May 21, 2010

To: Evan Haines
ICF International

CC: Una Song
United States Environmental Protection Agency

Re: Hewlett-Packard Response to *ENERGY STAR® Program Requirements for Computer Servers* Draft 1 Version 2.0

From: Hewlett-Packard Company,
Enterprise Servers, Storage and Networking Business Unit

This document may be published on the ENERGY STAR website.

Hewlett-Packard (HP) has a long-standing association with the *ENERGY STAR®* program and HP welcomes this opportunity to participate as a valued stakeholder in the process of creating Version 2.0 of *ENERGY STAR® Program Requirements for Computer Servers.*

The consistent feedback that should be apparent in this feedback document is the need to enable ENERGY STAR partner companies to swiftly, accurately and efficiently qualify as ENERGY STAR all servers. It is also highly important that server manufacturers be allowed to do in-house qualification tests. The following commentary provides the means to meet these goals and to foster future improvements in server energy efficiency.

1. Partner Commitment Issues

HP would like to highlight some partner commitment section issues.

1.1. Comments on the EPA Response to Inspector General Report

HP has a few comments on the ENERGY STAR program reaction to the US EPA Office of the Inspector General evaluation report titled *ENERGY STAR Program Integrity Can Be Enhanced Through Expanded Product Testing,* Report No. 10-P-0040, November 30, 2009 (IG Report).

1.1.1. Accredited Test Lab Requirements

Without all specifics being available, it is not possible to fully comment on the process that will be required to accredit HP's in-house test labs and allow us to do our own ENERGY STAR testing. HP strenuously objects to any requirement that would preclude self-certification by HP in-house test labs. Any requirement to use 3rd party test houses for ENERGY STAR compliance would be fervently opposed. Independence, accuracy and integrity can still be accomplished with accredited, in-house test labs that have proven their independence with safety and EMI certifications. FCC and other certifications

are done in-house, and HP has a track record in running ethical and competent test labs that pass audits, and should be trusted to do ENERGY STAR testing for HP products. HP supports certification of test labs via ISO/IEC 17025 or equivalent accreditation.

1.1.1.1. Relevance

The three product types evaluated in the IG Report are dramatically different than a computer server. DVD players, computer monitors and printers are relatively fixed functionality products whose features are not growing nearly as fast in performance and energy efficiency as server products. The products in the report are incapable of scaling their features, so conclusions drawn from the report are less valid for server products.

1.1.2. Energy Star Compliance Process Issues

The IG Report recommendations fail to adequately address a far more serious issue that is pointed out by that report. Page 3 of the IG Report states that "…selected non-ENERGY STAR products performed comparably to, and in some cases better than, ENERGY STAR products." This statement by the inspector general is an unflattering assessment of the success of the ENERGY STAR qualification process, and it highlights the fact that there are several procedural impediments for companies and those impediments block products from receiving the official ENERGY STAR recognition they would otherwise have earned. It is likely that thousands of servers from partner companies like HP, that could pass ENERGY STAR test requirements, are unable to receive official ENERGY STAR compliance due to unnecessary and/or costly complications included in the qualification process.

Currently, the ENERGY STAR label is not an accurate indicator of server energy efficiency excellence because a supermajority of today's servers that are produced by ENERGY STAR partner companies cannot adequately match their supply chains and configurable SKU structures to ENERGY STAR program requirements. Adding a significant testing process burden makes this issue worse, not better.

1.1.3. Inclusion of Non-compliant products in setting pass/fail thresholds

While it is true that non-compliance can be for either procedural or functional issues, the IG Report's recommendation to include non-compliant ENERGY STAR products in the data collection methodology could easily lead to unfair comparisons.

