



Standard Performance Evaluation Corporation (SPEC)

Server Efficiency Rating Tool (SERT) User Guide

1.1.1

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

Table of Contents

1. Overview	4
1.1. Summary.....	4
1.2. Setup Requirements	4
1.3. SUT Configuration Requirements	5
1.3.1. Platform Requirements.....	5
1.3.2. Storage Requirements	6
1.3.3. Memory Requirements.....	6
1.3.4. Network Requirements.....	7
1.3.5. Software Requirements	7
1.4. SERT Installation.....	7
2. Starting SERT via the Graphical User Interface (GUI).....	8
2.1. Prerequisites for Using the GUI	8
2.1.1. On the SUT and Controller system:	8
2.1.2. On the SUT only:	8
2.2. Invoking the GUI.....	8
2.3. Interacting with the GUI	8
2.3.1. Start	8
2.3.2. SUT Configuration	10
2.3.3. SUT Discovery	11
2.3.4. Test Environment.....	12
2.3.5. PTDaemon Configuration.....	14
2.3.6. JVM Configuration	19
2.3.7. JVM Review.....	20
2.3.8. Choose Test Suite.....	21
2.3.9. Launch Test	22
2.3.10. Results.....	24
2.3.11. Help.....	25
3. Platform Specific Settings	26
3.1. Microsoft Windows	26
3.1.1. Disabling the storage write cache on Microsoft Windows	26
3.1.2. Disabling the Windows firewall	27
3.2. Linux	28
3.2.1. Prerequisite Software	28
3.2.2. Enabling Transparent Huge Pages (THP).....	28
3.2.3. Disabling the storage write cache on Linux	28
3.2.4. Disabling the firewall	28
3.2.5. Linux on IBM Power Systems.....	29

3.3. IBM AIX	29
3.3.1. Disabling the AIX firewall.....	29
3.4. 32-bit ARM Platform	30
3.5. Oracle Solaris	30
3.5.1. Disabling the storage write cache on Oracle Solaris.....	30
3.5.2. Disabling the Solaris firewall.....	31
3.6. Lower Performance Platforms	31
4. Power Analyzer Range Settings	31
4.1. Automating Range Settings using the GUI	32
4.2. Manual Range Settings – Optional	32
4.3. Extracting Range Settings from Full Runs - Optional.....	32
5. Using the Collector Scripts	33
6. SERT Command Line Usage	34
6.1. SERT Configuration and Start Procedure.....	34
6.2. Generate report files with the reporter scripts	38
7. Coefficient of Variation (CV)	39
7.1. Host CV	39
7.2. Client CV	39
8. SUT Minimum Memory Requirement	39
9. Known Issues	40
10. Trademark	40
11. Copyright Notice	40

SVN Revision: 2023

SVN Date: 2016/01/27 17:53:45

1. Overview

1.1. Summary

Various government agencies around the world are currently working on Energy Efficiency Regulatory Programs for servers. The SPECpower Committee designed, implemented, and delivered the Server Efficiency Rating Tool (SERT)[™], a next-generation tool set for measuring and evaluating the energy efficiency of computer servers.

This User Guide is intended for both novice and experienced SERT testers. It is structured with the basics to get a novice user up and running quickly and includes a considerable amount of additional details for more experienced testers.

The SERT is composed of several elements including:

- The test harness (called **Chauffeur**) handles the logistical side of measuring and recording power data along with controlling the software installed on the SUT and controller system itself.
- The **Director** instructs the System under Test (SUT) to execute the workload.
 - The SUT can be a single or multi node server
- The **SERT Workload Suite** exercises the SUT while the test harness collects the power and temperature data.
- The SPEC Power and Temperature Daemon (**PTDaemon**) connects to the power analyzer and temperature sensor and gathers their readings while the workload executes.
- The **Reporter** gathers the environmental, power, and performance data after a run is complete and compiles it into an easy-to-read format.
- A Graphical User Interface (**GUI**) to simplify setting up and executing the SERT.

Additional background information describing the design and implementation of the SERT may be found in the SERT Design Document (<http://www.spec.org/sert/docs/designdocument.pdf>).

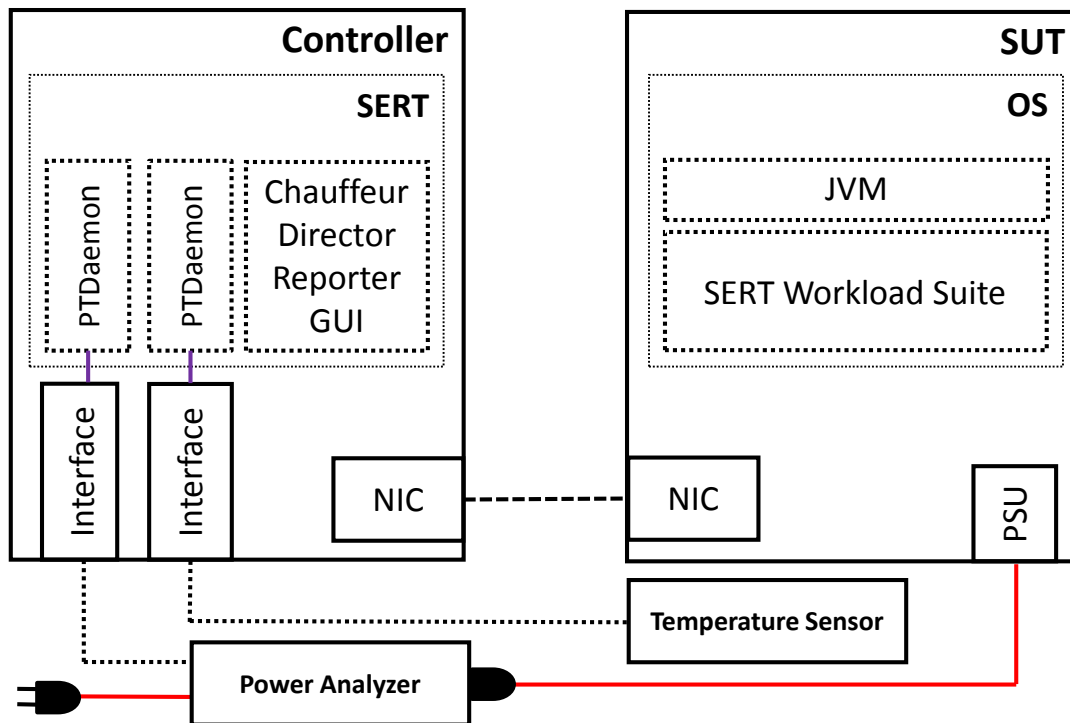
1.2. Setup Requirements

For the most basic SERT hardware measurement setup, one of each of the following is required:

- System under Test (**SUT**): the actual system for which the measurements are being taken
- **Controller** (e.g. a server, desktop PC, or laptop): the system to which the power analyzer, temperature sensor, and the SUT are connected
 - The Controller and SUT are connected to each other via an Ethernet connection.
 - The **Power Analyzer** and **Temperature Sensor** are connected to the Controller via device specific means.
 - Hardware requirements:
 - A hardware configuration capable of running a supported OS and JVM combination.
 - At least one wired Network Interface Controller (**NIC**) with a minimal transfer rate of 1Gbit/s.
 - At least two communication interfaces for controlling a power analyzer and a temperature sensor.
 - Software requirements:
 - OS (64 bit - x86): Current versions of Microsoft Windows OS and most current Linux distributions.
 - A current Java 7 or 8 run time environment (64 bit).
 - Vendor specific drivers supporting special power analyzer and / or temperature sensor connectors, if required.
- **SUT and Controller**: The date, time, and time zone must be set correctly.

- Power analyzer and temperature sensor:
 - **Note:** Before connecting the power analyzer and the temperature sensor to the Controller and the SUT, please read the “Power and Temperature Measurement Setup Guide” in the Documentation directory on the SERT media carefully! The latest version of this document can be found on the SPEC web site (http://www.spec.org/power/docs/SPEC-Power_Measurement_Setup_Guide.pdf). This practical guide explains how to set up and run various power analyzers and temperature sensors with the SPEC PTDaemon.
 - The power analyzer is connected to the Controller and used to measure the power consumption of the SUT, whilst the temperature sensor is also connected to the Controller and used to measure the ambient temperature in front of the SUT’s main airflow inlet.
 - **Note:** This release of the SERT is only intended to be used with servers running on AC power. DC power measurement is currently not supported.
 - Only a power analyzer and temperature sensor from the list of accepted measurement devices (http://www.spec.org/power/docs/SPECpower-Device_List.html) can be used with SERT.
 - The power analyzer must have been calibrated within 12 months before a SERT run is performed.

The following block diagram represents a simple SERT environment.



1.3. SUT Configuration Requirements

For a compliant SERT result, a basic set of hardware and software components is needed. In multi node test setups, all nodes must be equipped with exactly the same hardware and software components, i.e. only homogeneous configurations are allowed.

1.3.1. Platform Requirements

For the list of supported platforms, please see the SERT Client configuration guide: http://www.spec.org/sert/docs/SERT-JVM_Options-110.html.

1.3.2. Storage Requirements

- At least one storage device
- For a single drive, a minimum of two partitions are required
 - OS partition(s) and one SERT Data partition (min. 20 GB)
- For multiple drives, drives should be partitioned as follows:
 - First drive: OS partition(s) and one SERT Data partition (min. 20 GB)
 - Any additional drive: SERT Data drive (entire drive capacity)
- For all storage devices the device cache must be disabled and each device must be configured in “Write Through” mode; see section 3 for platform specific instructions.
- SERT Data partitions may only be hosted on filesystems (e.g., ntfs, ext4, zfs, jfs) that support direct I/O.
- RAID arrays of more than one device are not allowed, i.e. each individual storage device must be configured as JBOD or single disk RAID 0 array.
- The content of the storage devices being tested will not be erased. For meaningful results however, it is strongly recommended to reformat the filesystem on the storage devices before the SERT storage worklets are started for the first time. The test data files created during the first test run can be reused in subsequent runs.

1.3.3. Memory Requirements

The following table shows the minimum memory requirements by total system hardware threads.

Minimum System Memory to Run SERT		
Total System Hardware Threads (Number of CPUs * Cores/CPU * Hardware Threads / Core)	1-23	≥ 8 GB
	24-37	≥ 12 GB
	38-51	≥ 16 GB
	52-78	≥ 24 GB
	79-105	≥ 32 GB
	106-132	≥ 40 GB
	133-159	≥ 48 GB
	160-187	≥ 56 GB
	188-214	≥ 64 GB
	215-268	≥ 80 GB
	269-323	≥ 96 GB
	324-377	≥ 112 GB
	378-431	≥ 128 GB
	432-486	≥ 144 GB

Note: Some systems may be able to run with less memory than indicated in this table, but performance may be reduced. For exact minimum memory sizes within each range, please see chapter 8.

Memory should be evenly distributed between processors for best performance; e.g. sockets or physical packages. In the case of a two socket/processor system, the memory should be equally distributed between the sockets. You may encounter WARNING messages stating that the coefficient of variation between clients was greater than 5% if the memory is not evenly distributed.

1.3.4. Network Requirements

At least one Network Interface Controller (NIC) with a minimal transfer rate of 1 Gbit/s

1.3.5. Software Requirements

Required software components per server (host) are:

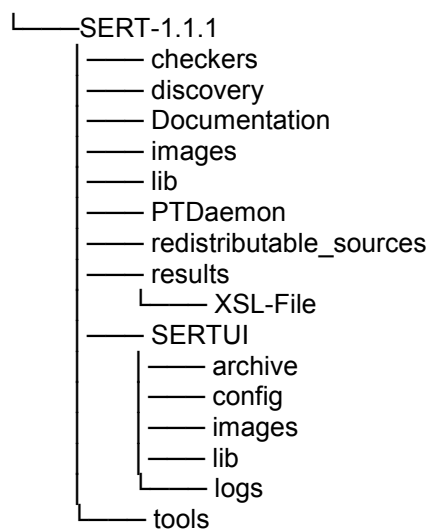
- Exactly one operating system instance (with one user space), including all modules that are installed during the default installation process. The OS must be from the list of supported software platforms (http://www.spec.org/sert/docs/SERT-JVM_Options-110.html).
- If the system uses a hypervisor by 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.
 - A version of a Java run time environment from the list of supported software platforms.
- For the purposes of running the SERT please ensure that the user is logged on as an Administrator (on Windows) or as root (on other platforms).
- On Linux distributions such as Ubuntu where root login is not normally available, the SERT processes should be launched using 'sudo'.

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

1.4. SERT Installation

Ensure that a Java runtime environment is already installed on the SUT and the Controller. Start the SERT setup via `java -jar SERT-1.1.1-setup.jar` and follow instructions on the screen. Ensure that SERT is installed on the OS partition for both SUT and Controller.

