SPECpower Committee

Server Efficiency Rating Tool (SERT) ™

Draft Design Document

Standard Performance Evaluation Corporation

Table of Contents

1 Introduction

1.1 Summary

The EPA's Energy Star Development Team is currently working on Version 2.0 of their Computer Server Specification[1]. Version 2.0 aims to evolve the program by adding a means to measure the overall efficiency of the server while it is performing actual computing work via an Active Mode Efficiency Rating Tool.

The SPECpower committee is currently working on the design, implementation and delivery of the Server Efficiency Rating Tool (SERT)™ next generation tool set that will measure and evaluate the energy efficiency of computer servers.

This first public draft outlines the design of SERT for review by EPA stakeholders and their associates.

1.2 About SPEC

The Standard Performance Evaluation Corporation (SPEC) was formed by computer industry participants in 1988 to establish industry standards for measuring compute performance. SPEC has since become the largest and most influential benchmark consortium. Its mission is to ensure that the marketplace has a fair and useful set of metrics to analyze the newest generation of IT Equipment.

The SPEC community has developed more than 30 industry-standard benchmarks for system performance evaluation in a variety of application areas and provided thousands of benchmark licenses to companies, resource centers, and educational institutions globally. Organizations using these benchmarks have published more than 20,000 peer-reviewed performance reports.

SPEC has a long history of designing, developing, and releasing industry-standard computer system performance benchmarks in a range of industry segments, plus peer-reviewing the results of benchmark runs. Performance benchmarking and the necessary work to develop and release new benchmarks can lead to disagreements among participants. Therefore, SPEC has developed an operating philosophy and range of normative behaviors that encourage cooperation and fairness amongst diverse and competitive organizations.

The increasing demand for energy-efficient IT Equipment has resulted in the need for power and performance benchmarks. In response, the SPEC community established SPECpower, an initiative to augment existing industry standard benchmarks with a power/energy measurement. Leading engineers and scientists in the fields of benchmark development and energy efficiency made a commitment to tackle this task. The development of the first industry-standard benchmark that measures the power and performance characteristics of server-class compute equipment started on January 26th 2006. In December of 2007, SPECpower_ssj2008 was released, which exercises the CPUs, caches, memory hierarchy and the scalability of shared memory processors on multiple load-levels. The benchmark runs on a wide variety of operating systems and hardware architectures. In version 1.10, which was released on April 15th 2009, SPEC augmented SPECpower_ssj2008 with multi-node support (e.g., blade-support).

1.2.1 SPEC Membership

SPEC membership is open to any interested company or entity. OAG members and associates are entitled to licensed copies of all released OSG benchmarks and unlimited publication of results on SPEC's public website. An initiation fee and annual fees are due for members. Nonprofit organizations and educational institutions have a reduced annual fee structure. Further details on membership information can be found on http://www.spec.org/osg/joining.html or requested at info@spec.org. Also a current list of SPEC members can be found here: http://www.spec.org/spec/membership.html.

1.3 The EPA's Energy Star for Computer Server Specification and SPEC

SPEC applauds the EPA for its goal to drive toward greater energy efficiency in IT Equipment, and SPEC considers the EPA Energy Star Program an industry partner in this effort. The development of

[1] US Environmental Protection Agency – Energy Star Program Requirements for Computer Servers http://www.energystar.gov/index.cfm?c=revisions.computer_servers

an Active Mode Efficiency Rating Tool is an essential component in the ongoing effort to reduce worldwide energy consumption and paves the way for a successful Energy Star for Computer Servers program that has the potential to harmonize energy efficiency programs worldwide.

SPEC welcomes this opportunity to work with the EPA on the Server Efficiency Rating Tool (SERT)™ in support of the Energy Star Specification for Computer Server and is proudly looking forward to continuing our long-standing association with the EPA Energy Star Development Team.

1.4 SPEC's General Development Guidelines

SPEC's philosophy and standards of participation are the basis for the development of SERT. The tool is being developed cooperatively by a committee representing diverse and competitive companies. The following guides the committee in the development of a tool that will be useful and widely adopted by the industry:

- Decisions are reached by consensus. Motions require a qualified majority to carry.
- Decisions are based on reality. Experimental results carry more weight than opinions. Data and demonstration overrule assertion.
- Fair benchmarks allow competition among all industry participants in a transparent market.
- Tools and benchmarks should be architecture-neutral and portable.
- All who are willing to contribute may participate. Wide availability of results on the range of available solutions allows markets to determine winners and losers.

Similar guidelines have resulted in the success and wide use of SPEC benchmarks in the performance and power/performance industry and are essential to the success of SERT.

