required outcomes:

- i) All test points shall result in successful results with no errors.
- ii) Test points shall be analysed for expected trends, and any anomalies shall be investigated.

EXAMPLE If two CPUs from the same CPU architecture are compared, and one is expected to have significantly more compute capability but consume more power on this metric, the results are verified to ensure they follow this trend.

2) Software functionality testing

Run the test method, including power and temperature collection, on a representative cross-section of applicable OSs, software components, and software versions. Include, at a minimum, the following:

- a) one OS from each major OS variant,
- b) if test method is delivered as source code, two variants of the compiler.

Required outcomes:

- i) All test points shall result in successful results with no errors.
- ii) Test points shall be analysed for expected trends, and any anomalies shall be investigated and fixed as appropriate.

3) Software variance testing

Run the test method, including power and temperature measurement, on the same server with the same software configuration 10 times.

Required outcome:

The coefficient of variation (CV) of the overall energy effectiveness score of 10 runs shall be less than 10 % and should be less than 5 %.

7.9.3 Real-world energy savings testing

One of the main purposes of this document is to ensure SEEM compliant regulations or programmes incentivize real-world energy savings. When creating a new metric, considerable time should be spent analysing and verifying this factor. For the SERTv2 metric, a number of years were spent developing the mix of worklets, then SPEC and The Green Grid spent many months analysing the final metric weighting (see NOTE in 5.3). It is strongly encouraged to perform additional analysis, but as a minimum, the following analysis shall be completed for all “implementer-specified” active state metrics.

- 1) Deployed Power Charting (See [D.3.1](#), Charting deployed power)
- 2) Deployed Power Ranking Table (See [D.3.2](#), Creating a deployed power ranking table)

8 Reporting SEEM

8.1 Server categories

8.1.1 General

Servers are designed in different form factors and for many different purposes. Grouping servers into categories based on form factor or functional similarities, and then assigning different thresholds per category, reduces the inherent disadvantage for servers with extra capabilities not measured by the applicable metric.

For the SERTv2 test method, the categories of [Table 1](#) are designed to effectively remove inequities, such as the disadvantage for the extra reliability, availability, and serviceability (RAS) features in resilient servers or the advantage of shared hardware in blade servers.

Alternately, some implementations of SEEM may focus on energy effectiveness improvements for products relative to their predecessors instead of competitors. In this case, a category system which groups servers by CPU architecture classes is appropriate, as specified in [Table 2](#).

For SERTv2, the implementer shall choose either the [Table 1](#) or [Table 2](#) category systems. After selection, the implementer may choose to combine two or more categories into one group (for example, combine A1 and B1), only if the implementer has data demonstrating the differences between the combined categories does not have a material impact on energy effectiveness. This combination shall not change the category names; the new combined category name shall be a concatenation of the letters of all combined categories followed by the number (e.g. if combining A1, E1, and J1, the new category name is AEJ1) and shall be documented by the implementer.

For implementer defined metrics, a new category system per metric shall be developed by the implementer which meets the requirements of [8.1.3](#).

8.1.2 SERTv2 server category definitions

8.1.2.1 General

Categories for SERTv2 shall be one of the following options, as selected by the implementer. The threshold values for each category of the selected option shall be determined by the implementer and meet the requirements of [6.5](#).

Table 1 — SERTv2 server categories by form factor and function (Option 1)

Category	Server type
One installed CPU	
A1	Rack
B1	Tower
C1	Blade or Multi-Node
D1	Resilient
Two installed CPUs	
E1	Rack
F1	Tower
G1	Blade or Multi-Node
H1	Resilient
Three or four installed CPUs	
I1	Rack
J1	Blade or Multi-Node
K1	Resilient

Table 2 — SERTv2 server categories by CPU architecture class (Option 2)

Category	CPU architecture class
One installed CPU	
A2	x86
B2	Oracle SPARC®
C2	IBM Power®
Two installed CPUs	
D2	x86
E2	Oracle SPARC®
F2	IBM Power®
Three or four installed CPUs	
G2	x86
H2	Oracle SPARC®
I2	IBM Power®

8.1.2.2 Resilient servers**8.1.2.2.1 General**

For a resilient server to be included in one of the resilient server categories (Table 1, categories D1, H1, and K1), the requirements of the 8.1.2.2.2 to 8.1.2.2.5 shall be satisfied. A resilient server which does not meet all the following requirements may be included in one of the other categories if it meets all the requirements of that category.

8.1.2.2.2 Processor RAS

The processor shall have capabilities to detect, correct and contain data errors, as described by all of the following.

- 1) Error recovery by means of instruction retry for certain processor faults.

- 2) Error detection on L1 caches, directories, and address translation buffers using parity protection.
- 3) Single bit error correction (or better) on caches that can contain modified data. Corrected data is delivered to the recipient as part of the request completion.

8.1.2.2.3 System recovery and resiliency

No fewer than six of the following characteristics shall be present in the server.

- 1) Error recovery and containment by means of:
 - a) data poison indication (tagging) and propagation which includes a mechanism to notify the OS or hypervisor to contain the error, thereby reducing the need for system reboots, and
 - b) containment of address/command errors by preventing possibly contaminated data from being committed to permanent storage.
- 2) The processor technology is designed to provide additional capability and functionality without additional chipsets, enabling the design into systems with four or more processor sockets.
- 3) Memory Mirroring: A portion of available memory can be proactively partitioned such that a duplicate set may be utilized upon non-correctable memory errors. This can be implemented at the granularity of DIMMs or logical memory blocks.
- 4) Memory Sparing: A portion of available memory may be pre-allocated to a spare function such that data may be migrated to the spare upon a perceived impending failure in an active memory device.
- 5) Support for making additional resources available without the need for a system restart. This may be achieved either by processor (cores, memory, I/O) on-lining support, or by dynamic allocation/deallocation of processor cores, memory, and I/O to a partition.
- 6) Support of redundant I/O devices (storage controllers, networking controllers).
- 7) Has I/O adapters or storage devices that are hot-swappable.
- 8) Can identify failing processor-to-processor lane(s) and dynamically reduce the width of the link in order to use only non-failing lanes or provide a spare lane for failover without disruption.
- 9) Capability to partition the system such that it enables running instances of the OS or hypervisor in separate partitions. Partition isolation is enforced by the platform and/or hypervisor and each partition is capable of independently booting.
- 10) Uses memory buffers for connection of higher speed processor-memory links to DIMMs attached to lower speed DDR channels. Memory buffer can be a separate, standalone buffer chip which is integrated on the system board or integrated on custom-built memory cards.

8.1.2.2.4 Power supply RAS

All power supplies installed or shipped with the server shall be redundant power supplies and concurrently maintainable. The redundant and repairable components may also be housed within a single physical power supply but shall be repairable without requiring the system to be powered down. Support shall be present to operate the system in a degraded mode.

8.1.2.2.5 Thermal and cooling RAS

All active cooling components shall be redundant and concurrently maintainable. The processor complex shall have mechanisms to allow it to be throttled under thermal emergencies. Support shall be present to operate the system in a degraded mode when thermal emergencies are detected in the system components.

8.1.3 “Implementer-specified” metric server categories

For each new metric, the implementer shall determine if there are groups of servers whose results are significantly influenced by factors not intended to be measured by the test method (see examples in [8.1.1](#)). For such determined factors, the implementer shall choose to either specify special test procedures to remove such bias (see [7.8.4](#)), or group such servers into a separate category, or both. Once categories are determined, the metric result for all servers in each category shall be mainly differentiated by the server’s energy consumption and performance while executing the test method. Once such categories have been determined, the implementer may specify additional categories for the purpose of focusing on self-improvement (see [Table 2](#) as an example). The use of one category shall only be used if, for the metric, there are no applicable servers with results which are significantly influenced by factors not intended to be measured by the test method.