Installing the SERT distribution media will create a directory tree of files and subdirectories.



This document can be found in the `Documentation` directory, while the various script and batch files discussed in the following sections are present in the `SERT` directory or the `PTDaemon` sub-directory. If the user will be using multiple machines to perform a SERT test sequence, it will be necessary to unpack (or install) all the files onto the storage devices which hold the operating system image on each system.

2. Starting SERT via the Graphical User Interface (GUI)

The SERT GUI provides an intuitive way for the user to configure and run a test, and to view and archive the results of those tests.

2.1. Prerequisites for Using the GUI

The instructions below assume that the GUI will be running on the Controller system, and that the SUT is a separate host.

2.1.1. On the SUT and Controller system:

- Ensure that the firewall is disabled; see section 3 for platform specific instructions.

2.1.2. On the SUT only:

- If applicable, please follow the platform specific settings in section 3.
- Set the `JAVA` variable in the `serthost.sh/.bat` script to the java executable SERT should use. For example:
 - For `serthost.bat`:
 - `set JAVA="C:\Program Files\Java\jdk1.7.0\bin\java.exe"`
 - For `serthost.sh`:
 - `JAVA=/usr/java/jdk1.7.0_71/bin/java`
- Run `[Install directory]\SERT-1.1.1\serthost.bat` or `serthost.sh`

2.2. Invoking the GUI

Open a cmd window and run one of the scripts `SERT-1.1.1\SERTUI\run-sertui.bat` or `run-sertui.sh` depending on your operating system.

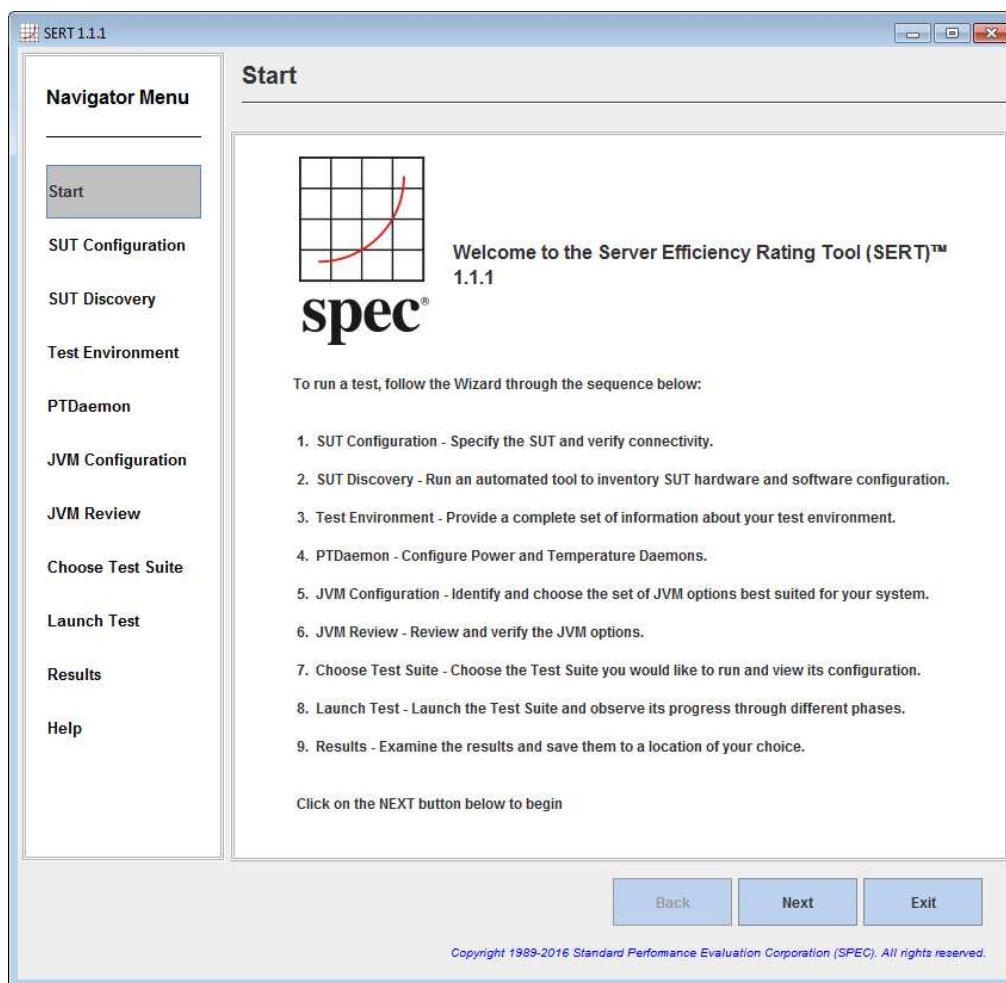
```
cd [Install directory]\SERT-1.1.1\SERTUI
run-sertui.bat
```

The command script doesn't return and will leave the window open.

2.3. Interacting with the GUI

2.3.1. Start

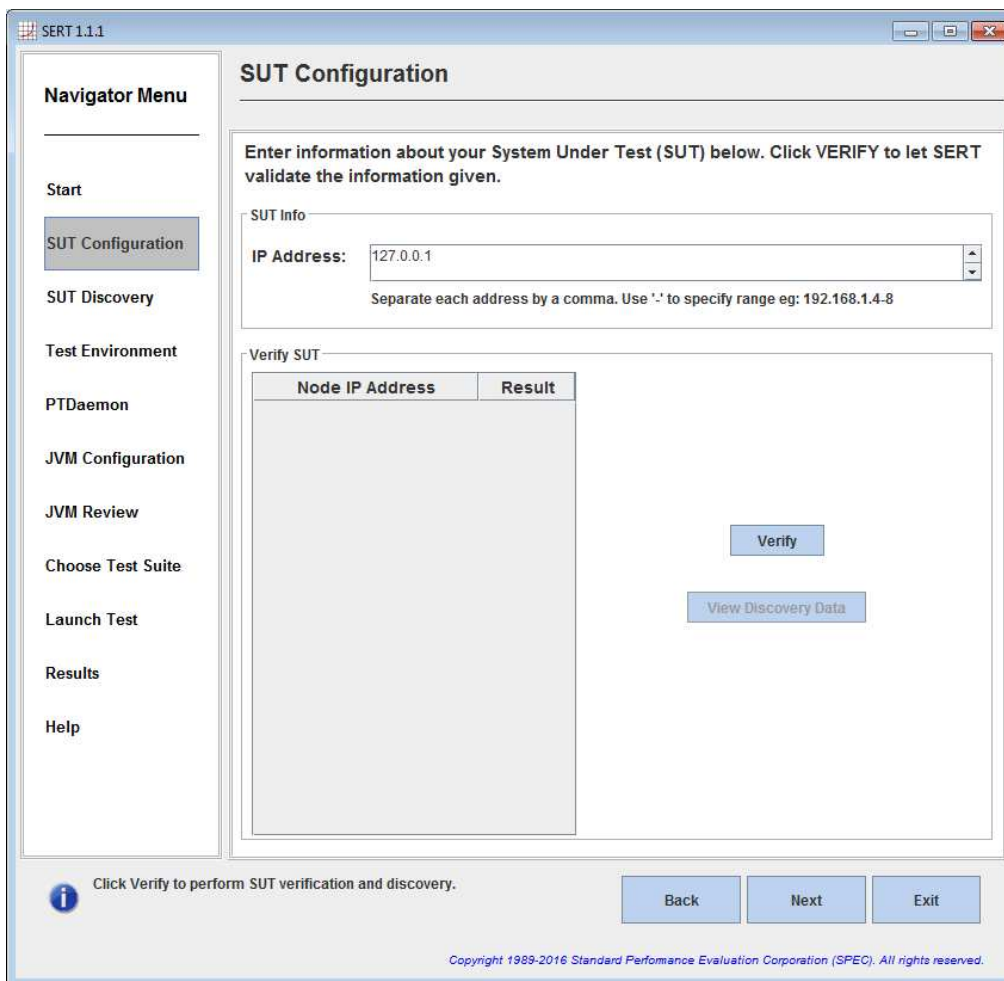
The **Start** panel of the GUI shows the sequence of steps necessary to run the SERT.



Click **Next** to continue, or **Exit** to terminate the GUI.

2.3.2. SUT Configuration

The *SUT Configuration* panel allows the user to configure the IP Address(es) of the SUT.



Provide the **IP Address**, multiple IP addresses can be provided via a comma-separated list (e.g. 10.0.0.101, 10.0.0.110) or by specifying an IP Address range (e.g. 10.0.0.101-120).

Click **Verify**, in order for the GUI to verify the connection to the remote service running on the SUT(s). If a failure occurs, an error message will appear in the lower left corner of the GUI.

When verification is successful, click **Next** to continue.

2.3.3. SUT Discovery

The **SUT Discovery** panel shows software and hardware configuration retrieved from the SUT. Please wait for the process to complete, after which a table of discovery items will be shown. Select the **Save** check box and click **Next** to save the configuration values into the test environment file (`test-environment.xml`). On subsequent runs, uncheck the **Save** box to bypass overwriting any configuration changes that might have been made.

If the test configuration is already correctly described in `test-environment.xml` you should deselect the **Save** check box to prevent this information being overwritten by SERTUI.

SUT Discovery

SUT Info

Network Address : 10.0.0.101
System Type : Linux
SERT Install Directory : /home/SERT-1.1.1

WARNING: Discovery items are not being saved.

The following values have been automatically detected and will be copied into the test environment configuration. Deselect the below check box if you would like to keep your test environment configuration as-is, or select individual items to save.

Save Discovery items to Test Environment Configuration file.

Toggle ALL items

Save	Discover Item	Result
<input type="checkbox"/>	CpuName	Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz
<input type="checkbox"/>	CpuHardwareThreadsPerC...	2
<input type="checkbox"/>	CpuNumberOfChips	2
<input type="checkbox"/>	CpuFrequencyMhz	1200
<input type="checkbox"/>	CpuCores	36
<input type="checkbox"/>	CpuActiveCores	36
<input type="checkbox"/>	BiosVersion	V5.0.0.9 R1.24.0.SR.1 for D3289-A1x
<input type="checkbox"/>	SystemVendor	FUJITSU
<input type="checkbox"/>	SystemModel	PRIMERGY RX2540 M1
<input type="checkbox"/>	MemorySlots	24
<input type="checkbox"/>	MemorySizeMB	258210
<input type="checkbox"/>	OSVendor	Red Hat
<input type="checkbox"/>	OSName	Red Hat Enterprise Linux Server
<input type="checkbox"/>	OSVersion	6.6 (Santiago)
<input type="checkbox"/>	NumaNodes	2
<input type="checkbox"/>	NicDescription	Emulex Corporation OneConnect NIC (Skyhawk) (rev 10)
<input type="checkbox"/>	NicDescription	Emulex Corporation OneConnect NIC (Skyhawk) (rev 10)
<input type="checkbox"/>	NicSpeed	0
<input type="checkbox"/>	NicSpeed	4000

Discovery items are not being saved. Click NEXT to continue.

Back Next Exit

Copyright 1989-2016 Standard Performance Evaluation Corporation (SPEC). All rights reserved.

2.3.4. Test Environment

The **Test Environment** panel allows the user to provide detailed information about the hardware and software setup used to obtain the SERT results. The information as defined in the Test Environment fields will appear in the Results Report at the end of the test run. Initially, the default template is populated with sample fields to provide an example configuration. The discovery process will also populate some fields in certain sections of the test environment file. If the status of any specific section is listed as “Defaults in use” displayed with a yellow background, this means that at least one of the fields listed within this section is still using a default value which does not reflect the actual system being measured. If the status of any specific section is listed as “Incomplete” displayed with a red background, the discovery process was unable to obtain any information. In both cases the user is required to provide the specific value for this field. If a section is listed as “Complete” with a green background, this indicates that all fields within this section have been updated either through discovery value population or manual editing.

Test Environment

Test environment information is required to document the hardware and software used in a SERT Test Suite.

- Edit each section by double clicking on the row
- Load a previously saved environment configuration by clicking LOAD

Section	Description	Items	Status
Test Information	Information identifying the testers and their location.	7	Defaults in use
SUT: Basic Information	Basic information about the system and its devices.	18	Defaults in use
SUT: CPU Information	Information about system CPU(s).	15	Defaults in use
SUT: Memory Information	Information about system memory.	4	Defaults in use
SUT: Memory DIMM Details	Information about memory DIMMs.	3	Complete
SUT: Power	Information about system power.	8	Defaults in use
SUT: Power Supply Details	Information about power supplies.	4	Defaults in use
SUT: Disk Drives	Information on system disk drives.	9	Incomplete
SUT: Network Interfaces	Information on system network interfaces.	6	Defaults in use
SUT: Other	Information on other hardware installed on the system.	2	Defaults in use
Software: Operating System	Information about the OS installed on the system.	5	Defaults in use
Software: Other	Information about other software installed on the system.	2	Defaults in use
JVM	Information about the JVM and its configuration.	2	Complete
Firmware	Information about system BIOS and management firmw...	4	Defaults in use
Notes	Additional notes on system details and configuration.	1	Defaults in use

Load... Save Copy... Preview Report

Click NEXT to continue when editing is complete.

