Standard Performance Evaluation Corporation (SPEC)

Server Efficiency Rating Tool (SERT) Run and Reporting Rules 1.1.1

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

SPECpower Committee

Table of Contents

SVN Revision: 2001

SVN Date: 2016/01/13 02:28:22

1. Introduction

1.1. Summary

The Server Efficiency Rating Tool (SERT)™ was created by Standard Performance Evaluation Corporation (SPEC), the world's leading organization for benchmarking expertise, at the request of the US Environmental Protection Agency (EPA). It is intended to measure server energy efficiency, initially as part of the second generation of the US EPA ENERGY STAR for Computer Servers program. SPEC anticipates usage by other regulatory programs world-wide. Designed to be simple to configure and use via a comprehensive graphical user interface, the SERT uses a set of synthetic worklets to test discrete system components such as 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 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. These rules describe what it takes to create valid SERT scores, not necessarily what the regulatory agency requires to comply with their program. To check for possible updates to the Run and Reporting Rules, please see http://www.spec.org/sert/docs/SERT-Run_and_Reporting_rules.pdf.

1.2. Philosophy

SPEC believes the user community will benefit from an objective series of tool results, which can serve as a common reference and be considered as part of an evaluation process. SPEC expects that any public use of results from this tool suite must be for server configurations that are commercially available.

For results to be valid, SPEC requires:

- Proper use of the SERT as provided.
- Availability of an appropriate Full Disclosure Report (FDR).
- Availability of the hardware and software used.

1.2.1. 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 is used. However, the SERT workload is not representative of workstation and client workloads.

1.3. Optimizations

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

SPEC is aware that it is sometimes difficult to draw an exact line between legitimate optimizations and optimizations that specifically target the SERT. However, with the rules below, SPEC wants to increase the awareness of implementers and end users of issues of unwanted tool-specific optimizations:

- Software used must generate correct code.

- Hardware and/or software optimizations must improve AC power consumption and/or performance for a class of programs, where the class of programs must be larger than the SERT.
- The vendor encourages the implementation for general use.
- The implementation is generally available, documented, and supported by the providing vendor(s).

1.4. Caveats

SPEC reserves the right to investigate any case where it appears that these guidelines and the associated Run and Reporting Rules have not been followed. SPEC may request that the claim be withdrawn from any public forum in which it appears and that the tester correct any deficiency in the product or process before submitting or publishing future results.

SPEC reserves the right to adapt the SERT, workloads/worklets, and rules of the SERT as deemed necessary to preserve the goal of fair measurements. SPEC will notify members and licensees whenever it makes changes to the tool and may rename the metrics.

1.5. Research and Academic Usage

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

1.6. Trademark

SPEC and the names SERT, SPECpower_ssj, and SPEC PTDaemon are trademarks of the Standard Performance Evaluation Corporation. Additional product and service names mentioned herein may be the trademarks of their respective owners.

1.7. Copyright Notice

2. Run Rules

2.1. Measurement

The SERT must be used to run as defined in these procedures to produce measured valid results. The SERT scores are a function of the workload (see section 2.3), and the defined control parameters (see section 2.6). The SERT results are not comparable to power and performance metrics from any other application.

2.2. Initializing and Running

For guidance, please consult the latest User Guide (http://www.spec.org/sert/docs/SERT-User_Guide.pdf) and Measurement Setup Guide (http://www.spec.org/power/docs/SPECpower-Measurement_Setup_Guide.pdf).

2.3. Workload

The SERT exercises multiple worklets in a Java framework. A detailed description can be found in the latest version of the design document on SPEC's website (http://www.spec.org/sert/docs/SERT-Design_Document.pdf).

2.4. Manual Intervention

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

2.5. SUT Configuration Information

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

2.6. Tool Control Parameters

The "config-all.xml" file is the default file which includes the control parameters for the SERT; most of them cannot be changed without causing the result to be marked as invalid.

2.7. Validity Checks

The run output file is checked for conformance to the Run and Reporting Rules. Validity messages are displayed in the report for different levels of potential issues. Some messages mark the results non-compliant; others indicate correctable issues or are informational. Regardless, the tool will run to completion producing a report that may not be valid for publication. Automatic validity checks are not exhaustive. Compliance with all other rules in this document is still required.

2.8. Client Configuration

JVMs are capable of modifying their behavior based on flags. Only Client Configurations (JVM options and JVM count) that are documented for the measurement environment http://www.spec.org/sert/docs/SERT-JVM_Options-110.html are allowed for valid measurements. Influencing the behavior of the JVM through other means is strictly forbidden.

