May 20, 2011

Dear EPA,

Thank you for holding the March 2011 face to face review in Washington, DC and the subsequent discussions on the various proposals for Tier 2 of the ENERGY STAR for Computer Server specification. The discussion points were helpful in outlining the issues and pursuing solution options for Tier 2.

Despite delays in the program due to the enhancements in the certification and verification program and technical issues on SPEC's Server Efficiency Rating Tool (SERT™), we remain supportive of SERT and development of the Tier 2 specification on servers. We have attached a summary of the feedback on the proposed changes to the specification and points discussed since the March meeting. We recommend coordinating another teleconference in a few weeks to discuss industry inputs and review the status of the data collection activities. In addition, we appreciate the support of the international stakeholders (e.g. Japan, Korea, China/PRC and Latin America) in the evaluation of SERT. We hope their concerns can be captured in the tool and Tier 2 development activities. Addressing the international concerns will help harmonize evaluation methods with other worldwide energy efficiency programs.

If you have any questions please feel free to contact me at (408) 765-5047, henry.l.wong@intel.com.

Sincerely,

Henry ML Wong
Sr. Power Technologist
Eco Technology Program Office
Intel Corporation

Intel Corporation
2200 Mission College Blvd, SC12-507
Santa Clara, CA 95054

Commentary by Section

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.

Although ACPI compliant systems could use S0 as a state definition for idle, we recommend classifying system states using wake latencies. System states can be generically described by the latency to resume operations from inactivity. By describing Idle as the ability to respond within a fixed period of time (e.g. ≤1mS), the specification would become independent of any specific operating system or architecture. Describing the idle state based on the wake latency allows one to generically describe other system (sleep) states and allows flexibility for technology improvements in power management.

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.

Intel Corporation
2200 Mission College Blvd, SC12-507
Santa Clara, CA 95054

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 recommend adding resilient servers to the list of systems out of scope. We also recommend referring small scale servers 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. 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.

Intel Corporation
2200 Mission College Blvd, SC12-507
Santa Clara, CA 95054

Test Configurations

Section B Ambient Temperature
We recommend the temperature condition for the test be clarified as the ambient temperature be 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. As noted in 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. For the accuracy of the device under test we recommend following the European normative standards described by TC59, referenced by CENELEC and the (TC108/TC59) Joint Working Group on Standby and Off Mode Power Consumption Measurement for Energy Using Products (EUP).

Reporting

Section 7.2 Low Power Modes
We recommend clarifying the objective and types of modes to document the power levels. Servers may have lower power states in both the active and inactive operation regions. The lower power states in the active region are dynamic to save energy over a period of time. Documenting power levels for these states is time consuming and does not reflect the energy saved. Low power modes that may exist in the inactive state of the system are sleep and off. These features are only deployed in specialized conditions and not used in the present day data center operations. Given the additional test and reporting burden, we recommend providing clarity in the types of low power modes to provide and a rationale for the use of this data.

Conclusion

Intel appreciates the opportunity to provide the EPA with the comments and the collaborative work in developing the ENERGY STAR Program Requirements for Computer Servers v2.0 specification. We hope to continue the collaboration through workshops and teleconferences to help accelerate the release of the specifications.

Intel Corporation
2200 Mission College Blvd, SC12-507
Santa Clara, CA 95054