Back Next Exit

Copyright 1989-2016 Standard Performance Evaluation Corporation (SPEC). All rights reserved.

The SUT Discovery from the previous step only partially identifies the configuration. Every section needs to be checked by the user, and entries must be edited manually as appropriate. These elements must be double-checked for correctness and examined for proper formatting. Highlight a row in the table and double click to edit the information associated with that section; e.g. select the **Test Information** section.

To make changes, double click the field to be edited in the right column, change the value directly, and hit the **Return key**. On completion, click **OK** to save.

Besides changing the available entries, some test environment sections support adding new groups of items. Currently this feature is available for **Memory DIMM Details** and for **Disk Drives**. Click on the **Add Group/Item** button to add a new group of items. The example in the picture below shows three new detail items for DIMM group [1]. The information in the value fields must be edited as described above. Use the **Delete** button to delete the highlighted items. Use the **OK** button to store the updated information and return to the main **Test Environment** panel.

The screenshot shows a window titled "Test Environment Section Editor" with a sub-section "Section: SUT: Memory DIMM Details". Inside the window is a table with two columns: "Item" and "value". The table contains the following data:

Item	value
Number of Dimms [0]	4
Size of Dimm (MB) [0]	16384
Description of Dimm [0]	16GB 2Rx4 PC4-2133P-R; slots 1, 3, 5, and 7 populated
Number of Dimms [1]	4
Size of Dimm (MB) [1]	8192
Description of Dimm [1]	8GB 2Rx4 PC4-2133P-R; slots 2, 4, 6, and 8 populated

Below the table, there are four buttons: "Add Group/Item", "Delete", "OK", and "Cancel".

Note that all sections of the test environment information must be complete before continuing further. This is indicated in the right column of the table showing a green color for all sections.

By clicking on the **Save Copy** button the user can save a fully populated template for subsequent re-use. Use the **Load** button to re-load an existing template in place of the current default template.

After editing test environment fields you should preview the System Under Test section of the result report file by clicking on the **Preview Report** button.

Click **Next** to continue.

Note: For the purposes of informal testing, the user may choose not to edit any of the test environment information provided by the default template. Simply click **Next** to continue.

All status fields showing the state "**Complete**" are highlighted green to represent that all information fields have been changed from their default values. It does not necessarily indicate that the information within these fields is correct and valid for publication. The tester has to verify the correctness of all the information manually and make the required changes.

The default parameters are defined with a leading underscore to clearly mark them as placeholder values that must be edited to produce valid results. The reporter recognizes default placeholders and highlights these fields with a yellow background in the result file.

2.3.5. PTDaemon Configuration

The **PTDaemon Configuration** panel allows the user to view and edit settings associated with temperature and power instances of PTDaemon. There is a tab for the **Power Analyzer** and for the **Temperature Sensor**. Additional devices can be added with the **Add** button or removed with the **Remove** button; this is a convenient way to retain configuration information in a lab environment where multiple devices are available.

For the **Power Analyzer** tab, there are four sections:

- **Measurement Device:** Choose which power analyzer is being used from the drop down menu. Once a device has been selected, only the options appropriate for that device will be displayed in the remaining sections.
- **Interface:** Choose how the power analyzer is connected to the controller, and specify the interface address.
- **Calibration:** Specify the date of last calibration for the power analyzer.
- **Logging:** If needed, save a log and/or debug output file by clicking on the checkbox. This should not normally be used except under guidance from SERT support personnel.

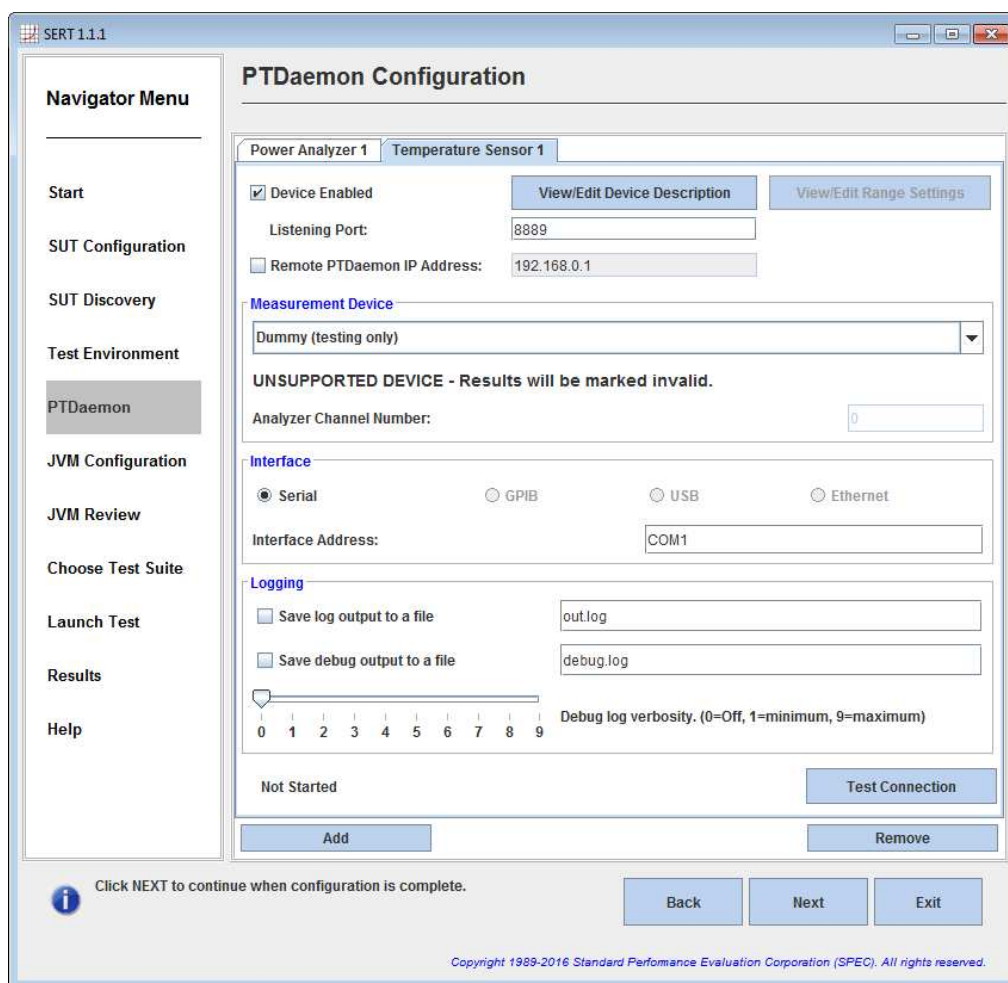
After making your choices, click on **Test Connection** to start the connection between the power analyzer and the controller. Be patient, as this process may take 30 seconds or more.

Power for the different SERT worklets and load levels may vary significantly. This may result in current readings which fall outside the limits of the given ampere range. This will lead to invalid power readings or high uncertainty values (>1%), in both cases the result will be marked invalid. In order to avoid such conditions, power ranges must be adjusted dynamically during a SERT test run. Also, the power analyzer voltage range must be set to match the standard voltage of the test environment. Proper setting of the voltage range is a prerequisite for correct calculation of uncertainties.

Both voltage and current ranges can be specified individually per workload, worklet and load level in the `listeners.xml` file or under **View/Edit Range Settings** in the PTDaemon configuration panel.

Range settings from previous runs will be imported from `listeners.xml` by the SERT UI and only need to be edited after configuration changes. It is strongly recommended though to check the settings for correctness before starting a test run.

Starting with SERT version 1.1.0, the GUI offers a method to partially automate the specification of range values by running a brief test prior to a full SERT run. See Section 4 for instructions on performing a range-setting run.



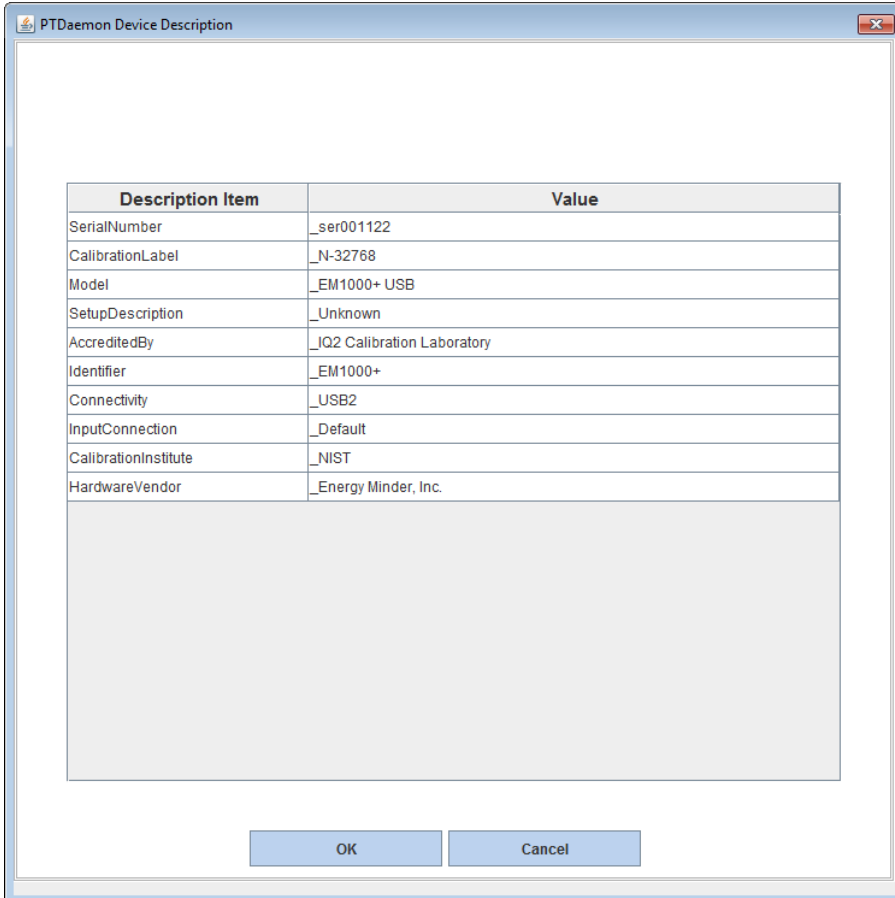
For the **Temperature Sensor** tab, there are three sections:

- **Measurement Device:** Choose which temperature sensor is being used from the drop down menu.
- **Interface:** Choose how the temperature sensor is connected to the controller, and specify the interface address.
- **Logging:** If needed, a log and/or debug output file can be saved by clicking on the checkbox.

After making your choices, click on **Test Connection** to start the connection between the temperature sensor and the controller.

2.3.5.1 View/Edit Device Description

The following figure shows an example of a power analyzer device description. Similar description panel is available for temperature sensor devices. Device descriptions should always be checked for correctness, especially if the test environment description was imported from an existing description file using the Load function of the Test Environment panel, see description above. Imported device information will only be shown here if the PTDaemonPort field in the test environment file does match the PTDaemon Listening Port configured in the Power Analyzer tab of the PTDaemon panel.