1.4.1 Differences from Conventional Benchmarks

Performance benchmarks and energy efficiency benchmarks tend to focus on capabilities of computer servers in specific business models or application areas. SERT is focused on providing a first order of approximation[2] of energy efficiency across a broad range of application environments.

- The absolute score is less relevant for the end user, because it will not reflect specific application capabilities.
- A rating tool that provides a pass-fail or a [Level 1/Level 2/Level 3] pass-fail rating is a better fit for EPA's Energy Star Environment for Computer Servers than a typical benchmark result with multiple digits of precision in the metric.
- Marketing of the absolute scores will be disallowed, in order to encourage more participation in the program

Benchmarks tend to focus on optimal conditions, including tuning options to customize the configuration and software to the application of the benchmark business model. The need to achieve competitive benchmark results often causes significant investment in the benchmark process. SERT is designed to be more economical and easier to use, requiring minimal equipment and skills though:

- Highly automated processes and leveraging existing SPEC methods
- Focus on as-shipped default settings for the server
- Free from super-tuning

SERT development is planned to include input from a broad spectrum of industry experts using partnerships with EPA Industry stakeholders to provide ongoing feed-back during the development process

- Collaborate on workload, metric and all other requirements of the EPA's Version 2.0 Framework

[2] Andrew Fanara, Evan Haines, Arthur Howard http://www.energystar.gov/ia/partners/prod_development/downloads/State_of_Energy_and_Performance_Benchmarking_for_Enterprise_Servers_Final.pdf

- The feedback about SERT's design will be collected, sorted, prioritized and passed on to SPEC Development team via EPA's associates (servers@energystar.gov)

Where a benchmark represents a fixed reference point, Energy Star programs are designed to foster continuous improvement, with thresholds for success rising as the industry progresses. SERT will be designed to match this paradigm, including:

- Quick adoption of new computing technologies
- Rapid turn-around for SERT version updates

1.5 Scope

The current scope of Version 2.0 Energy Star for Computer Servers includes servers with 1-4 processor sockets with a stated goal to expand to include blade technologies of similar scope. A design goal of SERT is to accommodate these and larger technologies. However, schedule constraints and the complexity of testing and validating the tool for larger configurations may preclude their inclusion in an Energy Star program.

Among the issues involved with support of larger systems are the overall capacity of the system to complete work, the ability to design a workload that scales with the inclusion of additional processors, memory, network interface cards, disk drives, etc. Different workload characteristics are required to demonstrate effectiveness for each of these components. Providing a workload that fairly represents their presence while not unfairly representing their absence is a challenge. These issues are more prevalent with larger systems that have more expansion capabilities than smaller servers.

For these areas where the EPA and their supporting partners conclude that the tool does not adequately represent the value of a component compared to its power requirements, the tool will be designed to accommodate the inclusion of "configuration power/performance adders". A design goal is to automatically include this additional information in the computation of the Energy Star qualification results, including detailed documentation that this was done. Note that these will only be included if, through the data collection process the EPA determines that "configuration adders" are necessary.

It is also possible that the workload mix that is defined for smaller systems will not scale well when examining larger systems.

1.6 Overview Summary

The following table summarizes some of the design goals that SERT will and will not provide.

IS	IS NOT
Rating Tool for overall energy efficiency	Benchmark
Measuring tool for power, performance and inlet-temperature	Measuring tool for Airflow, Air pressure, outlet-temperature
General compute-environment measure	Specific application benchmark measure
Support of AC- powered servers	Support of DC-powered servers (see also 4.5)
Used in single OS instance per server environments	Used in Virtualization environment
Energy Star Rating Tool	Marketing Tool

1.7 Scope / Schedule tradeoffs

SPEC benchmarks are developed with the goal to generate results which are directly comparable for multiple hardware and software architectures to the extent this is possible. The same basic goal directs the design of SERT as specified in this document.

Even though SERT is designed with the goal of being architecture agnostic, code needs to be implemented for each of the workloads and the tool harness on all supported architectures. Furthermore this code must be tested intensively on all architectures in order to ensure a functionally equivalent set of binaries, which generate fair and comparable results. Simply using a portable programming language will not be sufficient to achieve these goals. Consequently significant complexity is added to the development process.

Given that SERT is designed as a first order approximation rating tool, comparability may be handled differently than with benchmarks (second order approximation tools) which are used for competitive marketing. Nevertheless it's essential to ensure a minimal level of comparability.

The resources available in the SPECpower committee are limited and a timely development of the tool for a single architecture will be challenging. Support for additional architectures will remove resources from the development of the basic test routines because they will be needed for porting the code. Furthermore additive testing effort is required not only for the new architectures but for the original implementation as well in order to ensure comparability. Therefore each extra architecture will add a currently undetermined amount of time to the schedule.

