March 14, 2008

Dear EPA,

Please find enclosed Intel comments and feedback on the Energy Star for Servers Specification Draft v1.0, dated 2/14/08.

As noted in a previous teleconferences on the topic since the publication we welcome and are willing to assist with any subsequent workshops or development discussions on the topic.

We have also worked 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

President
Climate Savers Computing Initiative

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Note: if using Microsoft® Word, please use View Print Layout in order for graphs and charts to be visible.

Intel understands and appreciates the leadership shown by the EPA toward reducing the power 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 1 specification. As much as possible, we have attempted to calibrate our recommendations based on marketplace data, from various sources, manufacturers, end users, and industry analysts, such as IDC. Using these data outside of the direct conversation between the EPA and Intel would require additional permissions yet to be granted by IDC. Please feel free to contact Intel or sources identified to ask questions or investigate requests for permission to release the information.

We have provided comments to each section of the draft specification, consistent with the draft document's numbering scheme. Each section's comments are also titled per the draft specification 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 approached these recommendations with the philosophy of establishing definitions encompassing the broadest portion of the marketplace as related to power consumption and power management features. Focused this way, Server definitions should impact the overwhelming majority of the marketplace without undertaking numerous exceptions and complexities for special-purpose / low volume products. We focus definitions away from those segments of the market which present unique technology implementations and/or are inconsequential to overall market volume, and therefore ability to truly effect total power consumption.

Summary Recommendations-

1) Intel recommends that the EPA defines servers as "Designed and capable of having one to four processor sockets on its base board." Four-way or fewer 'processor sockets' define the server market. Socket counts above this are very low volumes of special purpose computing platforms with unique features and power consumption profiles which justify their own definitions and energy metrics. Also, a significant number of uniprocessor servers are still deployed, especially in the broader market of small and medium businesses.

2) Intel recommends that the EPA define servers as having "Support for up to 16 GB of ECC and/or buffered memory." 16 GB or less of main memory defines the volume server market as seen in IDC 2007* market data. Above this are very low volumes of special purpose computing platforms with unique features and power consumption profiles which justify their own definitions and energy metrics.

3) Intel recommends that the EPA focus the Tier 1 metrics on Pedestal and Rack Servers, while developing specification for blade servers. Most servers are rack form factor. Pedestals are a distant 2nd. Blade Servers are a unique form factor with their own power/thermal profiles and justify unique Energy metrics.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

4) Intel recommends that ENERGY Star approvals track real world operating power efficiencies not solely limited to idle state. We strongly recommend ENERGY Star power efficiency ratings be specified at 20%, 30%, 40%, and 50% utilization levels. Intel additionally advocates augmenting the SPECPower benchmarks to extend beyond java-oriented workloads to include Transactional and Decision Support workloads. We believe the SPEC community as well as the TPC forum are both aligned to achieve this goal within the year.

   Per IDC 2007 Server market analysis, the Server virtualization ramp is accelerating faster than previously projected. >38% of all 4-way servers were virtualized in 2007. AND the average utilization of virtualized servers is 55%, trending toward 60%!

   We respect that the EPA has previously been on a path to focus efficiency goals on power consumed by servers in an idle state. However, the IDC 2007* market data clearly shows that servers are increasingly being utilized at much higher rates than previously thought and that the trend is on the rise. For this reason, setting power standards for servers solely based on idle power consumption would be counterproductive to the overall goals of reducing the energy consumed.

   Additionally, modern day servers all incorporate monitoring and control points within multistage power supply units (PSUs) which throttle power consumption in proportion to workloads. i.e. they have functions which promote power saving when a processor is not doing work. Intel believes it is more directly in line with the goal of reducing overall server power consumption by focusing ENERGY Star requirement on efficiencies under real-world loads AND including these power-saving features into the components of the server, especially into its processors, firmware, and chipsets.

*IDC 2007 Server Market Study references data available through IDC Corporation, a leading industry analysis firm.