1.2. Recommended Qualification Process Streamlining

1.2.1. A Proposed ENERGY STAR Vision Statement

To ensure that ENERGY STAR is a relevant mark of server energy efficiency, and to prevent another unflattering program audit by the EPA's inspector general or the GAO, there should be a stated vision for the ENERGY STAR program to ensure that "all partner products that will pass audit requirements can bear the ENERGY STAR mark." While it is understandable that companies that choose not to be ENERGY STAR partners are precluded from the program, it is not acceptable for the program to create impediments to partners' swiftly and economically qualifying their highly configurable products.

1.2.2. Proposed ENERGY STAR Process Improvements

1.2.2.1. Test Simplification

Enable manufacturers to qualify and fully characterize "'Base Configuration' models" as ENERGY STAR and then allow SKUs based upon that well-documented, thoroughly-tested Base Configuration model to be ENERGY STAR compliant if they prescriptively add options that meet the power budgets for those types of options. This proposal is similar to FCC rules where separately certified products can be configured together without requiring additional testing to certify the resulting configurations. As a result, the number of manufacturers' tests would be reduced to a manageable quantity that can be done in a local test lab. This proposal also aligns better with the reality of how supply chains uniquely configure-to-order most products to customers' requirements, that would otherwise require the testing of millions of possible product families. Federal governments around the world primarily purchase "configure-to-order" server SKUs and configurations vary from agency to agency.

1.2.2.2. Compliance Accountability

Any product that a manufacturer claims as ENERGY STAR is subject to the threat of an audit at any time. As with benchmarking inaccuracies, competitors will alert the ENERGY STAR program of the need to do audits far more quickly and less expensively than random or periodic audits. The threat of an expensive and time-consuming audit will keep manufacturers honest. Therefore, any SKU that a manufacturer believes can pass an audit (due to its inherent design and configuration characteristics) should be able to claim ENERGY STAR compliance if testing has been documented for a similar subset configuration of the SKU.

1.2.2.3. EPA-Specified Energy Star Configurations

If a simplified process to enable widespread business use of ENERGY STAR in servers cannot be crafted, then there is an acceptable alternative solution. Since it is government procurement offices that plan to require ENERGY STAR server product purchases, it would be far easier for government CIOs to specify the detailed product configurations that they must all buy, and then let manufacturers focus ENERGY STAR testing on those specific configurations (with no additions or deletions of features allowed). This alternative gives the EPA control over exactly what gets tested and compared, takes the guesswork out of shooting at moving technological targets, dramatically cuts the test burden for manufacturers, and gives governments a competitive purchasing program.

1.3. Logos

HP wants to reiterate our support of the Version 1.0 and Version 2.0 decision to not require physical ENERGY STAR logos on server products.

2. Qualifying Products Eligibility Criteria Issues

The sections below discuss issues, changes needed, and suggested solutions to creating Eligibility Criteria that can help all eligible servers receive their earned recognition and ENERGY STAR status.

2.1. Section 1 Definitions Issues

2.1.1. Line 173-174. Some lower-end servers might not have ECC or buffered memory. Please delete these lines.

2.1.2. Line 175. The reference to "ac-ac" power supplies is a typo. This should be "ac-dc".

2.1.3. Lines 188-192. In the **Managed Server** definition: a server's management capabilities are not necessarily tied to its availability features. While redundant power supplies are often featured in a managed server, they are there for RAS reasons and not for system management reasons. Please delete line 191.

2.1.4. Lines 201-203. The definition for **Blade Server** is not clear enough to exclude servers that scale up their compute capacity by using processor/memory modules (e.g. 1 processor, 2 processor or 4 processor modules) to achieve larger processor counts that may subsequently be partitioned. It must be stated more clearly that "each Blade Server must be an independent Computer Server that shares a common infrastructure for

power, cooling and Blade Chassis management" and that the definition is not intended to describe processor/memory modules that may be used to scale-up a server to larger processor counts.

2.1.5. Lines 204-205. **Multi-bay Blade Servers** should clarify that a Multi-bay Blade Server may be implemented as either one blade that covers multiple chassis slots or as multiple blades that may be federated into a coherent processor/memory domain.