The resource and schedule problems recur with the support of multiple operating systems. SERT will be initially implemented on selected Operating Systems (OS) per HW architecture.

If support for additional architectures or OSs is desired, then active participation from requesting entities is mandatory. Companies dedicating additional resources to the SPECpower committee for development of SERT would relax the schedule constraints.

The combination of several architectures and various operating systems would result in a significant increase of testing efforts and postpone the release date considerably. An estimated schedule can be created once we have decided on all design details.

2 Server Efficiency Rating Tool (SERT)

2.1 Overview

SERT is composed of multiple components including a test harness, workloads, and a system under test (SUT). The test harness handles the logistical side of measuring and recording power data and includes the measurement hardware, control system, and the controlling software installed on the SUT and controller system. The workloads exercise the SUT while the test harness collects the power data.

2.2 SPEC PTDaemon

SPEC's power/temperature daemon (also known as PTDaemon, PTD or ptd) is used by 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.

The SERT harness connects to PTDaemon by opening a TCP port and using the simple commands detailed in the API section of this document. For larger configurations, multiple IP/port combinations can be used to control multiple devices.

PTDaemon can connect to multiple analyzer and sensor types, via protocols and interfaces specific to each device type. The device type is specified by a parameter passed locally on the command line on initial invocation of the daemon.

The communication protocol between the SUT and PTDaemon does not change regardless of device type. This allows SERT to be developed independently of the device types to be supported.

2.3 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.

The following environmental conditions need to be met:

- Ambient temperature range: 20°C or above
- Elevation: within documented operating specification of SUT
- Humidity: within documented operating specification of SUT

2.4 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 benchmark:

- Logging - The sensor must have an interface that allows its measurements to be read by the benchmark harness. The reading rate supported by the sensor must be at least 4 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 benchmark run.

2.5 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.

2.5.1 Power Analyzers and Temperature Sensors

SERT will utilize SPEC's accepted measurement devices list and SPEC PTDaemon update process. See Device List (http://www.spec.org/power_ssj2008/docs/device-list.html) for a list of currently supported (by the SPEC PTDaemon) and compliant (in specifications) power analyzers and temperature sensors.

2.6 Run Time

The run time will depend on the agreed set of worklets. The target run time is less than 2 hours.

2.7 Test Harness

SERT will require a test harness configuration very similar to that of SPECpower_ssj2008, with a controller system, power analyzer, temperature sensor, and with the likely addition of a networked client system to generate different types of workload traffic for the server under test (SUT).