8.2 Documentation

8.2.1 General

Documentation is an important element of SEEM and it provides transparency into the testing configuration and settings, allows for independent reproduction and verification of results and provides SEEM regulations or programmes with important information. [8.2](#) lays out the minimal documentation requirements to reduce the time it takes to collect and report documentation, especially for SEEM end users who participate in several SEEM compliant programmes. It includes a description of the documentation requirements for SEEM implementers and SEEM end users.

8.2.2 Implementation documentation

8.2.2.1 General

This subclause provides details of the compliance reports all implementers shall write for SEEM implementations, and the additional documentation required for “implementer-specified” metrics. Reports shall be made easily available to the SEEM end users.

8.2.2.2 Implementer’s compliance report for SERTv2

Implementations of SEEM which include SERTv2 shall provide a document to demonstrate SEEM compliance for their regulation or programme. The implementer shall list each SEEM requirement and provide an explanation describing how the implementer’s selection satisfies the requirement for each of the following items:

- 1) Power supply efficiency requirement ([5.2](#))
- 2) Thresholds per category ([6.5](#))
- 3) Categories ([8.1.2](#))

8.2.2.3 Implementer’s compliance report for “implementer-specified” metrics

Implementers of SEEM which include an “implementer-specified” metric, shall, for each “implementer-specified” metric, provide a document describing how the metric is SEEM compliant. To demonstrate SEEM compliance, the implementer shall list each SEEM requirement and provide an explanation describing how the implementer’s selection satisfies the requirement for each of the following items:

- 1) Applicability ([4.3](#))
- 2) Power supply efficiency requirement ([5.2](#))
- 3) Metric Requirements ([5.4](#))
- 4) Thresholds per category ([6.5](#))

- 5) Power requirements ([7.4](#))
- 6) Technical support ([7.7.1](#))
- 7) Product family and special configurations ([7.8.4](#))
- 8) Testing requirements ([7.9](#))
- 9) Categories ([8.1.3](#))
- 10) SEEM end user documentation ([8.2.3.4](#))
- 11) Minor update requirements ([9.4](#))

8.2.2.4 SERTv2 metric reporting specification

Implementers of SEEM which include SERTv2, shall provide a document to SEEM end users with the implementer's selection of all steps in [6.3](#) and extra documentation fields if required (see [8.2.3.2](#)).

8.2.2.5 "Implementer-specified" metric reporting specification

Implementers of SEEM which include an "implementer-specified" metric shall, for each "implementer-specified" metric, provide a document with the implementer's selection of each of the following items, in a similar way to how these items are documented for SERTv2 in this document.

- 1) Applicability ([4.3](#))
- 2) Power supply efficiency requirement ([5.2](#))
- 3) Thresholds per category ([6.5](#))
- 4) Power requirements ([7.4](#))
- 5) Technical support ([7.7.1](#))
- 6) Software Testing Variance ([7.6.2.2](#))
- 7) Product family and special configurations ([7.8.4](#))
- 8) Testing requirements ([7.9](#))
- 9) Categories ([8.1.3](#))
- 10) SEEM end user documentation ([8.2.3.4](#))
- 11) Minor update requirements ([9.4](#))

Sufficient information shall be included so a SEEM end user of the implementation is able to submit a server for certification to the regulation or program.

8.2.3 SEEM end user documentation

8.2.3.1 General

This subclause describes the documentation that SEEM end users shall include in submissions to SEEM compliant regulations or programmes. For "implementer-specified" metrics, the required documentation is specified by the implementer (see [8.2.3.4](#)). For compliance reports, a link to a valid compliance report shall be the only documentation required to demonstrate SEEM compliance.

8.2.3.2 SEEM end user documentation

SEEM results shall include documentation of all information required to reproduce the result. In addition, the following documentation, related to the testing performed, shall be included in submissions to the regulation or programme:

- 1) Server category (see [8.1](#))
- 2) SEEM active state metric
- 3) SEEM idle state metric (if required by implementer)
- 4) Version of SEEM standard
- 5) Name of SEEM compliant regulation or programme
- 6) Version of SEEM implementation
- 7) Per power supply, 80 PLUS efficiency level (see [5.2](#))
- 8) SEEM compliance report (or link) for power analysers not listed in [Annex B](#)
- 9) SEEM compliance report (or link) for temperature sensors not listed in [Annex B](#)

The implementer may specify additional required documentation fields per metric.

NOTE For power analysers and temperature sensors tested by SPEC, SEEM compliance reports are found on the *SPEC PTDaemon Website* [\[8\]](#).

8.2.3.3 SERTv2 end user documentation

SEEM implementations which include SERTv2 shall require the collection and reporting of all of the following.

- 1) Results.xml file from the measurement run. This file includes all measurement information, such as performance and power measurements at all load levels of all worklets, the CPU, Storage, and Memory Efficiency scores, and the Overall Efficiency score. See [Figure 1](#) and [Figure 2](#) for examples of SERTv2 result output formats that SERTv2 generates and that can be regenerated from a results.xml file.
- 2) Accurate documentation of all SERTv2 GUI documentation fields listed in [Annex C](#).
- 3) Measured idle state power, if required by implementer.
- 4) SEEM compliance report (or link) for SUT CPU Architecture, JVM, and OS, if not listed in [Annex A](#).
- 5) If not testing with the initial version (see 2)), the SEEM compliance report (or link) for the version of SERTv2 used.
- 6) If not testing with the initial version (see 2)), the SEEM compliance report (or link) for the version of PTDaemon used.
- 7) Submitted as a product family or single server.
- 8) Model name and number for every expansion APA included in product family or single server testing.
- 9) For multi-node servers, number of nodes.
- 10) Total number of memory channels server supports.
- 11) For blade servers, number of blade servers tested in the half chassis population (see 3)).

12) Any additional documentation which would be required to independently reproduce the SERTv2 result.

13) Any additional documentation fields specified by the implementer

NOTE For new versions of SERTv2, it is possible that SPEC will make SEEM compliance reports available on the *SPEC SERT 2.x.x Website*^[9]. For new versions of PTDaemon, it is possible that SPEC will make SEEM compliance reports available on the *SPEC PTDaemon Website*^[8].




Figure 1 — Example SERTv2 HTML Result Output

SERT (TM) Report
Copyright (C) 2007-2017 Standard Performance Evaluation Corporation (SPEC).
All rights reserved.
The SERT license agreement prohibits the use of numerical information from this report to make public comparisons promoting the use of one product over another.

HEWLETT PACKARD ENTERPRISE PROLIANT DL120 G7 (HISTORICAL SYSTEM)			
Test Sponsor	Hewlett Packard Enterprise	Software Availability	Apr-2016
Tested By	HPE Performance Engineering	Hardware/Firmware Availability	Jul-2016
SPEC License #	3	System Source	Single Supplier
Test Location	Houston, TX, USA	Test Date	Mar 25, 2017
SERT 2 Efficiency Score	6.4		

Summary

AGGREGATE SUT DATA	
Model	Hewlett Packard Enterprise ProLiant DL120 G7 (Historical System)
Number of Nodes	1
CPU Name	Intel(R) Xeon(R) CPU E3-1240 @ 3.30GHz
Total Number of Processors	1
Total Number of Cores	4
Total Number of Threads	8
Total Physical Memory	8.0 GB
Total Number of Memory DIMMs	2
Total Number of Storage Devices	1

Workload Efficiency Score			Idle Watts
CPU	Storage	Memory	
5.6	14.0	7.2	50.1

Figure 2 — Example SERTv2 Text Result Output

8.2.3.4 “Implementer-specified” metric SEEM end user documentation

For each “implementer-specified” metric, the implementer shall provide a document to SEEM end users including all required documentation. Required documentation shall be sufficient to independently reproduce server results using the metric’s test method. For examples, see [Annex C](#) and [8.2.3.3](#).

