



September 23, 2008

Dear EPA,

Please find enclosed Intel comments and feedback on the Energy Star for Servers Specification Draft v2.0, dated 8/15/08.

As noted by our participation in industry, Intel, and EPA-initiated forums, we believe the product specifications and resulting Energy Star program supporting it, is very important to the future energy efficiency movement in the enterprise server market. We view the criteria as a key world-wide consensus guideline to improving energy efficiency and productivity in the data center. We welcome further discussions and detailed development forums, by which the industry, Intel, and the US EPA may be able to finalize these specifications and guidelines.

We will continue to work extensively with our industry colleagues in Standard Performance Evaluation Council (SPEC), The Green Grid (TGG), Climate Savers Computing Initiative (CSCI), IT Information Council (ITIC), and Storage Network Information Association (SNIA), to support the goals of the Energy Star for servers program.

If you have any questions please feel free to contact myself or Henry L Wong, henry.l.wong@intel.com.

Sincerely yours,

Lorie Wigle
General Manager
Eco-Technology Program Office

Intel understands and appreciates the leadership shown by the EPA toward reducing the energy required by computer servers in doing real work. By establishing definitions and standards together with companies such as Intel and our customers- the computer server manufacturers- the EPA can cause the industry to implement server features that increasingly save energy while still delivering the productivity and greater positive environmental impact the global economy depends upon.

Consequently, Intel is pleased to provide the United States EPA with the following response to the ENERGY STAR Program Requirements for Computer Servers Draft 2 specification. The additional concerns, details and changes to the draft 1 were helpful to delineate the specific areas and approach the EPA desires. The comments listed below are specific to draft 2 and should be considered additional to Intel's feedback on draft 1.

As in previous feedback documents, we've provided comments to each section of the draft specification. Each section's comments are titled per the draft topics to ensure proper context. We would like the opportunity to review these comments with you and answer any questions the EPA team may have.

Section 1 and 2: Server Definition and Eligibility

We agree with the scope restriction to single, dual and quad-socket defined systems. These socket categories are simplistic groupings representing a more complex set of characteristics of systems within each category. Ideally, the classifications should be broken down further to address the energy profile of characteristics which are critical to the productivity of these systems. We also understand how prohibitive it may be to accommodate all of these characteristics. Instead, we advise that the Energy Star specifications recognize these characteristics and limitations placed on the specifications. The resulting specifications should indicate that given these limitations, the definitions, and eligibility, designation as Energy Star, is meant to highlight and incentivize relative improvement in energy efficiency. Given the limitations in market scope and detail, the specifications and indices are relative scales and do not necessarily comprehend energy efficiency or productivity in any particular usage or supported-industry.

Additional detailed recommendations:

1) Definition

A. Computer Server:

- Recommend to include telecom sector for server usage:

[131] "data centers and office/corporate environments" --> "data centers, telecommunications central offices, and office/corporate environments"

[138] "user-installed enterprise applications" --> "user-installed enterprise or communications applications"

[141] Some servers do support un-buffered DIMM, recommend to remove "and/or buffered memory" limitation.

C. Blade Chassis:

[149] "resources include" --> "resources may include"

D. Blade Storage:

[154] To be consistent with 1B, "Blade storage units are incapable" --> "Blade storage units are designed to be installed in a blade chassis and are incapable"

F. Server Appliance:

[163] "web interface" --> "web interface or command line interface"

I. High Redundancy Servers:

J. Standard Redundancy Servers:

Redundancy levels on section I and J seem vague as to whether these apply to simply power supplies or other characteristics in the system. We recommend providing some examples that may help clarify.

K. Computer Server Power Supply:

2) Qualifying Products

[217] Recommend clarifications on what requirements apply to computer server and blade server respectively.

Section 3 Efficiency Requirements for Qualifying Products

Section 3a Power Supply Efficiency

Intel recommends the EPA review the Climate Saver Computing Initiative (CSCI) recommendations and apply them to ENERGY STAR for Servers. We encourage the application of these recommendations based on the draft 2 feedback to support industry harmonization around power supply efficiency specifications. CSCI has been purposeful in working with and enhancing the industry PSU measurement processes used by the 80plus program and posted at www.efficientpowersupplies.org. We continue to work for harmonization around CSCI's common targets and methodology, which would enable aggressive PSU efficiency goals to be achieved industry wide.