2.1.6. Lines 210-212. In the **Blade Storage** definition: the first sentence should more accurately read "A storage subsystem that is designed to be installed in a blade chassis." Blade Storage may be "shared storage" that is available to devices inside and outside of the chassis, so it shouldn't presume that storage is only for servers within the chassis.

2.1.7. Line 213: Regarding **Blade Network Equipment**: There may be more than one type of networking blade in a blade chassis.

- One general type of networking blade is meant primarily to provide connectivity between blades in the chassis, with one or more uplinks to the network(s) outside of the chassis. Typically this type of blade would not be plugged into the unused blade slots that usually have server blades.
- A second type of networking blade would leverage the chassis shared resources and primarily provide network connectivity to devices outside of the chassis. This second type of networking blade would be more likely to use the unused blade slots that usually have blade servers.
- Network equipment may include networks other than Ethernet (IEEE 802.3).

2.1.8. Lines 249-250. Some multi-node servers are hot swappable, so the definition can't preclude hot swap-ability. Hot swap-ability cannot be used as a definitional difference between blade and multi-node.

2.1.9. Lines 281-283. For clarity, the Pedestal Server category should mention that it includes servers that are sold as "tower servers" by HP and others.

2.1.10. Lines 291-295. Some clarity may need to be added on the definition of DC-DC Power Supplies. Many DC-DC power supplies are sourced by a distribution voltage of -48VDC. There should be no confusion between these sorts of power supplies and regulators inside of a server that might also convert ±48VDC to lower voltages. ±48VDC is one of several typical backplane voltages for a blade server. The definitive difference between a DC-DC power supply and a DC-DC regulator is how much current or

power they can handle. As a suggestion, define DC-DC power supplies as >TBD watts, to define the difference from DC-DC power regulators.

2.1.11. Line 322. On small servers, I/O might be embedded on the motherboard and have no I/O expansion slots.

2.1.12. Lines 355-362. As further clarification to the Idle and Active state definitions, it may be helpful to also define or reference the ACPI states: Working (S0), Sleep with CPU On (S1), Sleep with CPU Off (S2), Standby (S3) Hibernate (S4), and Off (S5) so that there is no confusion about whether Idle is (or is not) entering an ACPI state other than S0 (it should NOT). Servers typically do not support the Standby (S3) state, since operating systems have reliability issues in S3. Hibernate (S4) is typically supported by servers, but it takes a long time to enter or to recover from an S4 hibernation, so this would be an infrequent event. S5 is also infrequent.

2.1.13. Lines 398-399. **"**Server Processor Utilization"…needs more details. We support The Green Grid's recommendation to change this section to read "The percentage estimate of the server's compute activities relative to the full operational voltage and frequency of the processor(s)."

2.1.14. In the System Configuration section, Maximum Configuration and Minimum Configuration are defined as the max and min configurations relative to "active mode efficiency". It is unclear whether configurations that scale up functionality beyond the "Maximum Configuration" (and still meet test specifications) could be included in the product family. Maxima and minima don't always occur at the max and min configurations.

2.1.15. Table 1 raises several issues for Motherboards. Instead of a single "Motherboard", servers may have assemblies of several printed circuit boards, so a single "part number" for a Motherboard is often not possible. An "assembly number" may be a more appropriate nomenclature for the collection of boards in each unique model or product family. To simplify the table, these differences could be noted with an asterisk.

2.1.16. In Table 1, HP requests that I/O devices NOT require the same technical and power specifications in all product family configurations. This would enable us to include any I/O device that fits a power profile and not have to re-test for every possible option.

2.1.17. Additionally on Table 1, it is common practice for bills of materials for printed circuit boards and assemblies to have "acceptable alternate" components. It would not be possible to test every conceivable permutation of substitute parts. We must be able to test just a single representative

assembly of parts and trust that other permutations will pass "by specification" and "by design". Processors and Power Supplies should be allowed to have different part numbers and specifications as long as they meet the prescriptive power requirements for the qualified product family.

