



Standard Performance Evaluation Corporation (SPEC®)

Server Efficiency Rating Tool (SERT™) Run and Reporting Rules

7001 Heritage Village Plaza, Suite 225
Gainesville, VA 20155,
USA

Table of Contents

1. Introduction	4
1.1. Overview	4
1.1.1. Applicability.....	4
1.2. Trademark and Copyright Notice	4
2. Run Rules	5
2.1. Manual Intervention.....	5
2.2. SUT Configuration Information	5
2.3. Client Configuration.....	5
2.4. Testbed Configuration	5
2.4.1. General Availability.....	5
2.4.2. System(s) Under Test (SUT).....	6
2.4.3. Memory	6
2.4.4. Network Interfaces	6
2.4.5. Storage.....	6
2.4.6. Software.....	7
2.5. Java Specifications	7
2.5.1. Tool Binaries and Recompilation.....	7
2.6. Validity Checks.....	7
2.7. Line Voltage Source	8
2.8. Environmental Conditions	9
2.9. Power and Temperature Measurements	9
2.9.1. Supported and Compliant Devices	9
2.9.2. Power Analyzer Setup.....	9
2.9.3. DC Line Voltage	9
2.9.4. Power Measurement Exclusion.....	9
3. Reporting Rules.....	10
3.1. Publication and Disclosure Requirement	10
3.1.1. Research and Academic Usage.....	10
3.2. Reproducibility	10
3.3. Testbed Configuration Disclosure	10
3.4. General Availability Dates.....	10
3.5. Disclosure Notes	10
4. Platform Support	11
4.1. Platform Enablement.....	11
4.2. Client Configuration Acceptance.....	12
5. SERT Update Process.....	12

6. Recommendation for Government Energy Efficiency Programs13

SVN Revision: 2713
SVN Date: 2019/11/28 15:30:42
This document version: <https://www.spec.org/sert2/SERT-runrules-20190815.pdf>
Latest SERT 2.x.x version: <https://www.spec.org/sert2/SERT-runrules.pdf>
Previous version: <https://www.spec.org/sert2/SERT-runrules-20190301.pdf>
Sections marked in italic are references to ISO/IEC 21836 (<https://www.iso.org/standard/71926.html>)

1. Introduction

1.1. Overview

The Server Efficiency Rating Tool (SERT™) was created by the Standard Performance Evaluation Corporation (SPEC®), the world's leading organization for benchmarking expertise. The SPECpower Committee designed, implemented, and delivered the SERT suite, a next-generation tool set for measuring and evaluating the energy efficiency of servers. The SERT suite was created with the input of various global energy-efficiency programs in order to incorporate their regional program requirements. Designed to be simple to configure and use via a comprehensive graphical user interface. The SERT suite uses a set of synthetic worklets to test discrete system components such as CPU, memory, and storage, providing detailed power consumption data at different load levels. Results are provided in both machine- and human-readable forms, enabling automatic submission to government-sponsored certification programs as well as both summary and detail reports for use by potential customers.

These rules abide by the norms laid down by SPEC in order to ensure that results generated with this tool are meaningful, comparable to other generated results, and repeatable, with documentation covering factors pertinent to reproducing the results. Per the SPEC license agreement, all results must adhere to these Run and Reporting Rules in order to be considered valid. Only valid results may be submitted for regulatory programs.

1.1.1. Applicability

SPEC intends that this tool measures the AC power and performance of computer servers for use by various energy efficiency programs.

The AC power consumption measured by this tool should not be assumed to represent the AC power consumption of other applications on the same hardware.

While this tool was designed to be a measure of computer servers, SPEC acknowledges that it may also be used to measure other classes of computing devices. Given the speed of technology advances in the industry, SPEC does not arbitrarily restrict the type of system on which the SERT suite is used. However, the SERT workload is not representative of workstation and client workloads.