2.9. Testbed Configuration

These requirements apply to all hardware and software components used in producing a SERT result, including the SUT, network, and Controller.

- 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 Controller must be a physically different system than the SUT.
- The JVM director and all instances of the SPEC PTDaemon must run on the Controller.

2.9.1. 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.

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.

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.

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.

Information must be provided in the disclosure to identify any component that is no longer orderable by ordinary customers. See http://www.spec.org/osg/policy.html#AppendixC – OSG Policy / Appendix C - Guidelines for General Availability.

2.9.2. General Availability for Historical Systems

Please see OSG Policy section 2.3.5 on SUT Availability for Historical Systems
http://www.spec.org/osg/policy.html#s2.3.5.

2.9.3. System(s) Under Test (SUT)

The SUT may be a single stand-alone server or a multi-node set of servers as described below in the following sections.

Only identical servers are allowed in a multi-node SUT configuration; each must be identically configured. This requirement is for servers that execute the workload of the tool, and it does not include components that support these servers, e.g., storage-blades, controllers, and shared appliances, which must be included in the power measurement.

All installed server-nodes must run the SERT, e.g., a multi-node SUT with eight installed servers must run the tool code on all eight nodes.

All software required to run the SERT must be installed on and executed from a nonvolatile storage device which is considered part of the SUT.

Storage external to the enclosure is allowed only if no other means of storage is available, e.g., server internal storage, storage blade, or enclosure storage. The power consumption of this external storage must be measured as part of the SUT.

2.9.4. Electrical Equivalence

Each SERT result submitted or made public must be for an actual run of the tool on the SUT named in the result. Electrically equivalent results are not allowed.

2.9.5. Hardware

Any hardware configuration of one or more systems and supporting components that is sufficient to install, start, and run the tool to completion in compliance with these run rules must be considered a compliant configuration.

External devices required for initial setup or maintenance of the SUT, but not required for normal operation or for running the tool (e.g., an external optical drive used for OS installation), may be removed prior to the tool being started.

If the model name or product number implies a specific hardware configuration, these specific components cannot be removed from the hardware configuration but may be upgraded. Any upgrades are subject to the support, availability, and reporting requirements of this document. For example, if the SUT is available from the vendor only with dual power supplies, both supplies must be installed and measured during the tool run. The power supplies may be upgraded if the vendor offers and supports such an upgrade, and the upgrade must be documented in the tool disclosure report.

A video monitor, if configured, may be powered by a separate power source and need not be included in the power measurement of the SUT.

The components are required to be:

- specified using customer-recognizable names,
- documented and supported by the providing vendor, and
- of production quality.

Any tuning or deviation from the default installation or configuration of hardware components, which is not expressly allowed in this document, is disallowed. This includes BIOS settings, power saving options in the system board management, or upgrade of default components.

2.9.6. Network Interfaces

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

2.9.7. Storage

If available, write caches in the storage subsystems (e.g., disk and RAID controllers) must be disabled or set to write-through mode. For systems where disabling either of the caches is not supported, a statement must be added to the SUT notes in the FDR.

2.9.8. 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 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. Full support for virtualized environments is considered for the future.
- A Java run time environment including one or more instances of a Java Virtual Machine (JVM).
- A 64-bit JVM is required on 64-bit platforms, and a 32-bit JVM is required on 32-bit platforms .

Optional power management software, when installed, 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

2.10. Java Specifications

The tested SUT must be run with one of the JVMs specified in http://www.spec.org/sert/docs/SERT-JVM_Options-110.html that adhere to the Java SE 7 (or higher) specification. Any tested system must include an implementation of the Java Virtual Machine as described by the following references, or as amended by SPEC for later Java versions:

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

2.10.1. Tool Binaries and Recompilation

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

2.11. 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:

- Frequency: (50Hz or 60Hz) ± 1%
- Voltage in rms: (100V, 110V, 120V, 200V, 208V, 220V, 230V, 240V or 400V) ± 5%

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 2.13.5. 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.

2.12. Environmental Conditions

Power measurements need to be taken in an environment representative of the majority of usage environments. The intent is to discourage extreme environments that may artificially impact power consumption or performance of the server, before and during the SERT run.
The following environmental conditions need to be met:

- Ambient temperature lower limit: 20°C
- Ambient temperature upper limit: within documented operating specification of the SUT
- Elevation: within documented operating specification of the SUT
- Humidity: within documented operating specification of the SUT
- Overtly directing air flow in the vicinity of the measured equipment in a way that would be inconsistent with normal data center practices is not allowed.

