IBM Comments:
SERT Draft Design Document

IBM has the following comments on the SERT Design Document:

Page 7, Lines 140 to 148; Page 5, Line 67: We support SPEC and EPA's statements that the SERT® Metric needs to be both hardware and operating system agnostic, with the ability to test and report on the key server architectures available in the marketplace using at least two available and agreed to operating systems. As we are an active participant in the development of the SERT Metric, we do not agree with the statements SPEC has made in the document suggesting that it will be necessary to limit the applicable architectures in order to deliver the SERT metric on a reasonable schedule that will align with publication of a final Tier 2 specification. We recognize that only those architectures and OS combinations for which there are volunteer contributors to develop and test the implementation can be completed, but we believe that there has been a sufficient commitment of resource to meet SPEC's and EPA's goal to be architecture and operating system agnostic.

Page 7, Lines 151-153: The design document should list the operating systems being used in the development of the SERT metric. It is our understanding that the two chosen operating systems are Windows and AIX and we think this is sufficient to support and test the currently available hardware architectures.

Page 9: lines 250-262 and Section 2.8: We support the concept of worklets, or component specific workloads, that stress various computer components in specific ways so as to demonstrate their value in a configured server. It is important to select functions that are common across architectures and operating systems. A routine that exercises a specific X86 instruction set is not acceptable. It is acceptable, where necessary, to develop worklets or work routines that can be run entirely in hardware on X86 but requires a combination of hardware and firmware to run on architecture, as long as both architectures deliver similar results. The important point is to provide a system stressor that exercises specific server function regardless of the architecture or operating system. It is also important to note that the nature of the worklets may be such that the definition of scaling from 0-100% might not include scaling the processor to 100% on a specific worklet or routine, since some worklets may be constrained by memory or bandwidth rather than by processor capacity. However, these limitations can be addressed in the scoring functions. We believe this is a good course of action, if it can be properly addressed when the final composite rating is created.

Overall we support the course and direction of the SERT development and believe it has the potential to deliver a workable metric for evaluating system performance/power characteristics for the ENERGY STAR® server requirements.

Page 11 to 12, lines 346 to 356: We do not think that an "adder" system to compensate for the presence of certain components is a foregone conclusion. We believe it is important to complete the metric and test various server configurations to determine how best to apply the metric to determining the performance power capabilities of different servers. We believe that there are at least three options for how the metric might be applied:

1. The metric may normalize for the different power/performance characteristics of different server configurations, allowing for a single rating system for all servers of 4 processors or less.
2. The collection of data for the SERT metric may show that are distinct categories of servers, either the current product differentiations or categories established by the number of processor sockets and rack or blade configurations, as well as possible differentiations within specific product categories, which should be rated by the SERT metric.
3. It may be necessary to create power or performance factors for certain component configurations to truly normalize the scoring.
4. Some other approach which cannot be identified until metric data is collected from actual systems.

We do support SPEC's efforts to enable the SERT metric to invoke adders if they are deemed necessary. However, we encourage EPA and the industry to develop sufficient metric data to determine the simplest and most elegant product categorizations to simplify the execution and understanding of the results for the SERT metric evaluations.

Page 12-13: The SPEC Draft Design Document provides very little detail about the computation of the SERT metric. We support the need for a single, combined metric from all workloads exercised by the tool. While we recognize that additional work is required to finalize the algorithms, we think it is important to publicize the statistical methods and algorithms chosen to process the raw data and generate the final scores for each worklet or workload routine and the composite system score.

Until the worklets are defined and a mass of data is accumulated, we believe that it is premature to select the geometric mean as the combinatory method. The geometric mean has the effect of making each component equally important in the measure - for example the CPU component would be equal to the network I-O component, even though the processors might account for 70% of the total power of the system and the network components only 10%. Other methods such as arithmetic mean, weighted arithmetic mean, average ratio to a reference machine, etc. should be considered as the method is developed. It may be appropriate to delay the selection of the statistical basis of the final scoring metric until test data can be gathered to allow assessment of the scoring options against their ability to generate or conform to defined product groupings.

It will be important for all stakeholders to have an opportunity to assess the chosen statistical approaches and the robustness of the scoring system before it is finalized within the metric.

Page 14, Section 2.12: We do not think that the "live CD" approach is a critical component of this effort and if it represents addition to the timeline, it should be removed from the scope. Server systems are sufficiently complex and there are sufficient differences between manufacturer's hardware and firmware that will necessitate some set-up by IT professionals responsible for the testing.

Page 15, Section 3 Logistics: This section does not clearly define the availability of the SERT model for use by other groups or jurisdictions. Given the recent efforts by several jurisdictions, spawned by the release of the ENERGY STAR Tier 1 server requirements, to set voluntary or regulatory server energy efficiency requirements, we think it's important that the SERT tool be available for us in other jurisdictions to allow for deployment of a single performance power metric for servers. While this is not a direct concern of EPA, the ubiquity of the ENERGY STAR program and its use to underpin other regulatory approaches makes it important to be able to extend the use of the SERT metric beyond the ENERGY STAR program.

Thoughts on Idle Metric Incorporation into SERT:

1. It is our (IBM's) opinion that the SERT metric is the right tool to integrate, performance, idle power, and power use over the operating range into a single metric.
2. We think that the metric can eliminate or reduce the need for categorization of servers, though we believe that will have to be proven out by the data collected from the metric run on different systems. I noted that we were specifically pushing the resilient server definition because we felt that the overhead inherent in a server with a lot of redundancy have different performance/power attributes than a non-redundant server which may justify a separate category.
3. I indicated that we believed that one way to incorporate idle into the SERT metric would be to weight the score for the idle state or the idle state and 10% utilization at 20 or 30% of the total score to give more weight to idle power draw, or more specifically the percentage reduction in power use between full power and idle power.
4. I also discussed with Evan the fact that what you are looking for in a SERT metric over the range of utilization of a server is a concave curve. So you are looking for higher power use at higher utilizations and a lower power use (a pretty flat change in power use) at lower utilization. I suggested that there are two indicators here - what is the difference in power use between 0% and 100% of power use and what is the slope of the curve from 0-20 (or 30%) utilization and 80% (or 70%) to 100% utilization. You are interested in a flatter curve at low utilization and a steeper curve at higher utilizations which may enable you to consider some form of metric using those three values.