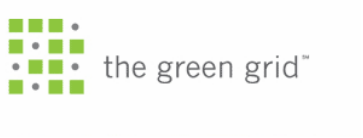

Response from The Green Grid to: ENERGY STAR® Program Requirements for Computer Servers Draft 1 (2/22/08)

The Green Grid Association, a consortium of industry leading companies welcomes the opportunity to comment on an early draft of topics under consideration for the ENERGY STAR for Computer Servers Specification. Some member companies of The Green Grid Association may in addition have provided additional considerations highlighted by their industry or company's particular perspective. Some members may have also provided their inputs through the Information Technology Industry Council (ITIC).

Introduction

A consortium of information technology providers, consumers and other stakeholders, The Green Grid Association seeks to improve the energy efficiency of data centers around the globe. The organization takes a holistic and comprehensive approach to data center efficiency and understands that addressing this challenge requires a high-level view of the entire data center and cooperation among a wide range of industry principals. Participants in The Green Grid include such diverse companies as major server and storage equipment manufacturers, major software providers, and large end-users / data center owners.

Overall

The Green Grid's feedback is referenced against the section numbers listed in the draft and is further identified by the title to place the comments in context. We applaud the logical approach in the outline and offer to provide additional support in the future development of the specification. We recognize the need to develop industry wide working groups and forums to assist in the technical development and are volunteering to support establishing these technical forums to facilitate this process. If you encounter any questions and/or wish to proceed with these activities, please feel free to contact Henry ML Wong, henry.l.wong@intel.com or Don Tilton, dtilton@spraycool.com.

Detailed Comments and Recommendation

Section 1. Definitions

The Green Grid (TGG) recommends the specification section be amended to substitute, edit, or add the following system definitions:

Computer Server: A computer that provides various processing, storage, and communication services in response to requests that originate and/or are mediated by other client computers and server computer systems. Computer servers have varying degrees of the following characteristics:

- Reliability, Availability, Serviceability, and Manageability (RASM) features
- Designed and certified to run Server Operating System and/or Hypervisors, and are intended to run general-purpose applications
- Baseboard management controller, service processor, or fully compliant WakeOnLAN capability.
- Include some type of network communication capability
- Designed to operate in a commercial data center environment (Electro-Magnetic Compatibility [EMC] and environmental rating).

Blade Server
A blade server (or any computer server) does not necessarily require a traditional hard drive or a spinning platen data storage device.

We request the addition of the following definitions:

Storage System: A system designed to provide data storage and may be part of data storage and archival process. While it may contain an embedded processor, this processor is not made available to run general-purpose applications. Please refer to the definitions provided by the industry consortia, Storage Networking Industry Association (SNIA).

Volume Server: A computer server as defined above packaged in either a 1U or 2U high rack-mount chassis. Although these dimensions may allow for number of socketed processors, volume servers are generally limited to 4 or less based on physical constraints. As implied by the rack mounting, these systems are generally purchased as part of larger computing structure.

Network Server: A computer server whose function is dedicated to service of network layer services (e.g. DNS*, Security, VPN*, Firewall, etc.) and not available for general compute function for applications. This class of server cannot have enabled sleep or standby features.

Telecom Server: A computer server whose function is dedicated to service of Voice services (e.g. VoIP* call control, Voicemail, etc.) and not available for general compute function for applications. This class of server cannot have enabled sleep or standby features.

Desktop Derived Server: DDS is a computer type specifically defined and covered under the ENERGY STAR v4.0 Computer Specification.

Direct Current (DC) Server: A computer server is driven by a direct current (DC) voltage power input. The AC to DC conversion may occur outside of the server providing DC power to multiple servers.

* Voice-over-Internet-Protocol (VoIP), Domain-Name-System (DNS), Virtual-Private-Network (VPN) are several commonly used acronyms describing networking technologies being used today.

Section 2. Qualifying Products

The ENERGY STAR Program for Computer Servers should apply to products meeting the definition of a *Volume Server* that contain a *Computer Server Power Supply*. A focus on *Volume Server*s is consistent with energy efficiency targets in EPA's report on server and data center energy efficiency*.

Note: Direct Current-powered servers, Blade Systems, Blade Chassis, Storage Systems, Networking, and Telecom systems should not be part of the Tier 1 compliance consideration of the ENERGY STAR Program Requirements for Computer Servers. Tier 1, as defined by draft 1 as the first phase of the program.