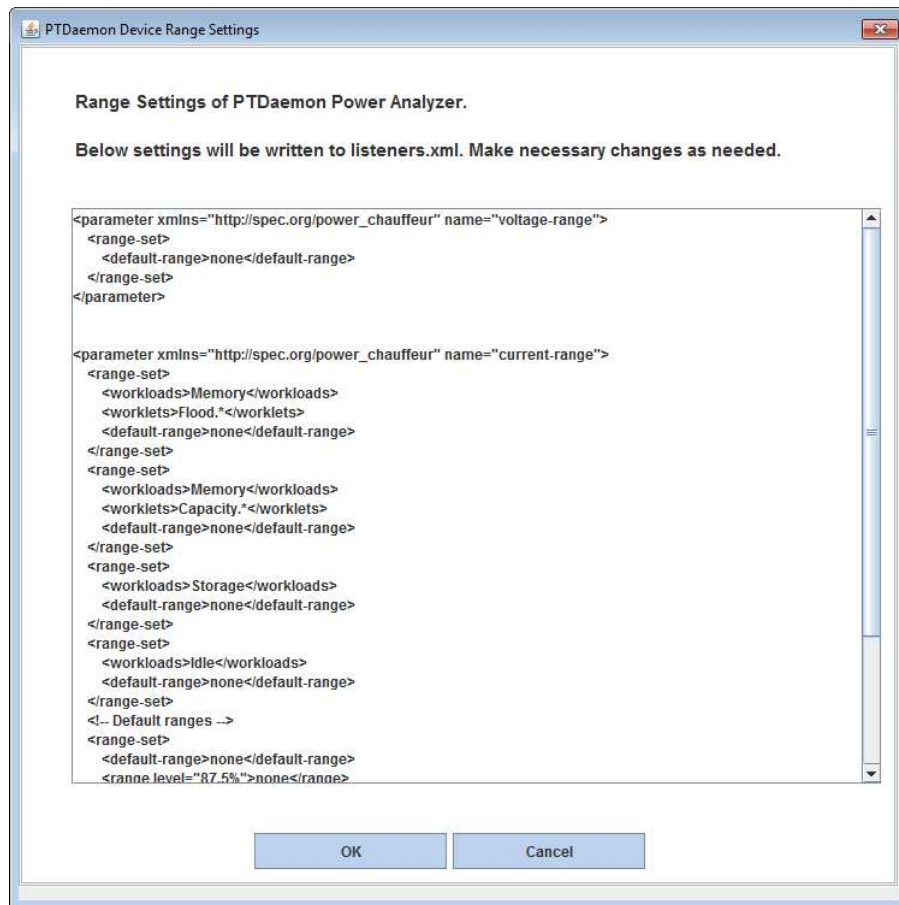
Description Item	Value
SerialNumber	_ser001122
CalibrationLabel	_N-32768
Model	_EM1000+ USB
SetupDescription	_Unknown
AccreditedBy	_IQ2 Calibration Laboratory
Identifier	_EM1000+
Connectivity	_USB2
InputConnection	_Default
CalibrationInstitute	_NIST
HardwareVendor	_Energy Minder, Inc.

Information shown here is imported from the `test-environment.xml` file. Each field can be edited individually to match the given configuration. Specifically you need to ensure that the power analyzer has been calibrated within one year before running the test. The corresponding CalibrationLabel must be specified in the appropriate field. A field for specification of the calibration date is available in the Power Analyzer tabs of the PTDaemon Configuration panel, see description above.

The device description is copied to the `test-environment.xml` file for use in subsequent SERT runs.

2.3.5.2 View/Edit Range Settings

Voltage and current ranges can be specified in the **PTDaemon Device Range Settings** panel, see picture below.




The default ranges (**none**) shown in the example above must be replaced by a suitable current or voltage value. Range definitions for other workloads, worklets and/or intervals can be added here as well.

The new information is copied to the `listeners.xml` file and used for subsequent SERT test runs.

The process to find suitable range settings for a given configuration is described in more detail in chapter [4 Power Analyzer Range Settings](#). Power analyzer specific information is given in the [SPECpower-Measurement Setup Guide.pdf](#) document.

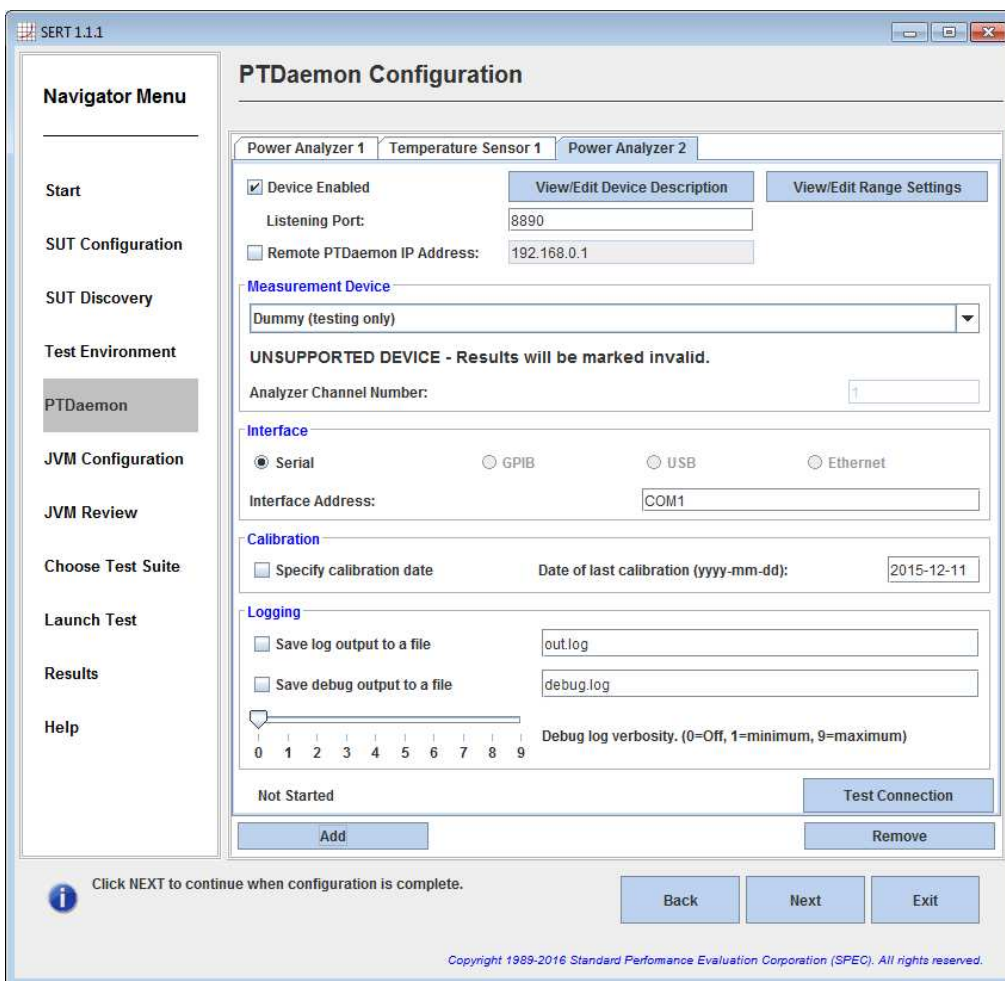
2.3.5.3 Additional Power Analyzer and Temperature Sensor Device

Additional device can be included by pressing the Add button.



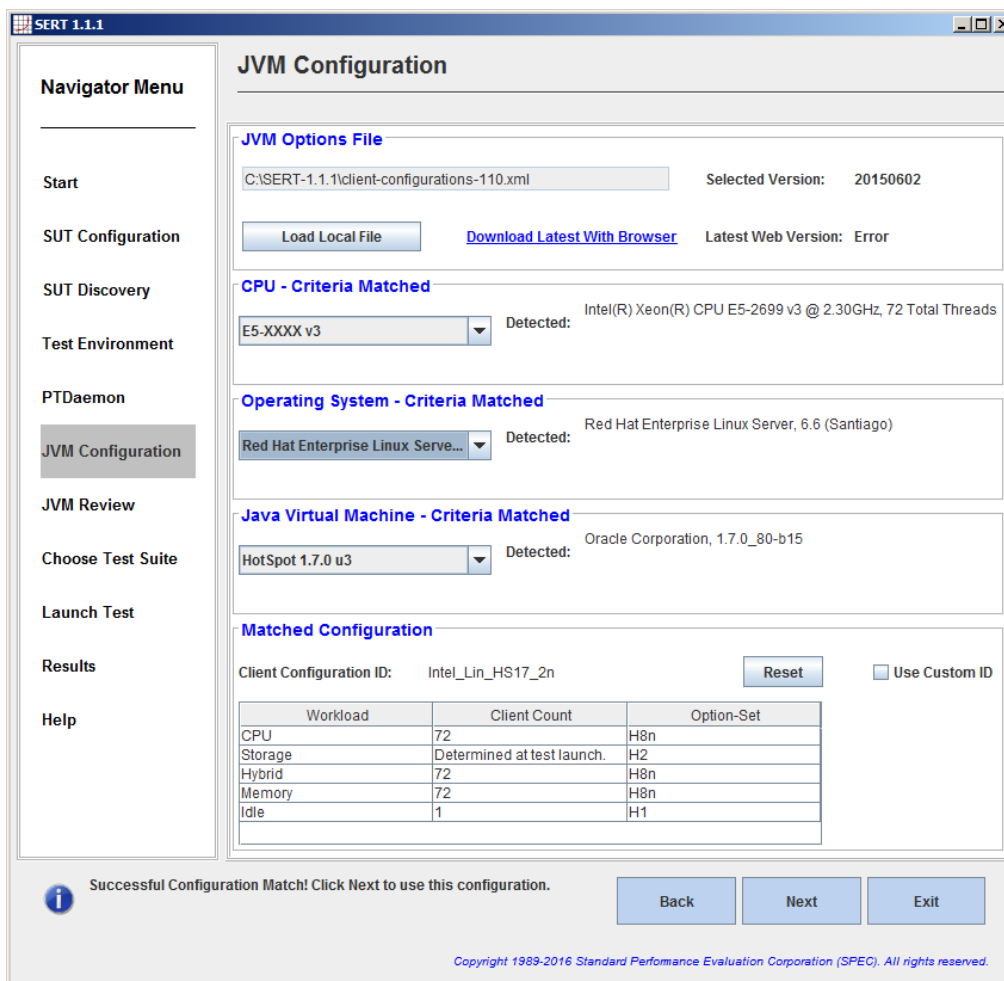
Select the device type and modify that setting to match your device. Please note to change the listing port as well (one port / device).

Example: Power Analyzer added



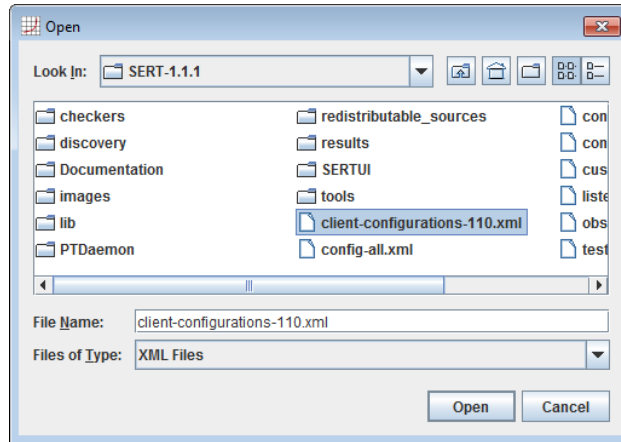
2.3.6. JVM Configuration

The **JVM Configuration** panel allows the user to select the parameters of the tested SUT configuration in order to identify the corresponding Client Configuration ID in the local `client-configurations-110.xml` file.



Note: Before selecting configuration parameters make sure that the local copy of the `client-configurations-110.xml` file matches the official version of this file which is available on SPEC's web site at <http://www.spec.org/sert/docs/client-configurations-110.xml>. The GUI will check this for you automatically; it is reported as the "Latest Web Version". If necessary, you can set a web proxy in the GUI startup scripts (`run-sertui.bat` or `run-sertui.sh`).