2.1.18. Line 442. We request that the requirement for storage change to "…HDD, SSD or other local storage device…"

2.1.19. Lines 446-447. It is unclear why a "Base Configuration" should differ from a "Minimum Configuration" and the currently listed definition for Base Configuration makes no sense. HP would like to see a Base Configuration or Minimum Configuration used as the simplified test platform for an entire family of SKUs.

2.2. Qualifying Products Issues

2.2.1. HP supports the inclusion of blade servers and multi-node servers in participation for ENERGY STAR version 2.0 for Computer Servers.

2.2.2. HP does not support the inclusion of resilient servers in this version of the specification, but we would like to support the investigation into how best to include this class of server in future specifications.

2.2.3. HP supports the Green Grid's basic approach to qualifying blade servers. Other sections of this feedback document offer additional ways to simplify server and blade testing, and subsequent audits.

2.3. Energy Efficiency Criteria Issues

2.3.1. Section 3.1 of the Draft specification uniformly discusses power supply units (PSUs) as being "in the chassis". A more generic and preferable approach would be to require that PSUs supplying power to a chassis must meet the requirements of Table 2. This change would support power supply solutions that are aggregated at a multi-server level, but might not be classifiable as either multi-node or blade servers.

2.3.2. Section 3.1 of the Draft specification needs to include dual-node and multi-node servers, which also use power supplies similar to pedestal, tower and/or rack mount servers.

2.3.3. Lines 517-522. HP supports the abandonment of the NPL (Net Power Loss) efficiency methodology. NPL is an example of a theoretical concept that doesn't align with supply chain realities and adds no discernable efficiency benefits.

2.3.4. Lines 538-567. The best way to test blade server solutions is going to depend on how many different product families are going to be required by the complexity of the ENERGY STAR qualification process.

- If every possible permutation of options has to be tested, then the testing environment must be simplified.
- If we only need to test Base Configurations, then a highly populated chassis becomes more affordable.
- To reiterate other portions of this feedback document, it is HP's position, that all ENERGY STAR Partner servers that can pass an audit should receive the ENERGY STAR qualification, and the ability to use the ENERGY STAR logo on products should not be overburdened by the testing process. See feedback paragraphs 2.3.6, 2.3.7 and 2.3.8 for additional suggestions for blade solution qualification.

2.3.5. Lines 536-537. The way this section reads seems to require divulging proprietary information. We support listing power management modes, but not the underlying "techniques" and technology used to yield the results.

2.3.6. Line 558-572. Table 4 needs to comprehend dual-node and/or multi-node chassis with removable-nodes that have similarities to both rack-mount and blade servers.

2.3.7. A blade chassis should receive an ENERGY STAR logo if it has power supplies that meet the requirements of Table 2, and has power management and other basic features that may be similar to a Version 1.0 four socket server. Documentation would accompany it that divulge the configuration it was tested and approved for compliance, but once a chassis is qualified it should be able to be ENERGY STAR regardless of what is plugged into it.

2.3.8. Table 4 should define idle power limits that are "divide by N" equivalents to rack-mount server equivalents.

- The minimum number for "N" may be different for every manufacturer and every chassis+blades combination, and should be allowed to be the minimum number of homogeneous servers in a chassis+blades solution that will pass the test threshold.
- Documentation would be provided that states that to achieve the expected energy efficiency levels a chassis must have the equivalent of N blades of this type installed for the chassis+blades. A smaller N is better.
- Blades and chassis must be able to be sold separately with their logos intact and documentation for how similar solutions were tested during the qualification.

2.3.9. Lines 577-582. We support the concept of shipping an ENERGY STAR qualified server blade and documenting the situations under which it been qualified for the ENERGY STAR logo. However, the environment where it is installed must only be required to match the basic chassis characteristics (e.g. power supply type), since the actual installed chassis will be likely to contain blades, storage and/or network equipment that were not included in the qualification test environment.

