ACTIVE STATE EFFICIENCY CRITERIA (Section 3.5)

The Green Grid SERT™ Analysis Working Group (WG), in collaboration with the SPECpower Committee, has assessed the options for combining the performance and power data generated by the SERT™. This assessment is based on the analysis of an extensive data set consisting of SERT data from ENERGY STAR® certified systems and the SPECpower Committee's dataset. ITI, the SPECpower Committee and the TGG WG recommend that a single metric be used for assessing the server efficiency. The recommendation has the following key points:

1. The server efficiency metric is created by using the measurements collected from the SERT by combining the performance and power interval data into worklet efficiency scores, combining the worklet efficiency scores by workload type (CPU, memory, storage), and combining the three workload scores with an applied weighting. All combinations of the interval data should be accomplished using the geometric mean (geomean) function; see the actual calculations in "Creating a Server Efficiency Metric.pdf"

2. For thresholding purposes, two configurations – a low-end, and a high-end configuration should be tested and reported. Specific requirements should be set for the CPU capability, memory capacity, and storage type of the low-end and high-end configurations and any active state efficiency thresholds should be set for only those two configurations.

3. Idle power is a poor indicator of server efficiency and idle power thresholds should not be set in Server ENERGY STAR V3 if active efficiency thresholds are established. The worklet efficiency scores, which include utilization and power use as low as a 12.5% server utilization level, represent that a server with a larger difference between the maximum and minimum active power measurements may achieve a better worklet and overall server efficiency score given the same relative performance.

Recommended Combined, Single Metric:

ITI, the SPECpower Committee, and TGG are jointly proposing that the ENERGY STAR adopt a single server efficiency metric, which uses the geomean to combine the SERT normalized performance and power interval data into individual worklet performance/power efficiency scores. It then combines the worklet efficiency scores as the geomean of Workload types (CPU, memory, and storage) and then uses a weighted geomean to calculate the combined component/workload scores into a single metric. The details of the combinatory equations are provided in the document "Creating a Server Efficiency Metric.pdf" which is provided with these comments.

Mathematical Combination Method: The TGG SERT Analysis Working Group (WG) and the SPECpower Committee evaluated the relative merits of the arithmetic mean, geomean, and harmonic mean (see document "Mean Calculation Methodology Analysis.pdf" included with these comments) and determined that the geomean offers the best methodology to combine the performance and power interval data and the resulting worklet efficiency scores. Using the geomean prevents any single

performance score or load level from unduly influencing the combined metric. Of the three means methods considered only the geomean provides the same answer when dealing with ratios (i.e. the mean of a ratio = the ratio of means). As the server efficiency metric will inherently be a ratio of performance to power, the geomean is the only method that will provide a consistent efficiency value when combining both ratio metric efficiency values and performance and power values.

Weighting of the workload types:

The worklets combined for each workload type are:

CPU: Compress, CryptoAES, LU, SOR, XMLValidate, SHA256, SORT, and Hybrid ssj

Memory: Flood2 and Capacity2

Storage: Sequential and Random

The WG and the SPECpower Committee are continuing to evaluate the SERT data to refine their worklet and workload weighting recommendations to find agreement on an appropriate weighting. The TGG WG and the SPECpower Committee are confident the required analysis can be completed within the planned schedule for the development of the V3 requirements.

Product Categories and Configuration Types:

ITI, the SPECpower committee, and the TGG WG recommend adding a separate category for storage servers, as these servers have unique properties which are not evaluated properly by the single combined metric proposal discussed above.

For active state and idle power evaluation, it is recommended that a low-end, and a high-end configuration be defined and required to be tested, with the SERT tool and the .xml report submitted and made publically available. If ENERGY STAR sets minimum efficiency thresholds, the data indicates that a separate threshold should be set for the low-end and high-end configurations.

Low-end configuration: Lowest available performance processor, as characterized by lowest product of the core count and CPU nominal frequency, at least the minimum memory capacity required by SERT (SERT Users Guide – Memory Requirements), and one specified drive. Accommodation will need to be made for servers that do not support HDDs or require a drive for operation.

High-end configuration: Highest performance processor as characterized by the highest product of the core count and CPU nominal frequency, available for the server, three to four times minimum SERT memory capacity requirement, and two specified drives.

In addition to the two configurations that would be evaluated for conformance to an efficiency level, the typical configuration should also be tested under SERT and the SERT results reported via the .xml report. Because the typical configurations will not be consistently configured, an active efficiency threshold should not be established for this configuration. The current definition of the typical configuration should be maintained.


ITI


the green grid®
get connected to efficient IT