1.2. Trademark and Copyright Notice

SPEC, the SPEC logo, and the name PTDaemon are registered trademarks of the Standard Performance Evaluation Corporation (SPEC) and SERT is a trademark of SPEC. Additional product and service names mentioned herein may be the trademarks of their respective owners. Copyright © 1988-2019 Standard Performance Evaluation Corporation (SPEC). All rights reserved.

2. Run Rules

2.1. Manual Intervention

No manual intervention or optimization of the Controller, SUT, or its internal and external environment is allowed during the measurement run.

2.2. SUT Configuration Information

The configuration information used to generate the final report must be populated accurately by the tester.

2.3. Client Configuration

Only Client Configurations that have been accepted by SPEC and are documented at https://www.spec.org/sert2/SERT-JVM_Options-2.0.html are allowed for valid measurements.

SERT will automatically pick the correct JVM options for each workload from the specified set and start the client JVMs with these options. The option sets for all supported platform / OS / JVM combinations have been defined by the platform manufacturer and must be used for compliant SERT runs. If the CPU, Operating System, or Java Virtual Machine are not detected properly, the drop-down boxes can be used to select the correct values; however, selecting the incorrect values would result in a non-compliant run. If the values are all set correctly but no matching configuration is found, this platform is not currently supported. Please check http://www.spec.org/sert2/SERT-JVM_Options-2.0.html for a list of supported platforms. You may be able to use a different JVM or Operating System instead. If the processor is not on the supported list, you may wish to contact your processor vendor to encourage them to complete the platform acceptance process so it will be supported in SERT.

JVMs are capable of modifying their behavior based on flags. Configuring the JVM through other means is strictly forbidden.

2.4. Testbed Configuration

The SERT suite is developed for out-of-the-box configurations. Limited tuning has been implemented to ensure consistency across architectures.

- Any deviations from the standard default configuration for the testbed components must be documented so an independent party would be able to reproduce the configuration and the result without any further assistance.
- The JVM director and all instances of the SPEC PTDaemon® must run on system(s) other than the SUT.

2.4.1. General Availability

Availability requirements are defined by the regulatory agency in charge of the specific energy-efficiency program that is using data from the SERT suite.

The entire SUT must be comprised of components that are generally available on or before the date of publication, or must be generally available within three months of the first publication of these results.

World-wide availability is not a requirement. However, the entire SUT must be available in a region that supports the current and voltage sources used in the measurement of the SERT suite.

Products are considered generally available if they are orderable by ordinary customers and ship within a reasonable time frame. This time frame is a function of the product size and classification and common practice. Some limited quantity of the product must have shipped on or before the close of the stated availability window. Shipped products do not have to match the tested configuration in terms of CPU count, memory size, and disk count or size, but the tested configuration must be available to ordinary customers. The availability of support and

documentation of the products coincide with the release of the products. Documentation must be provided in a language that is appropriate for the region where the SUT is claimed to be available.

Hardware products that are still supported by their original or primary vendor may be used if their original general availability date was within the last five years. The five-year limit is waived for hardware used in controller systems.

Software products that are still supported by their original or primary vendor may be used if their original general availability date was within the last three years.

2.4.2. System(s) Under Test (SUT)

Each SERT result submitted or made public must be for an actual run of the tool on the SUT named in the result. For multi node configurations, all nodes must be equipped with identical hardware and software components, i.e. only homogeneous configurations are allowed.

2.4.3. Memory

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 Equation 4). If this server configuration does not meet these requirements, then SERTv2 was not designed to function on this server and this server configuration cannot be tested in this configuration.

$$SERTv2 \text{ Minimum System Memory Size} = 2GB + 496MB * \text{System Threads}$$

(Equation 4)

where System Threads is calculated in Equation 5

[SOURCE: ISO/IEC 21836:2020 7.8.2 Single Configuration Testing 2)]

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

(Equation 5)

[SOURCE: ISO/IEC 21836:2020 7.8.3.2 High-End Configuration]

2.4.4. Network Interfaces

At least one port of the SUT's fastest network interface controller must be connected to the Controller

