

October 15, 2009

Environmental Protection Agency
ENERGY STAR Program

We appreciate the EPA's efforts at standardizing energy efficiency and consumption of computing devices, as reflected in this Preliminary Draft Version 1.0 Tier 2. We are grateful to submit our comment for your review.

Sincerely yours,

Jay Taylor (512) 728-3777
Sr Engineer Global Strategist
Dell, Inc. MS PS4-30
One Dell Way
Round Rock, TX 78682

We have reviewed the Preliminary Draft Version 1.0 Tier 2 Specification and have the following comment.

Dell considers cooperation and support of the Energy Star standards to be an essential component in the ongoing efforts to reduce the impact of climate change. Dell is still concerned with the developments of the Energy Star standard targeted at Servers. We consider the EPA and specifically the Energy Star program our partner in this effort.

Dell Detailed Comments and Recommendations for the Draft Document

Page 6

K. High Performance Computing System

HPC Systems are sometimes aggregated from other base systems, usually in clusters. The architecture of these systems are focused upon delivery of an outcome data set of values or capture of those values for later transformation. As such these systems are usually custom in design and architecture. They should be excluded from consideration.

Page 9

CC. Base Configuration

The base configuration definition is irrelevant to the standard, and would be revised continually in the standards. The addition of this definition adds complexity to the standard unless there is a specific deliverable related to this definition. We recommend the definition be removed unless a criteria is to be developed from it.

Page 10

EPA intends to review additional product types for inclusion in the Tier 2 specification as previously communicated. EPA's intent remains to have widest reasonable/feasible scope under the general definition of Computer Server, allowing the manufacturer community to qualify a diverse array of products that fit primary definition. Specific server types noted for further investigation include *systems with greater than four sockets, Blade Systems, Fully Fault Tolerant Servers, Server Appliances, and Multi-Node Servers.*

Stakeholders have communicated to EPA that the majority of the market is covered by 1S, 2S, and 4S servers in rack, pedestal, and blade configurations. This is largely the same scope as present in Tier 1, with the addition of blades.

Because broad coverage remains ideal for the program, EPA will consider information on server types outside of this range before determining a final scope for Tier 2.

Monolithic servers Pedestal, 1U, 2U and 4U rack servers remain the majority of the market. Additional servers identified represent a small portion of the marketplace.

Since blade servers have been identified as an area of interest; it should be noted that significant architecture differences exist among the market participants. Areas of commonality do exist, they all have power supplies, features in the chassis versus the blade, number of blade servers per chassis, additional components such as storage and networking devices located in the chassis or in a blade slot indicate a diversity of architecture that does not allow a ready comparison among manufacturers. Dell's concern is that specifying one or more performance characteristics of a blade server may reward certain architectures that are not the most efficient at delivering total or specific performance capabilities. Indeed the architecture of blade servers is such that the configuration complexity of the systems delivers critical performance characteristics due to and by architecture.

We propose that; like the Tier 1 specification for 3 and 4 socket server systems, PSU efficiency, and systems measurements be required for blade server systems, however Idle and benchmark performance measurements are reported on the power and performance data sheets. We do not consider the option for short term evaluation methods for an intercept with Tier 1 to be a viable option. The complexity of the specification coupled with the significant work required to deliver the methods and measurements are still incomplete.

For Dell servers that are greater than 4 sockets, fully fault tolerant, multi-node, and resilient and HPC represent a small fraction of the total populations of servers available or shipped by Dell.

Page 13

Net Power Loss

Net power loss is a term not utilized in the industry in any context outside of definitions delivered by the Energy Star program. This to be developed criterion would require the industry to identify value, explain and educate customers in an already complex sales environment. We continue to recommend dropping this approach.

Page 16

Power Supply Efficiency

Power supplies are long lead items for servers, usually custom designed and requiring extensive worldwide safety and regulatory certifications. The expense of these power

supplies is such that the continuous revision of the power supply efficiency requirement in the standard is a cost burden for industry. We recommend keeping the existing power supply efficiency limits for the next Tier.

The multi output power supply efficiency should line up with the Energy Star 5.0 requirements for small scale servers.

