May 20, 2008

Dear EPA,

Please find enclosed Intel comments and feedback on the ENERGY STAR® for Computer Servers Definitions and Qualifying Products memo, dated 4/25/08.

As noted previously we welcome and are willing to assist with any subsequent workshops or development discussions on the topic, including the upcoming workshop in Seattle in early July'08.

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 (ITIC), and Storage Network Information Association (SNIA), to support the goals of the ENERGY STAR for servers program. As such, you may find some similarities and consistencies in the concerns being addressed. We hope this consistency aids in highlighting and focus on key considerations in the development of the specifications.

This document contains information that should remain with the EPA and not be published publicly. We have provided a post-able version referencing the information disclosed in this document.

If you have any questions please feel free to contact me at henry.l.wong@intel.com.

Sincerely yours,

Henry M Wong
Sr. Staff Technologist
Eco Technology Office
Intel Corporation

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Summary Response (Computer Server Definition and Qualifying Products)

Computer Server market segmentation into impact-oriented categories of product with like requirements and capabilities is helpful. A matching segmentation of certification criteria may also be required to achieve the EPA's overall objective of reduced energy consumption. The definition of servers should provide that level of segmentation, reviewing both the capabilities and usage targets. We should also indicate that a means to gauge that segment's efficiency and impact needs to be established and that different categories may require different assessment tools.

Definitions:

- ENERGY STAR definitions should separate a 4 Processor Socket systems (4S), highlighting that the usage modes, lower-level resources and capabilities of these systems are different than a single or dual processor system. As observed in the IDC surveys, it is on 4S systems that server consolidation seems to be occurring along with the resulting reduction in overall power consumption. The capacity and capability assessment tools to gauge the efficiency of these system are, however, not available as yet.
- These platforms are using virtualization technologies toward this goal.
    - The optimal memory and disk configurations of 4S virtualization platforms will necessarily consume more idle- power than 1P or 2P solutions, though the 4S configurations are critical to consolidation and data center wide efficiency.
    - A 'one size fits all' certification spec centered only on idle power may actually encourage the purchase of a higher number of less capable servers. This would result in an increase in overall power consumption rather than drive reduction through consolidation using better configured, more energy-efficient server virtualization platforms.
    - 4S should eventually adopt an increased utilization / virtualization benchmark as part of its ENERGY STAR credentials as these metrics become available.
- A recommended solution may be to add the following sub categories of servers (in addition to blades and DC servers) :
    - Volume server as a 1-2 or more processor socket server
    - Medium scale server as a 4 or more processor socket server, with ≥ 16GB of main system memory
    - High-End server as a 4 or more processor socket server, with ≥ 32GB of main system memory
- Blade chassis and blade servers are preconfigured and proprietary vendor by vendor. Consequently, it will be problematic to define common metrics as these proprietary implementations offer unique functionality and/or differences in energy performance
- As a definition, listing DC/DC based servers is fine. We would like more clarity on the EPA's strategy regarding DC/DC based servers. Currently DC/DC servers have been used in telecom environments and in other compact data center configurations. The building infrastructure impact, cooling considerations, and system configurations are different than rack based systems we observe in other data centers. We recommend a more detailed study on its configuration and usage model prior to creating an ENERGY STAR program for this class of machine.
- For the EMC classification question, the reason we should not include class B devices, is that class B is designed for consumer goods not meant to be maintained and serviced as are servers in a data center.
- In some places there are typographical errors : e.g. "Ac" should be "AC" and "Dc" should be "DC.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

More specifics on the definitions

Although Intel believes that the overwhelming majority of Servers are defined by single and dual socket processor platforms, Intel recognizes that 4 Socket server platforms (4S) are an important and emerging category. (* reference: IDC 2007 Worldwide Server Forcasts by Wayness)

The 4S platform is demonstrating itself in today's marketplace as the 'workhorse' for server consolidation though virtualization. As such, 4S platforms are increasing Server processor utilization and getting more "work" done for a given amount of power, primarily because applications on multiple smaller servers can be consolidated onto one larger server and the smaller servers can be retired.

Intel encourages the ENERGY STAR® Specifications to include 4S servers but strongly suggests that the EPA develop ENERGY STAR certification criteria that takes into consideration performance under varying load levels including, but not limited to Idle power.