9 Minor update acceptance testing

9.1 General

In the evolving world of computer servers, software updates are needed to enable support for new hardware and software platforms, to address software bugs, or to enhance usability. However, in order to accurately compare SEEM results, consistency in components which affect the measurement is required. Additionally, when SEEM is implemented, it is imperative that changes which affect result comparability are not introduced for the duration of the regulation or programme. To accommodate these requirements, the following test plans have been developed to maintain comparability of measurement results while allowing updates. For “implementer-specified” metrics, a set of update requirements shall be created, as described in [9.4](#).

9.2 SEEM minor update testing and requirements

9.2.1 SEEM component versions

For SEEM, only the following component versions shall be used:

- SPEC PTDaemon™ version 1.8.1 (initial version), 1.9.0, 1.9.1, or later versions that comply with [9.3.3](#)⁴⁾,
- Power analysers: Power analyser models listed in [Annex B](#), or a power analyser which complies with all the requirements of [9.2.2](#).
- Temperature sensors: Temperature sensor models listed in [Annex B](#), or any temperature sensor which complies with all the requirements of [9.2.3](#).

9.2.2 Power analysers

9.2.2.1 General

A list of power analysers which meet the requirements of [9.2.2](#) is available in [Annex B](#). The subclauses of [9.2.2](#) contain a set of general requirements, two test plans to verify a power analyser’s accuracy, and an additional test plan related to the power requirements of a given SUT to be tested.

In order to use a new power analyser with SERTv2, software support for the new power analyser shall be added to the SPEC PTDaemon tool. The SPEC process to add this support is included in the *SPEC Power Analyzer and Temperature Sensor Acceptance Process Document*^[10].

The SPEC adoption process may include additional requirements to those described in [9.2.2](#).

9.2.2.2 General power analyser requirements

The list of power analysers that have already been shown to comply with these requirements are listed in [Annex B](#). Additional power analysers may be used only if they have been shown to satisfy all of the following requirements.

- 1) Measurements — the analyser shall report true root mean square (RMS) power in watts and at least two of voltage, amps or power factor.

4) PTDaemon is a trademark of the Standard Performance Evaluation Corporation. References to PTDaemon do not constitute an endorsement by ISO/IEC.

- 2) Measurements shall be reported by the analyser with an overall uncertainty of 1 % or better during all sections of the pulse acceptance test, specified in [9.2.2.4](#).
- 3) Calibration — the analyser shall be able to be calibrated by a standard traceable to the National Institute of Standards and Technology (U.S.A) or a counterpart national metrology institute in other countries.
- 4) Crest Factor — the analyser shall provide a current crest factor of a minimum value of 3. For power analysers which do not specify the crest factor, the analyser shall be capable of measuring an amperage spike of at least 3 times the maximum amperage used in the pulse acceptance test, specified in [9.2.2.4](#).
- 5) Logging — the analyser shall support reading data in real time, using a documented software interface. If a software package is required to access this interface, this software shall be available to all users of the device, or freely available and redistributable. Analysers shall support a reading rate of at least 1 set of measurements per second. The measurement set shall include watts and at least two of volts, amps, and power factor. The data averaging interval of the analyser shall be 1 (preferred) or 2 times the reading interval.

9.2.2.3 Simultaneous measurements test

To verify the measurement accuracy of a power analyser not included in [Annex B](#), the following test shall be performed and shall achieve the specified result.

- 1) Select an SUT, a power analyser from the list in [Annex B](#) and a new power analyser to be tested.
- 2) Ensure the new power analyser satisfies all the requirements of [9.2.2.2](#).
- 3) Connect the new power analyser and the power analyser from [Annex B](#), in series, between the SUT and the power source (in either order).
- 4) Configure and run the SSJ and Idle state worklets of SERTv2.
- 5) Reverse the order of the power analysers, so the opposite analyser is now connected to the SUT. This step is intended to eliminate any measurement differences due to voltage drops or energy consumption of the analysers themselves.

Repeat Step 4).

Required outcome:

For all intervals of the measurement phase of each SERTv2 worklet, the difference between the average power values from the two power analysers shall not be larger than the sum of the uncertainties of each power analyser at each measured power level.

NOTE As a convenience to the tester, the uncertainties for each worklet are calculated by PTDaemon, and reported in the SERTv2 results-detail files in the "Power Measurement Uncertainty" column of the "Power Data" section.

9.2.2.4 Pulse acceptance test

To verify the measurement accuracy of a power analyser not included in the list in [Annex B](#), the following test shall be performed and shall achieve the specified result.

- 1) Configure a hardware or software programmable load generator to generate 10 pulses each of 50 ms, then 10 pulses each of 100 ms to 900 ms in 100 ms steps (10 total steps), starting every 5 000 ms. Baseline load shall be between 40 W and 150 W within one of the voltage ranges specified in [7.4](#). The pulses shall reach at least 20 % higher than the baseline.
- 2) Measure the pulses with the power analyser being evaluated.

- 3) All measured power values shall correspond to the pulse widths generated by the load generator, within a tolerance of $\pm 25\%$. All pulses shall be captured by the power analyser.
- 4) Reconfigure the load generator to generate 100, 50 ms pulses starting every 5 010 ms, followed by 100, 100 ms pulses starting every 5 010 ms.
- 5) Measure the pulses with the power analyser being evaluated.
- 6) The power analyser shall indicate all pulses, with two exceptions. It is allowed to indicate up to two duplicate sample readings and two missing sample readings. A duplicate sample reading is when the power analyser indicates two pulses occurred when the load generator only generated one pulse.

9.2.2.5 Server-specific power analyser requirements

There are many different sizes and types of servers which have differing power requirements. In addition to the general power analyser requirements of [9.2.2.2](#), there are also requirements that the power analyser has sufficient accuracy in the ranges used by the SUT.

When running SERTv2 on a SUT, using a frequency and voltage described in [7.4](#), the power analyser shall provide an uncertainty under 1 % and a crest factor of 3 or more, for loads with a power factor between 0,8 and 1,0, and wattages from 20W to the maximum used in a SERTv2 run on the SUT.

9.2.3 Temperature sensors

A list of temperature sensors which meet the requirements of this subclause is available in [Annex B](#). In order to use a new temperature sensor with SERTv2, it is required that software support for the new sensor be added to the PTDaemon tool. The SPEC process to add this support is available in the *SPEC Power Analyzer and Temperature Sensor Acceptance Process Document*^[10]. The SPEC adoption process may include additional requirements to those described in this subclause.

All temperature sensors used in SEEM shall meet all of the following requirements.

- 1) Logging — the sensor shall support reading data in real time, using a documented software interface. If a software package is required to access this interface, this software shall be available to all users of the device, or freely available and redistributable. The reading rate supported by the sensor shall be at least 4 samples per minute.
- 2) Accuracy — measurements shall be reported by the sensor with an overall accuracy of $\pm 0,5\text{ }^{\circ}\text{C}$ or better for temperatures equal to and between $20\text{ }^{\circ}\text{C}$ and $50\text{ }^{\circ}\text{C}$.

9.3 SERTv2 minor update testing and requirements

9.3.1 General

In order to ensure accuracy and comparability of results for a specific version of SEEM, [9.3.1](#) specifies requirements which ensure result comparability between all testing software and hardware components.

