



ENERGY STAR® for Computer Servers Tier 2 Development Feedback and Recommendations (May 2011)

The Green Grid, a consortium of industry-leading companies, welcomes the opportunity to comment on topics under consideration for the ENERGY STAR for Computer Servers specification.

Introduction

A consortium of information technology providers, consumers, and other stakeholders, The Green Grid seeks to improve the energy efficiency of data centers around the globe. The association takes a holistic and comprehensive approach to data center efficiency. The Green Grid understands the need to quickly proceed with data collection, develop an efficiency evaluation tool and expand the scope of the ENERGY STAR® program to cover blades and other servers. Participants in The Green Grid include such diverse companies as major server and storage equipment manufacturers, major software providers, and large data center end users/owners. Because the consortium is comprised of representatives from throughout the IT ecosystem, we offer a unique perspective and ability to support a comprehensive specification development.

Summary

The Green Grid appreciates the coordination of the March and subsequent development dialogs on product family and test plans for Tier 2 of the server specification. We have summarized the inputs and recommendations to ensure other ENERGY STAR stakeholders can review and assist in the development of the specification.

The key areas of concern regarding EPA's proposal are:

- **Computer Server Definitions:** ECC is a standard across most servers. We recommend keeping this definition. Systems without ECC should be considered small scale servers covered under the Computer V5/6 requirements and not part of the server specification, due to their tolerance for memory failures.
- **Product family description:** The 5 point family test definition is good but slightly more than needed. We recommend a 3 point test which should be sufficient coverage without overly burdening large extensive configurations.
- **Form factor and exceptions:** Blades should include form factor descriptions such as single wide, double wide and half-height as sub-categories of blades. One member company has tested single and double wide blades with 2 processor sockets and found significant differences in the power profile. The blade form factor contributes to the blade capacity and cooling requirements of a populated system.
- **Data Collection plans and schedule:** Tier 2 development will likely take 12-14 months based on previous activities. To meet that schedule we recommend allowing system manufacturers to use their power calculators to provide volume data for statistical distributions, after the manufacturers have demonstrated the relationship between actual system test data, such as data supplied for ENERGY STAR® qualification and power use calculated from the power calculator. The calculators will utilize the extensive characterization and development work employed by each manufacturer to support the performance specifications on their products. This will also reduce the uncertainty in data collection across the large number of configurations.
- **Test configurations:** Should explicitly indicate that these describe the conditions of the test equipment and setup. The accuracy and resolution are not applicable to the system or device under test. We recommend aligning with international standards, where CENELEC, TC59, and the JWG have defined device testing accuracy for standby and off power modes.

The Green Grid continues to actively support the efforts of ENERGY STAR and the EPA in promoting energy-efficient products and maintaining the program's integrity. As indicated in our comments on the partner agreement and program section, we hope to review any details in an upcoming teleconference. If any of the ENERGY STAR team has specific questions on these comments, please feel free to contact our task force lead, Henry M. Wong (408-765-5047).

Detail comments:

Computer Server Definition

Error correcting code, ECC is a generic characteristic of most computer servers. The dynamic memory size and speed require data error detection and correction to meet the quality of service in most data center operations. Computers operating as servers that do not need error correction are systems that have a small memory footprint and can tolerate occasional errors. These small scale servers are already covered in the ENERGY STAR for Computers specifications. We recommend leaving ECC as part of the definition of all servers covered by the ENERGY STAR for Computer Server specification.

Adding examples in the Resilient Server definition is a welcome addition. As noted above, basic ECC should not be listed specific to resilient servers. The fault isolation and extensive correction techniques described in this section are sufficient to differentiate a resilient server from the basic ECC capabilities found in typical servers.

For ACPI compliant systems, S0 is an industry accepted (standard) system state definition for idle. For non-ACPI compliant systems and architectures, we recommend classifying their idle state by using wake latencies, i.e. response to OS request. Though there is no industry definition of idle state for these machines, their Idle state can be described by the latency to resume operations from inactivity. By describing Idle for non-ACPI compliant systems as the ability to respond within a fixed period of time (e.g. $\leq 1\text{mS}$), the specification would preserve idle state expectations across different architectures. TGG members with non-x86 systems are evaluating and will present a specific recommendation of an alternative approach for their systems.