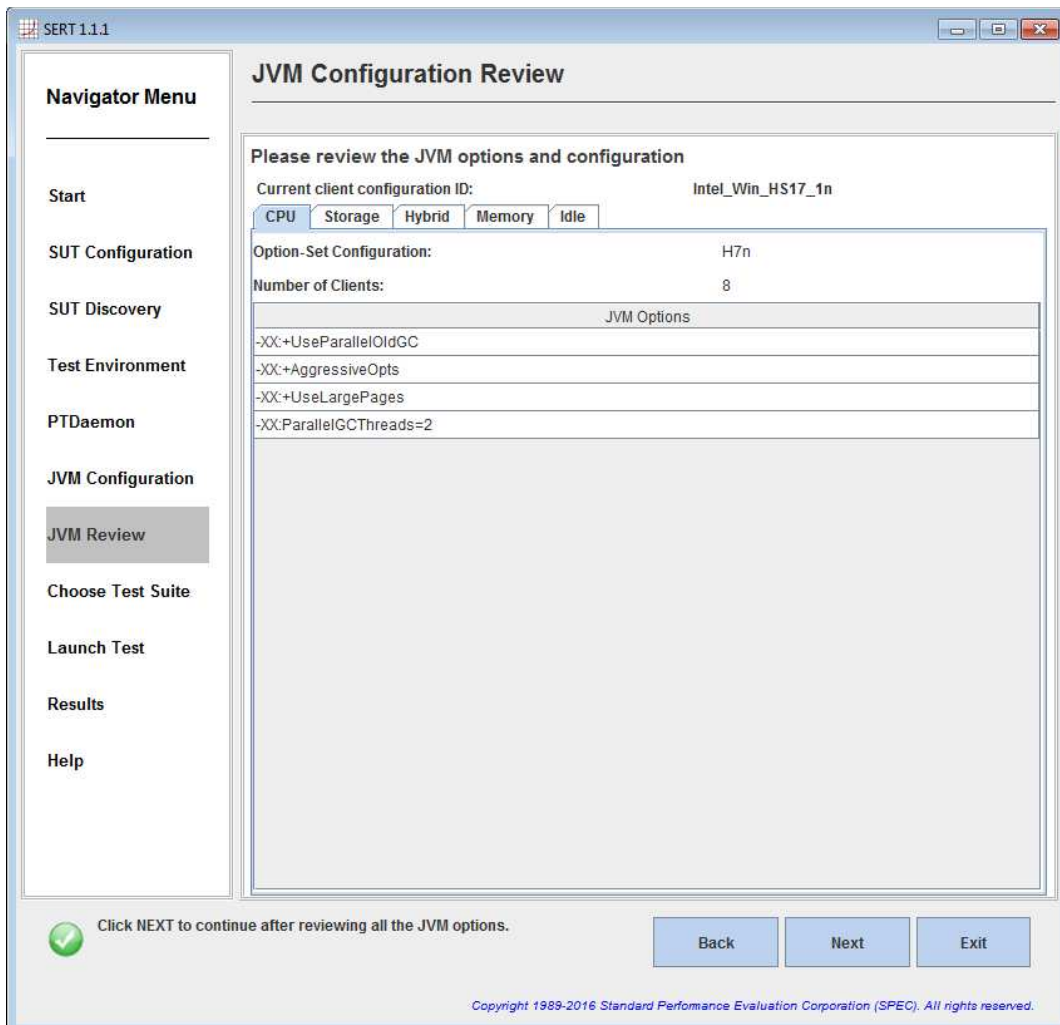
If the version strings don't match, you can download the latest version by clicking on the **Download Latest With Browser** hyperlink. After you have downloaded the latest version, click on the **Load Local File** button for opening this version in SERTUI (see picture below). In case the controller doesn't have internet access you have to find another system with internet access for copying the file on a removable storage medium and then transferring it to a folder on the controller system.



The code 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.

2.3.7. JVM Review

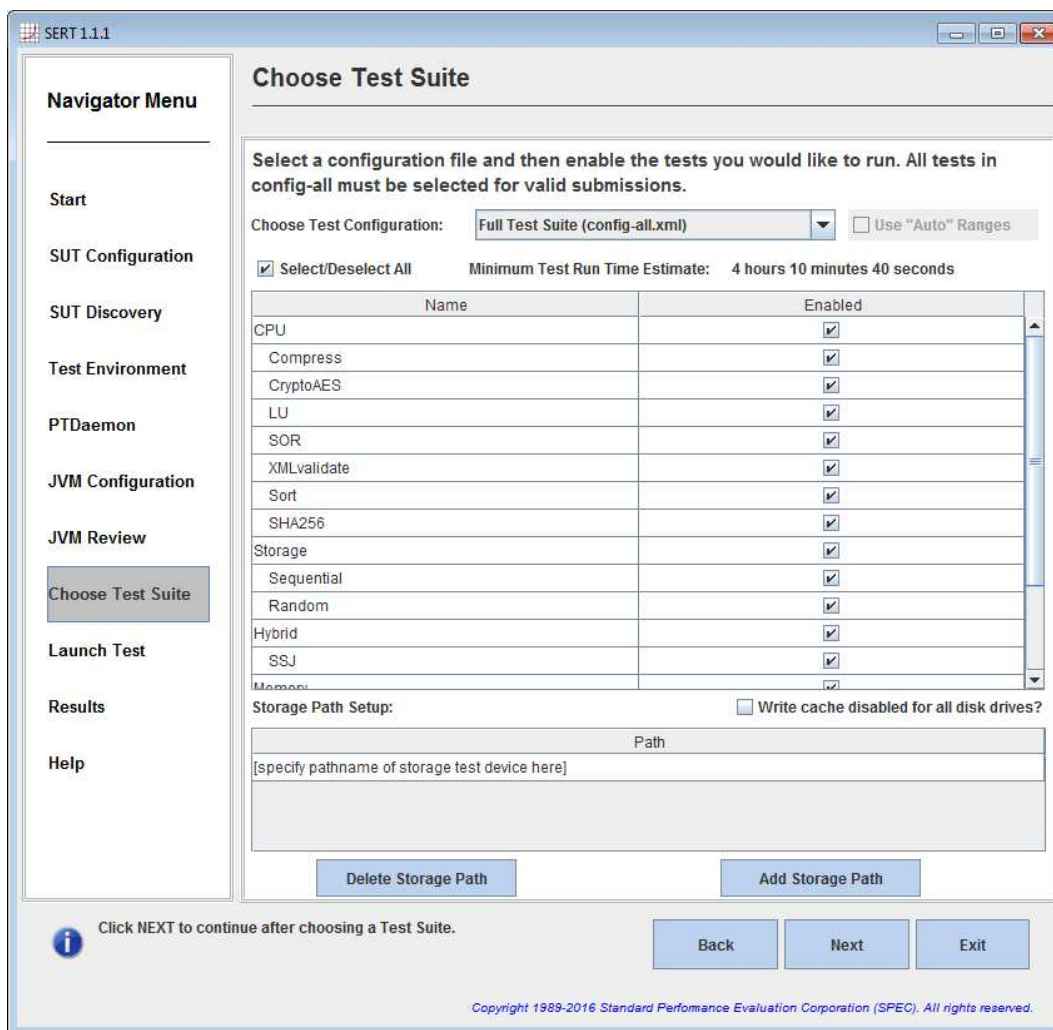
The **JVM Configuration Review** panel allows the user to review the JVM options and the number of client JVMs used for each workload during the measurement.



2.3.8. Choose Test Suite

The **Choose Test Suite** panel allows the user to choose the test to run, to configure key launch variables, and to view configuration information about the test. For a compliant run the “Full Test Suite (config-all.xml)” configuration must be selected.

Note: For the purposes of informal testing, the user may choose to select individual worklets. For a valid run, all worklets must be enabled.



In the Storage Path Setup section, specify the paths of all configured storage devices under test. Separate pathname fields are required for each storage device. More devices can be added using the **Add Storage Path** button. Each time the **Add Storage Path** button is clicked, a new pathname field is created at the bottom of the existing list of pathnames as shown in the following figure. Clicking on the **Delete Storage Path** button will delete the currently selected pathname field. You must enter a pathname referencing a filesystem on a storage device here.

The following requirements apply to pathnames:

- The system partition, e.g. “C:” under Windows must not be specified here. Instead only the pathname for the additional test partition must be given here.
- It is recommended to use descriptive pathnames for mounting the test storage devices and to avoid drive letters for pathname specification. Example: “/SAS-10K_0” instead of “K:” under Windows OS.

Note: According to the SERT run rules, write caches must be disabled for all storage devices. You must select the **Write cache disabled** check box to confirm the storage write caches is disabled. If this check box isn't selected a warning will be displayed later when the test is launched.

Click **Next** to continue after making the selections.

Note: If you start SERT and the specified storage path name does not exist, the following message will appear:

```
Storage verification failed on SUT <SUTNAME>. Storage path <PATHNAME> does not exist. Would you like to continue launch of test suite?
```

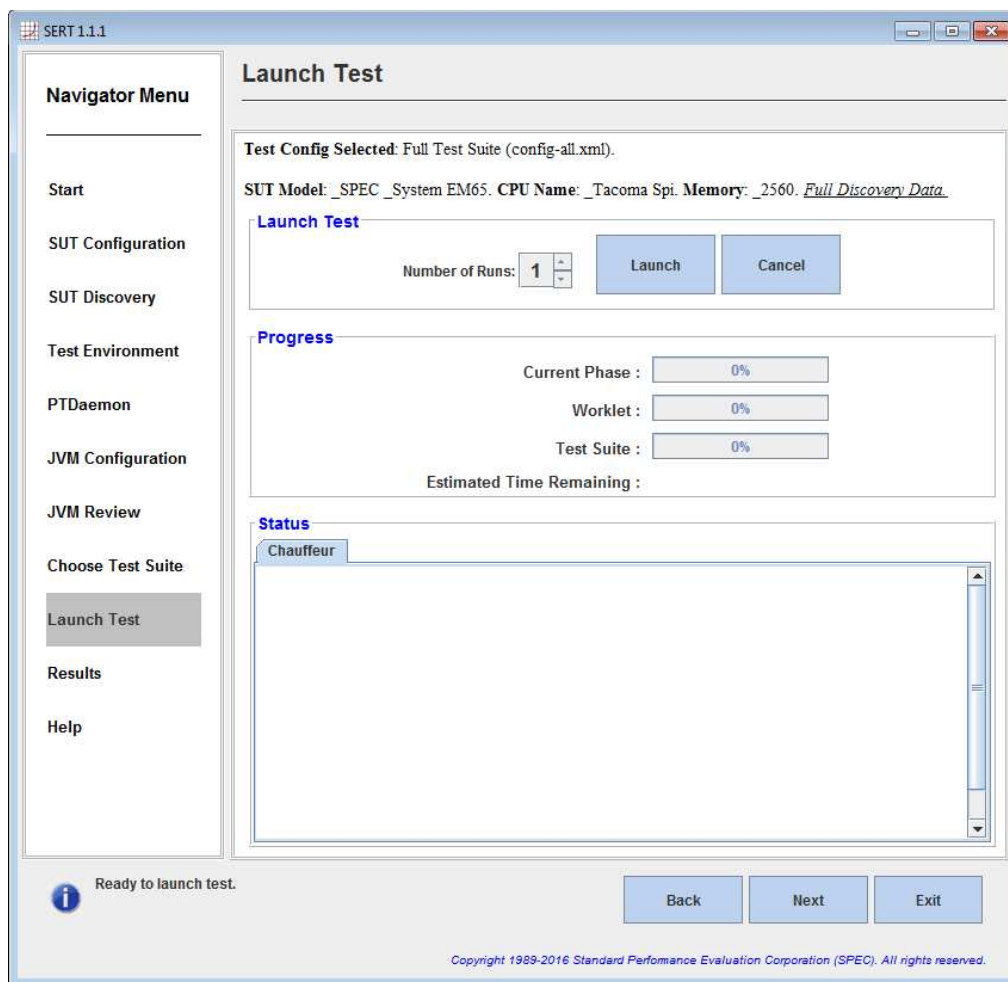
Note: If there is insufficient free space on the specified storage device(s), the following message will appear:

```
Storage verification failed on SUT <SUTNAME>, not enough free space available. Would you like to continue launch of test suite?
```

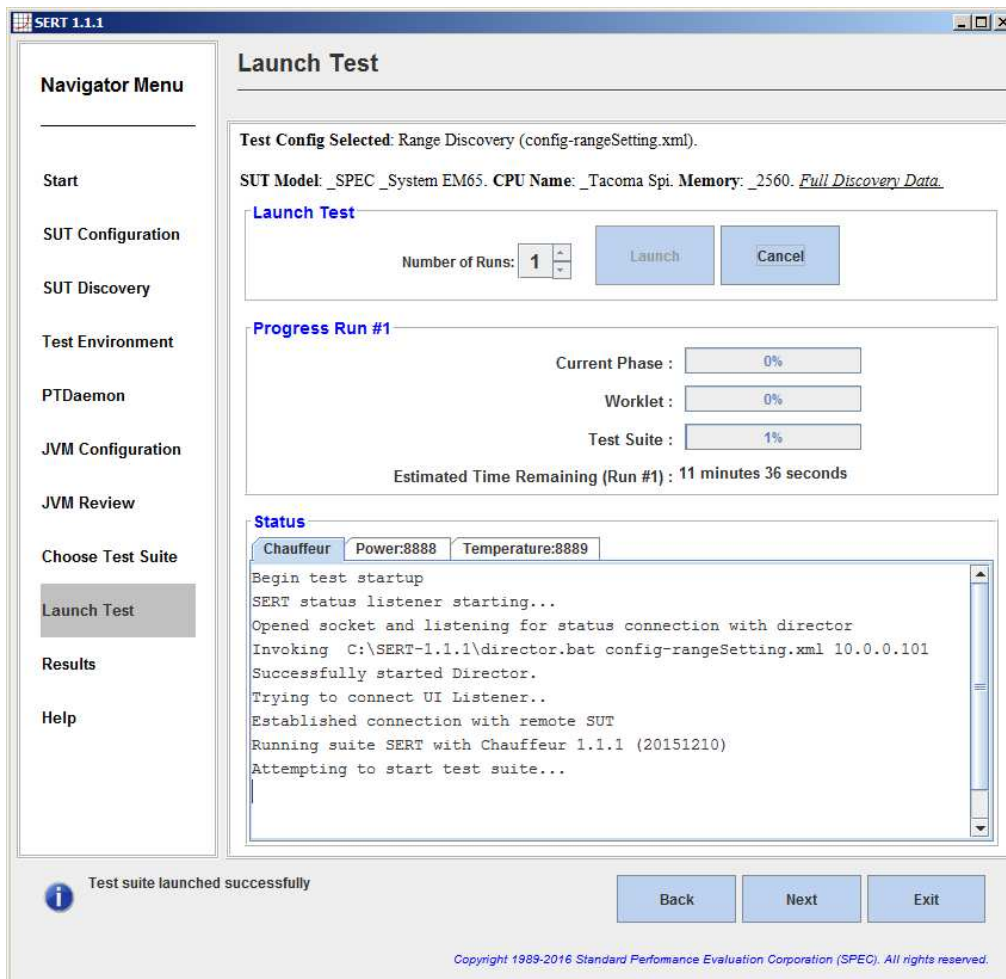
If either error appears, it is highly recommended that the test be terminated immediately and the reported error be corrected before restarting the SERT. If the test is not terminated immediately, it will eventually fail.