9.3.2 SERTv2 Component versions

For SERTv2, only the following component versions shall be used:

- SPEC Server Efficiency Rating Tool (SERT) version 2.0.1 (initial version), 2.0.2, 2.0.3, or later 2.x.x versions which comply with [9.3.3](#).
- SPEC SERT Run and Reporting Rules Version 2.x.x (20190815)^[16] only.
- SERT Client Configuration XML: CPU architecture, OS, and JVM version sets listed in [Annex A](#), or which comply with [9.3.5](#).

9.3.3 SERTv2

Updated minor versions of SERTv2 shall be tested with the following test plan and meet the stated requirements.

- 1) Select two servers which are currently available for sale from [Annex A](#), or that have been shown to be compliant with SEEM. Select two OSs which are developed by different companies, and two CPU architectures developed by different companies.
- 2) Configure each server as specified in [7.1](#).
- 3) For each configuration, run a baseline compliant SERTv2 run using the initial version of SERTv2.
- 4) For each configuration, run a compliant SERTv2 run using the new version of SERTv2.
- 5) Per configuration, the overall SERTv2 efficiency score shall not differ more than 3 % between the original and new version of SERTv2. Additionally, per configuration, the 100 % performance score of each CPU and Memory SERTv2 worklet shall not differ by more than 10 %.

9.3.4 PTDaemon

Updated versions of PTDaemon shall be tested with the following test plan and meet the stated requirements.

- 1) Select at least two different controller system OSes, two different servers which are still available for sale, two temperature sensors and two power analysers. The selected power analysers shall be manufactured by different companies, and it is preferred, but not required, that the temperature sensors are of different models and manufacturers. All selected power analysers and temperature sensors shall be selected from [Annex B](#) or have been shown to be SEEM compliant.
- 2) Select and configure (as specified in [7.1](#)) one of the two servers to run SERTv2. Select and configure one of the two controller systems with one of the selected OSs. Connect one power analyser and temperature sensor to the SUT.
- 3) Install the initial version of PTDaemon (as specified in [9.2.1](#)).
- 4) Run the SERTv2 test.
- 5) Install and configure the new version of PTDaemon.
- 6) Run the SERTv2 test.
- 7) Repeat steps 1 through 6 once more using the second server, power analyser, and temperature sensor. Also, reconfigure the controller system with a second OS.
- 8) For each hardware configuration, per SERTv2 worklet, the sum of the average-watts measured with the initial version and new version of PTDaemon shall differ by less than 5 % and the average temperature difference per worklet shall be less than 3 °C.

9.3.5 SERTv2 Client Configuration XML

9.3.5.1 General

The *SPEC SERT 2.x.x Client Configuration XML*^[11] is a file within SERTv2 which consists of the sets of CPUs, Operating Systems, and JVMs which can be used with SERTv2. An HTML version^[12] is also available. [Annex A](#) includes a list of sets of CPUs, OSes, and JVMs which are preapproved for use with SEEM.

For CPU, OS and JVM sets not listed in [Annex A](#), compliance can be obtained by following the compliance test plans of [9.3.5.4.3](#) and working with SPEC to add support in SERTv2. Creation of a new client configuration requires selection of a set of JVM optimization flags which comply with the requirements of [9.3.5.2](#).

Finally, software support for the new OS and JVM flag set is required within the SERTv2 tool. SPEC has a process for this support to be added, which is described in Section 4 of the *SPEC SERT Run and Reporting Rules* [16]. This SPEC acceptance process may include additional requirements to those described in [Clause 9](#). Once accepted, SPEC typically posts an updated client configuration HTML and XML file for use with SERTv2.

9.3.5.2 JVM flag requirements

Compliant configurations shall only use JVM flags which meet the following requirements.

- 1) Described in publicly available documentation from the JVM vendor.
- 2) Supported for use in production environments.
- 3) Does not disable RAS features that are enabled by default in the JVM.
- 4) Does not cause applications to behave in ways contrary to the Java Virtual Machine specification. For example, this requirement precludes options that relax the level of precision required by mathematical operations.
- 5) Does not enable optimizations that target the specific code sequences executed by the SERTv2 suite. Only optimizations which are expected to improve performance or reduce energy consumption of a wider class of applications than the SERTv2 suite shall be selected.

9.3.5.3 SERTv2 Invalid messages

9.3.5.3.1 General

SERTv2 includes a comprehensive error checking system which produces Invalid and Warning messages for many reasons. When completing compliance testing, the tester is expected to encounter some Invalid messages, such as those stating the configuration being tested is not included in the Client Configuration XML. This section specifies two groups of Invalid messages, and some sections of the test plan state the situations where encountering these messages is allowed. These two groups do not include all possible SERTv2 Invalid messages.

9.3.5.3.2 Group A Invalid messages

The following is the list of all the Invalid message in Group A. The symbols <> denote a field that SERTv2 fills in dynamically, such as a version number or a field name.

- 1) This result was obtained with a development build <> and may not be used for compliant results.
- 2) This result was obtained with a trial version of SERT <> and may not be used for compliant results.
- 3) The specified client configuration <> could not be found in the client configurations file.
- 4) The client configuration <> does not contain a client specification for <>.
- 5) The client count for <> in client configuration <> does not match the definition from the client configurations file.
- 6) The option set for <> in client configuration <> does not match the definition from the client configurations file.
- 7) No signature found for <>.
- 8) The <> file has been modified.
- 9) The signature for <> is invalid.
- 10) One or more TestEnvironment entries (including <>) uses a default value (beginning with "_").

- 11) The node Quantity (<>) is invalid. It must be an integer value.
- 12) The <> date is formatted incorrectly. The correct format is YYYY-MM (actual date was <>).
- 13) No power analyzer was used for this run.
- 14) The power analyzer <> is not an accepted device.
- 15) No temperature sensor was used for this run.
- 16) The temperature sensor <> is not an accepted device.

9.3.5.3.3 Group B Invalid messages

The following is the list of all the Invalid message in Group B. The symbols <> denote a field that SERTv2 fills in dynamically, such as a numeric test value.

- 1) Transactions per second was <> of target (threshold is at least <>).
- 2) Transactions per second was <> of target (threshold is at most <>).
- 3) The coefficient of variation among hosts was <>, which is greater than the threshold of <>.

9.3.5.4 CPU and software compliance test plans

9.3.5.4.1 General

Once software support for a new CPU architecture, OS, and JVM has been added to SERTv2 through the SPEC process, the applicable test plan shall be successfully completed for the configuration to be compliant with SEEM. The intent of these test plans is to verify that SERTv2 runs successfully on different configurations, and that the score changes align with expectations, such as a higher core count CPU obtaining a higher CPU Worklet performance score, or higher frequency memory obtaining a higher Memory Worklet performance score. On a given CPU architecture, when a major OS or JVM has been shown to be compliant, all other minor updates to that OS or JVM are also considered compliant without retesting.

Unless otherwise noted, only a single run of each test point is required. Where identical testing is required in multiple sections of a test plan, the test only needs to be run once. It is allowed to use normalized results when demonstrating compliance. For all test cases and all selected configurations, the JVM, OS and CPUs of the architecture being tested for compliance, shall be used and all required outcomes shall be satisfied.

For the three CPU compliance test plans of [9.3.5.4.3.4](#), [9.3.5.4.3.5](#), and [9.3.5.4.3.6](#), the tester may choose to obtain compliance for the entire CPU architecture, or only a set of CPU models or socket types within a CPU architecture, referred to as a CPU architecture portion. For instance, if a new CPU architecture supports 1, 2, and 4 socket configurations, the tester may choose to only obtain compliance for 1 and 2 socket configurations. Later, the tester may obtain compliance for 4 socket configurations with additional testing. If a tester is only obtaining support for a CPU architecture portion, this shall be specified in the testing report, and the scope of the test plan will be reduced to the portion being tested to obtain compliance.

