



**TO: Andrew Fanara, EPA
Becky Duff, ICF**

**FROM: Ken Salaets
Director**

DATE: March 25, 2009

SUBJECT: Draft 4, ENERGY STAR® Computer Server Specification

ITI's Energy Star Working Group has reviewed the ENERGY STAR Draft 4 Computer Server Specification (hereafter, "specification") and appreciates the opportunity to offer the following comments and recommendations after the initial deadline. As always, we welcome any questions that you may have relative to this comment or any other issue.

EFFECTIVE DATE

As we have stated in the past, the proposed implementation date for introduction of the ENERGY STAR for Server specification is insufficient for deploying a standard of this magnitude and complexity across the server industry. Although EPA cites an immediate effective date for new product categories as "standard practice," we are not aware of any previous IT office specification being implemented without giving the market sufficient time to adapt to new ENERGY STAR specifications. In addition, we are unaware of any compelling reason to accelerate the effective date of the server specification ahead of the 270 day "appropriate lead time" cited in the of Energy Policy Act of 2005 "for a new or a significant revision to a product category, specification, or criterion..."

Clearly, the proposed specification qualifies as a "new or significant revision of a product category." Accelerating its effective date will have a detrimental impact on the manufacturing and distribution process, particularly in light of the Federal procurement mandate for ENERGY STAR listed products, and could create the same competitive disadvantages cited by EPA in delaying implementation of the Thin Client specification. For these and other reasons, we urge you to delay implementation of the server specification until a date much closer to the 270 day metric established by Congress.

We are also concerned about the possibility of as many as three additional revisions before October 2010. This would place an immense and unreasonable burden on manufacturers, and would prove disruptive to development cycles of future server families and corresponding silicon architectures.

IDLE MODE

We have concerns with the agency's emphasis on limiting "idle mode" consumption without a performance context, as it has the potential to drive the IT industry and server manufacturers to the lowest configuration for delivering the ENERGY STAR limits. EPA has mitigated this concern somewhat with the processor power management criteria for 4 socket servers, as it enables manufacturers to provide high performance, high capacity ENERGY STAR® qualified solutions to its clients. The 4 socket server solutions will offer our clients a system capable of virtualization and high utilization with reduced power when the equipment is idle; delivering the most workload per unit of power applied. This will allow IT uses to tradeoff the benefits of more simplified ENERGY STAR® systems against fewer higher performance servers as documented in the Green Grid Members Meeting Presentation given in February by Henry Wong and Jay Taylor. In a data center containing hundreds to thousands of servers, a higher percentage of lower-configured ENERGY STAR listed servers could translate into higher energy consumption when compared to the use of highly utilized 4 processor servers with power management..

Additional recommendations:

Blades. Treat Blades like 4S systems (Power Mgt, reporting data, and no idle limit). The different architectures and overhead distribution practiced by blade manufacturers, the complexity of a fully populated, "multi-node" blade chassis, the difficulty of the proposed EPA blade testing method, and the fact that a subject of this complexity will take 2 or 3 additional iterations to reach a consensus idle criteria, suggests that requiring enabled processor power management on blade seems is a workable solution to address blades in the final draft. This would also limit the amount of discussions and development requests that could divert industry resources from, and slow down, the Tier 2 energy efficient performance metric(s) development.

SAS Controllers. Based on existing data on SAS controllers (needed for performance enhancement), we recommend a 10W SAS controller (add-in card) adder.

ENERGY STAR DOCUMENTATION

In an effort to reduce impact to the environment, all materials such as data sheets should be provided in a paperless requirement or other deliverables, such as soft copy or web delivery. For the purpose of efficiency, reduced cost and maintaining the greenest of standards, ITI recommends that all ENERGY STAR-related marketing and technical materials be distributed in electronic format. For example, a manufacturer could establish a central, self-hosted web site, or web pages on its corporate site, with all information pertaining to ENERGY STAR compliance.

PROCESS IMPROVEMENT

We note with that the text in the various notes of Draft 4 exceed the lines of the draft specification itself. This observation speaks to the significant complexity of the subject. We propose that EPA consider adopting a roadmap approach to improve the process of developing or revising specifications. This roadmap would provide focus or targets for improvement, and would allow if properly implemented for ENERGY STAR features to be included in feature sets identified by marketing for implementation by the development teams. A roadmap approach would allow for less churn and cost when implementing the ENERGY STAR requirements.

Several technologies or approaches have been proposed for Tier 2. Of specific note are the proposals on Energy Efficient Ethernet and Net Power Loss. We offer the following comments on these two proposals:

Energy Efficient Ethernet. We are unaware of any devices on the horizon that will be on the market within the term of the proposed server specification. The silicon and technology would need to be available now in order to meet the suggested October 2010 implementation timeframe. Accordingly, we question the need to include such a specification.

Net Power Loss. As a practical perspective, net power loss is an effort to quantify the “right sizing” of a power supply to the server. In general, the IT industry agrees with right sizing power supplies to the load. The challenge is explaining these requirements or approaches to customers of ENERGY STAR systems. We submit that “appropriate sizing” is a better approach, where the power supply rating is matched with the potential load. For example, the metric could state that a power supply unit is not to exceed 120% of the maximum ENERGY STAR configuration power of the system. We believe this would be a more practical approach both in terms of customer requirements and energy efficiency.