- A Blade Chassis must be allowed to have a rich set of blades installed into it without a particular blade losing its ENERGY STAR qualification.
- Similarly, a Blade Server must be allowed to be installed into a richly configured and heterogeneous chassis without losing its ENERGY STAR qualification.

2.3.10. Lines 696-705. **Support for Energy Efficient Ethernet IEEE 802.az?** We support providing documentation transparency of whether a server does or does not have IEEE 802.3az, HOWEVER making IEEE 802.az a requirement for ENERGY STAR would not make sense in any foreseeable future versions of the specification. The definition of a Computer Server in the ENERGY STAR specification does not (and should not) require that Ethernet connections are provided. Until IEEE 802.az is tested, adopted and widely available, any thought of making it an ENERGY STAR requirement would be an unnecessary burden on qualifying servers and will limit the government's ability to purchase ENERGY STAR qualified servers.

2.3.11. If idle power test criteria are defined for components, the idle power for basic components like memory and hard drives have not changed since the release of version 1.0 and so if idle power adders are defined, they should not be any lower than those defined by version 1.0.

2.3.12. A significant issue with version 1.0 idle power criteria exists for two-socket servers that ship with one processor. Our supply chain ships basic two-socket models with only one processor, and so the lowest power two-socket servers are barred from getting an ENERGY STAR logo, since the power threshold is too low for that case. We need to reduce confusion for customers. Due to 1P/2S servers having to have the same power budget as 1P/1S servers in Version 1.0, HP is unable to qualify 1P/2S "energy efficient SKUs" for ENERGY STAR, but we can qualify 2P/2S "performance SKUs".

2.3.13. If idle power criteria are included in version 2.0, HP requests that idle power pass/fail thresholds be added for four socket servers. The disparity caused

by 2P servers having Idle power limits and 4S servers having no limits creates an unfair competitive issue, which is especially apparent in the resilient server category.

2.3.14. HP applauds the EPA's collaboration with the Standard Performance Evaluation Corporation (SPEC) on the Server Efficiency Rating Tool (SERT).

2.4. Standard Information Reporting Issues

2.4.1. As discussed elsewhere in this feedback document, HP proposes simplifying the Product Family definitions and testing processes, so that the number of Product Families are bounded and based upon a few common Base Configurations. We support supplying documentation informing customers of the differences between features in the test environment versus what is actually shipped, but we do not support having to test the thousands of possible configure-to-order configurations for each server model. The current Product Family definition doesn't currently provide a means to adequately overcome the testing complexity.

2.5. Standard Performance Data Measurement and Output Issues

2.5.1. Lines 812-815: A sampling rate of ≥1 Hertz is unnecessary and problematic. A 10 second minimum sample period is more reasonable, but guaranteeing the accuracy of power, utilization and air temperature measurements in that sample period might not be possible with current state-of-the-art "commercial off-the-shelf" (COTS) technology.

2.6. Testing Issues

2.6.1. Multi-node Servers and Blade Servers: any direct comparisons of power or performance should be done on a "divide by N" basis, since multi-node and blade servers are actually "N" independent physical servers with various levels of shared infrastructure.

2.7. Product Qualification Issues

2.7.1. As previously discussed, the current definition of Product Families precludes having more than a few tested solutions and the implications of Table 1 cause re-testing every time a slight change happens in the supply chain. Implications of this may include the need to send compliance testing off-shore for the thousands of configurations and changes that are needed, and

will certainly preclude testing the vast majority of servers that should qualify to bear the ENERGY STAR logo.

2.8. Effective Date Issues

2.8.1. HP expects a 9 month transition period from Version 1.0 to Version 2.0. Products that meet Version 2.0 will want to be able to qualify for the new specification as soon as it is released, however, especially those products that were excluded from Version 1.0.

2.9. Appendix A Issues

2.9.1. Lines 1005-1010. It doesn't mention where the power is being measured, but the correct place to measure Blade Server Power is where the AC or external DC power enters the enclosure (e.g. at the power supplies that convert data center distribution power to chassis distribution power).