While executing all test plans of [9.3.5.4](#), it is expected and allowed to encounter SERTv2 Warning messages as well as SERTv2 Group A Invalid messages (see [9.3.5.3.2](#)). No SERTv2 Invalid messages are allowed during testing performed for use in SEEM government regulations or programmes.

All configurations shall follow the configuration requirements of [7.1](#) as well as use balanced memory configurations, with memory evenly distributed between processors and at least one DIMM populated per memory channel. As stated in [4.2](#), SERTv2 does not apply to servers which do not support balanced and evenly distributed memory population.

The five compliance test plans include a set of compliance tests and the number of test points (SERTv2 runs) per test, as summarized in [Table 3](#). Links to each compliance test plan and compliance test are included in the [Table 3](#) titles for convenience.

Table 3 — Tests and Number of Test Points per Compliance Test Plan

		Compliance Test Plans (see 9.3.5.4.3)				
		Software minor (see 9.3.5.4.3.2)	Software (see 9.3.5.4.3.3)	CPU model (see 9.3.5.4.3.4)	CPU architecture (see 9.3.5.4.3.5)	CPU (see 9.3.5.4.3.6)
Compliance Tests (see 9.3.5.4.2)	<i>Validity & variance</i> (see 9.3.5.4.2.2)	5	5	5	5	5
	<i>CPU socket</i> (see 9.3.5.4.2.3)		2		Largest, all partial	All, incl. partial
	<i>CPU frequency</i> (see 9.3.5.4.2.4)					2
	<i>CPU core count</i> (see 9.3.5.4.2.5)				2	2
	<i>Memory size</i> (see 9.3.5.4.2.6)		2		2	2
	<i>Memory frequency</i> (see 9.3.5.4.2.7)					2
	<i>Storage quantity</i> (see 9.3.5.4.2.8)		2			2
	<i>Storage technology</i> (see 9.3.5.4.2.9)		2			2

NOTE In [Table 3](#), *partial* refers to server configurations where some but not all of the CPU sockets are populated, such as a 4P server with only 2 CPUs installed.

9.3.5.4.2 Compliance tests

9.3.5.4.2.1 General

The SEEM compliance test plans of [9.3.5.4.3](#) are composed of a set of the compliance tests defined in this subclause. The compliance test plans define which test(s) are required, the number of test points, and some configuration information, whereas the compliance tests of this subclause define the test steps and their required outcome(s).

9.3.5.4.2.2 Validity and variance test

Run a series of 5 consecutive SERTv2 runs on the configuration specified in the compliance test plan.

Required outcomes:

- At most, 1 of the 5 runs can include SERTv2 Group B Invalid message(s) (see [9.3.5.3.3](#)).
- The coefficient of variation (CV) of the performance of the highest load level of all SERTv2 worklets shall be less than 10 %.

9.3.5.4.2.3 CPU socket test

To analyse CPU socket scaling, all hardware and software configurations that affect CPU performance, other than the number of CPUs, CPU sockets, and memory size, shall remain constant, except if a

difference is required for functionality. If possible, CPUs of the same model shall be used for all tests. If CPUs of the same model are not supported in all tested CPU socket counts, CPUs with the most similar performance characteristics, such as core count, nominal core frequency, and cache size, shall be used. When possible, the system memory size divided by number of CPUs shall remain constant between all tests, and the memory shall be distributed evenly between CPUs. If the same amount of memory per CPU is not possible, the most similar amount of memory per CPU possible shall be used.

- 1) Run SERTv2 on the socket counts specified in the compliance test plan. All server CPU sockets shall be populated with CPUs. If a compliance test plan specifies testing more than one socket count, but the CPU architecture or CPU architecture portion only supports one socket count, test only one socket count.
- 2) If servers of this CPU architecture or CPU architecture portion are also sold in configurations with unpopulated CPU sockets, run SERTv2 on one example of each unpopulated CPU count per socket type. For example, if a 4-socket server is sold in configurations with only 1 CPU installed as well as configurations with only 2 CPUs installed, test one configuration with 1 CPU installed and one configuration with 2 CPUs installed.

Required outcomes:

- i) For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).
- ii) For each server with more CPUs than the least supported, the 100 % load level performance score of all SERTv2 CPU worklets shall be larger than the corresponding score from the run with fewer of the same model CPUs installed.

9.3.5.4.2.4 CPU frequency test

Run SERTv2 on the number of CPUs with different core frequencies specified in the compliance test plan. To enable CPU core frequency scaling analysis, select CPU models which have different CPU nominal frequencies, but are as similar as possible in other CPU performance relevant characteristics, such as core count and cache size. Differences in storage configurations between tests are allowed. If there are not the specified number of CPU models of this architecture or architecture portion sold at different frequencies, run with only the most different CPU frequencies possible. One of the CPU frequencies selected shall be the maximum frequency sold of this architecture or architecture portion.

Required outcomes:

- i) For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).
- ii) Each of the 100 % load level performance scores of the SERTv2 CPU worklets, divided by system threads, as calculated in [Formula \(5\)](#), for server configurations with higher frequency CPU(s) shall be higher than the corresponding score for server configurations with lower frequency CPU(s).

9.3.5.4.2.5 CPU core count test

Run SERTv2 on the number of CPUs with different core counts specified in the compliance test plan. To enable CPU core count scaling analysis, select CPU models which have different core counts, but are as similar as possible with other CPU performance relevant characteristics, such as core frequency and cache size. For comparability, it is strongly preferred that core count tests be performed using the same server and software configuration, with only the CPU changing between tests. It is allowed, however, for different servers to be used for testing, if all hardware and software components which affect CPU performance, except for the CPU model, are the same. Differences in storage configurations between test points are allowed. If there are no CPU models of this architecture, or architecture portion, sold with the specified number of different core counts, only run with the most different core counts possible. One of the selected core counts shall be the maximum core count sold of this architecture or architecture portion.

Required outcomes:

- i) For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).
- ii) Each of the 100 % load level performance scores of the SERTv2 CPU worklets for server configurations with higher core count CPU(s) shall be higher than the corresponding score for server configurations with lower core count CPU(s).

9.3.5.4.2.6 Memory size test

To enable memory size scaling analysis, select one software and hardware configuration and test the different memory sizes specified in the compliance test plan. If no platform of this CPU architecture supports the number of different memory sizes specified, test the most different memory sizes supported. The balanced memory requirement of [9.3.5.4.1](#) shall be followed for all tested memory configurations.

For the purposes of this test plan, smallest effective memory size is the largest of:

- 1) 8 GB,
- 2) SERTv2 Minimum System Memory Size (see [Formula \(4\)](#)), and
- 3) all memory channels populated with a single DIMM of the smallest size sold with the server.

Required outcomes:

- i) For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).
- ii) The SERTv2 Memory Capacity³ Max performance score of the server configuration with more memory shall be higher than the corresponding score of the server configuration with less memory.

9.3.5.4.2.7 Memory frequency test

To enable memory frequency analysis scaling, select one software and hardware configuration and run SERTv2 on the specified number of different memory frequencies. If no platform of this CPU architecture, or architecture portion, support the specified number of different memory frequencies, run with the most different memory frequencies possible.

Required outcomes:

- i) For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).
- ii) The SERTv2 Flood³ Full performance score of server configurations with higher frequency memory shall be higher than the corresponding score of server configurations with lower frequency memory.