2.3.9. Launch Test

The **Launch Test** panel is used to initiate a SERT test, and to monitor its progress. Click **Launch** to begin the test. **Cancel** is available and may be used as necessary.




After a few seconds, Status messages and Progress information will be displayed and the test will proceed as intended. Estimated Time Remaining is approximate and might be significantly longer for high-end configurations, especially multi-node configurations. In some cases, the Progress Bar and Estimated Time Remaining may appear to not show any progress, even though the SERT test is functioning correctly.



When the test completes, click **Next** to continue.

2.3.10. Results

The **Results** panel allows the user to view and archive the results from the test pass that has just completed.



Click **View** to load the results in the default Web browser. Please notice any Warning or Invalid messages on the top of the report. If power and temperature data is shown as zero or invalid after a test run please verify your connections and power cycle the Power Analyzer, as the SERT cannot currently repair this misconfiguration.

Click **Save As** to save the results folder to a location of choice.

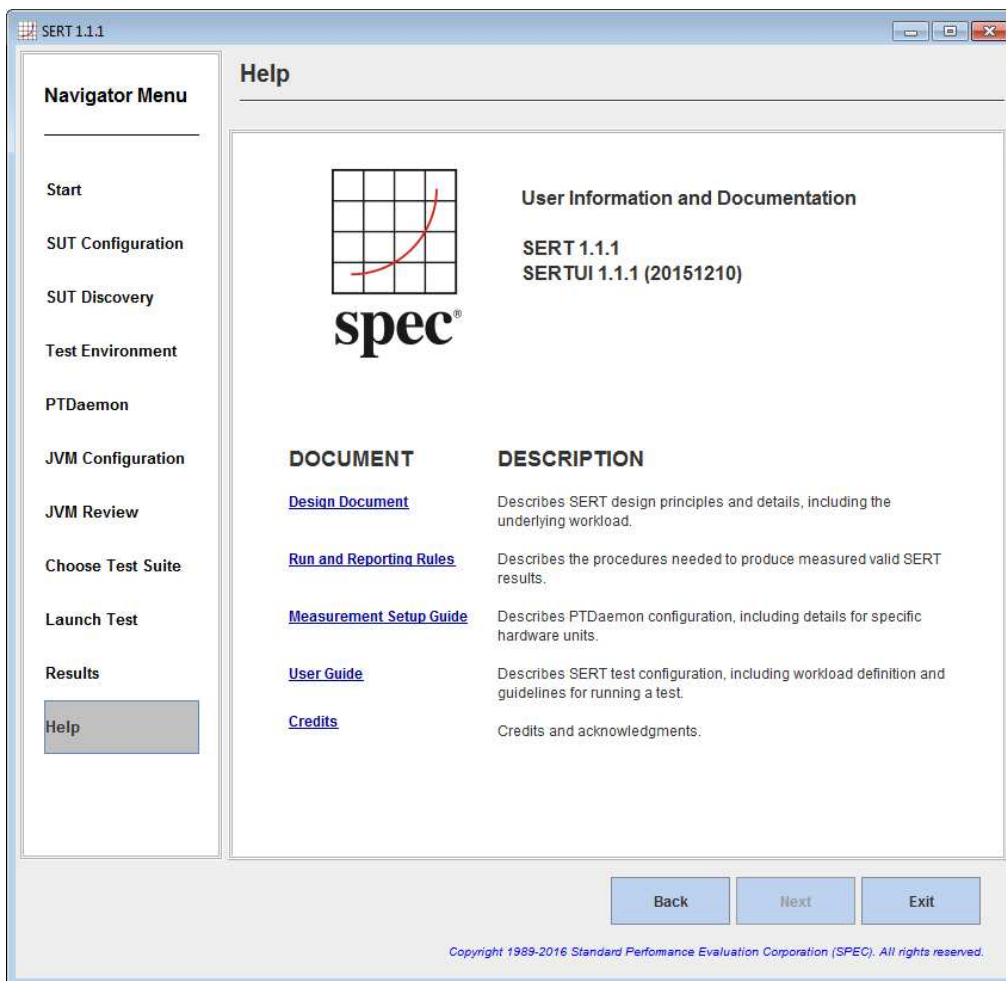
Click **Archive** to name, annotate, and save the results folder to the default "Archive" directory and enter them into the table provided.

Once placed in the archive, the results can be browsed at any time by clicking on **View**.

The test is now complete. Exit the GUI by clicking on **Exit**.

2.3.11. Help

The **Help** panel provides release information and links to various documents made available to the user.



3. Platform Specific Settings

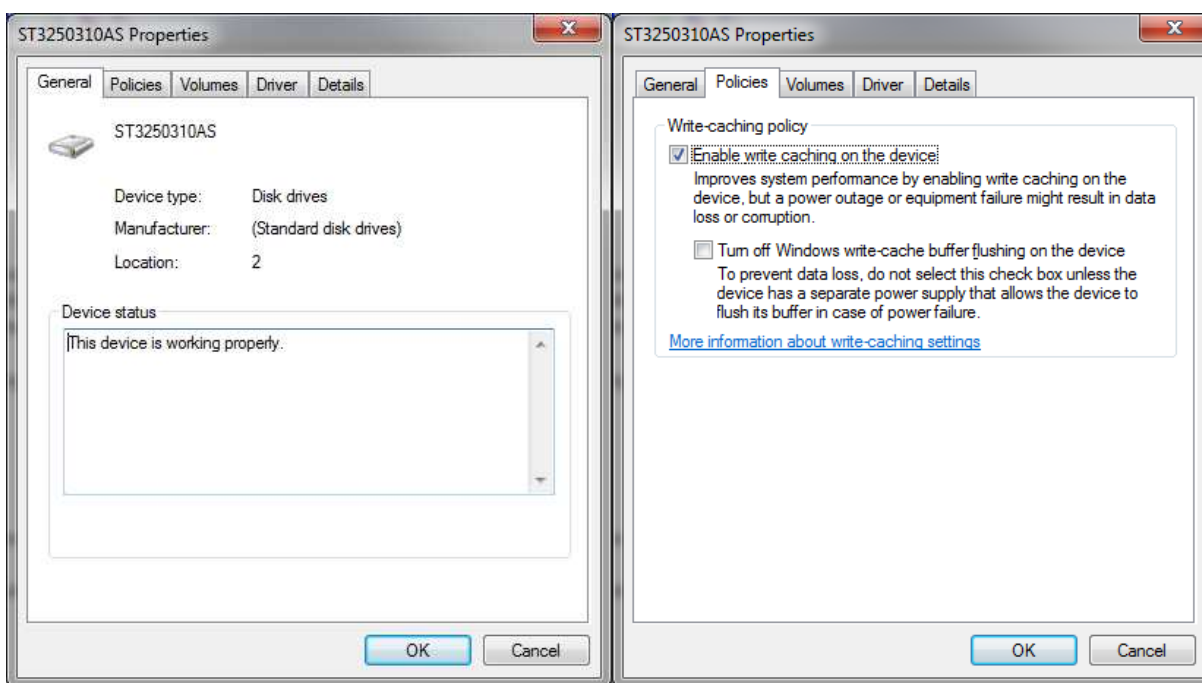
3.1. Microsoft Windows

For a 64-bit environment, it is necessary to enable the “Lock Pages in Memory” user rights policy, which may be done as follows:

- On the **Start** menu, click **Run**. In the **Open** box, type **gpedit.msc**. The **Group Policy** dialog box opens. The user account running this tool needs to be a system administrator.
- On the **Group Policy** console, expand **Computer Configuration**, and then expand **Windows Settings**.
- Expand **Security Settings**, and then expand **Local Policies**.
- Select the **User Rights Assignment** folder. The policies will be displayed in the right hand details pane.
- In the pane, double-click **Lock pages in memory**.
- In the **Local Security Policy Setting** dialog box, click **Add**.
- In the **Select Users or Groups** dialog box, add the user account that will be used to run the SERT (typically “Administrator”).
- **Restart the machine** for these changes to take effect.

3.1.1. Disabling the storage write cache on Microsoft Windows

For systems with Windows Server installed, each storage device will usually appear in the Device Manager. To disable write caching, select the desired storage device in the Device Manager, then either double click, or right click and select **Properties**. Select the **Policies** tab to display the interface that allows the disabling of write caching:




This (right) shows write caching enabled, click on the checkbox to disable it; e.g. the tick will disappear.


3.1.2. Disabling the Windows firewall

The firewall may be disabled by the following:

Open the Windows Control Panel and then click on the Windows Firewall applet.



Select Turn Windows Firewall on or off from the left hand side of the Firewall applet.



Click on the Turn off Windows Firewall for the relevant network(s); e.g. Domain, Private or Public, then click on OK.

3.2. Linux

3.2.1. Prerequisite Software

In addition to the base operating system and Java Virtual Machine, the SERT requires the numactl command. If this is not installed by default on your Linux distribution, use the appropriate commands to install the numactl package. For example:

```
yum install numactl (RHEL)

sudo apt-get install numactl (Ubuntu)
```

3.2.2. Enabling Transparent Huge Pages (THP)

Transparent Huge Pages (THP) must be enabled when running the SERT in a 64-bit Linux environment where THP is supported. Earlier versions of RHEL and SLES do not enable this feature by default, although other untested distributions may do so. To determine whether Transparent Large Pages are enabled for the system, run the following command:

```
"cat /sys/kernel/mm/redhat_transparent_hugepage/enabled" (RHEL)
-or-
"cat /sys/kernel/mm/transparent_hugepage/enabled" (SLES11 SP2 / Ubuntu):
```

The user should see the following output: "[always] never", with the always in brackets indicating that Transparent Large Pages are enabled by default on the system. If this file exists, but the output shows [never], or [mmap], please run the following command at startup to enable Transparent Large Pages:

```
"echo always > /sys/kernel/mm/redhat_transparent_hugepage/enabled" (for RHEL)
-or-
"echo always > /sys/kernel/mm/transparent_hugepage/enabled" (SLES11 SP2 / Ubuntu).
```

3.2.3. Disabling the storage write cache on Linux

For instructions on how to disable the different write caches, please read the manuals provided by the manufacturers of the system, controller and storage devices. Different versions of Linux provide different mechanisms for disabling disk caches, so it may be necessary to contact the manufacturer of your storage device(s) if neither the manufacturer's manuals nor the operating system documentation explain how to do this.

3.2.4. Disabling the firewall

Different Linux distributions use different commands to disable the firewall. In RHEL 6.x and SLES, issue the following command as root:

```
chkconfig iptables off
```

For RHEL 7 and other Linux distributions that use systemctl, issue this command as root:

```
systemctl disable firewalld
```

For Ubuntu 14.04 LTS, use:

```
sudo ufw disable
```

A reboot may be required after running these commands in order for them to take effect.

3.2.5. Linux on IBM Power Systems

Transparent Huge Pages are not supported on IBM Power systems running RHEL 6.x or SLES 11.x. For these distributions, large pages must be configured manually on systems with more than 32 GB of memory.

Sign on as the root user.

Download the newest libhugetlbfs library from the project home, at the following URL:

<http://sourceforge.net/projects/libhugetlbfs/files/>

Execute the command:

```
cat /proc/meminfo | grep MemTotal
```

This returns the total memory available to the operating system in Kilobytes [Value].

Compute the target large pages as follows:

```
targetLP = Truncate(Value/18204)
```

(Note: 18204 is a factor designed to allocate approximately 90% of available memory for large pages).

