March 25, 2009

Dear EPA,

Please find enclosed Intel comments and feedback on the Energy Star for Servers Specification v1.0 Draft 4, dated 2/20/09.

Intel remains committed and supportive of the US EPA's efforts to define energy efficiency goals and targets across the spectrum of computer products including the current proposal for Energy Star for Servers. We are encouraged with the progress on Tier1 and welcome the opportunity to assist in meeting the EPA's target release of the specification by May 2009. Intel is actively involved with industry stakeholders on the remaining implementation concerns with Tier1 and plans for an Energy Efficient Performance solution for Tier 2.

We 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 (ITI), Alliance for Telecommunications Industry Solutions (ATIS) and Storage Network Information Association (SNIA), in addition to supporting the Energy Star for servers program to deliver increasing energy efficiency.

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

Sincerely,

Lorie Wigle
General Manager
Eco-Technology Program Office

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Intel welcomes the opportunity to work with the EPA in driving toward greater energy efficiency in enterprises. Energy Star for Servers is an aggressive program that could be used to harmonize energy efficiency programs world wide if it is written to achieve the stated energy saving goals.

Intel appreciates the opportunity to provide the EPA with the following response to draft 4 of the ENERGY STAR Program Requirements for Computer Servers v1.0 specification. Specifically, draft 4 represents significant progress in creating an introductory specification for computer servers. The adjustments in categorizations of servers between managed and non-managed systems, accommodating the complexity of 4-Socket systems configurations, and recognition of fault tolerant systems will help finalize the specification. Intel believes that the Tier 1 specifications should be adjusted and finalized as quickly as possible to better focus the industry's development efforts in creating energy efficient performance metric(s) for use as an alternative to system Idle power for the Tier 2 specifications.

As with our feedback on previous drafts, the response is organized per section. We've also included a general commentary reflecting the updates and changes in draft 4. We would like to continue our practice of having the opportunity to review these comments with the extended EPA team to answer any questions you may have. The comments listed below are specific to draft 4 and should be considered as additional to Intel's feedback on the previous drafts.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Overall Summary

Intel appreciates EPA's recognition of the complexity in 4-socket configurations and addressing the unintended consequences with applying idle criteria to this category of product. The statistical analysis used to set a more scalable adder system on 1- and 2-Socket systems are consistent with our findings and addresses the systemic error highlighted in our responses to draft 3. We feel there remain a few areas to quickly close on Tier 1. We recommend that the priority for Tier 1 completion should be to minimize the implementation complexity. The industry's development focus could then be applied to Tier 2 energy efficient performance metric(s).

Intel recommends for Tier 1:

+ Treat Blades similarly to 4Socket systems (Power Mgt, Reporting data, and No idle limit). The blade category is complex due to variability in chassis dependencies, testing, and a wide range of configurations. Attempting to conduct several rounds of data collection and review will take a significant amount of time and could negatively impact the development of Tier 2.

+ As referenced by CSCI, low load efficiency and Power Factor Correction (PFC)'s should be changed on low (input) power PSU's. We recommend removing the low load (≤ 20%) testing requirements (PFC and efficiency) for PSU's ≤500W. Removal will simplify the testing, has little impact to energy savings, and reduces the possible unintended consequences of over-sizing supplies.

+ Two-Socket systems configured with 1 processor should be treated as a 2-Socket system (i.e. 2P limit) with a 15W idle power subtraction. The associated system limits would be consistent with those customers purchasing the features enabled these systems, and allows for incremental upgrades when the projected needs increase.

+ Based on SAS storage power data and reflective of the performance requested many customers, we recommend a 10W SAS (add-in card) Idle power adder.

+ EPA's product family qualification simplifies the management of ENERGY STAR products for EPA. However, we recommend that the

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

testing be limited to maximum and minimum configurations which comply with the Energy Star limits. We also recommend that processor speed or SKU's should also be allowed to vary within the product family, due to the predictability of idle contributions within a processor model type. We recommend reviews with system manufacturers to use in line manufacturing data as an alternative to demonstrate compliance within a product family.

+ For DC-DC systems (i.e. -48Vdc), the test voltage should be -53Vdc ± 1Vdc. Although these systems are quoted as -48Vdc, with a large operating range, most installations are running at -52Vdc to -54Vdc. This testing condition is consistent with the industry standards conducted by ATIS.

+ With respect to processor utilization, thermal and power data reporting, we recommend removing the 1second sampling requirement. The sampling requirement may restrict innovations in sampling and controls of the lower subsystems. The request that currently exists in draft4 for an accurate status at 30second intervals should be sufficient for the server level controls without impacting the subsystems.