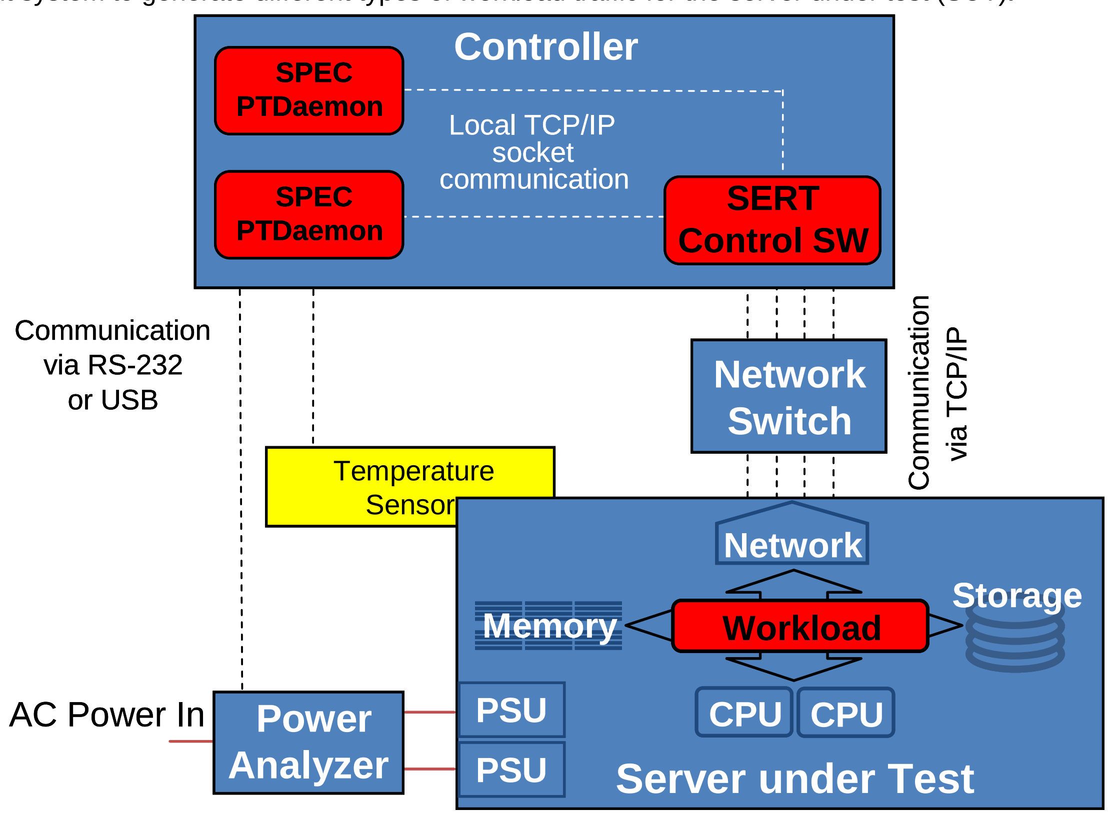

2.7.1 Graphical User Interface

One of the stretch goals is the incorporation of a graphical user interface (GUI) to facilitate configuration and setup of test runs, allow real-time monitoring of test runs and to review the results. The SERT GUI will lead the user through the steps of detecting or entering the hardware and software configuration, setting up a trial run or a valid test, displaying results reports and other functions common to the testing environment.

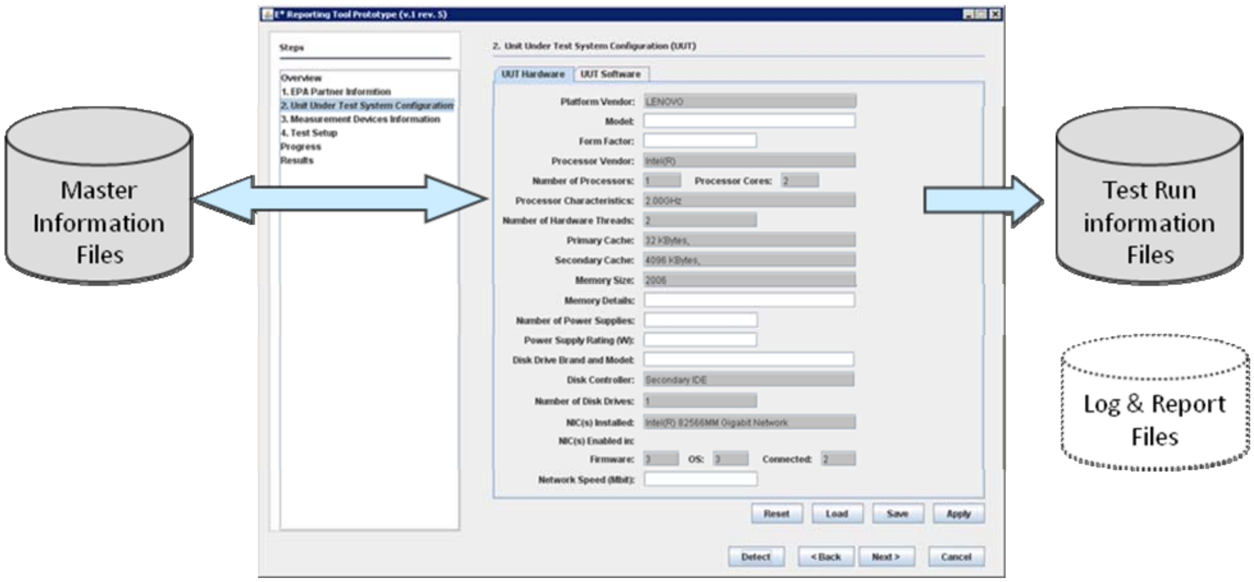

The SERT GUI will include several features to enable SERT testing with minimal training and enhance the accuracy of results:

- Easy Navigation with Tabbed Screens
- How to Use (in-line usage guidance and help)
- Configuration Discovery (Detect function) will automatically populate most fields about SUT and Controller hardware and software.
- The GUI will display, allow entry of and store required information about the test environment
    - For use in reports: e.g. Company Info, Platform Config, Run-Time parameters, etc.
    - Master and Test Run information files can be stored, enabling reuse, saving time with multiple platforms.
- Test Setup, Execution and Progress Display
    - Start measurements; Choose type of run (trial or final)
    - Display progress, warnings and errors.
- Display results and enable printing and capture of reports
- Provisions for redundant components and power and performance adders.

2.8 Workload

The existing SPEC benchmarks are mainly based on tailored versions of real world applications representing a typical workload for one application area or a synthetic workload derived from the analysis of existing server implementations. These benchmarks are suitable to evaluate different sub-areas of the overall server performance or efficiency if power measurements are included. They are not designed to give a representative assessment of the overall server performance or efficiency.

