September 19, 2008

To: Rebecca Duff
ICF International

CC: Arthur Howard
ICF International

Andrew Fanara
United States Environmental Protection Agency

Re: Hewlett-Packard Response to the *ENERGY STAR® Program Requirements for Computer Servers* **Draft 2** Feedback

From: Hewlett-Packard Company (HP), Enterprise Storage and Servers Business Unit

This document may be published on the Energy Star website.

Hewlett-Packard welcomes this opportunity to again provide comments on *ENERGY STAR® Program Requirements for Computer Servers* (hereafter in this feedback document called "Energy Star for Computer Servers") — Draft 2. HP is proud to continue our long standing association with the *Energy Star* program. The comments and issues in the Draft 2 document fall into two categories and are covered in the two major sections of this review:

1. *Draft 2 Energy Star* Partner Requirements
2. *Draft 2 Energy Star* Product Eligibility Requirements

Several key issues are detailed on the pages below:

1. Draft 2, Energy Star: Partner Commitments

1.1. The requirement in lines **33-34** stating that "The ENERGY STAR mark must be clearly displayed on the front or side of the product…" is in direct conflict with statements made by Andrew Fanara at the Redmond, WA stakeholders' meeting. Hewlett-Packard and other manufacturers in attendance were all clear in their opposition to this requirement at the stakeholders' meeting, and Andrew's statement led us to believe that this issue was closed.

There is no room on the front of servers for physical Energy Star labels and putting a label on the side would be invisible to users. Due to the configurability of servers, some SKUs can't be Energy Star, so the factory costs of adding a label only to servers SKUs with certain parts installed would be unmanageable.

1.2. On line **30** the phrase "When a Partner qualifies **the** product, …" should say "When a Partner qualified **a** product, …"

September 19, 2008

2. Draft 2, Energy Star: Product Eligibility Requirements

The following is a compiled list of comments on the line numbers listed in the document. Comments are broken into two general lists. The *Substantive Feedback* section lists substantial changes to the approach documented in Draft 2, while the *Content Feedback* section details issues that are important to clarify the intent of the document.

2.1. Substantive Feedback

2.1.1. HP supports the idea of having different pass/fail criteria for low redundancy and higher redundancy servers. However, lines **181-186** are inadequate to describe the rich diversity of features that are possible while delivering a wide range of Reliability, Availability, Serviceability and Manageability (RASM) features.

Describing servers with <u>no</u> redundancy as "Standard Redundancy Servers" makes no sense. "High Redundancy Servers" have no single definition. RASM features in servers cannot be provided for zero idle power, and the amount of RASM features varies widely. RASM features may include RAID disk controllers (and the required number of hard drives to deploy the chosen RAID level), hot plug hard drives and support circuitry, system management controllers, hot plug I/O cards, ECC memory, hot-plug and/or redundant fans, hot plug and/or redundant power supplies, et al.

2.1.2. Lines **236** and **237**: The 1000W choice in Draft 2 for a division of high wattage vs. low wattage single-voltage power supplies seems arbitrarily chosen and should be increased to at least • 1200 Watts and >1200 Watts. From a technology perspective, 1500 watts is the approximate design crossover point for large and small power supplies.

2.1.3. Line **236**: Allowing multi-voltage power supplies to have an easier set of thresholds, without bounds for wattage, provides incentives for server manufacturers to move from single-voltage power supplies to less-efficient multi-voltage power supplies.

2.1.4. Lines **236** and **237**: We reiterate HP's previous objection to the additional and unnecessary testing required for efficiency and power factor at both 10% and 20% power supply loads. 20%/50%/100% loads are the industry-standard test procedure to verify efficiency curves for a range of loads.

2.1.5. Lines **236** and **237**: If the 10% load requirements are retained, then (similar to multi-voltage power supplies) the 10% pass/fail threshold for single voltage power supplies should <u>not be required</u> if the power supply is placed into a server that has the ability to prevent power supply loading from going below 20%. This prevention of <20% loads can be done by: 1) "right-sizing" and/or 2) via power supply technology that shifts loads and turns off unnecessary power supplies and/or 3) by shipping without n+n power supply redundancy.

2.1.6. Line **237**: Power factors specified at 10% loads are difficult to meet. To meet the strict efficiency standards at light loads, larger power supplies require operating in a burst mode. This mode degrades the ability to correct power factor. This tradeoff makes meeting both high power factor and high efficiency requirements very difficult. If you reassess your data, you will find that a smaller percentage of the power supplies in your sample are able to handle both the power efficiency and power factor requirements in Draft 2.