* Report to Congress on Server and Data Center Energy Efficiency, US EPA, Aug. 2, 2007

Section 3A. Power Supply Efficiency Requirements

TGG recommends that the EPA adopt the Climate Savers Computing Initiative's (CSCI) power supply (PSU) efficiency targets. Please refer to CSCI for the updated targets that has been based on the research conducted.

Note: The CSCI specification on power supplie units does not currently include any efficiency targets at load levels below 20%. We also expect the treatment of fans and fan power to follow the current testing standards established by the Climate Savers Computing Initiatives (CSCI).

Section 3B. Idle Power

Strictly assessing "Idle Power" on a server, does not provide any measure of efficiency or indeed the consumption level to expect from the server in deployment (even in an inactive state). TGG recommends that Idle Power be removed as a stand alone criterion in an ENERGY STAR program for servers for the following reasons:

a) Idle Power level on a server can not be consistently defined
b) Varies based on resident software for the target application suite
c) Varies based on the hardware configuration- so targeted as to support the SLA
d) Dependent on the firmware and operating system image
e) Dependent on resident security and maintenance activity

Unlike a client based computer, a server's idle power levels are highly configuration and connection dependent. A server's energy consumption includes expenditures to support a more sophisticated and complex computing network.

Section 3C. Standard Information Reporting Requirements

While TGG understands that EPA wants to aid consumers in the comparison of otherwise functionally equivalent server systems, SPEC contains specific publishing guidelines associated with the data on any system. The specific use of SPECpower™, highlighted in the draft is problematic since this may conflict with the reporting rules and obligations that SPEC has

designated. TGG recommends that that for any energy efficiency benchmark, the appropriate benchmarking organization be consulted and its publishing guidelines supported. Many of these guidelines were agreed upon to ensure accuracy and to promote proper representation of the data. Consistency and agreements with these organizations on reporting rules would also ensure continued maintenance and support of these benchmarks.

As noted in other feedback, SPECpower is limited in scope. TGG recommends establishing and hosting a work group to include:

a) the EPA,
b) the EPA consultants,
c) industry standards bodies,
d) other key stakeholders

to plan additional benchmarks and workloads. TGG does not believe that solely publishing SPECpower information without comprehension of additional workloads would fairly represent energy efficiency in those applications.

Section 3D. POWER AND TEMPERATURE MEASUREMENT REQUIREMENTS

Real Time Reporting: We believe that it is important that equipment be capable of reporting server power and thermal measurements in a standard format and we support EPA's desire to have all ENERGY STAR systems report AC power consumption and inlet air temperature.

We do not support specifying a specific data management protocol, as many IT and equipment manufacturers are marketing or developing systems to manage this data across a full range of OEM equipment. We believe that requesting that the information is accessible for 3rd party access and monitoring is sufficient. A real-time access to such information in an xml, csv, or other open format would allow any existing or future system to make use of the information.

Section 3E. Power Management and Virtualization

Many computer servers contain virtualization and power management features. Such features may include:

Dynamic Voltage/Frequency Scaling
Dropping into lower inactive power levels as workload reduces.
Processor Sleep
Variable Speed Fan Control based on power or thermal readings
Memory Power down
Lower Power Memory States
Lower Power I/O Interfaces
Export power and utilization to System administrators for action.
Power Capping
Virtualization

Tier 2 Requirements: (please also refer to section 5 below regarding schedules)
Any second phase should encompass adequate review of information established from Tier 1. The information includes the scope of the server categories, the benchmarking data and power and thermal sampling information supported. The minimum (fastest) server refresh cycle in a data center is 3-5yrs. Changes in this market are likely to require that length to observe changes in the installed base, and the value in the data collected.

Section 4. Test Criteria

Please refer to the Climate Savers Computing Initiative for the updated test methodology for power supply efficiency. It should also be noted that for functional testing, System manufacturers be allowed to use the Operating-System (OS) or Hypervisor of their choosing as long as it's a recognized enterprise system used by the product line. Many systems are shipped without an OS, pending an on-site end user approved OS image is installed prior to the integration and deployment of the system.

Section 5. Effective Date

A. Tier 1 best case would indicate an approximately Q4'08 implementation. A Q4'08 schedule is obviously dependent on the scope and testing details that are indicated in the specification.

B. Tier 2 should accommodate lifecycles for servers, which appear to be a minimum of 3-5yrs based on industry analyst studies.