The design goal for the SERT workload however is to include all major aspects of server architecture, thus avoiding any preference for specific architectural features which might make a server look good under one workload and show disadvantages with another workload.

The SERT workload will instead take advantage of different server capabilities by using various load patterns, which are intended to stress all major components of a server uniformly.

If some components cannot be stressed adequately by the respective load pattern this can be compensated by adjusting the threshold for these components, e.g. increasing the power allowance for additional components which are not used by the load pattern.

It is highly unlikely that a single workload can be designed which achieves the goals outlined above, especially given the time constraints of the schedule targeted for Energy Star for Servers Version 2.0 by the EPA. Therefore the SERT workload will consist of several different worklets each stressing specific capabilities of a server. This approach furthermore supports generating individual efficiency scores for the server components besides the overall system score.

2.8.1 Worklets

Developing the workload in the traditional SPEC way based on real world applications would result in complex test environments and high run times, especially for the IO intensive workloads, e.g. many client systems would be required for network IO and large disk sub systems for disk IO. The resulting costs for running such tests could be prohibitive for a rating tool. Therefore the SERT workload will be a collection of synthetic worklets for a variety of different load scenarios.

In order to achieve consistent results from all worklets and a broad coverage of technologies the following guidelines need to be observed:

- Each worklet must be adjustable to different performance levels, e.g. some predefined levels between 100% (maximum load) and 0% (idle)
- Each worklet must calibrate to maximum performance level by itself, i.e. no definition of the 100% level by the test user
- Multiple programming languages may be used
- Precompiled binaries of the test programs should be used where possible.
- Each worklet should scale with the available hardware resources. More resources should result in a higher performance score, e.g. more processor/memory/disk capacity or additional processor/memory/disk modules yield a better result in the performance component of the efficiency rating.

In order to follow these guidelines the workloads will probably be based on batches of discrete work, where each batch constitutes a transaction. The different load levels will be achieved by scheduling the required number of transactions.

Each worklet should perform specific tasks as in the following examples. The definition of such building blocks is not finalized and needs further investigation. Regardless, they should cover a large variety of functions from low-level program primitives to more complex transactions.

2.8.2 Worklet Examples

- **CPU**
  floating point matrix multiply, business logic simulation, data compression / encryption
- **Memory**
  throughput, latency, read, write
- **Network IO**
  throughput, latency, streaming, packet read/write
- **Storage IO**
  throughput, database access pattern simulation, file read/write

2.8.3 Workload terms

	Example
Workload	All CPU, Memory, Disk and NIC worklets at all load levels
Worklet	Disk bandwidth test at all load levels
Transaction	memory read + modify + memory write
Operation	Disk write

The SERT workload consists of multiple worklets. Each worklet will execute a series of transactions which are comprised out of various operations.

2.8.4 Worklet Execution

In the current design the worklets will run consecutively, each in its own phase. This allows generating independent scores at different load levels which can be combined to an overall metric.

Selected worklets or additional worklets specifically designed for concurrent execution may be run simultaneously in an extra phase. Concurrent execution of worklets will introduce more realistic task-switching, especially using IO load modules.

2.8.5 Load Levels

Multiple load levels are a desired goal of SERT and the design will include support for multiple levels. The active idle load level as well as a 100% workload level (not max power) are already good candidates. Prototype testing will show which levels will be included and if any weighting will be necessary.

2.8.6 Scaling

Since the server efficiency rating of a given server is the primary objective of SERT, one of the main design goals for SERT is to be able to scale the performance on the system in proportion to the system configuration.

As more components (processors, memory, disk storage, network interfaces) are added to the server, the workloads should utilize the additional resources so that the resultant performance is higher when compared to the performance on the same server with a lesser configuration. Similarly, for a given server, when the components are upgraded with faster counterparts, the performance should scale accordingly.

This is a very important aspect of the tool since adding and upgrading components typically increases the total power consumed by the server which will affect the overall efficiency result of the server.

Creating a tool that scales performance based on the number/speed of CPUs is most readily achievable – for the other components, the complexity of implementing such a tool increases substantially.

While SERT will be designed to scale performance with additional hardware resources of the SUT, if there are performance bottlenecks in system components unrelated to the added hardware the SUT itself may not be able to sustain higher performance. In such cases the addition of components to the SUT will normally result in higher power consumption without a commensurate increase in performance.

2.9 Server Options and Expansion capabilities

A server may have many optional features that are designed to increase the breadth of applications. These features not only absorb additional power, but also require more capacity in the power supplies and cooling system. Some SERT workload components will be designed to demonstrate the enhanced capabilities that these features provide. However, while the tool needs to credit these capabilities for the expanded workloads that they will accommodate, it cannot penalize efficient servers that are not designed with substantial expansion options. A balance must be struck between providing enhanced ratings for enhanced configurations and avoiding easy qualification of servers by simply adding features that may not be needed in all situations.