9.3.5.4.2.8 Storage quantity test

Run SERTv2 on the specified number of different storage device quantities. All tested storage devices shall be of a similar model and the same technology (HDD, SSD, etc). It is allowed for the storage device where the OS is installed to be of a different technology than the rest of the storage devices. Changes in memory or CPU configuration between different tests are allowed if they do not significantly affect storage performance. If no platform of this CPU architecture or architecture portion supports the specified number of different storage device quantities, run with the most different storage device quantities supported. It is recommended to run with the maximum number of storage devices supported to ensure functionality across all storage controllers and connectors.

Required outcomes:

- i) For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).
- ii) The SERTv2 Storage worklets 100 % performance score of server configuration with more storage devices shall be higher than the corresponding score of the server configurations with fewer storage devices.

9.3.5.4.2.9 Storage technology test

Run SERTv2 with the specified number of storage devices of different technologies, for example HDDs and SSDs. Changes in memory or CPU configuration are allowed if they do not significantly affect storage performance. If no platform of this CPU architecture, or architecture portion, supports the specified number of different types of storage technologies, run with most different types of storage devices supported.

Required outcome:

For every test point, a SERTv2 result shall be completed with no Group B Invalid messages (see [9.3.5.3.3](#)).

9.3.5.4.3 Compliance test plans

9.3.5.4.3.1 General

This subclause defines the five compliance test plans used to demonstrate SEEM compliance for new server CPU, OS, and JVM configurations. Each compliance test plan includes a set of tests defined in [9.3.5.4.2](#). The compliance test plans of this subclause define which test(s) are required, the number of test points, and some configuration information, whereas the compliance tests define the test steps and their required outcome(s).

9.3.5.4.3.2 Software minor update compliance test plan

Minor updates to compliant major versions of OSs or JVMs are already considered compliant.

This test plan is intended for situations where changes to JVM flags or other settings are desired for a new minor update to a compliant OS or JVM.

New major versions of OSs cannot obtain compliance with this test plan. This test plan only applies when there are no changes to the JVM version or flags, no software updates to SERTv2, and the only required changes are to the Client Configuration XML file.

Compliance for a new OS or JVM version may only be obtained with this test plan if one of the following is applicable.

- 1) Adding support for a new minor JVM or OS version when the major JVM or OS version is already supported on the desired OS and CPU architecture.
- 2) Adding support for a new similar OS variant which is already supported for the same CPU architecture. A common example is a Linux distribution which has many similar components and a similar kernel version to a compliant Linux distribution.

Complete the validity and variance compliance test ([9.3.5.4.2.2](#)) with a configuration including the new software components obtaining compliance.

9.3.5.4.3.3 Software compliance test plan

Compliance for an OS and/or JVM version on a server with a compliant CPU architecture or CPU architecture portion can be shown to be compliant by successfully completing the following test plan.

If compliance for a new OS and/or JVM on a CPU architecture that has not been shown to be compliant is desired, the correct CPU compliance test plan of [9.3.5.4.3.4](#), [9.3.5.4.3.5](#), and [9.3.5.4.3.6](#) shall be completed instead of this test plan. In some situations, the reduced test plan of subclause [9.3.5.4.3](#) is applicable and can be used instead of this test plan. All test points shall include the new software components obtaining compliance.

Complete the following compliance tests using the specified test parameters:

- 1) Validity and variance test ([9.3.5.4.2.2](#))
- 2) CPU socket test ([9.3.5.4.2.3](#)) with two socket counts, one being the largest socket count which is supported by the CPU architecture or CPU architecture portion.
- 3) Memory size test ([9.3.5.4.2.6](#)) with two supported memory sizes, the smallest effective memory size (see [9.3.5.4.2.6](#)), and a different system memory size equal to, or larger than, 25 % of the max memory size supported.
- 4) Storage quantity test ([9.3.5.4.2.8](#)) with two different quantities of storage devices.
- 5) Storage technology test ([9.3.5.4.2.9](#)) with two different storage technologies.

9.3.5.4.3.4 CPU model compliance test plan

Compliance for new CPU models can be obtained with this test plan only if they are adding support for new CPU models within a compliant CPU architecture. New socket counts may be shown to be compliant using this test plan. Only JVM and OS versions already compliant with this CPU architecture may be used.

When testing new CPU models for compliance, the tester shall identify which socket counts are being tested for compliance. This selection will affect the scope of the test plan and determine which CPUs and socket counts receive an entry in the SERT Client Configuration XML file.

All test points shall include the new CPU model obtaining compliance.

Complete the Validity and variance compliance test ([9.3.5.4.2.2](#)) with a configuration that meets all of the following requirements:

- 1) The largest CPU socket count supported by the new CPU models being tested.
- 2) The largest CPU core count of the new CPU models being tested.
- 3) A memory size equal to or larger than 25 % of the maximum memory size supported by the CPU models being tested.

9.3.5.4.3.5 CPU architecture compliance test plan

Compliance for a new CPU architecture or CPU architecture portion can be obtained with this test plan, only if adding support for a new CPU architecture within a compliant CPU architecture class.

JVM and OS versions that have not yet been shown to be compliant for this CPU architecture class may be used during this testing, and the new CPU architecture, JVM, and OS will obtain compliance.

When testing a new CPU architecture for compliance, the tester shall identify if all or only a portion of the CPU models sold of that architecture are being tested for compliance. Additionally, if the tester only desires to obtain compliance for certain socket counts, this shall be specified. This selection will affect the scope of the test plan and determine which CPUs and socket counts receive an entry in the SERT Client Configuration XML file.

All test points shall include the new CPU architecture and software obtaining compliance.

Complete the following compliance tests using the specified test parameters:

- 1) Validity and variance test ([9.3.5.4.2.2](#))
- 2) CPU socket test ([9.3.5.4.2.3](#)) with the largest socket count which is desired to be made compliant for this architecture or architecture portion. If the CPU architecture supports more than one socket type, test the largest socket count per CPU socket type.
- 3) CPU core count test ([9.3.5.4.2.5](#)) with two CPU models with different core counts. One of the core counts selected shall be the maximum core count sold of this architecture, or architecture portion.
- 4) Memory size test ([9.3.5.4.2.6](#)) with two supported memory sizes, the smallest effective memory size (see [9.3.5.4.2.6](#)), and a different system memory size equal to, or larger than, 25 % of the max memory size supported.

9.3.5.4.3.6 CPU compliance test plan

Compliance for any new CPU architecture, CPU model, OS, or JVM can be obtained if the following test plan is successfully completed, even if no other CPU architectures in the same CPU architecture class have been shown to be compliant. In some situations, the reduced test plans of [9.3.5.4.3.4](#) or [9.3.5.4.3.5](#) may be applicable and can be used instead of this test plan.

When testing a new CPU architecture for compliance, the tester shall identify if the intent is to obtain compliance for all or only a portion of the CPU models sold of that architecture. Additionally, if the tester only desires to obtain compliance for certain socket counts, this shall be specified. This selection will affect the scope of the test plan and determine which CPUs and socket counts receive an entry in the SERT Client Configuration XML file.

All test points shall include the new hardware and software obtaining compliance.

Complete the following compliance tests using the specified test parameters:

- 1) Validity and variance test ([9.3.5.4.2.2](#))
- 2) CPU socket test ([9.3.5.4.2.3](#)) on all socket counts which are desired to be made compliant for this architecture or architecture portion.
- 3) CPU frequency test ([9.3.5.4.2.4](#)) with two CPUs which are sold at different nominal CPU core frequencies. One of the frequencies selected shall be the maximum frequency sold of this architecture, or architecture portion.
- 4) CPU core count test ([9.3.5.4.2.5](#)) on two CPU models with different core counts. One of the core counts selected shall be the maximum core count sold of this architecture, or architecture portion.
- 5) Memory size test ([9.3.5.4.2.6](#)) with two supported memory sizes. First, the smallest effective memory size (see [9.3.5.4.2.6](#)), and secondly, a different system memory size equal to, or larger than, 50 % of the max memory size supported.
- 6) Memory frequency test ([9.3.5.4.2.7](#)) with two different memory frequencies.
- 7) Storage quantity test ([9.3.5.4.2.8](#)) with two different quantities of storage devices.
- 8) Storage technology test ([9.3.5.4.2.9](#)) with two different storage technologies.