Section 3a Power Supply Efficiency

PSU Efficiency:

Intel recommends that ENERGY STAR for Servers align to Climate Saver Computing Initiative (CSCI) power supply efficiency specifications. CSCI is an industry-wide organization working with ingredient (power supply) manufacturers, system integrators, and researchers, and is optimally supported to provide comprehensive evaluation methods in constantly advancing system design. CSCI has been purposeful in aligning to existing industry PSU measurement processes used by the 80plus program and posted at www.efficientpowersupplies.org. Harmonization around a CSCI's common targets and methodology would enable PSU goals to be achieved industry wide.

CSCI PSU specifications recognize that there are several PSU architectures that may be used in Servers; most notably, Multi-output and Single-output designs. Therefore, it is important that the ENERGY STAR specifications include targets for both multi-output and single output PSU's. Recommended efficiency targets are:

Percentage of Rated Output Power	Test Input Voltage	20%	50%	100%
Minimum Efficiency Requirement - Single voltage	230VAC	85%	89%	85%
Minimum Efficiency Requirement - Multi-Voltage	115VAC	82%	85%	82%
Power Factor			0.9	

The multi-output PSU efficiency targets are consistent with the current DT Tier II draft spec targets and align to CSCI year 2 requirements. The single voltage PSU efficiency targets also align to year 2 CSCI requirements and will be in effect at the time of the rev 1.0 Server ENERGY STAR spec release.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Efficiency Test Conditions:

- Intel recommends that efficiency targets be met using the test protocol at www.efficientpowersupplies.org. Under this protocol, Single Voltage PSU's are tested at 230VAC input, and Multi-Voltage PSU's tested at 115VAC input voltage.
- Fan power should be included in multi output efficiency testing and not included single output PSU efficiency testing. Quite often, single output PSUs have system fans cooling the PSU as well as other component in the system. Multi-output PSUs usually do not have this situation and the PSU fan is typically only used to cool the PSU. In addition, this remains consistent with the industry standard PSU efficiency test protocol (www.efficientpowersupplies.org).
- Intel recommends that PSU efficiency acceptance criteria align to CSCI's acceptance method. CSCI's method sets the "target efficiency" as the Mean, or average efficiency of the population of PSU's. The requirement states that in a sample size of 30 pcs, the "mean efficiency" be greater than or equal to the "target efficiency", and that the standard deviation of the efficiency be less than 1.0%. Smaller sample sizes may be tested with an offset in the measurement required due to the statisitcal confidence of the smaller sample size. This criteria aligns to broadly used acceptance criteria for statistical analysis, and applies it to PSU efficiency. The acceptance criteria and white paper describing the method are posted at www.climatesaverscomputing.org.

Further, Intel strongly recommends that efficiency requirements at 10% load not be included in the Server ENERGY STAR specification:

- The recommended CSCI targets have very aggressive efficiency levels at the 20% load point. Given the improvement of efficiency at 20%, the 10% load point will improve by definition.
- OEM's are innovating redundant PSU solutions and not allowing PSU's to operate at the 10% load level. Spec'ing a 10% load efficiency requirement will penalize those that are innovating to optimize power delivery efficiency across the load range and at realistic system loads.
- OEM's are also looking to "right size" PSU's so that they more efficiently deliver the power across the load range to the system. A stringent 10% load requirement will negatively impact efficiencies at higher load points.

Further, Intel strongly recommends that "PSU right-sizing" not be included in the Server ENERGY STAR specification:

- PSU manufactures are working toward a flat efficiency curve across the load range. This strategy allows OEM's to re-use PSU's across several platforms and take advantage of economies of scale. By attempting to spec right-sizing, OEM's will be penalized for working to optimize PSU solutions across product lines
- The complex nature of Server taxonomy means that by definition, right-sizing is very difficult to define. Intel does not see a clear way to define what a "right-sized" PSU is.

Based on initial CSCI estimates, power supply improvements to the highest efficiencies could recover up to 21% of the energy expended in servers today.