2.4.5. Storage

All the following storage device requirements shall be followed:

- For a single storage device configuration, a minimum of two partitions are required
 - OS partition(s) and one SERT Data partition (min. 20 GB)
- For a multiple storage device configuration, drives must be partitioned as follows:
 - First storage device: OS partition(s) and one SERT Data partition (min. 20 GB)
 - Any additional storage device: SERT Data drive (entire drive capacity)
- For all storage devices, the device write cache must be disabled and the "Write Through" mode on RAID controllers must be set if possible; see the SERT User Guide for platform specific instructions. For systems where disabling write caches or setting "Write Through" mode is not supported, a statement must be added to the SUT notes in the FDR.
- SERT Data partitions may only be hosted on filesystems (e.g., ntfs, ext4, zfs, jfs) that support direct I/O. Note: USB Hard drives that don't support direct I/O will not work with the SERT suite.

- RAID arrays of more than one device are not allowed, i.e. each individual storage device must be configured as JBOD or a single disk RAID 0 array.

Note: The content of the storage devices being tested will not be erased. However, for meaningful results, it is strongly recommended to format the filesystem before running the SERT storage worklets for the first time. The test data files created during the initial run can be reused in subsequent runs.

2.4.6. Software

Required software components per server (host) are:

- Exactly one operating system instance (with one user space), including all modules that are installed during the default installation process.
- If the system uses a hypervisor by default, the SERT suite may run in a guest instance (virtual machine) of that hypervisor. Only one guest may run on each host. The guest must be configured to use all hardware resources that are not reserved by the hypervisor.
- A 64-bit Java run time environment including one or more instances of a Java Virtual Machine (JVM).

Optional power management software, when installed by default, must be reported. The operating system must be in a state sufficient to execute a class of server applications larger than the tool alone. Disabling operating system services is forbidden.

Transactions per second was xx% of target

For end user testing, the following values in the config-all.xml file may be set to any integer between 1 and 10 to provide additional warmup time to allow systems to reach steady state before measurements start. For platform acceptance testing, any integer between 1 and 5 may be set. This should help in cases where runs are invalid due to 100% throughput being too high or too low, e.g. INVALID: Sequential worklet, interval 100%: Transactions per second was 104.7% of target (threshold is at most 104.0%).

- `<interval-count id="cpuWarmupIntervalCount">3</interval-count>`
- `<interval-count id="memCapacityWarmupIntervalCount">1</interval-count>`
- `<interval-count id="storageWarmupIntervalCount">1</interval-count>`
- `<interval-count id="hybridWarmupIntervalCount">3</interval-count>`

2.5. Java Specifications

The SERT 2 suite requires Java 7 or later and its implementation must follow the following specification:

- Java Virtual Machine Specification (ISBN-13: 978-0133260441)

2.5.1. Tool Binaries and Recompile

The SERT suite binaries are provided in jar files containing the Java classes. Valid runs must use the SPEC provided jar files and these files must not be updated or modified in any way.

2.6. Validity Checks

At the end of the run, the below checks are performed on the results file. The run will be marked invalid if the configuration or results do not adhere to the below requirements. "INVALID" messages will appear in the report to explain the reasons why the result was marked invalid. In addition, warning messages may be included in the report to alert the reader about possible issues. Note, these validation checks are not a substitute for a thorough review of SERT run rule compliance as it is not possible to automatically verify compliance with all SERT run rules.