9.4 “Implementer-specified” metric minor update requirements

Once a SEEM regulation or programme is finalized, it is critical that the test method does not change in a way which materially affects the metric. However, some updates to the test method are needed for compatibility with new software or hardware.

For “implementer-specified” metrics, for each test method component, other than power analysers and temperature sensors, the implementer shall specify if updates are allowed. For all components for which

updates are allowed, the implementer shall specify a test plan with quantitative passing requirements, which only allows component updates which have an acceptably small effect on the metric (see [9.3](#) as an example).

NOTE Power analysers and temperature sensors are exceptions because their update requirements are specified for all SEEM implementations in [9.2](#).

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 21836:2020

Annex A (normative)

SERTv2 pre-approved CPU architectures, operating systems, and Java Virtual Machines (JVMs)

[Table A.1](#) is a list of sets of CPUs, Operating Systems, and JVMs that have been vetted for use with SERTv2 prior to the writing of this document and shall be approved for use with SEEM without the need to complete the test plans of [Clause 9](#). New configurations need to be shown to be compliant by completing the test plans of [Clause 9](#). Only CPUs, OSs and JVMs from the same set (a row in the below table) are approved for use together in SEEM. Unless otherwise noted, OSs and JVM minor updates and patches of the same version stated in the table, are also approved. Additional CPU, OS, JVM sets can only be used with SEEM if they meet all the requirements of [9.3.5](#). Detailed settings for each configuration are available on the *SPEC SERT 2.x.x Client Configuration Website*^[12]. Additionally, the settings for supported configurations are included in the SERTv2 software package and selectable in the SERTv2 GUI.

Table A.1 — Pre-approved CPU Architectures, Operating Systems, and Java Virtual Machines

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	E3-XXXX	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX	Windows Server 2012	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX	Windows Server 2012	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v2	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v2	Windows Server 2012	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v2	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX v2	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX v2	Windows Server 2012	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX v2	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v3	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v3	Windows Server 2012	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v3	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX v3	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX v3	Windows Server 2012	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E5-XXXX v3	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v4	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v5	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_1n
Intel	E3-XXXX v3	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_1
Intel	E3-XXXX v3	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_1
Intel	E5-XXXX v3	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_1
Intel	E5-XXXX v3	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_1
Intel	E3-XXXX v4	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_2
Intel	E3-XXXX v4	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_2
Intel	E5-XXXX v4	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_3
Intel	E7-XXXX v4	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_3
Intel	E5-XXXX v4	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_4

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	E7-XXXX v4	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_4
Intel	Xeon Platinum	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Platinum	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Platinum	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Platinum	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Platinum	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Platinum	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Gold	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Gold	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Gold	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Gold	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Gold	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Gold	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Silver	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Silver	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Silver	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Silver	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Silver	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Silver	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Bronze	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Bronze	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Bronze	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Bronze	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Bronze	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Bronze	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Platinum Gen2	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Platinum Gen2	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Platinum Gen2	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Platinum Gen2	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Platinum Gen2	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Platinum Gen2	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Gold Gen2	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Gold Gen2	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Gold Gen2	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Gold Gen2	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Gold Gen2	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Gold Gen2	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Silver Gen2	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Silver Gen2	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Silver Gen2	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Silver Gen2	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Silver Gen2	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Silver Gen2	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Bronze Gen2	Windows Server 2012 R2	HotSpot 1.8.0 pre121	Intel_Win_HS18_5
Intel	Xeon Bronze Gen2	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Bronze Gen2	Windows Server 2012 R2	OpenJDK 11	Intel_Win_HS18_5
Intel	Xeon Bronze Gen2	Windows Server 2016	HotSpot 1.8.0 pre121	Intel_Win_HS18_5

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	Xeon Bronze Gen2	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5
Intel	Xeon Bronze Gen2	Windows Server 2016	OpenJDK 11	Intel_Win_HS18_5
Intel	E3-XXXX v6	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	E3-XXXX v6	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	E-2XXX	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	E-2XXX	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Celeron G4XXX	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Celeron G4XXX	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Core i3-8XXX	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Core i3-8XXX	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Core i3-9XXX	Windows Server 2012 R2	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Core i3-9XXX	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5a
Intel	Pentium Gold G5XXX	Windows Server 2016	HotSpot 1.8.0 post120	Intel_Win_HS18_5b
Intel	E3-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	Intel_Win_J917_1dyn
Intel	E3-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	Intel_Win_J917_1dyn
Intel	E5-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	Intel_Win_J917_1dyn
Intel	E5-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	Intel_Win_J917_1dyn
Intel	E3-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	Intel_Win_J917_1dyn
Intel	E3-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	Intel_Win_J917_1dyn
Intel	E5-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	Intel_Win_J917_1dyn
Intel	E5-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	Intel_Win_J917_1dyn
Intel	E3-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_1dyn
Intel	E3-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_1dyn
Intel	E5-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_1dyn
Intel	E5-XXXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_3dyn
Intel	E3-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_3dyn
Intel	E5-XXXX v2	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX v3	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_3dyn
Intel	E3-XXXX v3	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX v3	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_3dyn
Intel	E5-XXXX v3	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX	Windows Server 2012	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX	Windows Server 2012	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX v2	Windows Server 2012	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX v2	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX v2	Windows Server 2012	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX v2	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX v3	Windows Server 2012	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX v3	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX v3	Windows Server 2012	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E5-XXXX v3	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E3-XXXX	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	Intel_Lin_HS17_1n