Section 3b Idle Power

Using idle power as the server Energy Star metric may not have the desired effect to reduce overall energy consumption; it could actually have the opposite effect and lead to higher overall energy consumption. An idle power specification would only serve to reduce efficiency in the data center and impede the consolidation efforts implied in the EPA's Data Center report published last year.

Reason:
Server idle power has little correlation to energy efficiency. If idle power is chosen as the server Energy Star metric, this may drive the wrong purchase decisions, leading to higher overall energy consumption.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Servers are deployed to complete a certain task or service a particular workload. Idle power is driven primarily by the server hardware configuration (# of processors, # of DIMMs, # of disks, # of power supplies and power supply efficiency rating). If idle power is chosen as the server Energy star metric, the type of servers that are expected to qualify for the Energy Star metric are going to be lower performance servers. There is a potential that procurement departments start mandating the purchase of Energy Star servers (which is a good thing), but to meet a certain performance target, a higher quantity of lower-performing servers will be needed instead of purchasing fewer higher performance servers. Higher performance servers that consume more power at idle wouldn't qualify for Energy Star. The resulting impact will likely become an increase power and thermal burden, due to additional machines to an already constrained data center. This will not be a motivating specification for consolidation and would impede industry progress.

In addition, there are clear industry trends towards virtualization and server consolation which, over time, the duration servers spend in idle state will be reduced.

Difference Between Client Computers and Servers:

Idle power was selected as the Energy Star metric for client computers, where there is often a direct correlation between the number of persons and the number of client computers purchased. On the other hand, server purchase decisions are based on the amount of work that needs to be completed, and is not a linear relation like client computers.

Food for Thought:

As an example, to meet a particular performance target, is it better to purchase qty 1000 Energy Star servers (that have less performance) or qty 1 high performance server that does not meet Energy Star? Which solution uses more overall energy at idle and under load? Which purchase decision will occur when non-technical procurement departments mandate the purchase of Energy Star servers?

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Details:

Below are some sample system idle power measurements. Idle power is driven primarily by the server hardware configuration (# of processors, # of DIMMs, # of disks, # of power supplies and power supply efficiency rating).

System Vendor	Model	Processors	Memory	PSU	Idle Power	SPECpower score	Comments
FSC*	PRIMERGY TX150 S5	1* Intel 3070	4GB: 2*2GB	1*410W	91W	356	1 socket rack mount server
Dell*	PowerEdge 1950 III*	2*Intel E5440	8GB: 4*2GB	1*670W	143W	712	2 socket rack mount server
Dell*	PowerEdge 2950 III*	2*Intel E5440	16GB: 4*4GB	1*750W	147W	682	2 socket rack mount server
Fujitsu Siemens Computers*	PRIMERGY RX300 S4**	2*Intel E5440	16GB: 4*4GB	1*700W	166W	690	2 socket rack mount server
Intel	SE7520AF2 Server Board	2*Intel Xeon 3.6GHz	4GB: 4*1GB	1*600W	159W	87.4	
HP*	DL580 G5*	4*Intel L7345	16GB: 4*4GB	2*1200 W	271W	546	4 socket rack mount server
IBM*	X3850	4*Intel X7350	16GB: 8*2GB	2*1440 W	350W	Not measured	4 socket rack mount server; idle power is Intel internal measurement
IBM*	X3850	4*Intel X7350	64GB: 32*2GB	2*1440 W	431W	Not measured	4 socket rack mount server; idle power is Intel internal measurement

http://www.spec.org/power_ssj2008/results/power_ssj2008.html as of March 3, 2008 except where noted.
* Other names and brands may be claimed as the property of others.

In the above table, the servers with the lowest idle power are not the most energy efficiency as shown by the SPECpower* score. If idle power is chosen as the Energy Star metric, a server with a single processor and with low memory capacity will generally have the lowest idle power and qualify for Energy Star rating. This would inherently penalize higher performance servers that have more capability (more compute power, more memory, more I/O, better redundancy) because these servers have higher idle power. In addition, higher performance servers are often purchased because of the RAS (reliability, availability and serviceability) characteristics, which don't necessarily increase server performance, but can increase system power.