The following are required for a valid run and are automatically checked:

ResultModificationValidator	Verifies that the results.xml file was not modified after the run, aside from in the section of the file that is explicitly marked where editorial changes are allowed
CompleteRunValidator	Ensures that all of the configured measurement intervals were actually included in the result
TransactionFailuresValidator	Confirms that all transactions completed successfully
TargetThroughputValidator	Verifies that the target throughput for each load level was achieved (within acceptable thresholds)
ThroughputVariationValidator	Ensures that each host in multi-node runs has similar performance
VoltageValidator	Verifies that the measured voltage is within 5% of the line standard.
PowerAnalyzerDeviceValidator	Confirms that power analyzer are configured properly and on the accepted device list
PowerErrorReadingsValidator	Validates that there are no more than 1% error readings of power data and no more than 2% error readings for power-factor, voltage, and current
PowerUncertaintyReadingsValidator	Confirms that power analyzers report no more than 5% of readings with uncertainty greater than 1%, and no more than 1% of readings that have unknown uncertainty
TemperatureSensorDeviceValidator	Confirms that temperature sensors are configured properly and on the accepted device list
MinimumTemperatureValidator	Verifies that the minimum reported temperature is at least 20.0 degrees C throughout the run

2.7. Line Voltage Source

The preferred line voltage source used for measurements is the main AC power as provided by local utility companies. Power generated from other sources often has unwanted harmonics which are incapable of being measured correctly by many power analyzers, and thus, would generate inaccurate results.

The AC Line Voltage Source needs to meet the following characteristics:

AC Frequency: $\pm 1\%$ of 50 Hz or 60 Hz

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

[SOURCE: ISO/IEC 21836:2020 7.4 Power Requirements]

The usage of an uninterruptible power source (UPS) as the line voltage source is allowed, but the voltage output must be a pure sine-wave. For placement of the UPS, see 2.9.2. This usage must be specified in the Notes section of the FDR.

Systems that are designed to be able to run normal operations without an external source of power cannot be used to produce valid results. Some examples of disallowed systems are notebook computers, hand-held computers/communication devices, and servers that are designed to frequently operate on integrated batteries without external power.

Systems with batteries intended to preserve operations during a temporary lapse of external power, or to maintain data integrity during an orderly shutdown when power is lost, can be used to produce valid tool results. For SUT components that have an integrated battery, the battery must be fully charged at the end of each of the measurement intervals, and proof must be provided that it is charged at least to the level of charge at the beginning of the interval.

Note that integrated batteries that are intended to maintain such things as durable cache in a storage controller can be assumed to remain fully charged. The above paragraph is intended to address “system” batteries that can provide primary power for the SUT.

If an unlisted AC line voltage source is used, a reference to the standard must be provided to SPEC. DC line voltage sources are currently not supported.

For situations in which the appropriate voltages are not provided by local utility companies (e.g., measuring a server in the United States which is configured for European markets, or measuring a server in a location where the local utility line voltage does not meet the required characteristics), an AC power source may be used, and the power source must be specified in the Notes section of the disclosure report. In such situations, the following requirements must be met, and the relevant measurements or power source specifications disclosed in the general notes section of the disclosure report:

- Total Harmonic Distortion of source voltage (loaded), based on IEC standards: < 5%
- The AC power source needs to meet the frequency and voltage characteristics previously listed in this section.
- The AC power source must not manipulate its output in a way that would alter the power measurements compared to a measurement made using a compliant line voltage source without the power source.

The intent is that the AC power source does not interfere with measurements such as power factor by trying to adjust its output power to improve the power factor of the load.

2.8. Environmental Conditions

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 above 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 may artificially impact energy consumption or performance of the server, before and during the measurement.

[SOURCE: ISO/IEC 21836:2020 7.2 Environment]

2.9. Power and Temperature Measurements

2.9.1. Supported and Compliant Devices

For SERT, only the power analyzers and temperature sensors compliant with ISO/IEC 21836 sub clause 9.2.1 shall be used. The list of current ISO complaint devices can be found here: https://www.spec.org/power/docs/SPECpower-Device_List.html.

2.9.2. Power Analyzer Setup

The power analyzer must be located between the AC line voltage source (or UPS) and the SUT. No other active components are allowed between the AC line voltage source and the SUT. Power analyzer configuration settings that are set by the SPEC PTDaemon must not be manually overridden. The power analyzer must be calibrated within the past 12 months for a compliant SERT result.