2.1.7. Line **238**: Regarding inclusion or exclusion of DC-DC power supplies, industry sales of DC-DC power supplies are less than 1% and there is an industry trend moving DC power from ±48 VDC to 380 or 400 VDC. Additional concerns have to do with large efficiency differences between ±48VDC power supplies that either: 1) comply with telco (NEMA) requirements to handle wide voltage swings in the -48 VDC battery-backed power distribution, or 2) expect their ±48 VDC to be well-regulated (non-telco power). While power supply efficiencies in the latter case are likely to be better, the environments are not the same.

2.1.8. Line **238**: HP suggests that it be possible to amend the "Generalized Test Protocol for Calculating the Energy Efficiency of Internal AC-DC and DC-DC Power Supplies", without revising Energy Star v1.0 for Servers. Adding a 380 VC or 400 VDC test methodology are two of many possible additions that may be needed during the life of this Energy Star specification.

2.1.9. Lines **240-248**: HP objects to Energy Star specifying pass/fail idle power thresholds for servers and prefers that idle power only be used as an item listed in the Qualified Product Information standard reporting sheet.

Pushing for specific idle thresholds will cause the unintended consequences of either 1) an unbounded number of thresholds that have to be accurately specified and tested for all possible server features, or 2) only a few de-featured models will be able to meet the fixed idle power thresholds described in Draft 2.

Any single idle power threshold is meaningless without a defined set of specific RASM features. Idle power thresholds are equally meaningless if there is no correlation to a measurement of peak performance on a valid benchmark for that server type and application environment. **Servers with more performance capability should get proportionally higher idle power thresholds.**

Tables 3 and 4 are gross oversimplifications of the rich diversity possible in 1P, 2P and 4P servers. 1P, 2P and 4P categories bear no direct relation to the compute capacity of the server. It is also very unclear whether the column headings are talking about capabilities for memory expandability or installed memory capacity, and regardless of the answer, direct comparisons and wattages must be for specific memory sizes. Memory requirements grow as the compute capacity grows, and compute capacity grows when servers have more performance, more CPU cores, and/or more processor sockets.

You should not equate 1P, 2P or 4P processors that have different inherent performance levels. In a direct comparison of "like" processors, frequency differences, # of processor cores per socket, and low-power screening provides extreme variations in power and performance, so (e.g.) all "2P" servers cannot be equated in an idle power threshold. On dissimilar processors (including generational differences from the same processor vendor), the number of sockets, the frequency and server performance have no correlation.

"Power per DIMM" is a non-linear equation. E.g., adding one pair of DIMMs may require more wattage per DIMM than adding two pairs of DIMMs.

2.1.10. Lines **247-248**: While we support the concept of larger memory installations getting more power allocations, by picking • 16GB and <16GB as the only two choices of memory size, this gives unfair bias against servers that need memory sizes other than those that are optimized by those two choices. E.g. 32GB servers could not pass with thresholds set for a 16GB server.

2.1.11. Lines **240-248**: The methodology for including blade servers and blade enclosures in idle power thresholds is undefined in Draft 2. Since an enclosure full of blade servers often provides performance and RASM equivalent to an equal number of rack servers, rather than putting blades into a different class of product, it would be fair to let blades compete directly with rack-mount servers.

The power associated with blade server solutions are generally lower power (per server) than their rack-mount server equivalent. This enables blade servers to compare directly to rack-mount server thresholds. Encouraging blades is a good way to create incentives for customers to move to more efficient blade solutions, rather than penalizing blades' power efficiency excellence with tougher standards.

Complicating some direct comparisons is that blade enclosures often include network and/or storage switches and other functionality. Rack servers connect to external LAN and SAN switches and routers. The power of the switches in the enclosure is less than if those switches were implemented externally. Power threshold allowances are needed to support the extra functionality (like network or storage switches) that may be in a blade enclosure.

2.1.12. Line **283**: AC-DC multi-voltage power supply testing should also be allowed at 230 VAC, but with elevated power supply efficiencies to reflect the better efficiency expected at 230 VAC. A large majority of multi-voltage server power supplies (and most desktop/notebook PC power supplies) can handle a range of voltages from 100-240 VAC ±10%. With 230 VAC as the only voltage available in most of the world (e.g. Energy Star's European Union partner), it makes sense to add a set of 230 VAC thresholds for multi-voltage power supplies.