+ For Tier 2, we recommend removing the energy efficient network specification and requirements and the storage targets from Tier 2. The IEEE802.3az standard is preliminary and will be limited in adoption by the Tier 2 timeframe. IEEE802.3az may be a consideration after Tier 2, when technical issues such as interoperability across the network technologies can be resolved. As discussed with the Storage and Network Industry Association (SNIA), the differing characteristics of energy efficiency in storage, warrants the development of a separate specification for this group of systems for now.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Review by Section
Section 1 & 2: Server Definition and Eligibility

Server Definitions
We agree with establishing definitions for "fully fault tolerant" and "dual-node" systems. The definitions may be further clarified by providing examples. Fault tolerant systems can be described with examples such as lock-step operation, with duplicate compute resources dedicated to a given operation such as the CPU and local memory. For "dual-node" servers, one may note that these do not include multimode systems, which would cause a system to emulate an 8-way or higher computing cluster.

Blades
Intel agrees that including blades in the program is desirable. We do, however, believe that blades and their hardware dependencies to the chassis and shared resources add complexity similar to the 4-Socket systems.

Enterprise Server Purchasing and Integration
Intel supports the concept of using the product "family" approach noted in draft 4. We believe that only a maximum and minimum configuration confirmation is necessary, where the "typical" configuration may be too arbitrary and market dependent to define. It is also unclear how a "typical" configuration data could be used, given the variety of serviced markets. A typical configuration for a database managed service may more heavily depend on broader I/O and local memory, whereas a typical system for transactions may rely on more compute capabilities. By containing the compliance configuration boundaries, variations interpreting "typical" can be avoided. Within the product family range, existing in line monitors on subcomponents should be sufficient to demonstrate compliance within the family. This method is consistent with the manufacturing process of these systems, consistent with the methods used to ensure compliance to the manufacturers' datasheet, and precludes the costly expenditures for additional end product sampling. As noted in previous comments, final product assemblies pulled and evaluated can no longer be sold as new, thus increasing the cost burden without providing additional assurances of compliance. We recommend the EPA discuss the possibility of utilizing manufacturers' in-line monitors as an alternative to sampling in order to demonstrate compliance within a family.

With respect to the processor speed or SKU restriction, manufacturers should be allowed to characterize and submit maximum and minimum processors SKU's to demonstrate compliance across the processor family chosen. The processors within a product family may vary in speed and features "enabled" however, the idle power contribution across that processor product family is predictable. System manufacturers can reasonably ensure compliance in the family even though they may populate the system with a lower speed processor from the same family. As noted above, the in-line monitors, whether provided by the subcomponent vendors or sampled, can and should be used as a proxy for within system family compliance.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Section 3 Efficiency Requirements for Qualifying Products

Section 3a Power Supply Efficiency

Intel recommends the EPA to continue its excellent engagement with Climate Savers Computing Initiative to resolve efficiency limit issues with low load efficiency specifically, the efficiency and power factor correction limits on low capacity power supplies appear to be more for documentation completeness rather than a true need. The additional testing requested will increase costs, and may cause a mixed incentive to oversize the power supply. Given that with the variability in this range for smaller sized power supplies, the limited energy savings and mixed incentive counter to right sizing, we recommend that for the <=500W power supplies the efficiency and PFC levels of 20% loading and below be removed.

Section 3b Idle Power

Intel appreciates the EPA's recognition of the additional complexity of 4-socket systems. We agree that the information disclosure on performance, maximum power, idle power and performance will help consumers in an energy efficient configuration choice for this category.

For blade servers, we believe the complexity in configuration, and variety in dependencies on chassis, such as airflow, double wide slots, and backplane configurations, create a similar situation to 4-Socket systems. The available idle data and hardware dependencies on these structures are extremely limited. With the desire to focus on an energy efficient performance metric for Tier 2, the industry may need to duplicate the data collection and analysis process on blade servers after the metric(s) development. We believe this extends the development work on Tier 1 and potentially impacts Tier 2 schedules as well. We recommend that for Tier 1, blade servers be treated similar to 4-Socket servers. The intercept of an energy efficient performance metric(s) for this category would be part of the Tier 2 development.

Partially populated 2-Socket platforms

We understand the desire to place limits strictly on what is populated in the system. As Intel has previously commented, the socket designation highlights the category and group of product features supported on the platform. The system categories contain more features than simply an incremental processor. Therefore, although one may populate a 2-Socket system with a single processor and half the capable memory, the system management features, platform consistency, and optional upgrade may be required by those customers. We recommend that a subtraction of a processor idle contribution to a 2-Socket system be employed, as opposed to considering these systems with the 1socket, 1 processor designed systems. One should also note that twice as many 1processor systems do not incorporate the platform features mentioned for a 2-Socket system. We estimate that an aggressive AC platform idle subtraction of 15W from a 2socket system would fairly address the partially populated case.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