Conclusion:

Using idle power by itself as a standalone metric to establish the server Energy Star metric is not recommended. Servers do spend time at idle, but the absence of a calibration factor between servers that have different performance characteristics is the main concern with using idle power as the only metric. We recommend that server performance, power at peak performance, along with idle power be used to establish the server Energy Star metric.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Section 3C. Standard information reporting

We appreciate EPA's proposal requesting a standard framework to communicate performance and power information. As noted in the draft, SPECpower is an ideal benchmark to begin providing performance, power, and energy information in a common communication vehicle. However, the details associated with SPECpower™ and reporting requirements were not highlighted and reviewed with SPEC. There has also been some industry complaints associated with the limited scope of SPECpower benchmark as well. To address these concerns in the industry, we recommend that:

1) The EPA convene with SPEC to comprehend the documentation requirements manufacturers must meet to provide SPECpower data

2) The EPA investigate the application of SPEC's power assessment protocols to support peak_power, energy, and active_idle power data for other performance benchmarks such as Linpack, SPECmail, and SPECweb.

By introducing a framework with multiple data points, the focus would be to provide a standard means of communicating performance, energy, and power information to IT managers, in addition to establishing a quality data base for the EPA with regards to energy efficiency benchmarks. The data submitted would be a wide range of manufacturers. We believe IT managers would have an additional reason to value an ENERGY STAR product, that key planning information is provided in a standard consistent manner.

Section 3d Standard Temperature and Power Measurement

Intel agrees that promoting the availability of power and temperature information is important to allow systems to monitoring the condition of the server. The number of existing methods prevent realistically calling for a standard that may not be ready. The ability to regularly and openly provide inlet air temperature and power consumption should not be impeded by the adoption of a fixed standard.

Intel recommends that ENERGY STAR for servers specify the type of information to monitor, eg. Inlet air temp and power, and that such data be accessible to system management software in an open format such as xml, csv, or other text. Intel agrees with the EPA data set of power use, utilization, and server inlet temperatures as the key parameters for monitoring to gauge server efficiency and data center health.

Specifically, SMASH is a middle-ware type implementation that may be ideal for some mid-tier enterprise data centers, but could also cause others to bring on unneeded resources to carry the additional impact of a SMASH implementation, actually making the data center less efficient due to more equipment not doing useful (computational) work. Intel supports SMASH but only in the right application. Alternatively, new manageability capabilities are regularly being developed. For instance DCMI (with largely the same adopters as SMASH) will potentially be released before the next revision of the E* document. This specification, which would use fewer resources than SMASH, is intended for large data centers and would more efficiently provide the same data requested.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Section 3e. Power Management and Virtualization

Power Management. Energy Efficient Servers will take advantage of features built into the silicon and system board to reduce the power of components under active workload. Since idle periods in servers can be very short, optimum energy efficiency requires advanced features be integrated directly into the system design to achieve maximum speed and benefit. Below is a table of power management features recommended for energy star requirements:

ITEM	DESCRIPTION
P-states	Processor clock speed and power actively matched to workload (Intel EIST/DBS and AMD PowerNow!™)
T-States	Thermal management states. Actively managed reduction of processor power
Processor Core State Power Management	Ability to manage system package power on both a per-core and per package basis (as in Intel® Xeon® Quad Core Processors and AMD Barcelona processors).
Remote Power Management	Ability to limit system power by remote agent control.

Similarly, we believe ENERGY STAR should promote the use of virtualization as it obviously contributes to energy savings in the server consolidation usage model. However, we do not have a robust pass/fail test to suggest that would qualify only some servers. Over time as the SPEC virtualization work completes, there could be a metric to use to evaluate consolidation capacity.

Section 4. Testing Criteria

For PSU testing, please see Climate Savers Computing Initiative. The test procedures already address the fan and fan power considerations.

For System efficiency testing or benchmark testing, please consult with SPEC.org or the associated benchmarking organization. This will ensure proper, comparative and consistent methods of assessment.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124