2.1.13. Line **295** (and elsewhere) states that Energy Star will recognize the top 25% of models "currently available" in the marketplace. The power supply efficiency and power factor thresholds chosen in Draft 2 make the 25% goal an impossibility on January 1, 2009. Additionally, there is low confidence that idle thresholds will be defined in such a way as to recognize the actual top quartile of server models. The pass/fail idle criteria give unfair advantage to de-featured servers and inaccurately presumes widespread availability of ultra-high-efficiency power supplies in that timeframe.

2.1.14. Lines **306-323**: Appendix A does not adequately comprehend the inclusion of Blades in Tier I. Under "System Characteristics" There should be a comprehension that blade enclosures and blade servers disaggregate some aspects of a rack-mount server and may have additional features that are unlike rack-mount servers.

Blade enclosure parameters should include:

- Number of blade slots available for servers
- Number and type of blade servers installed to yield the certified results
- Capacity of shared storage in the enclosure (storage <u>not</u> in blade servers)
- Number of switch blade slots
- Number and type of switches installed (10 Gigabit Ethernet, Fibre Channel, etc.)
- Number of enclosure management processor modules installed

September 19, 2008

2.2. Content Feedback

2.2.1. Lines **139** and **140** are unnecessary. The only server type excluded by this paragraph is a server with zero "processor sockets", which is impossible, unless the intent is to exclude systems that solder down their processors without the need for a removable "socket".

2.2.2. The term "processor socket" on lines **139, 219, 220, 224, 245** and **308** is not defined and begs the question as to whether this term applies to a processor that is soldered down and not in a removable socket.

2.2.3. Similar to 2.2.2, there is a lack of definition of the term "processor board" on line **140**. A processor board plugs into a type of socket, just as some single-chip processors may plug into a different kind of socket.

2.2.4. Lines **156-158**: The term Direct Current Server is defined but not used elsewhere in the specification. This type of server is still a server as defined by the term "Computer Server" in lines **128-142**. Deciding to include or exclude on the basis of a DC-DC power supply doesn't require the definition of a unique server type. A Blade Chassis might also have a DC-DC power supply.

2.2.5. Similar to 2.2.4, Lines **181-186** define two sub-types of Computer Servers. Since Lines **216-217** say that only Computer Servers, Blade Servers and Blade Chassis are eligible for Energy Star, then it begs the question whether servers defined in 1E, 1I and 1J are eligible or ineligible.

2.2.6. The definition (and exclusion in Tier I) of "desktop-derived servers", as another sub-type of Computer Servers, should be considered for addition to the list of defined server types. Spelling out the exclusion of desktop-derived servers needs to be included on line **217**.

2.2.7. HP approves of the move to align with Climate Savers Computing Initiative (CSCI) efficiency and thresholds and test methodologies. HP prefers a complete alignment with CSCI power supply efficiencies and power factors (e.g. removing 10% load requirements). The Net Power Loss efficiency threshold idea mentioned on line **238** would be very difficult to implement.

2.2.8. Line **237**: If DC-DC power supplies are added, then power factor is obviously not applicable to them.

2.2.9. Line **238**: HP approves of the decision to exclude fan power for the purposes of efficiency testing. This is the only fair way to compare power supplies that may or may not have power supply fans.

September 19, 2008

2.2.10. Lines **236** and **237**: The accuracy of the efficiency percentages and power factors are not clearly identified. Without clear guidance, rounding rules would allow power supplies to pass that are actually a half percent worse than the specification. Please list the required number of digits of accuracy for compliance measurements.

2.2.11. Line **250**: Use of SPECpower_ssj2008 to define how to measure the "active idle" for a specific server configuration is an acceptable methodology. Displaying this active idle information on the Qualified Product Information standard reporting sheet is the best use of idle power information.

2.2.12. Line **283**: DC-DC power supplies should not be referred to as "DC-DC (All)". In this instance it should be listed as "DC-DC ±48V" if that is the DC-DC power supply assumed. Requiring a ±48VDC input voltage automatically excludes power supplies designed (e.g.) for the 380-400VDC input voltage range. Obviously, 400VDC power supplies cannot be tested at 48VDC.

While a 48VDC ±X% power supply should be tested specifically at 48VDC, there are also (at least) two variants of 48VDC power supplies that should be differentiated. NEMA-rated (telco) power supplies require a power supply to respond to the wide range of DC voltages that result from the battery charge/discharge range in a telco environment. Non-NEMA 48VDC power supplies might expect a more regulated 48VDC supply voltage and thus could yield higher power supply efficiencies. NEMA-rated power supplies should not be required to meet non-NEMA efficiency capabilities.

Additional lines should be added for higher voltage DC-DC input voltage power supplies. 380 VDC and 400 VDC distribution voltages appear to be emerging in a few large data centers.