Execute the following command to set the number of large pages:

```
echo [targetLP] > /proc/sys/vm/nr_hugepages
```

Create the necessary pseudo-filesystem by executing the following commands

```
mkdir /libhugetlbfs
groupadd libhuge
chgrp libhuge /libhugetlbfs
chmod 770 /libhugetlbfs
usermod root -G libhuge (this assumes "root" is the SERT user on the system)
mount -t hugetlbfs hugetlbfs /libhugetlbfs
```

Edit `/etc/fstab` and add the following line

```
hugetlbfs /libhugetlbfs hugetlbfs mode=0770,gid=1000 0 0
```

Execute the following command to mount the huge table space:

```
hugeadm --create-global-mounts
```

3.3. IBM AIX

It is necessary to setup large pages correctly.

Sign on as the root user.

Execute the command:

```
bootinfo -r
```

This returns the total memory available to the operating system in kilobytes [Value].

Compute the target large pages as follows:

```
targetLP = Truncate(Value/18780)
```

(Note: 18780 is a factor designed to allocate approximately 87.25% of available memory for large pages).

Execute the following commands:

```
vmo -r -o lpgg_regions=[targetLP] -o lpgg_size=16777216 -o v_pinshm=1
bosboot -aD
shutdown -rF
```

This will cause the system to be rebooted with the large pages set up correctly for the SERT.

3.3.1. Disabling the AIX firewall

AIX does not have a firewall enabled by default. For information on configuring or disabling IPSec, see the AIX Knowledge Center:

https://www-01.ibm.com/support/knowledgecenter/ssw_aix_71/com.ibm.aix.security/ipsec.htm

3.4. 32-bit ARM Platform

The following items need to be considered for a 32-bit ARM platform:

The variable `ARM_JAVA_FP_ABI` in `serthost.sh` must be set appropriately for your Java version. The currently supported HotSpot 7 JRE uses the soft-float ABI so the variable should be set to "armel" in that case. For other JVMs please consult the vendor's documentation.

Since the supported Linux versions use the ARM hard-float ABI, while HotSpot 7 uses the soft-float ABI, several non-default libraries may need to be installed in Linux so that HotSpot will function properly. This may be accomplished via the command "`sudo apt-get install -y libc6-armel libsfgcc1`".

For processors running at speeds less than 2GHz, some of the worklets may require extra warmup time in order to reach steady-state performance before the measured intervals begin. Default warmup configurations on all worklets are a single 30-second warmup, but up to ten 30-second warmup periods are allowed. These values can be modified manually in `config-all.xml`; their names are `cpuWarmupIntervalCount`, `memCapacityWarmupIntervalCount`, `storageWarmupIntervalCount`, and `hybridWarmupIntervalCount`. 5 warmup periods should be sufficient for processors running at speeds over 1GHz.

3.5. Oracle Solaris

If the Host process is not run as root, discovery data will be incomplete, as some commands that the discovery process runs require elevated privileges. The `prset(1M)` command is used to effect process affinity, and that requires either root privileges or the `PRIV_SYS_RES_CONFIG` privilege. See `privileges(5)` and `usermod(1M)` for instructions on assigning that privilege to a normal user.

The Solaris virtual memory system will not overcommit memory, so it is necessary to configure more swap space than is default in order for the SERT Memory workload to run properly. This swap space will not actually be used, and will therefore not adversely affect scores; it is only necessary for very large memory allocations to succeed. Large memory systems may require swap space equal to the size of main memory, though on small memory systems it may be sufficient to configure swap space equal to about half the size of main memory. Disks used exclusively for swap space must also be included in the configuration for the Storage workload. It is recommended that those disks be partitioned with a small 20GiB UFS partition for SERT use, with the rest of the disk dedicated to swap.

The SERT Storage workload does not support measuring any kind of RAID or other disk aggregation. This includes ZFS in any multi-disk configuration. All ZFS pools on the system must consist of only a single disk. ZFS filesystems used for SERT Storage workload data set must have the "primarycache" and "secondarycache" properties set to either "none" or "metadata". The easiest way to accomplish this is to create a separate filesystem for the SERT Storage workload data set with a command such as:

```
zfs create -o mountpoint=/sert/data0 -o primarycache=metadata -o
secondarycache=metadata rpool/sert_data
```

When possible, UFS should be preferred as the filesystem to use for the SERT Storage workload. It provides significantly better performance than ZFS in no-cache conditions.

3.5.1. Disabling the storage write cache on Oracle Solaris

The cache on individual storage devices may be disabled using the expert mode of the `format(1M)` command. This command requires root privileges to run. For each disk attached to the system, run

```
format -e <path to disk>
```

Enter the 'cache' command, then the 'read_cache' subcommand, and then select 'disable'. Repeat for the 'write_cache' subcommand.

There is a shell script in the `tools` subdirectory called `solaris-disable-disk-caches.sh` which will perform these actions automatically for all disks attached to the system. It also reports whether the cache was successfully disabled for each disk.

Many Solaris-based systems ship with hardware RAID controllers with on-board cache. If it is possible to disable that cache, it will be done with an HBA-specific configuration utility such as MegaCli for LSI-based RAID controllers. Consult the documentation for your specific server and HBA model to determine how the cache can be disabled. If it is not possible to disable the cache on the hardware RAID controller, it may not be used for a compliant SERT run.

3.5.2. Disabling the Solaris firewall

The Solaris firewall may be disabled by issuing the following command as root:

```
sudo svcadm disable svc:/network/ipfilter:default
```

The state of the firewall can be verified using the `'sudo ipfstat -io'` command. When the firewall is disabled, its output will look like

```
empty list for ipfilter(out)
```

```
empty list for ipfilter(in)
```

3.6. Lower Performance Platforms

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 lower performing platforms to reach steady state before measurements start.

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

4. Power Analyzer Range Settings

Most power analyzers can be set to multiple current ranges. In order to achieve accurate results, the range settings on the device must be set appropriately for the load being measured. Most devices support an auto-ranging mode which will adjust the device automatically as the load changes; however, use of this mode is not allowed for measurement because it takes time for the range to adjust (during which the measured data may not be accurate), and because PTDaemon cannot determine the accuracy of the data while the device is in auto-ranging mode.

The SERT allows you to configure separate ranges for each power analyzer and can adjust the ranges as needed for each measurement interval. These settings appear in the `listeners.xml` file. Each `PowerAnalyzerListener` includes `"voltage-range"` and `"current-range"` parameters. The default configuration includes multiple current range settings for different load levels, but all are set to the same value: `"none"`. This indicates that no change will be made to the settings on the device during a test run.

The range settings may need to be adjusted when making changes to the system under test, the worklets being tested, or changes to the configuration of the operating system. The validation performed at the end of the run will determine if power measurement accuracy thresholds are not being met. If your result includes warnings regarding the power analyzer uncertainty, adjusting the current ranges may resolve the issue.

4.1. Automating Range Settings using the GUI

The **Choose Test Suite** panel gives you the option to select a “Range Discovery” run rather than the standard SERT full run. By selecting this option, the tool will run an abbreviated subset of the SERT worklets which should be sufficient to create range settings that will provide valid measurements in a majority of cases.

The steps to complete such a run are as follows:

1. Once at the “Choose Test Suite” panel, select the “Range Discovery” choice from the dropdown menu.
2. For most power analyzers, it is safe to check the “Use Auto Ranges” box and let the analyzer select the range choices for the run.
3. Complete the remaining settings on this page as documented previously in the User Guide.
4. Proceed to the launch page and launch the run.
5. When the run completes, you will see a message in the lower portion of the panel stating that “range settings have been updated”
6. Return to the PTDaemon configuration panel, stop the PTDaemon instance, and then review the range settings.
7. Return to the “Choose Test Suite” panel, select “Full Test Suite”, and continue on to make a full run.

Some power analyzers may be unable to select appropriate ranges automatically, and the range setting run will result in invalid values for ranges. In such a case, you may need to edit range settings from the PTDaemon panel, setting all range values to “none”, and then manually set the analyzer to a range that will measure the system at its highest load and rerun the Range Discovery test with “Use Auto Ranges” disabled.

The Range Discovery run should give working settings in most cases, but it is impossible to cover all system types and analyzer types. If range values generated by the run do not give valid measurement results for a full run, you may need to use one of the manual processes described below.

4.2. Manual Range Settings – Optional

To configure the ranges manually, replace each of the `<voltage-range>` and `<current-range>` entries in the `listeners.xml` file with a value appropriate for your system under test. The range setting should generally be slightly higher than the maximum expected reading during that interval. Each model of power analyzer will support a specific set of ranges; however, the configuration can use any value, and PTDaemon will round the range settings up to the next supported range.

4.3. Extracting Range Settings from Full Runs - Optional

The SERT also includes a process for extracting appropriate ranges for your configuration after a full run has been completed. After a run completes, the results files will be stored in a numbered directory in the results directory, for example `results/sert-0007/results.xml`.

Use the directory name from that test and run the following command:

```
reporter.bat -s -p -r results/sert-0007/results.xml -x
org/spec/power/sert/resources/rangeSettings.xsl -o rangeSettings.txt
-or-
./reporter.sh -s -p -r results/sert-0007/results.xml -x
org/spec/power/sert/resources/rangeSettings.xsl -o rangeSettings.txt
```

This will generate a `rangeSettings.txt` in the same directory as `results.xml`. This file will contain an XML listener specification for each power analyzer used in the range setting run. Open the file in a text editor and follow the instructions (which will be to copy some text from that file to the `listeners.xml`).

5. Using the Collector Scripts

The SERT package is distributed with two log collection scripts to assist in packaging up various log files and other critical configuration files for troubleshooting. The collector scripts are intended to be run on both the SUT as well as the controller system.

The two files are Operating System-specific and may be found in the base SERT installation directory:

`collectlogfiles.bat` is intended to be executed on a Microsoft Windows Operating System.

`collectlogfiles.sh` is intended to be executed on a Linux or AIX system.

In order to help SPEC diagnose a failed run, please follow these steps:

1. Ensure the run has completed and the GUI has been stopped on the Controller side.
2. On the SUT, open up a command line window, change directories to the base SERT installation directory, and execute either `collectlogfiles.bat` (for Windows systems) or `collectlogfiles.sh` (for Linux systems).

If you required to share your results files, please use the `/r` (Windows) or `-r` (other platforms) option, e.g. `collectlogfiles.bat /r` in order to package up the latest results files, various log files and other critical configuration.

If either of these script files are executed when the SERT has not yet been executed, when no results have been created or when some other part of the test pass has failed then you may see one or more (up to five) instances of the following message: The system cannot find the file specified.

On non-Windows systems where `bash` is not the default shell, it may be necessary to run "`bash collectlogfiles.sh`".

3. Once the script completes, navigate to the `SERTlog` subdirectory and copy the latest archive file (.zip for Windows, tar.gz for Linux) to the system for submission to SPEC for analysis. The files will be created in the following format:

`<hostname>_YYYY-MM-DD_HHMM.<zip/tar.gz>` (where YYYY is the year, MM is the month, DD is the day, and HHMM represents the hours and minutes count).

Always ensure the latest log archive file is captured, as a new version will be generated in the `SERTlog` subdirectory every time the `collectlogfiles` script is run on that system.

4. Follow the above procedure for the controller system, ensuring the correct script is executed for the Operating System running on the controller.
5. Use the form on <http://www.spec.org/sert/feedback/issuereport.html> in order to submit the files.

Note: Technical support requiring involvement of specialists is done by volunteers from our member institutions (<http://www.spec.org/spec/membership.html>). Please do not send us any proprietary information in your queries. If you feel you cannot clearly ask your question without reference to proprietary information please clearly mark such communications and we will not involve our SERT specialists, but rather will answer your query only if SPEC's own staff SERT generalists are able to do so.

6. SERT Command Line Usage

6.1. SERT Configuration and Start Procedure

This section provides details to enable a successful test with a simple server configuration as an alternative to running the GUI. Further sections provide additional details for some of the configuration settings that will be required in order to perform a successful test pass. It is recommended to read this section as well as the SERT Design Guide to understand more of the details behind SERT, and how to configure and run it. Please note that in the following instructions there are typically two script file options: *.bat for the Windows OS and *.sh for UNIX and Linux OSes.