2.13. Power and Temperature Measurements

The SERT provides the ability to automatically gather measurement data from accepted power analyzers and temperature sensors and integrate that data into the SERT scores. It will be required that the analyzers and sensors must be supported by the measurement framework, and must be compliant with the specifications in this section.

2.13.1. Power Analyzer Requirements

To ensure comparability and repeatability of power measurements, the following attributes for the power measurement device are required for the SERT. Please note that a power analyzer may meet these requirements

when used in some power ranges, but not in others, due to the dynamic nature of power analyzer Accuracy and Crest Factor. The usage of power analyzer's auto-ranging function is not permitted.

- **Measurements** - The analyzer must report true RMS power (watts) and at least two of the following measurement units: voltage, amperes, and power factor.
- **Accuracy** - Measurements must be reported by the analyzer with an overall uncertainty of 1% or better for the ranges measured during the benchmark run. Overall uncertainty means the sum of all specified analyzer uncertainties for the measurements made during the benchmark run.
- **Calibration** - The analyzer must be able to be calibrated by a standard traceable to NIST (U.S.A.) (http://nist.gov) or a counterpart national metrology institute in other countries. The analyzer must have been calibrated within the past year.
- **Crest Factor** - The analyzer must provide a current crest factor of a minimum value of 3. For analyzers which do not specify the crest factor, the analyzer must be capable of measuring an amperage spike of at least three times the maximum amperage measured during any one-second sample of the benchmark run.
- **Logging** - The analyzer must have an interface that allows its measurements to be read by the SPEC PTDaemon. The reading rate supported by the analyzer must be at least one set of measurements per second, where set is defined as watts, and at least two of the following readings: voltage, amperes, and power factor. The data averaging interval of the analyzer must be either one (preferred) or two times the reading interval. "Data averaging interval" is defined as the time period over which all samples captured by the high-speed sampling electronics of the analyzer are averaged to provide the measurement set.

Examples:
An analyzer with a vendor-specified accuracy of +/- 0.5% of reading +/- 4 digits, used in a test with a maximum power value of 200W, would have an "overall" accuracy of (((0.5%*200W)+0.4W)=1.4W/200W) or 0.7% at 200W.

An analyzer with a wattage range 20-400W, with a vendor-specified accuracy of +/- 0.25% of range +/- 4 digits, used in a test with a maximum power value of 200W, would have an "overall" accuracy of (((0.25%*400W)+0.4W)=1.4W/200W) or 0.7% at 200W.

2.13.2. Temperature Sensor Specifications

Temperature must be measured no more than 50mm in front of (upwind of) the main airflow inlet of the SUT. To ensure comparability and repeatability of temperature measurements, SPEC requires the following attributes for the temperature measurement device used during the SERT run:

- Logging - The sensor must have an interface that allows its measurements to be read by the SERT harness. The reading rate supported by the sensor must be at least four samples per minute.
- Accuracy - Measurements must be reported by the sensor with an overall accuracy of +/- 0.5 degrees Celsius or better for the ranges measured during the SERT run.

2.13.3. SPEC PTDaemon

SPEC PTDaemon (also known as power/temperature daemon, PTD, or ptd) is used by the SERT to offload the work of controlling a power analyzer or temperature sensor during measurement intervals to a system other than the SUT. It hides the details of different power analyzer interface protocols and behaviors from the SERT software, presenting a common TCP-IP-based interface that can be readily integrated into different benchmark harnesses.

2.13.4. Supported and Compliant Devices

The SERT will utilize SPEC's accepted measurement devices list and SPEC PTDaemon update process. See Device List at http://www.spec.org/power/docs/SPECpower-Device_List.html for a list of currently supported (by the SPEC PTDaemon) and compliant (in specifications) power analyzers and temperature sensors.

The process to add software support for a power analyzer or temperature sensor to the infrastructure can be found on the Power Analyzer Acceptance Process page at http://www.spec.org/power/docs/SPEC-Power_Analyzer_Acceptance_Process.pdf.

2.13.5. Power Analyzer Setup

The power analyzer must be located between the AC line voltage source 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.

2.13.6. DC Line Voltage

SPEC PTDaemon is neither supported nor tested with DC loads today and currently no resources are devoted to including this support. The SPECpower Committee is in favor of including DC support if new resources from companies whose focus is DC computing become available to the SPECpower Committee to address the development and support opportunities.