SERT's goal is to avoid unnecessarily penalizing servers that are designed for low expandability, while crediting servers with greater expandability. For example a configuration with four I/O adapters in PCI slots may execute the workload of the tool more effectively than a configuration with only one such adapter. On the other hand it may only run the workload of the tool as effectively as a configuration with two network adapters. Because the configuration with four adapters may run some real workloads more effectively than configurations with only two adapters, the EPA may elect to allow for some form of "configuration adders" to provide credit for the power infrastructure needed to support the additional PCI slots.,

While the decision to include this form of power credit and the value of any such credit will be in the hands of the EPA, the tool will be designed and tested to ensure that, should "configuration power-performance adders" credits be included, the tool will accommodate them.

2.9.1 Redundancy

Many servers have redundancy built in for power supplies and cooling fans. Some servers include different levels of redundancy for memory, disk, and even processors. A design goal is to include accommodation for redundant components, although no specific tests are planned for energy measurement under fault tolerant conditions when one of a redundant set of components is disabled.

2.9.2 IO Component

Disk and Network IO component are strongly desired to provide a better-rounded picture of system performance and power than a CPU-centric test. SPEC is in the early stages of evaluating IO workloads for SERT, so this section provides many discussion points but not necessarily conclusions.

SPEC recognizes that some of the items in the next two sections may not be reasonable or practical to test or measure in a meaningful way. In those cases we would suggest the use of "configuration power-performance adders" to compensate for the extra power draw associated with extra functionality.

Other items under consideration include:

- Different types/quantities of IO for different server categories
- Self-calibrating performance measurements for the disk and network subsystem

2.9.3 Disk IO

Ideally the disk IO component of SERT would give credit for:

- Higher performance disk subsystems
- Larger capacity disk subsystems
- Reliability and availability features (RAID, battery backed cache, etc)

2.9.4 Network IO

Ideally the network IO component of SERT would give credit for:

- Higher performance Network Interfaces
- Larger transfer speed Network Interfaces
- Reliability and availability features

2.10 Metric/Score

While SERT is not intended to be a benchmark, nevertheless as a rating tool it must produce a metric or score indicative of the efficiency of the server under test. That metric must combine both the performance of the SUT as well as its power consumption in a way that allows comparison among all systems subjected to it. The desired outcome of that comparison is a quantitative measure of the relative power-performance efficiencies of the systems. The system which produces the higher metric should have greater power-performance efficiency than the system which produces the lower metric.

Since different architectures perform differently on different workloads, SERT may be composed of several discrete worklets to ensure architecture neutrality. Each worklet will produce a measure representing the performance achieved by the SUT, which then must be combined with the measures produced by the other worklets to yield a metric indicative of the overall performance of the SUT on all worklets used in the tool. SPEC recommends that the multiple performance measures produced in this manner be combined into a single metric as the geometric mean of the individual measures.

The geometric mean of individual worklet performance may be used whether the individual worklets are run sequentially or simultaneously. Depending on the worklets chosen and the magnitudes of their individual measures, we intend indexing the measures to a set of reference scores before combining them into the single metric as the geometric mean. These techniques have the advantages of rendering the single metric unit-less, and of keeping the scale of the individual measures within similar ranges, so that a worklet with large magnitude individual measure does not overwhelm the result from a workload with a smaller measure.

Once determined, the overall performance must be combined with the measured power consumption of the SUT in a way that demonstrates the power-performance efficiency of the system. This will be a complex calculation automatically performed by SERT to take into account the power-performance efficiency of the SUT at different utilization levels.

The metric that is produced by SERT is separate from the Energy Star rating. The EPA will determine criteria for Energy Star acceptance of which the SERT scores may be only a part. It's anticipated that the top 25% of tested units will achieve Energy Star qualification. A "gold-level" Energy Star qualification may be available for units achieving in the top 5% of results. Additionally the EU has proposed a system of graduated achievement in power-performance efficiency with levels A through F, for which they will determine the overall criteria.

Server under test may be placed in different categories by the EPA. The EPA will decide how to apply these categories and whether units in a particular category may be compared to units in another category.

2.11 Reporting and Output Files

SERT will produce two reports and a set of log files. The reports will be created in XML format, in order to reduce the effort for both EPA and the partner in displaying and or storing the desired information. We will take steps in order to ensure authenticity (e.g. encryption) of the reports.

2.11.1 Report 1: "Summary Report"

This report will contain a placeholder for a "pass or fail" notice for the tested platform, to be provided by the EPA. A test run is marked non-compliant if the test completes with technical errors. In such a case, error messages and/or warnings will be automatically included in the report. The information in this report is public and could be used for marketing purpose.