We note that 4S servers typically have a higher level of idle power than less capable servers. The 4S platforms tend to have more memory at higher memory densities, increased capability in the chipsets and more sophisticated cooling and power management systems. The additional capabilities aid in delivering energy-efficient performance across the range of server utilization, but consume more idle power than less capable system.. These systems are optimum efficiency selections, delivering more work, and serving larger user communities for less energy consumed than multiples of 2S and 1S platforms which they replace. We recommend that 4S platforms be treated separate to the 1-2Socket systems if included within the qualifying products targeted.

As a reference point, below is chart that compares idle power in relation to system performance. The data is from the same HP DL380G5 2S server where the only variable that was changed was memory configurations.

The key takeaway from this chart is that servers that have slightly higher idle power (because of increased memory capacity) provide much greater performance as compared to servers that are configured to consume the least amount of power at idle.

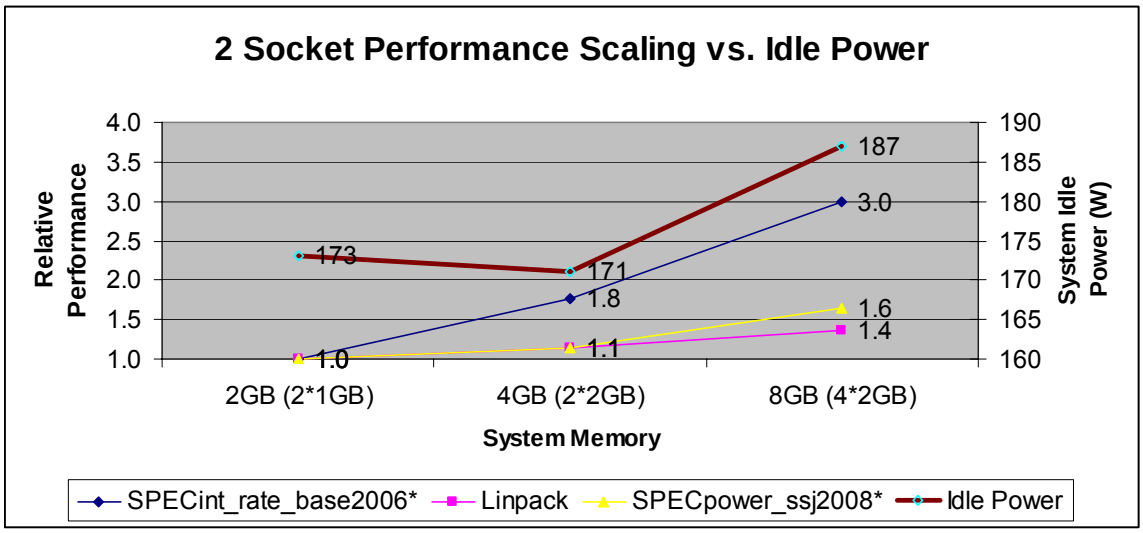

Note: The reason the idle power drops slightly from the 2GB to 4GB memory configuration is these are different DIMMs. Each DIMM vendor/model consumes slightly different amounts of power.

FCC Class B systems (designed for operation in the home) should not be included as the class of servers we're focused on is for the data center. A conglomerate of home based stand alone servers requires more cooling and power resources; and is rarely done in an enterprise environment. Class A devices are

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

designed to allow integration into rack or bladed systems, increase the compute density for efficiency and optimizing the use of the area's cooling and power delivery.

Qualifying Products

Products categories that would apply to the ENERGY STAR specification revisions should be prioritized. We recommend that those categories of servers be applied only when an appropriate assessment tool can be developed for that category of server. Based on the market volumes, available assessment tools, and energy impacts we recommend the following qualifying products priority:

- Volume server (i.e. 1-2 Processor Socket, 1-2U rack form-factor)
- Blades
- DC-DC server
- Medium scale server (i.e. ≥ 4+ Processor socket, ≥ 16GB main system memory)
- High-End server (i.e. ≥ 4+ Processor socket, ≥ 32GB main system memory)

Given several practical implementation concerns and the need to limit the scope of the ENERGY STAR program, we recommend that categories not included be explicitly stated. We hope to ensure that customers seeking ENERGY STAR server types that are not covered are able to continue their current purchasing activity until an assessment method exists and ENERGY STAR can include this category of machine in the program.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

Appendix

Document from EPA

ENERGY STAR® Revised Definitions for Computer Servers Based on Draft 1 Specification Comments

1) Definitions: Below are the definitions of the relevant terms in this document.