Additionally, directly comparing servers powered by AC against servers powered by DC is not appropriate, since the AC-DC conversion losses are not included in DC-powered server. Therefore, the SPECpower Committee recommends creating a separate category for DC-powered servers.

2.13.7. Power Measurement Exclusion

Switches (network or KVM) need not be included in the power measurement of the SUT for a multi-node configuration.

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. Reporting Metric and Result

While the SERT is not intended to be a benchmark, nevertheless as a rating tool, it must produce a metric or score(s) indicative of the efficiency of the SUT.

Since different architectures perform differently on different workloads, the SERT is composed of several discreet worklets to further ensure architecture neutrality. Each worklet will produce a measure representing the performance and power consumption achieved by the SUT.

In the 1.x.y release of the SERT, individual worklet measurements will be normalized relative to an arbitrarily selected baseline. No summation of metrics across multiple worklets is planned. The complexity of performance and power measures across components at multiple target load levels makes creation of a metric difficult. Therefore, it is recommended to implement a ~9-12 month reporting-only phase first. Once a sufficient level of data is collected from report-only submissions, SPEC plans to recommend a data-driven metric and scoring algorithm. The results that are produced by the SERT are separate from the rating of the different energy-efficiency regulatory programs. The SUT might be placed in different categories of the energy-efficiency regulatory programs.

3.1.1. Reporting and Output Files

The SERT produces multiple reports, and includes code that will ensure the authenticity of the reports.

In order to reduce the effort of displaying and/or storing the desired information, the primary report is generated in the XML format. This report contains all the information regarding the SERT run, including all hardware and software configuration information of the Controller, the SUT, and the SERT workloads. It also includes all pertinent information about the worklets such as JVM affinity, options, and other launch settings, along with the resulting performance, power, and efficiency results.

In addition, multiple human readable versions of the report are also generated, two HTML and two plain text. These reports are designed to be viewed on a screen with minimum width of 1280 pixels and contain a subset of data that is derived from the primary XML report. The names and brief descriptions of the reports are given below:

- The first HTML report, named "results.html", contains all pertinent hardware and software information about the test environment, including the SUT, and the resulting aggregate performance, power consumed, and efficiency score achieved by the different worklets.
- The second report, named "results.txt" contains the same information as the results.html file above, only it is in a plain text format.
- The third report is in HTML format as well, and is named "results-details.html" and contains all the information present in the results.html file above, along with a detailed breakdown of the performance and power at each load level for all the worklets. In addition, for each worklet, the run-time parameters such as number of CPUs, threads, and command line options are captured.
- The fourth report, named "results-details.txt" contains the same information as the results-details.html file above, only it is in a plain text format.

3.1.2. Configuration Power/Performance Modifier

These are "substitutions" for real measurements for items the SERT cannot measure or for which the performance cannot be determined (e.g., redundant power supplies) and need to be created during the *Metric and Level Proposal* phase. The design allows for modifiers; nevertheless, SERT 1.1.0 does not implement this feature yet. The intention is to build on data collected from the first version of the SERT to create modifier proposals for a future revision.

3.2. Publication

Any entity choosing to make statements using the SERT 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 http://www.spec.org/fairuse.html.

3.3. Disclosure Requirement

Any entity publically disclosing information associated with a SERT result must be willing to share the results.xml file with SPEC for evaluation purposes. Please see OSG Policy section 2.3.7 on Required Disclosure for Independently Published Results: http://www.spec.org/osg/policy.html.

3.3.1. Estimates

This rule is covered in SPEC Fair Use Rule: http://www.spec.org/fairuse.html.

3.3.2. Comparison to Other Tool Suites

This rule is covered in SPEC Fair Use Rule: http://www.spec.org/fairuse.html.

3.3.3. Addendum to OSG Fair Use Policy

This rule is covered in SPEC Fair Use Rule: http://www.spec.org/fairuse.html.

3.4. 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.

SPEC is aware that power or performance results for pre-production systems may sometimes be subject to change, for example, when a last-minute bug fix reduces the final performance.

If the sponsor becomes aware that the SERT score of a typical released system is more than 5% lower than that reported for the pre-release system, the tester is required to submit a new result for the production system, and the original result must be removed, corrected or marked non-compliant (NC).

3.5. 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.

3.6. 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. All the system, hardware, and software features are required to be available within three months of the first publication of these results. With multiple sub-components of the major components having different availability dates, the latest availability date must be listed for that major component. The SERT software components are not included in this date.

3.7. 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.

Please note that the method that started the tool needs to be disclosed in the “System under Test Notes” when no keyboard was used (e.g., Run was started via Remote Desktop).