Items included in this report are:

- EPA Partner name and EPA Partner ID
- EPA Energy Star Category of the tested platform
- Test Date and Location (plus "Tested by")
- Tested Platform Manufacturer and Model Number
- Placeholder for "Pass/Fail"
- Warnings or Error Notices if applicable
- System Configuration information (Redundant components to be marked appropriately):
    - form factor
    - number and type of processors
    - available processor sockets
    - memory size, type, # memory DIMMs, # DIMM Slots, Max Memory Capacity
    - available expansion slots
    - number of and make-model of power supply, output rating, min/max
    - Input power

- OS supported / OS used for test
- number of and make-model of storage controller
- number of and make-model of mass storage devices
- number of and make-model of network interface cards (NICs)
- Management Controller or Service Processor Installed? [Yes/No]
- Other Hardware Features / Accessories

2.11.2 Report 2: “Power and Performance Data Sheet”.

This report will contain all the information the EPA requires and that is deemed necessary by SPEC. The Power and Performance Data Sheet will be public, but marketing use is prohibited by SERT Fair Usage Rules. The information is intended to be delivered to the EPA in a form most expeditious for EPA review.

This report will contain all the data from the “Summary Report” with the following additional detail sections:

- Overall Result / Score
- All target load level results
- Hardware and Software Configuration
- Power Measurement Summary
- Environmental information

2.11.3 Log files

A set of log files will be produced for each test run.

- The information in the log files is intended to be “non-public”.
- These files will be identified by a run serial number such that multiple consecutive test runs produce multiple log file sets.
- Each log file will be a record of actions from the software during the various phases of the testing, including errors and warnings.
- The intent of the log files is for auditing and support purpose.
    - Problems or failures can be more easily resolved with this low level detail record. If any issues arise with regard to the accuracy or veracity of the partner reports, these log files (potentially encrypted) should be adequate to resolve most issues.
- Examples of log file content are:
    - Handshake validation messages among various components
    - Error or warning messages
    - State change messages/notifications.
    - ‘Transaction’ instantaneous/periodic summary information
    - ‘Transaction’ response times

The EPA may require that any or all of the above outputs be delivered prior to Energy Star qualification. Regardless, the partner must commit to archiving all output from any results submitted to the EPA.

2.12 Test Software

Another “stretch goal” of SERT is to enable a “Live CD” approach to tool installation, for some environments – such that the entire tool suite along with the underlying operating system could all be run from a single bootable CD or DVD with no other operating system installed on the SUT. This should provide increased ease of installation and improve the adoption rate of the tool.

Possible issues with this approach include the lack of specific hardware drivers for newer devices, the potential lack of vendor specific power management, licensing and availability issues for some operating systems. Alternatives include allowing additional drivers to be installed during setup, or providing separate test installers with binaries for use with a vendor's own as-shipped OS installation.

2.13 Validation / Verification

SERT software components will implement software checks wherever possible to increase information accuracy, verify user input, monitor run-time data collection, and validate results with the intent of improving accuracy and remedying user errors, preventing invalid reports to the EPA.

When conditions or results do not meet specific criteria, warnings will be displayed and error messages will appear in the SERT reports.

These features will make it easy for the EPA Partner to generate compliant results and prevent submission of erroneous reports to the EPA.

Examples of compliance checking are:

- Verify input properties (parameters) and run-time duration of load levels.
- Temperature out of range will be reported.
- Power and Temperature read errors must be under a chosen threshold.

All the SERT software components will perform validation checks within the domain of their functions, e.g. warnings of connection problems, log measurement errors and out-of-range conditions, warning the user of missing or incomplete information and check the validity of some entered data.

Other new validation methods will be considered as the SERT software design and implementation progresses.

3 Logistic

3.1 Ownership

SPEC retains ownership of the SERT code and has the final sign off on the code. SPEC is committed to working with the EPA during the development of SERT to ensure that code meets the industries requirements.

3.2 Support

The end user support logistics and scale involved with Energy Star programs are beyond what SPEC's volunteer development team can sustain and support. Thus, SPEC should only be considered as a second line of support reserved for fundamental SERT issues only. The EPA will be responsible for providing direct customer support for any Energy Star program which utilizes SERT.

3.3 Maintenance

SPEC will maintain the SERT code base.

3.4 Updates

The initial version of SERT is targeted to support EPA E* Tier 2 requirements. After the release of the initial version, SPEC will continue to develop SERT in two ways.

- Fixes for discovered bugs and feature enhancements or additions will be provided as revisions of the current version. Licensees of the current version of SERT are entitled to these bug-fix and/or enhancement revisions.
- SPEC will create new versions of SERT to satisfy the requirements of later E* Versions. End-user use of new versions requires license purchase.