A. Computer Server: A computer that provides services and manages networked resources for client devices such as: desktop computers, notebook computers, thin clients, wireless devices, other computer servers and other networked devices. Computer servers primarily respond to requests and are accessed via network connections, and not through direct user input devices such as a keyboard, mouse, etc. For purposes of this specification, computer servers ***must include all*** of the following characteristics:

- Marketed and sold as a server;
- Designed and capable of having at most four processors (i.e., 1 - 4 individual processor sockets);
- Support for error-correcting code (ECC) and/or buffered memory;
- Dedicated management controller, such as Baseboard Management Controller (BMC), service processor, or ability to detect Wake On LAN (WOL) packets (Magic Packet and/or Directed Packet Filtering) to wake or power on from low power states;
- Include at least two ports for network communication capability, e.g., Ethernet, Fibre Channel, etc. (both ports can be the same technology);
- Include Reliability, Availability, Serviceability, and Manageability (RASM) features;
- Designed for and listed as supporting Server Operating Systems and/or Hypervisors, and targeted to run user-installed enterprise applications; and
- Designed and placed on the market as a Class A product as per EN55022:1994 under the EMC Directive 89/336.

Note: The purpose of including these detailed characteristics in the definition above is to (1) clearly delineate the types of computers covered by this specification and (2) separate servers covered by this specification from those products being addressed in the Version 5.0 computer specification, which is currently under development. Therefore, any "low end" computer server that does not meet the definition above will continue to be covered by the computer specification. For example, a desktop-derived server targeted to run user installed enterprise applications and which meets all of the requirements above is eligible for qualification under this server specification. All other desktop-derived servers, such as "home" or "media" servers, will continue to be covered by the computer specification. Lastly, servers with greater than four processor sockets are considered outside the scope of this specification, due primarily to their increased complexity, but may be considered in future revisions.

In the definition above, EPA is particularly interested in comments on:

(1) The use of Wake On LAN (WOL) as an alternative to a service processor or BMC

(2) The relevance of the Class A EMC designation and whether this should also refer to Class B under the directive.

Once this computer server definition has been established, EPA's next step will be to determine which types of servers will require further delineation based on their unique functionality and/or differences in energy performance. EPA has initiated this process as illustrated in the definitions for blade and appliance servers provided below.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

B. Blade Chassis: An enclosure containing shared resources for the operation of blade servers and blade storage. These resources include power supply(s) for power conversion, dc power distribution, cooling, network hardware, and system management, and may also contain shared storage. A blade chassis contains multiple slots which can be populated with a number of blades of different sizes and are often capable of accommodating blades from different vendors.

C. Blade Server: A computer consisting of, at minimum, a processor and system memory that relies on certain shared resources (e.g., power supply, cooling, etc.). Blade servers are designed to be installed in a blade chassis and are incapable of operating independent of the chassis.

D. Blade Storage: A storage element that relies on certain shared resources, contained in a blade chassis. Blade storage units are incapable of operating independent of the blade chassis.

E. Direct Current (Dc) Server: A computer server designed to operate with a dc-dc power supply or a server which runs directly off dc voltage supplied to internal dc-dc converters from an external source.

F. Server Appliance: A self-contained server system bundled with a pre-installed operating system and application software that is used to perform a dedicated function or set of tightly coupled functions. Server appliances deliver services through one or more networks (e.g,. IP or SAN), and are typically locked down systems managed through a web interface. Server appliance hardware and software configurations are customized by the vendor to perform a specific task, and are not intended to execute user-supplied software. Example services that may be made available via a server appliance include: name services, firewall services, authentication services, and encryption services.

G. Storage Equipment: A system designed specifically to provide data storage external to the server and which may be part of an archival process. While storage equipment may contain an embedded processor, this processor is not made available to execute user- installed software applications, but may execute data specific applications, e.g. backup utilities, data compression, install agents, and other tasks.