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	E3-XXXX	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E3-XXXX	Oracle Linux 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E5-XXXX	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E5-XXXX	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E5-XXXX	Oracle Linux 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E3-XXXX v2	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E3-XXXX v2	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E3-XXXX v2	Oracle Linux 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E5-XXXX v2	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E5-XXXX v2	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E5-XXXX v2	Oracle Linux 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E3-XXXX	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E5-XXXX	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E3-XXXX v2	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E5-XXXX v2	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E3-XXXX v3	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E5-XXXX v3	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E7-XXXX v3	Red Hat Enterprise Linux Server 7	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E3-XXXX v3	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E5-XXXX v3	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_2n
Intel	E5-XXXX v4	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_1
Intel	E5-XXXX v4	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_1
Intel	E5-XXXX v4	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_1
Intel	E7-XXXX v4	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_1
Intel	E7-XXXX v4	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_1
Intel	E7-XXXX v4	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_1
Intel	Xeon Platinum	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Platinum	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Platinum	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Platinum	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	Xeon Gold	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Gold	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Gold	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Gold	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	Xeon Silver	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Silver	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Silver	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Silver	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	Xeon Bronze	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Bronze	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Bronze	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Bronze	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	Xeon Platinum Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Platinum Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Platinum Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Platinum Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	Xeon Gold Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Gold Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Gold Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Gold Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	Xeon Silver Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Silver Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Silver Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Silver Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	Xeon Bronze Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 pre121	Intel_Lin_HS18_2
Intel	Xeon Bronze Gen2	Red Hat Enterprise Linux Server 7	HotSpot 1.8.0 post120	Intel_Lin_HS18_2
Intel	Xeon Bronze Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	Intel_Lin_HS18_2
Intel	Xeon Bronze Gen2	Red Hat Enterprise Linux Server 7	OpenJDK 11	Intel_Lin_HS18_2
Intel	E3-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	E3-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E3-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E3-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E3-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E3-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E3-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E3-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E3-XXXX	Oracle Linux 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E3-XXXX	Oracle Linux 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E3-XXXX	Oracle Linux 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E3-XXXX	Oracle Linux 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E5-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E5-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E5-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E5-XXXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Oracle Linux 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Oracle Linux 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Oracle Linux 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E5-XXXX	Oracle Linux 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Oracle Linux 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Oracle Linux 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	E3-XXXX v2	Oracle Linux 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Oracle Linux 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	SUSE Linux Enterprise Server 11	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Oracle Linux 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Oracle Linux 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Oracle Linux 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E5-XXXX v2	Oracle Linux 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E7-4XXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E7-4XXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E7-4XXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E7-4XXX	SUSE Linux Enterprise Server 11	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Oracle Linux 6	J9 1.7.0 SR1	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Oracle Linux 6	J9 1.7.0 SR2	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Oracle Linux 6	J9 1.7.0 SR3	Intel_Lin_J917_1dyn
Intel	E7-4XXX	Oracle Linux 6	J9 1.7.0 SR6	Intel_Lin_J917_1dyn
Intel	E3-XXXX v2	Solaris 11	HotSpot 1.7.0	Intel_Sol_HS17_1
Intel	E5-XXXX v2	Solaris 11	HotSpot 1.7.0	Intel_Sol_HS17_1
Intel	E7-XXXX v2	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_2n
Intel	E7-4XXX	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_2n
Intel	E7-XXXX v3	Windows Server 2008 R2 SP1	HotSpot 1.7.0	Intel_Win_HS17_3n
Intel	E7-XXXX v3	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_3n
Intel	E5-XXXX v4	Windows Server 2012 R2	HotSpot 1.7.0	Intel_Win_HS17_3n
Intel	E7-4XXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	Intel_Win_J917_2dyn
Intel	E7-4XXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	Intel_Win_J917_2dyn
Intel	E7-4XXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR3	Intel_Win_J917_2dyn
Intel	E7-4XXX	Windows Server 2008 R2 SP1	J9 1.7.0 SR6	Intel_Win_J917_2dyn

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
Intel	E7-XXXX v2	Windows Server 2012 R2	J9 1.7.0 SR6	Intel_Win_J917_3dyn
Intel	E7-XXXX v2	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E7-XXXX v2	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E7-XXXX v2	Oracle Linux 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E7-4XXX	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E7-4XXX	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	Intel_Lin_HS17_1n
Intel	E7-4XXX	Oracle Linux 6	HotSpot 1.7.0	Intel_Lin_HS17_1n
AMD	EPYC 7xx1	Windows Server 2012 R2	HotSpot 1.8.0 post120	AMD_EPYC_Win_HS18_1
AMD	EPYC 7xx1	Windows Server 2016	HotSpot 1.8.0 post120	AMD_EPYC_Win_HS18_1
AMD	EPYC 7xx1	Ubuntu 18.04.x	OpenJDK 11	AMD_EPYC_Lin_OJDK11_1
AMD	EPYC 7xx2	Ubuntu 18.04.x	OpenJDK 11	AMD_EPYC_Lin_OJDK11_2
AMD	EPYC 7xx2	Ubuntu 19.04.x	OpenJDK 11	AMD_EPYC_Lin_OJDK11_2
AMD	EPYC 7xx2	Windows Server 2019	OpenJDK 11	AMD_EPYC_Win_OJDK11_1
AMD	Opteron 32xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 32xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 41xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 41xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 42xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 42xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 43xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 43xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 61xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 61xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 62xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 62xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 63xx	Windows Server 2008 R2 SP1	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 63xx	Windows Server 2012 R2	HotSpot 1.7.0	AMD_Win_HS17_1
AMD	Opteron 32xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 32xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 32xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 41xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 41xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 41xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 42xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 42xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 42xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 43xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 43xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 43xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
AMD	Opteron 61xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 61xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 61xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 62xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 62xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 62xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 63xx	Red Hat Enterprise Linux Server 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 63xx	SUSE Linux Enterprise Server 11	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 63xx	Oracle Linux 6	HotSpot 1.7.0	AMD_Lin_HS17_1
AMD	Opteron 32xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 32xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 41xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 41xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 42xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 42xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 43xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 43xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 61xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 61xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 62xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 62xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 63xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR1	AMD_Win_J917_1dyn
AMD	Opteron 63xx	Windows Server 2008 R2 SP1	J9 1.7.0 SR2	AMD_Win_J917_1dyn
AMD	Opteron 32xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 32xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 41xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 41xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 42xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 42xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 43xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 43xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 61xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 61xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 62xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 62xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 63xx	Windows Server 2012 R2	J9 1.7.0 SR1	AMD_Win_J917_2dyn
AMD	Opteron 63xx	Windows Server 2012 R2	J9 1.7.0 SR2	AMD_Win_J917_2dyn
AMD	Opteron 32xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 32xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 32xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 32xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
AMD	Opteron 32xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 32xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 41xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 41xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 41xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 41xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 41xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 41xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 42xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 42xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 42xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 42xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 42xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 42xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 43xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 43xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 43xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 43xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 43xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 43xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 61xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 61xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 61xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 61xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 61xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 61xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 62xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 62xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 62xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 62xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 62xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 62xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 63xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn

Table A.1 (continued)

CPU vendor	CPU model(s)	Operating system	JVM	Configuration
AMD	Opteron 63xx	Red Hat Enterprise Linux Server 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 63xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 63xx	SUSE Linux Enterprise Server 11	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron 63xx	Oracle Linux 6	J9 1.7.0 SR1	AMD_Lin_J917_1dyn
AMD	Opteron 63xx	Oracle Linux 6	J9 1.7.0 SR2	AMD_Lin_J917_1dyn
AMD	Opteron A11xx	openSUSE Leap 42.x	OpenJDK 1.8.0	AMD_AArch64_Linux_OJDK18_1
IBM	POWER8	AIX 7.1	J9 8.0 SR1	IBM_AIX_J980_1dyn
IBM	POWER8	AIX 7.1	J9 8.0 SR2	IBM_AIX_J980_1dyn
IBM	POWER8	AIX 7.1	J9 8.0 SR3+	IBM_AIX_J980_1dyn
IBM	POWER8	Red Hat Enterprise Linux Server 7	J9 8.0 SR1	IBM_Linux_J980_1dyn
IBM	POWER8	Red Hat Enterprise Linux Server 7	J9 8.0 SR2	IBM_Linux_J980_1dyn
IBM	POWER8	Red Hat Enterprise Linux Server 7	J9 8.0 SR3+	IBM_Linux_J980_1dyn
IBM	POWER8	Ubuntu 16.04.x	J9 8.0 SR3+	IBM_Linux_J980_1dyn
IBM	POWER9	AIX 7.2	J9 8.0 SR3+	IBM_AIX_J980_2dyn
IBM	POWER9	Red Hat Enterprise Linux Server 7	OpenJDK 1.8.0	IBM_Linux_OJ980_1dyn
IBM	POWER9	Red Hat Enterprise Linux Server 7	J9 8.0 SR3+	IBM_Linux_J980_1dyn
Oracle	SPARC T4	Solaris 11	HotSpot 1.7.0	SPARCT4_Sol_HS17_1
Fujitsu	SPARC64-X	Solaris 11	HotSpot 1.7.0	SPARC64_X_Sol_HS17_1
Marvell	ThunderX2	Red Hat Enterprise Linux Server 7	OpenJDK 11	Marvell_ThunderX2_Linux_OJDK11_1