2.9.3. DC Line Voltage

The SERT suite is neither supported nor tested with DC loads today.

2.9.4. Power Measurement Exclusion

Switches (network or KVM) can be excluded from the power measurement of the SUT.

3. Reporting Rules

In order to publicly disclose the SERT results, the tester must adhere to these reporting rules in addition to having followed the Run Rules above. The goal of the reporting rules is to ensure the system under test is sufficiently documented so that someone could reproduce the test and its results and to ensure that the tester has complied with the Run Rules.

3.1. Publication and Disclosure Requirement

Any entity choosing to make statements using the SERT suite must follow the SPEC Fair Use Rule. Fair Use: Consistency and fairness are guiding principles for SPEC. To help assure that these principles are met, any organization or individual who makes public use of SPEC tool results must do so in accordance with the SPEC Fair Use Rule, posted at <https://www.spec.org/fairuse.html>.

Any entity publically disclosing information associated with a SERT result must be willing to share the results.xml file with SPEC for evaluation purposes. Any SPEC member may request that a full disclosure of the result be forwarded to SPEC. The disclosure is expected to be complete, including the configuration description. The full disclosure must be made available to the relevant subcommittee no later than 10 working days after the request.

SPEC may take action: When public claims are made about SPEC results, by any party (including but not limited to members, licensees, the press, academic researchers), SPEC reserves the right to take action if the full disclosure is not made available, or shows different results than the tester's claim, has other rule violations, or violates SPEC policies (such as the Fair Use policy).

3.1.1. Research and Academic Usage

SPEC encourages the use of the SERT suite by the Research and Academic Community; please consult the SPEC Fair Use Rule for Research and Academic Usage at <https://www.spec.org/fairuse.html#Academic>.

3.2. Reproducibility

By submitting or publishing a tool disclosure (report), the test sponsor implicitly states that the system performance and power measured is representative of such systems.

3.3. Testbed Configuration Disclosure

The system configuration information that is required to reproduce published power and performance results must be reported. The principle is that if anything affects power or performance or is required to duplicate the results, it must be described. Any deviations from the standard, default configuration for the SUT must be documented so an independent party would be able to reproduce the result without any further assistance.

3.4. General Availability Dates

The dates of general customer availability must be listed for the major hardware components (including system firmware) and server software, by month and year.

3.5. Disclosure Notes

The Notes section is used to document additional important information required to reproduce the results from other reporting sections that require a larger text area.

4. Platform Support

The SERT suite has two levels of platform support:

Enablement: Enablement indicates that a processor architecture, operating system, or JVM has the necessary functional support implemented in the SERT suite, and has passed functional testing to ensure that SERT runs correctly.

Acceptance: Accepted “client configurations” are required to produce compliant results. A client configuration applies to a specific combination of processor architecture, operating system, and JVM, and specifies the runtime parameters that will be used with this configuration.

The sections below define the process for extending enablement and acceptance to new platforms.

4.1. Platform Enablement

Platform enablement primarily relates to functional support for a processor architecture, operating system, or JVM. Enablement is not sufficient for producing compliant runs, but is a prerequisite for accepting new platform configurations.

Enabling support for a new processor architecture, operating system, or JVM involves implementing any platform-specific code that is necessary for the SERT suite to run properly in that environment, and then producing test results to confirm that the SERT suite functions properly.

Specific aspects of the SERT suite that may require platform-specific code include:

- Affinity
- Discovery
- Storage worklet implementation
- Run scripts (serthost.sh / serthost.bat)
- Documentation updates (particularly the User's Guide); this includes any necessary steps for configuring large pages (if supported) and for deactivating write cache for storage devices

Implementing these functions requires SERT source code access; therefore, this work must be done by a SPEC member or supporting contributor. For further information please contact info@spec.org.

The SERT suite testing must be completed to confirm that the changes are working properly. The extent of testing required depends on the extent of code changes required. For example, if the only necessary changes are to the Discovery code, testing would focus on the accuracy of the discovery data for the new platform as well as regression tests on existing systems that share the same discovery scripts.