H. Network Equipment: A product whose primary function is to route Internet Protocol traffic among ports connected to it. In data centers this includes switches and routers.

Note: Definitions for blade servers and server appliances were created because they may require the development of separate requirements in future drafts based on differences in design and functionality compared to the larger server population covered by this specification. Additional types of servers may be defined in future drafts, as necessary. For example, it was suggested that EPA include definitions for network and telecom servers. EPA also created definitions for storage equipment, blade storage, and network equipment to ensure that these product types are properly excluded. Several of these definitions have been modified based on the Draft 1 comments. EPA is interested in stakeholder feedback on the above definitions, as well as feedback on specific product types (e.g., network servers, firewall servers, etc.) and where they fit best based on the definitions outlined above.

Computer Server Power Supplies

I. Computer Server Power Supply: A server component designed to convert high voltage input power to lower voltage dc output(s) for the purpose of powering the server. A computer server power supply must be separable from the main system and must connect to the system via a removable or hard-wired male/female electrical connection, cable, cord or other wiring.

J. Ac-Dc Power Supply: A power supply designed to convert line voltage ac input power into lower voltage dc output(s) for use by the server.

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124

K. Dc-Dc Power Supply: A power supply designed to convert dc input voltage from one level (including -48 V dc) into lower voltage dc output(s) for use by the server. Dc-to-dc converters (also known as voltage regulators) internal to the product and used to convert low dc voltage (e.g. 12 V dc) into other dc voltages for use by the individual server components are not considered dc-dc power supplies under this specification.

L. Single-Voltage Power Supply: A power supply designed to convert high voltage input power into only one single lower dc voltage output. Although single-voltage power supplies only provide one primary voltage for powering the server during operation, they may also include one or more lower voltage standby rails (typically 5 V dc) used only to initiate start-up when the main system is powered down.

M. Multi-Voltage Power Supply: A power supply designed to convert high voltage input power simultaneously into multiple different lower dc voltage outputs. There may be other low power DC output voltage(s) for standby power.

Note: This section has been revised to be more inclusive of dc-dc power supplies. EPA is continuing to investigate the technical feasibility of including servers that utilize dc-dc power supplies in this specification. The definition for single-voltage power supplies has been amended to indicate that power supplies with a dedicated standby rail are still considered single-voltage supplies for the purposes of this specification.

2) Qualifying Products: A computer server must meet the definition provided in Section 1.A, 1.B or 1.C above, to be eligible for ENERGY STAR qualification under this specification. Storage equipment, blade storage, and network equipment, as defined above, are not eligible for ENERGY STAR qualification under this specification.

Note: EPA plans to address storage equipment and network equipment under separate specification development initiatives. As a result, these product types are excluded from qualifying for ENERGY STAR under this specification.

When reviewing this document, stakeholders should consider the following important points:

- EPA's intention is to cover a broad array of server types that fit the technical definition provided in 1.A.
- EPA also intends to cover blade chassis and blade servers under this specification. However, since these server systems have a different approach and form factor from volume rack mount servers, EPA will have to determine which requirements might pertain to an ENERGY STAR qualified blade chassis (e.g., power supply efficiency) and which might pertain to a particular blade server itself (e.g., power management, idle, etc).
- EPA is also interested in the possibility of addressing servers using dc-dc power supplies in this specification. Feasibility for this approach will depend on the emergence of an appropriate and industry accepted test procedure that can be used to test these product types and determine comparable energy performance levels. If developed, EPA will consider using an amended power supply test procedure to accomplish this task.
- EPA will decide whether different types of servers may require separate energy performance criteria or levels. A decision to address these different types of servers will require the creation of separate "sub categories" of ENERGY STAR servers, with their own unique energy efficiency criteria and/or levels, and will be dependent on availability of energy performance data that supports this conclusion. Moreover, a justification for subcategories will be driven by significant differences in functionality and/or observable energy performance for a given workload.

Stakeholders are encouraged to provide feedback on this latest proposal to Rebecca Duff, ICF International, at rduff@icfi.com **by Friday, May 9, 2008.**

Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124