Section 3b Idle Power

The industry system manufacturers, users, industry stakeholders, and Intel have demonstrated that unqualified idle power as a metric is not an appropriate yardstick to determine if a server is energy efficient or not. In fact, it would encourage purchasing inefficient system configurations, and discourage advances in productivity and energy efficiency in the industry. In particular, technologies such as virtualization and system power management, which require more robust configurations and resources impacting idle power, are discouraged by an unqualified idle metric. We do however, agree that system idle power is a valuable piece of information and should be communicated in a standard and common method for servers. We recommend that the system Idle power become part of the standard data reporting section.

Our criteria recommendation is to configure the Tier 1 server specification to encourage and motivate advances in productivity and energy efficiency. Whereas a more detailed specification that cover these characteristics can be addressed in Tier 2.

Upon review of energy efficiency and productivity of servers in the data center, a highly efficient and productive configuration uses virtualization and ability to autonomously manage system resources dependent on policies for the environment. These policy engines exist at multiple levels of the software and hardware stack of the system. Systems which have been demonstrated in the market include system management tools provided by system manufacturers, Intel's Node Manager, and VirtualMachineManager's. These system resource management tools coordinate execution of tasks and employ hardware power management features, if available. It is this base level set of features and capabilities which allow power and energy savings

in the data center. Without these hardware level features, the dynamic range of the system and energy savings can not be achieved.

As an alternative to an Idle Power specification, we recommend the following:

- + Energy Star compliant systems contain the following power management features:
 - o *Mandatory:*
 - Dynamic Voltage and Frequency Scaling (DVFS, Speedstep™)
 - Memory Power Management or Low Power Memory States
 - Exports power and utilization to System interface.
 - o *Some number (X) of the following:*
 - Processor Sleep;
 - Power Capping;
 - Dropping into lower static power modes as workload reduces;
 - Variable Speed Fan Control based on power or thermal readings;
 - Lower Power I/O Interfaces
 - Rack level or processor level liquid cooling

The proposed alternative to idle promotes both a minimum set of information and features to allow the system interface to execute a power management policy (at specific levels of hardware); and, encourages the adoption of additional energy efficiency techniques. The system management sophistication is consistent with the “premier” energy efficient systems the Energy Star logo should promote.

The number of optional features, X, should be determined based on investigation of various server configurations currently being sold. We recommend a review of current server configurations and tool availability to determine the premier levels that meets the intent of the Energy Star designation.

Note that given the unique market servers provide, the purchasing methods are not the same as other compute systems. Some issues include, operating system (images) are not shipped with the system but are customer approved images installed on site. Inactive power levels, with or without resident software, can not be accurately emulated by the manufactures. When measuring system power, especially inactive states such as idle, a common configuration (software stack) must be used.

Additional detailed feedback to section 3

B. Idle Power

- Other than redundancy, # of sockets, and memory, there are different system configurations that contribute impact to idle power consumption such as I/O slots, HDD. Any power or energy review must either scale to or accommodate these valuable application features.
- The separation of 16GB and >16GB will make the higher capacity category inappropriate for main stream and high end servers. With main stream and high end servers vary in memory capacity from 16GB to few hundred GB, this will always lead to preferring the minimal capacity system in this category. Small memory configurations will actually hinder and impose limits to virtualization that can be accomplished. Metrics need to take scaling into account.
- The absolute values without a concept of performance encourage bad tradeoffs. These encourage lots of little servers, which may not be as efficient as a more powerful 1S or 2S server.

C Power and Temperature Measurement Requirements

[263] "on AC power" --> "on AC (or input DC) power"

Section 4 Test Criteria

Many servers are sold without an operating system (OS) installed. It is important to note that actual server idle power is much higher than without an OS installed because the OS is required to put the processor/system into the lowest power state. In order to accurately measure idle power, the OS must be installed. On systems that don't have the OS installed by default out of the factory, the idle power will be higher. Recommend the EPA define an OS configuration if testing prior to delivery to the end user is desired.

- [283] AC testing: In North America, 208Vac is much more common than 230Vac.
- [283] DC testing: Should be tested at -53VDC +/- 1VDC as noted above.

Appendix A (Standard Information Reporting)

- Include product dimensions
- [309] "Dimm" --> "DIMM"
- [310] Include type of I/O slot and speed/width (i.e., x16) of each
- Include any definition of on-board storage?
- [344] "nominal and at peak (35C ambient inlet)" --> "25C and 35C ambient inlet"
- Should include System Idle and Maximum Operational Power.

- Include information on monitoring of input feed voltage/amperage?
Include other voltage/current monitoring information?
- Do the SPECpower_ssj2008 results need to be officially posted or just available on demand?
- EPA should provide sample loading criteria.
- May want to list PSU max power required in what the SPECpower_ssj2008 100% power usage is. These are often quite different.