2.9.2. Line 1011. Related to comment 2.9.1, the data center distribution power input to the chassis for Blade Servers or for Multi-node Servers is the proper place to measure total power consumption of the entire assembly of blades or nodes. This total power may include power for infrastructure components, like Blade Network Equipment, that are excluded from participation in ENERGY STAR but are required for the proper operation of the total blade solution.

2.9.3. Line 1039. Telling us where to populate a blade pushes for an unnecessary conformity in how manufacturers prioritize cooling for blade slots. Each type of chassis has a different recommended order of population for blades. This line should read "when testing a single blade, install it in the location that is recommended by the manufacturer's documentation, or if not specified install the blade in a top corner in the chassis."

2.9.4. Lines 1043-1046. Similar to comment 2.9.3, each type of chassis has a different recommended order of population for blades, to achieve the optimal energy efficiency. The order of population shown in this section of the Draft Specification is unnecessarily detailed.

2.9.5. Lines 1043-1046. Prescribing a half full chassis is unnecessarily restrictive. A better approach would be to allow a manufacturer to install "N" blades in a chassis, where "N" is the number of blades that meet the "per blade" pass/fail criteria for that combination of chassis and blades. Then the manufacturer must document how many blades were required to meet the requirement, and manufacturers could compete on that basis. Choosing an

arbitrary "half-full" chassis creates an unnecessary barrier for compliance and there is less means for healthy competition.

2.9.6. Line 1052. An Ethernet port is not defined as a requirement in the definition of a Computer Server in section 1A. It is best if system requirements are not buried in an addendum.

2.9.7. Line 1076. Testing "as-shipped" is a consistent impediment for qualifying both product families and specific configurations.

- Specific "build to order" SKUs are shipped without hard drives and with some empty processor sockets, yet they are populated in the re-seller channel with those and other options. We need to be able to test a representative configuration that is a superset of what is shipped, and document how the results were derived. This should also enable a server shipped as a one processor-two socket server, to be tested as a two processor server.
- "Configure to order" SKUs should be able to be shipped by having a "Base Configuration" that is fully tested and options that are able to be added by prescription, i.e. the options each consume less power than the prescribed power budget for each device. This would solve the problem of ENERGY STAR requiring millions of tests to qualify all possible servers that would pass audit. This is also quite different from requiring ENERGY STAR servers to be tested "as shipped", yet it still yields servers that will pass audits.
- HP cannot violate its ethics policies. If the EPA requires testing "as-shipped" for the millions of permutations of server options, we are unfairly barred from markets for ENERGY STAR servers by the unacceptable testing cost and complexity caused by the unintended consequences of well-intended requirements.

2.10. Appendix B Issues

2.10.1. There is no content in this section, so there are no issues at this time.

3. Questions raised in the draft specification notes:

3.1. **Should there be a "volume server" category?** The HP answer would be "no". The term "volume server" does not have enough detail to differentiate between the many platforms with varying scalability and RASM features, and our highest sales volume servers have moderate to high levels of RASM features.

3.2. **Should the definition of "managed server" include remote power control, remote reset, hardware event logging, and/or remote cold start capability independent of OS-based management?** These seem like a reasonable subset of requirements for managed servers.

3.3. **Are there additional blade server definitions that need to be included in section 1.B.2.?** Blade slots are open to development of blades that are compatible with the electrical and architectural environment of the chassis. Much like I/O option slots that provide connectivity for new functions, there may be new types of blades that are not yet comprehended.

3.4. **Is the Availability Environment Classification system adequate to define availability metrics?** The Harvard Research Group's AE-0 to AE-4 availability levels are reasonable generalizations and classifications of end-user requirements for availability, but it is of grave concern to HP whether ENERGY STAR can adequately value the features it takes to deliver superior levels of availability. ENERGY STAR shouldn't be used as a way to level the playing field to the least common denominator, so there needs to be a way to reward increased value in superior solutions.