For Line 162 in the CPU definition, please remove "server" from the term "server motherboard". The motherboard may share many of the same features as those used for client systems. The qualifier is more appropriately used to describe the OS as a "server" operating system.

Family Definition

EPA's 5 point testing proposal of:

- I. maximum performance (power),
- II. maximum configuration,
- III. lowest configuration
- IV. lowest performance (power)
- V. typical

is a good list that recognizes characteristics that bound a product family. We believe, however, that 3 points are sufficient to represent the product family. Limiting the family definition to only necessary data points is a critical feature for very expensive configurations such as 4Socket and bladed systems. Given that a fully configured bladed system could be 24 blades and cost well over \$300,000, limiting the testing

to only necessary data points would help improve the adoption of the program. We recommend that the following 3 data points would be sufficient:

- I. max power,
- II. max configuration and
- III. minimum configuration.

The other data-points are not expected to change the compliance capability of the product family and would be within the envelope of the 3 data points.

Form factor and exceptions

The description of the scope of products included and exempt from the specifications should correspond to the description of servers in the beginning of the specifications. We also recommend referring small scale servers without ECC to the ENERGY STAR for Computers specifications and storage to the Data Center Storage specification development.

In the form factor description, we recommend listing the form factor of blades which may indicate the energy footprint of the system. For example single-wide, double-wide, and half height blades define different energy footprints. These descriptions may be a necessary identification of bladed systems certified as ENERGY STAR compliant.

Data collection plans

Schedule should reflect the testing requirements to configure and evaluate configuration sensitivities. Based on experience and reviews with colleagues in the industry, we believe the specification development will take approximately 12 to 14months, with the schedule beginning in March of 2011. Our estimated schedule break down is approximately:

Activity	Duration	Comment
Data Collection	4-5months	Use power calculators for quantity. Correlation supports the accuracy and configuration sensitivities.
Data Analysis	2 months	Review of base models, feature adders.
Draft revisions(x2)	~5months	2 draft revisions of the specifications, proposed limits, documentation requirements, and test configurations
Final Revision	2months	Final revision including CB requirements/training and international standards comparisons
Release	1month	EPA's previous requirement for interagency review and approval.

As we discussed in some of the teleconferences, the schedule assumes the use of system manufacturer power calculators for quantity information and test correlations supporting the models in the calculators. Using power calculators which were built from characterizations and analysis at system manufacturers, would reduce the time necessary in gathering a sufficient dataset. We suggest correlations be conducted against either new systems and/or data from previous ENERGY STAR submissions from that manufacturer.

As part of the data collection, we recommend highlighting the need to collect Idle power sensitivities to configurations enhancements on a fixed base model. The configuration enhancements should include memory technology & speed, I/O, storage drives, and redundant power supplies. Data collection of the feature enhancements on a fixed base model will help focus on the additive idle of the enhancements instead of base model variations.

With regard to I/O devices, we believe there are additional categories of devices that may be included in these systems such as co-processing units, memory structures, and various I/O protocol drivers. There exists co-processing devices in the market that may consume up to 60W. We recommend addressing these in the compliance test process by removing them from the systems during compliance testing or allocating a budget for devices that may not be removable from the system.

Test Configurations

Section B Ambient Temperature

We recommend the temperature condition for the test be clarified as the ambient temperature between 17 C and 27 C. The current wording can be mis-interpreted as requiring testing to cover the entire temperature range.

Section E testing accuracy

The accuracy requirement appears extreme for servers or any other device for energy testing requirements. The .01W accuracy to a 95% confidence interval appears to be a condition of the test equipment instead of the accuracy of the measurement of the system under test. It should be consistent with IEC 62301 Edition 2, aka EC 62301-1 (renamed to EN 50564), only the test equipment uncertainty can be expected to maintain this accuracy under a fixed load. By definition, the system under test would have measurement uncertainties that exceed these limits. We recommend clarifying that the testing accuracies are referenced to the test apparatus and not system measurements.