I/O Adders
The adders identified for memory and redundant supplies are in-line with our analysis of the database of systems used and our understanding of the architecture of these systems. We do believe that due to the storage performance requirements for some users and the serviceability requirements on larger 2-Socket systems, adders should be provided for SAS controllers and redundant fans. We find that those who require fast access to local bulk memory require SAS controller & drives, which add approximately 10-15W to system idle. A 10W adder for SAS would aggressively limit the idle contribution, yet recognize the productivity needs of this additive feature. For redundant fans, the redundancy is usually combined with redundant power to allow smooth transition upon either power or cooling failure. Therefore, for those applications where only power and cooling redundancy is required, as opposed to multi-machine backup, redundant fans offer the most efficient and serviceable alternative. Power redundancy does not also mean fan redundancy on the platform, so they should be treated independently. Cooling redundancies may also offer lower and finer levels of fan speed control to reduce fan power consumption consistent with compute demands on the system (beyond just high load and inactive). We recommend a 10W idle adder for redundant fans.

Section 3d Data Measurements and Reporting

The data measurements and reporting targets should not pre-determine the sampling (i.e. "1 second") that the system manufacturers use to export an external value (e.g. 30 second average noted in draft 4). The 30 second data reporting should be sufficient for monitoring and control systems, whereas the sampling and capture rate of 1 second is both an enormous amount of information and may conflict with subsystem monitors and sampling that may be occurring. We recommend that the prescriptive intermediate sampling rate be removed to allow customizations in the subsystem while providing actionable system data at the 30 second averaging interval.

Section 4 Test Criteria

As noted in previous comments, many computer servers are shipped without an operating system. In fact, the final configuration, OS image, system settings, and management tools are installed well after the hardware has been shipped from the system manufacturer. We recommend testing models based on their full scale capabilities, as a means by which to rank systems. The procedures should be reviewed and evaluated with the industry organizations previously defined (TGG, ITI, or CSCI) as part of the process of modulating the Energy Star practices to accommodate the procurement and integration methods for enterprise servers.

Testing system power levels are very dependent on efficiency and conversion that occurs at the power supply. As observed with the idle power limits, the number and type of conversion has a direct affect on the value obtained at idle. The input line voltage has a direct affect on these values and the resulting difference in platform power can not be controlled to provide a similar rating. We recommend a solution be derived with CSCI, such as settling on a fixed, worst-case line level to run the compliance testing.

48Vdc systems
Based on industry work with ATIS, we continue to encourage the EPA to reconsider establishing a new standard that differs from an existing telecommunications standard. We concur with the experience and recommendations from ATIS. We recommend a review of

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

the following aspect with ATIS or other industry groups familiar with this category of product.

The bulk of the 48Vdc systems are actually -48Vdc systems. Despite its name, -48Vdc systems almost NEVER (<1% of the time) operate at -48Vdc. The ANSI standard for -48Vdc is ATIS-0600315.2007, and this states a nominal voltage to be -53Vdc (this is a compromise between the two major battery types: VRLA-based systems typically operate at -54Vdc, while flooded (wet cell) battery systems typically operate at -52Vdc). There is an appreciable difference in operating efficiency of some equipment at -48Vdc vs. -53Vdc, and we encourage the industry to optimize their power converters for the voltage that is typically used rather than at some arbitrarily different voltage used for comparison. For the purpose Energy Star testing, we recommend a test voltage of -53Vdc +/- 1Vdc for "-48Vdc" systems.

Tier 2 Requirements Effective October 1, 2010

Developing energy efficient performance metric(s) will be a challenge to achieve within the timelines provided. We are convinced that the industry would be capable of developing the tools for energy efficient performance by the end of the 2009 calendar year, such that a late 2010 Tier 2 Energy Star for Computer Servers can be enabled.

We believe the incorporation of either Energy Efficient Ethernet (i.e. IEEE 802.3az) or a storage device specification to Tier 2 of the server specification would be premature. The Ethernet specification will have been released for a very short time. As with many newly developed standards, interoperability and availability of compliant Ethernet products are expected to be issues during 2011. To gain a robust product mix (e.g. 10Gbe and 1Gbe) of IEEE802.3az compliant devices and resolve potential interoperability issues, we recommend including this in a future revision beyond Tier 2. We believe, based on current technology projections that inclusion of IEEE802.3az into an Energy Star revision could occur in 2012. With regards to storage criteria, we agree with SNIA that storage parameters for energy efficiency or energy efficient performance are vastly different from compute servers. We recommend that Energy Star for Computer Data Storage continue separate from the computer server specification. Separation of the programs would allow each group to better define the tools needed to gauge efficiency for the product's use and market. We recommend revisiting the option to combine the specifications after Tier 2 of the Energy Star for Computer Server specification.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124