- Configure the analyzer. Edit the power analyzer PTDaemon (runpower.bat or runpower.sh) and temperature sensor PTDaemon (runtemp.bat or runtemp.sh) script files. Ensure the proper communication ports and network ports are used. (The batch and script files are present in the SERT\PTDaemon directory).
 - For each additional power analyzer: Create a copy of runpower.bat or runpower.sh and ensure that the proper communication ports and network ports are set.
 - In order to analyze problems reported by PTDaemon it may be helpful to generate log files. To do this, edit the script files to set the LOGFILE and optionally the EXTLOG variables. LOGFILE should be set to the name of the file which will hold the log output. Log files may not be shared, so the names given to each must be unique. If extended logging of range settings, warnings, errors, and other important events is desired, set the EXTLOG variable to "TRUE".
- Start all power analyzer PTDaemon(s) (runpower.bat/runpower.sh) and temperature sensor PTDaemon (runtemp.bat/runtemp.sh).
- Edit the test-environment.xml file (in the SERT directory), describing the hardware and software details of the test setup. The default values in this file need to be changed to match the real test environment.
 - For each additional power analyzer: Add the following section to the test-environment.xml in the <MeasurementDevices> section and change the default values to match the real test environment.

```
<PowerAnalyzer>
  <PTDaemonHostname>localhost</PTDaemonHostname>
  <PTDaemonPort>8888</PTDaemonPort>
  <HardwareVendor> Energy Minder, Inc.</HardwareVendor>
  <Model> EM1000+ USB</Model>
  <SerialNumber> ser001122</SerialNumber>
  <Connectivity> USB2</Connectivity>
  <InputConnection> Default</InputConnection>
  <CalibrationInstitute> NIST</CalibrationInstitute>
  <AccreditedBy> IQ2 CalibrationLaboratory</AccreditedBy>
  <CalibrationLabel> N-32768</CalibrationLabel>
  <DateOfCalibration>2010-01-01</DateOfCalibration>
  <SetupDescription> Unknown</SetupDescription>
</PowerAnalyzer>
```

- For multi-node environments: the node quantity needs to be set. An example for a 3 node environment


```
<Node>
  <Quantity>3</Quantity>
```

- If the default port numbers in the PTDaemon script files have been changed, the `listeners.xml` file (in the SERT directory) must be changed accordingly. Also any range settings for the power analyzer can be specified in the `listeners.xml` configuration file.
 - For each additional power analyzer: Add the following section to the `listeners.xml` and change the default values to match the real test environment (**Note:** Do not forget to change the port number).

```
<listener enabled="true">
  <type>PowerAnalyzerListener</type>
  <classpath>
    <path>lib/ptdaemonClientApi.jar</path>
  </classpath>
  <parameters>
    <parameter name="hostname">localhost</parameter>
    <parameter name="port">8888</parameter>
    <parameter name="current-range">
      <range-set>
        <workloads>Memory</workloads>
        <worklets>Flood.*</worklets>
        <default-range>none</default-range>
      </range-set>
      <range-set>
        <workloads>Memory</workloads>
        <worklets>Capacity.*</worklets>
        <default-range>none</default-range>
      </range-set>
      <range-set>
        <workloads>Storage</workloads>
        <default-range>none</default-range>
      </range-set>
      <range-set>
        <workloads>Idle</workloads>
        <default-range>none</default-range>
      </range-set>
      <range-set>
        <default-range>none</default-range>
        <range level="87.5%">none</range>
        <range level="75%">none</range>
        <range level="62.5%">none</range>
        <range level="50%">none</range>
        <range level="37.5%">none</range>
        <range level="25%">none</range>
        <range level="12.5%">none</range>
      </range-set>
    </parameter>
  </parameters>
</listener>
```

- On the system under test, edit the SERT host script file (`serthost.bat/serthost.sh` – both be found in the SERT directory) for the appropriate system configuration. For single runs, the `KEEPALIVE` variable can be cleared; for multiple consecutive runs, it should remain at the default setting `-keepalive`.

Check the following specifications for proper values matching the test system configuration:

JAVA = <for a non-default Java Virtual Machine (JVM), change the java command to the pathname of the java command you want to run>

The `serthost.sh/.bat` script files currently both include the default definition:

```
JAVA=java
```

- On the system under test, run the `serthost.bat/.sh` script file. On Windows, this should run as an Administrator. On Linux, it should run as root. With Ubuntu (or other Linux distributions where a root login is not normally available), use "`sudo serthost.sh`".
- Check `client-configurations-110.xml` is up-to-date, and download the latest version from <http://www.spec.org/ser/ser/docs/client-configurations-110.xml> if necessary and copy to base folder. Compare the `version="YYYYMMDD"` string in the header section of the file with your local copy, where the version string specifies the date this file was last changed in the format year month day, e.g. 20130529.
- On the Controller running the Director, edit the SERT configuration file `config-all.xml` (in the SERT directory):

- Search for `id="clientConfig"` and replace the string "`defaultConfig`" with the configuration name from `client-configurations-110.xml` matching your "processor / OS / JVM / memory-per-client" combination (Use the guidance on the Client Configuration page http://www.spec.org/ser/SERT-JVM_Options.html for the correct configuration name.). Chauffeur picks the correct JVM heap sizes (parameters `-Xms` and `-Xmx`), number of JVMs and Java parameters based on the specifications in `client-configurations-110.xml` for the selected reference configuration. The computed values must be used for a valid measurement.
- Search for "`diskioFilePaths`" in the `<definitions>` section at the beginning. Edit the following lines:

```
<parameter id="diskioFilePaths">
    <path> [specify pathname of storage test device here]</path>
</parameter>
```

Change/remove/add path names of the storage devices to be tested.

Each path name should be listed in its own `<path>` tag.

- Edit the `director.bat/.sh` script files for single runs or the `start_director_multiple.bat/.sh` script files to start the desired number of SERT runs. First, the host IP address must be specified in `director.bat/.sh`, then edit the following line:

```
HOSTS=_hostname_or_IP_of_Host_systems_
```

- For multi-SUT environments: Add each SUT IP address (separated by a comma) to the HOST argument. **Note:** Do not use a space behind the comma(s). An example for 3 SUTS:
`HOSTS=10.10.10.1,10.10.10.2,10.10.10.3`
- On the controller system, call the `director.bat/.sh` script file for a single run or the `start_director_multiple.bat/.sh` script file for multiple consecutive runs.
- The duration of the whole test sequence depends on the number of test cases defined above. The result files for each SERT test run (`results.html` and `results.xml`) are stored in a separate directory with a name beginning with `sert-nnnn` e.g. `sert-0000`.
- On current releases of the SERT, a typical test pass exercising all worklets (CPU, memory, and Storage I/O) on a typical two-socket server with 16GB or more of memory and at least two on-board storage devices, may take four hours or more. This duration is governed by the amount of hardware present in the SUT, so more richly populated platforms may take longer to execute.

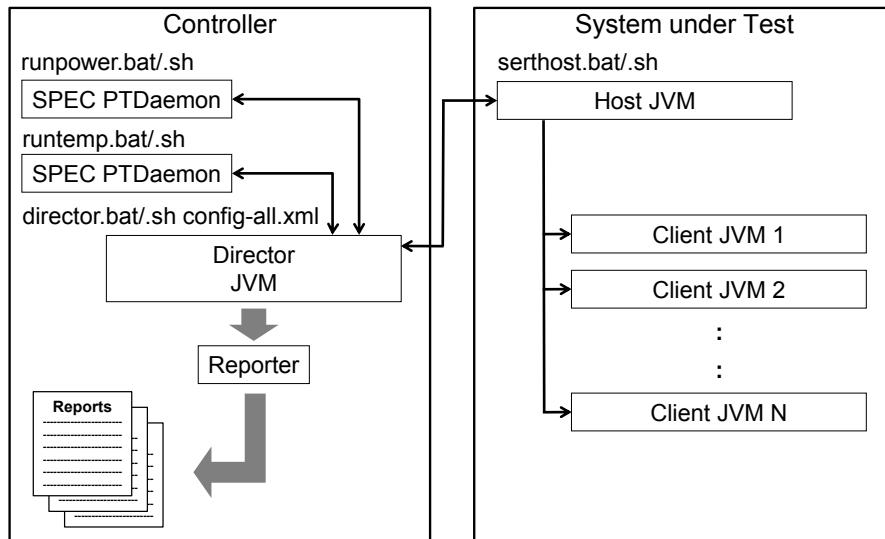


Figure 2: SERT Startup Procedure

6.2. Generate report files with the reporter scripts

In some cases the user might want to create a new .HTML, .CSV or .txt report from an existing SERT xml output file, for example after editing user-modifiable fields (configuration information in the TestEnvironment section).

In order to re-generate an .HTML, .CSV or .txt report from the raw files the following comment need to be used.

```
reporter.bat [options] {for Windows-based system}
reporter.sh  [options] {for Linux/UNIX-based system}
```

Required option:

-r [SERT xml results file]

Example: `reporter.bat -a -r results/sert-0002/results.xml`

Additional options:

-a	Generate all report formats (HTML, CSV, and Text) default: generate only HTML output
-c	Generate CSV report files
-o Output	Output file for generated report default: match the filename of the raw file, with a .txt, .csv, or .html extension (to match the report type). May not be used with the -a option
-p	Direct XSL transform (must be used with -x) default: generate a report from the transformed output
-s	Skip validation. Reports will automatically be marked invalid
-t	Plain ASCII text output default: generate HTML output
-x XslFile	Use a non-default XSL file to generate the report

Here are some examples of how to use the reporter scripts for regenerating default report files or for generating special output files:

- regenerate all HTML report files
`reporter.bat -r results\sert-0000\results.xml`
- regenerate all default report files
`reporter.bat -a -r results\sert-0000\results.xml`
- generate special CSV report files
`reporter.bat -s -c -x results\XSL-File\csv-detail.xsl -o results-
detail.csv -r results\sert-0000\results.xml`

Additional XSL scripts for generating other CSV report files are included in the `results/XSL-File` folder.

7. Coefficient of Variation (CV)

The coefficient of variation (CV) is defined as the ratio of the standard deviation σ to the mean μ :

$$c_v = \frac{\sigma}{\mu}$$

It represents the extent of variability in relation to mean of the population. The Host CV and Client CV are presented as percentage numbers in the SERT report.

7.1. Host CV

The Host CV is calculated from all individual node performance "Scores" for the whole set of nodes (population) is shown in the report for each interval. The Host CVs are calculated and displayed for multi node runs only, for single node runs a "0.0%" is displayed.

The Host CV is used by SERT to identify significant deviance of results of some nodes from the mean. Similar results per node are expected for a homogeneous set of nodes, i.e. the CV should be fairly small. High CV values indicate greatly different results for one or more nodes. CV values above certain thresholds will cause warnings or error messages in the SERT result files, and may lead to a run being invalid and thus unsuitable for submission.

7.2. Client CV

The Client CV is calculated from all individual client performance "Scores" for the whole set of clients (JVM instances) including all nodes for multi node runs (population) is shown in the SERT report for each interval.

The CV is used by SERT to identify significant deviance of results of some clients from the mean. Typically similar results per client are expected, i.e. the CV should be fairly small. High CV values indicate greatly different results for one or more clients. For unbalanced configurations, e.g. different types of storage devices (HDD and SSD) or dissimilar amounts of memory attached to different processors, high CV values are "normal". CV values above certain thresholds will cause warnings or error messages in the SERT result files, and may also lead to a run being invalid and thus unsuitable for submission. Warnings can be ignored for unbalanced configurations as high CV values are expected in this case.

8. SUT Minimum Memory Requirement

For the CPU and Hybrid worklets, Chauffeur allocates a maximum of 256 MB heap space per hardware thread, and reduces the heap if there isn't sufficient physical memory. Sufficient is defined using the following algorithm:

$$\text{heap space per thread (64bit)} = (\text{physicalMemoryGB} - 1) * 0.85 / \text{hardwareThreads}$$

$$\text{heap space per thread (32bit)} = (\text{physicalMemoryGB} * 0.875) * 0.85 / \text{hardwareThreads}$$

Example:

Configuration: 8GB physical memory, 32 hardware threads

$$\text{heap space per thread (64bit)} = (8 - 1) * 0.85 / 32 = 0.186 \text{ GB} = 190 \text{ MB}$$

In order to ensure that 256 MB will be used per JVM, we can solve the formula above for *physicalMemoryGB*, and we get:

$$\text{minimumMemory (64bit)} = 1 + (0.25 * \text{hardwareThreads} / 0.85)$$

$$\text{minimumMemory (32bit)} = (0.25 * \text{hardwareThreads} / 0.74375)$$

Example:

Configuration: 32 hardware threads

$$\text{minimumMemory (64bit)} = 1 + (0.25 * 32 / 0.85) = 10.4 \text{ GB}$$

Per thread heap space sizes below 256MB may result in lower performance for each client. The SERT harness only ensures that all clients will be using the same heap size to avoid performance differences between them.

The formulas given above don't pertain to the memory worklets. These worklets will use all available physical memory after subtracting the memory used by the operating system and other run time software (JVM etc.).

9. Known Issues

Please check for the latest updates <http://www.spec.org/sert/docs/SERT-FAQ.html>.

10. 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.

11. Copyright Notice

Copyright © 1988-2016 Standard Performance Evaluation Corporation (SPEC). All rights reserved.