Testing should also demonstrate that the new platform is capable of achieving compliant SERT results on one or more typical system configurations using the new platform support.

In some cases, adding enablement for a new platform may not require code changes, but just updates to documentation and testing.

Finally, in order for SPEC to officially add enablement support for a new platform, the submitting party must provide a commitment for future support and testing. Support for a particular platform may be withdrawn in a future update of the SERT suite if no party is able to provide adequate testing to ensure that the new update still runs properly on that platform.

4.2. Client Configuration Acceptance

ISO/IEC 21836 9.3.4 defines the process and should be references for acceptance of new client configurations. The introduction to this clause is as follows.

The SERT Client Configuration XML is a file within SERTv2 which consists of a list of sets of CPUs, Operating Systems, and JVMs which can be used with SERTv2. The HTML (posted on SPEC's website) and XML (version used for testing) of the SPEC SERT 2.x.x Client Configuration file are linked to in the Bibliography. 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 9.3.4.4.3 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 sub clause 9.3.4.2.

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. This SPEC acceptance process may include additional requirements to those described in these sub clauses. Once accepted, SPEC typically posts an updated client configuration HTML and XML file for use with SERTv2.

[SOURCE: ISO/IEC 21836:2020 9.3.4.1 General]

5. SERT Update Process

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 OSes 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 (v2.0.1).*
- 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 %.*

[SOURCE: ISO 21836:2020 9.3.2 SERTv2, adjusted "(specified in 2)" to v2.0.1]

6. Recommendation for Government Energy Efficiency Programs

SPEC applauds the global efforts of driving towards greater server efficiency by government program adding a means to measure the overall server. It is the SPECpower Committee's intent to work with its member companies and regulatory bodies to insure that SERT versions referenced in or used for regulatory programs have the necessary support. It will be important to undertake periodic updates in the regulatory requirements as the SERT suite is improved and updated to test the broader server capabilities as new technology generations are introduced.

The SERT suite must be used in the intended manner to ensure accurate and repeatable results; therefore, we recommend the inclusion of the following items in every Server Efficiency Regulatory Program:

- **Measurement** - The SERT suite kit must be used to run and produce measured SERT results. Valid runs must use the provided binary files and these files must not be updated or modified in any unsupported way.
- **Manual Intervention** - No manual intervention or optimization for the SUT or its internal and external environment is allowed during the SERT measurement.
- **General Availability** - The implementation of the system under test must be generally available, documented, and supported as of the HW and SW availability date stated in the result file in order to ensure that the systems and their hardware and software components actually represent real products that solve real business and computational problems.
- **Result File Collection** - It is strongly suggested that Server Efficiency Regulatory Programs are collecting the results.xml file of each SERT result.
- **Power management settings** for the SUT must be the default settings that are shipped with the SUT.
- DC line voltage sources are currently not supported, nonetheless if DC servers are included in a government program, **separate categories for alternating current (AC) AC and Direct Current (DC) power servers** need to be defined, since the AC-DC conversion losses are not included in DC-powered server.
- **Competitive comparisons** that promote the use of one product over another and use numeric data generated by the SERT are expressly disallowed.

The validation checks described in Section 2.6 are not a substitute for a thorough review of SERT run rule compliance as it is not possible to automatically verify compliance with all SERT run rules. SPEC recommends that Government Energy Efficiency Programs verify the following items before accepting any SERT 2 results:

- **Verify** that the client configuration supports the used OS, JVM, and CPU architecture.
- **Verify** that the Hardware and Software configuration is filled out correctly and no default values (e.g., "ReallyFast Java Accelerator Card", "Open Doors 2010 F500 Operating Instruction", "Controller cache disabled in EasySetup") were left in the full disclosure report.
- **Verify** that the full disclosure report and all fields as described in <https://www.spec.org/sert2/SERT-resultfilefields.html> are completed.