Page 17

Active Mode Power Requirements

The addition of active mode power requirements will be gated by a tool that delivers this criterion. Dell's concern remain that the tool be delivered soon enough in the process that the extensive testing required may be delivered without haste. Careful study of the impact this tool has on architecture must be performed and an understanding of architecture and component tradeoffs must be understood before moving this forward as a criterion. We remain concerned on the development of this tool and the timeline associated with delivering the Tier 2 specification.

The associated text and description for Table 3, needs clarity. The mix of minimum and base configuration references within the text conflicts with the table and text in grey.

Dell supports the current effort by SPEC Power to generate a Server Efficiency Rating Tool. This is currently offers the best approach to establishing an active mode rating for computer servers. The SPEC efficiency rating tool is an industry initiative with broad support and should provide a robust metric tool which focuses on the both active and idle mode energy using characteristics of a server. We propose support of the SPEC efficiency rating tool as a priority for the Tier 2 Energy Star for servers specification.

Additionally, Dell recommends the Energy Star program consider these modifications to enhance the applicability of the rating tool and simplify the overall evaluation of power use by computer servers:

1. The rating tool should focus on scaling performance and power capabilities of the functions common to all servers, including but not limited to; processors, both floating point and integer computations, disk, I/O, and memory access. A tool which focuses on functional energy efficiency performance offers the best approach to providing an effective comparison of like systems with a range of configurations and processors.
2. Idle power requirements should be integral to the evaluation tool, allowing the server to receive an overall score, incorporating active and idle power performance. It could be advantageous to provide a higher weighting to the idle power requirements in calculating a score from the tool, but including idle in the metric provides the opportunity to normalize the score for more capable processors which are penalized

in the current standard for the 1 and 2 socket “idle” only criteria. A separate idle power limit would not be necessary and not advised, as a standalone idle test would be redundant to the information delivered from the tool, increase costs, and complexity of compliance testing.

Page 20

Supervisor Power Management

This text should be clarified or removed. The reference is to supervisor power management is for the OS, hypervisor, or management controller? The previous row in the table already identified the Energy Star requirement for implemented power management features.

Page 20

Other Requirements – Energy Efficient Ethernet

The requirement of implementing this standard upon release does not consider the business impact of implementation. The standard is not released and not final, silicon and software impacts are not understood. We agree with adoption of new standards that offer the possibility of reduced energy consumption or improved efficiency as require by customers or implemented by industry. We disagree with requiring implementation of technology on behalf of customers (Federal or not) without a strong customer preference or requirement. In our conversations with federal customers we are not getting this technology as a demand.

Page 21

Standard Information Reporting Requirements

The Power and Performance Data Sheet and QPI forms require the same information in two different formats. (i.e PSU Efficiency and PFC vales) Also many of the entries are duplicated. A consolidated sheet should be considered. Also, as previously reported, for family declarations the QPI improperly calculates Maximum or Minimum configuration values for system memory or disk capacity. Families that have common form factor but different wattage power supplies are required to identify two families, for example.

The are literally only a few data sheets available for review from OEMs or the Energy Star web site. The section outlined by line 554 would seem to indicate that complaints have been registered about the documents currently available. Further clarification is necessary to understand what message is being delivered. In addition the complexity of reporting through the power and performance data sheet and the QPI is such that there are many opportunities to have misunderstanding or improperly listed products as a direct result of the programming in the sheets.

Page 23

Processor Utilization Measurements

Dell disagrees with the approach outlined. Adopting the measurement proposed implies that the operating systems, or hypervisor reporting of CPU utilization is somehow inaccurate or in need of improved accuracy for improved efficiency. Also, OEM's are not in control of the OS or hypervisor algorithms that deliver CPU utilization and would require concurrence of the major suppliers of operating system and hypervisors to achieve the Energy Star criteria.

Page 23

Sampling Requirements

Minimum reporting intervals or sampling intervals are needed to provide data on server operation, such as a minimum sample rate of one sample per second. Is the intent of the Energy Star program to issue a standard or methodology for averaging algorithms or measurement reporting algorithms for servers?

Dell supports revealing standardized data measurements. The system for collecting and reviewing the data is crucial. We request language that supports or acknowledges that the operating system reports processor utilization and that a system may report power consumption and air temperature; however latency in workload and network loading may delay the availability of revealed information.

Thank you for the opportunity to comment.

Jay Taylor (512) 728-3777
Sr Engineer Global Strategist
Dell, Inc. MS PS4-30
One Dell Way
Round Rock, TX 78682