SPEC will cooperate with the EPA on the logistics of distributing updates.

3.5 Licensee and price

SERT will have a SPEC license and fee structure for end user companies. As with other SPEC products, all SPEC OSG member organizations are entitled to a complimentary SERT license and free version and revision updates in consideration of their development costs and contributions.

Additional potential EPA E* partner companies are entitled to purchase their initial SERT license at a cost anticipated to be on the order of $1000. As new E* Versions are announced and SERT is updated, new versions of SERT which support the latest E* Version may be purchased by licensees of the immediately preceding version at the anticipated cost of $500.

When available, bug fix or enhancement revisions to the current version will be provided to current licensees free of charge.

3.6 Delivery

SPEC will delivery SERT via CD/DVD to the EPA Partners.

3.7 Submission Review Process

The EPA will be responsible for reviewing submissions for acceptance for the Energy Star for Server program. SPEC will not review Energy Star for Server submissions. SPEC is committed to automating SERT extensively with the end goal of minimizing the review process.

3.8 Trademark

Product and service names mentioned herein may be the trademarks of their respective owners.

4 SERT and EPA Energy Star for Server Version 2.0

In order to ensure that SERT is utilized in the intended matter, we recommend the inclusion of the following items in the Energy Star for Server Specification.

4.1 Tested as Shipped

To provide results that are representative of a customer environment, the goal is to test systems in an "as-shipped" state. No super tuning would be allowed, but rather a limited list for tuning/optimization might be permitted. This list would be agreed on before SERT release.

4.2 Fair Use of SERT information

A clear goal of the Energy Star program is to have the broadest possible participation among vendors. Experience in the computer industry's performance benchmark community demonstrates that when performance details become available for marketing purposes, only vendors with superior (at the time of publication) products are incented to publish results. To encourage broader participation across the industry, a set of strong rules must be in place that will restrict marketing use of any of the detailed information generated by the tool. No data besides the actual Energy Star qualification should be utilized in EPA Partners' marketing collateral. These rules will be stipulated in both the license for the tool and the EPA Partner agreement.

Note that, while these rules are not strictly a part of the tool "design", the existence of these rules are necessary to allow the flexibility of the design and the delivery of detailed consumer information that is desired.

4.2.1 Fair Use Rules

- The only information provided by the tool that can be used for marketing collateral is the Energy Star qualification of a server configuration or server family
- The only information provided by the tool that can be used for public comparison is the Energy Star qualification of a server configuration or server family All other publicly available information from the tool is made available to help to verify that the tests were run correctly and to allow consumers to better understand how well the configurations tested match their specific needs.
- If the tool is used for research to generate information outside of the Energy Star program, the information may not be compared to the Energy Star program results and competitive comparisons may not be made using the data generated.

- The EPA Energy Star Qualification is governed by EPA rules.

4.3 Accredited, Independent laboratory

The requirement to use accredited, independent laboratories may place a large burden on EPA Energy Star partners, especially smaller companies. We recommend the use of an independent laboratory as an option, but not implementing this as a requirement.

4.4 Supply Voltage tolerance

In order to use a voltage within a 1% difference, an extra voltage source is needed. This will unnecessarily increase the cost for the partner, especially smaller companies. We recommend the tolerance be set to ± 5%.

4.5 DC Line-Voltage

The SPEC PTDaemon does not support DC measurements and currently no resources are devoted to including this support. It is our opinion that comparing servers powered by AC against servers powered by DC is not fair, since the AC-DC conversion losses are not included in DC-powered server. We would recommend creating a separate category for DC-powered servers.

4.6 Airflow

Regarding the requested thermal flow rates of servers (Total Power Dissipation; Delta T at Exhaust at Peak Temperature; Airflow at Maximum Fan Speed at Peak Temperature; Airflow at Minimum Fan Speed at Peak Temperature). SPEC recommends deleting this requirement in order to enable smaller companies to become an Energy Star partner as well.

We believe that a unified reporting of the thermal flow rate across the industry could help facility managers to plan datacenters better with the additional thermal load of the servers in mind. Unfortunately, the volumetric flow depends on the configuration of the server and the EPA's concept of Product Family could not be used. Also we believe that the guideline described is not strict enough to ensure unified measurements and reporting across the industry and the additional cost for each partner needs to be considered as well. In order to qualify a server, a volumetric airflow bench (neither a simple wind tunnel nor an anemometer) is necessary (Cost ~$5000) as well as a temperature chamber (Cost ~$5000). We estimate the time spent to produce these results is around 8 hours for each configuration. SPEC recommends deleting this requirement in order to enable smaller companies to become an Energy Star partner as well."

