



**TO: Andrew Fanara
Energy Star Office
U.S. Environmental Protection Agency**

**FROM: Ken Salaets
ITI**

DATE: October 21, 2009

SUBJECT: Preliminary Draft ENERGY STAR[®] Tier 2 Server Specification

I apologize yet again for the delay in submitting comments. Our detailed comments and recommendations follow. As always, we welcome the opportunity to provide additional information or to respond to questions. You can reach me at 202.626.5752 or ksalaets@itic.org.

Page 6

K. High Performance Computing System

HPC Systems are often aggregated from other base server systems, and organized, usually in clusters. The architecture of these systems is focused upon delivery of an outcome data set of values, or capture of those values for data transformation and storage. As such, these systems are typically “custom” in design and architecture. Industry requests they 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 criterion 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 server Pedestal, 1U, 2U and 4U rack servers continue to account for the majority of products on the market. The additional servers identified represent a significantly smaller portion of the marketplace.

Blade servers have been identified as an area of interest; it should be noted that significant architecture differences exist among the market participants for blade systems. Areas of commonality do exist. For instance, they all have power supplies. However; features in the chassis versus on 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 all indicate a diversity of architecture that does not allow a ready comparison among manufacturers or against monolithic architectures. Industry's concern is that specifying one or more performance characteristics of a blade server may reward specific architectures that are not the most efficient at delivering total or specific performance capabilities. It is also not possible to anticipate this reward in advance. 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 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 industry.

Page 13

Net Power Loss

Net power loss is a term not utilized in the industry or by customers in any context outside of definitions delivered by the ENERGY STAR program. If EPA insists on mandating this as developed criterion, it will require the industry to identify the value, and then have to explain and educate customers in an already complex sales environment. We question the utility and necessity of this burdensome requirement, and continue to recommend dropping this approach.

Page 16

Power Supply Efficiency

Power supplies are long-lead items for servers. Typically, they are custom-designed and require 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

creates a significant burden for manufacturers, while returning minimal return on the investment. We recommend keeping the existing power supply efficiency limits for the next Tier. In addition, we recommend aligning the multi-output power supply efficiency 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 the tool that delivers this criterion. Industry's concern remains that the tool needs to be delivered soon enough in the process so that the extensive testing that will be required may be delivered without haste or churn. 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.

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

Page 20

Supervisor Power Management

This text should be clarified or removed. Is the reference to supervisor power management for the operating system, the hypervisor or management controller? The previous row in the table already identified the ENERGY STAR requirement for implemented power management features.

Other Requirements – Energy Efficient Ethernet

The requirement of implementing this standard upon release unreasonably fails to take into account the significant business impact of that implementation. The standard is not currently released or final, the impact on silicon and software is not fully understood, nor is the requirement currently incorporated in the silicon roadmaps for device delivery.

Industry supports the notion of adopting new standards that offer the possibility of reduced energy consumption or improved efficiency as required by customers or implemented by industry. We disagree, however, with requiring the implementation of technology on behalf of customers (Federal or otherwise) in the absence of a strong customer preference or requirement. In conversations with Federal customers, we are not hearing a demand for this technology.

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 values. In addition, 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. For example, families that have common form factors but different wattage power supplies are required to identify two families.

There 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 it creates numerous opportunities to have misunderstandings or improperly-listed products as a direct result of the programming in the sheets. Consolidation of the two formats can resolve these issues.

Page 23

Processor Utilization Measurements

Industry disagrees with the approach outlined in this section. Adoption of the proposed measurement implies that the operating system or hypervisor reporting of CPU utilization is somehow inaccurate or in need of improved accuracy for improved server or data center efficiency. Also, OEMs are not in control of the operating system or hypervisor algorithms that deliver CPU utilization, and would require the concurrence of the major suppliers of operating system and hypervisors in order to meet the ENERGY STAR criteria. The approach should be dropped.

Sampling Requirements

We request that the agency continue to allow a reporting of a 30 second average, sampled on 1 second intervals, every 30 seconds. A data center operator is unlikely to be analyzing power usage on a frequency smaller than 15 minutes, so use of a rolling average or a 30 second average reported every 30 seconds does not materially impact the data center operator and the use of the information. EPA should replicate the Tier 1 Sampling Requirements in Tier 2. Many, perhaps more scalable, high-level data collection tools cannot sample faster than 30 seconds and are often configured to sample more slowly.

Industry does support revealing standardized data measurements. However, the system for collecting and reviewing the data is crucial. We request the inclusion of 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.