POWER SUPPLIES AND POWER FACTOR CORRECTION (PFC)

In Draft 4, PFC is specified down to 50 watts, citing as a reference another ENERGY STAR specification that deals with single output power supplies. ITI questions the proposal to specify PFC under 100 watts. It is not a requirement for an external power supply (EPS, also a single output device) in the ENERGY STAR EPS specification. In addition to the Active Mode efficiency requirements found above, power supplies with greater than or equal to 100 watts *input* power must have a true power factor of 0.9 or greater at 100% of rated load when tested at 115 volts @ 60Hz.

ITI questions the proposal to specify PFC under 100 watts for a number of reasons. For example, the input power measurement is not achievable, as we cannot provide the 10% accuracy in the low range of the power supply operation (0 to 15% load). Further, a power meter on 2000KVA power supply 10% at full range is +/-200 watt. If run it at 1% load (20 watts +/- 2 watts). Typical power meters are specified for full range tolerance. This gives very poor tolerance or accuracy at lower loads. Due to this we have gone with stating that from 15-100%

range we have 10%. Below 15% power supply loading, the 10% accuracy level cannot be maintained. ITI recommends that this be acknowledged in the ENERGY STAR specification.

For power supplies, ITI proposes the >1000 watt power supply specification set at the Climate Savers Silver for the Tier 1 spec. We could support utilizing the Gold levels for Tier 2. ITI also proposes removing the rolling data requirement, since industry power supplies perform power measurements synchronized in time periods (an average generated at 30 seconds intervals) instead of a rolling average. Most systems cannot afford to poll more often than every 30 seconds, because the data collection becomes overwhelming and chews up significant register space, resulting in slower system performance. In addition, significant average movements will not be identified or tracked on a 30 second interval or faster interval by the operating system.

IDLE MODE CRITERIA

Clarity is needed in this area. Sockets may be installed with single or multiple processors, or cores. The idle mode limits from the specification do not allow for flexibility of configuration of cores; and will reward single core socket servers over multi core socket servers without regard to performance or the compute delivery requirements of the data center.

Significant investment by the IT industry and server manufacturers in time and effort is required to deliver minimum, typical and maximum values. We received feedback from our members that maximum power consumption is not always aligned with maximum configuration or maximum performance, and is influenced by application (benchmark tool), BIOS, operating systems or hypervisor settings, and by component suppliers.

In the case for the product family, and meeting the idle criteria at the minimum and maximum configuration, the EPA should accept that the other models are compliant if:

- a) All components for a given model are built to a specific power spec and the variation should not be more than +-5% across any number of manufacturers.
- b) The idle power curve for a given set of configurations will look like a power supply efficiency curve as the complexity of the configurations increase. The idle power will vary some around that curve, but in general the two end points will appropriately define the model and processor combination.
- c) The VAR delivery has created additional dilemmas; it leaves the VAR two choices:
 1. Keep its value add services within the configuration defined by the manufacturer data sheet, which limits the VARs ability to meet the customer's needs and the real purpose of the VAR.

2. The VAR will have to be an ENERGY STAR partner, provide shipment information (which may overlap with the manufacturers report as well, so there will be double counting and there will be too many partners for EPA to deal with).
3. The best solution is to designate the ENERGY STAR qualification at the OEM.

Blades are excluded from the scope of this document. However; ITI would propose that ENERGY STAR consider including blades within the scope of this specification, in a manner similar or identical to the 3 and 4 socket servers. That is power management requirements, reporting requirements and power supply efficiency limits, and posting of idle data. A concern ITI has is when identifying minimum, typical and maximum blade configurations, is the context referencing installed blades in chassis, or configurations of blades.

Our members indicate that blade chassis are almost never shipped or operated with a single blade server, although that would be a minimum configuration. Testing costs and time could be reduced by allowing a measurement based on a blade configuration in a chassis, multiplied by the blade count.

Language should be inserted into the power monitoring section to explicitly state that “accuracy requirements are limited to system operating conditions.”

ITI supports revealing standardized data measurements. The system for collecting and reviewing the data is also important. We recommend the inclusion of language that supports an understanding that the operating system reports processor utilization, and that a system may report power consumption and air temperature, while also acknowledging that latency in workload and network loading may delay the availability of revealed information.

ADDITIONAL RECOMMENDATIONS

To facilitate rapid closure to Tier 1 in order that EPA and industry can focus our collective efforts on developing energy efficient performance metric(s) for Tier 2, ITI offers the following additional recommendations:

Clarification on Idle limits. We recommend that the following text be added prior to the idle limit chart: “Idle limits are applicable to 1- and 2-sockets systems.” The clarification would avoid confusion with other product categories previously mentioned such as appliances, fully-fault tolerant, and 4-socket systems.

Classifications and edits. For classifications and testing, the industry appreciates EPA providing a product family qualification route to simply management of ENERGY STAR products for EPA, our customers and ourselves. Industry believes, however, that we can reach a better definition of the maximum and minimum configuration test points, the associated testing requirements and the qualification status of all of the configurations within a product family.

Tier 2 Requirements. For Tier 2, we are pleased to and anticipate working with the EPA in the development of energy efficient performance metric(s) in conjunction with other stakeholder organizations. We are, however, concerned with inclusion of the energy efficient network specification and requirements and the storage targets to Tier 2. Given the preliminary nature of specifications and products that could meet IEEE802.3az, and the differing characteristics of what storage, we do not recommend including these areas as part of the Tier 2 development.