



September 28, 2012

To: Robert J. Meyers
Project Manager, ENERGY STAR for Computer Servers
United States Environmental Protection Agency

Re: Hewlett-Packard Response to *ENERGY STAR[®] Program Requirements, Product Specification for Computer Servers*, Draft 3 of Version 2.0

From: Hewlett-Packard Company, Enterprise Group

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 the *ENERGY STAR[®] Program Requirements for Computer Servers*.

The consistent feedback that should be apparent in both this and previous replies is the need to enable ENERGY STAR partner companies to swiftly, accurately, economically and efficiently qualify and certify all server configurations that meet the ENERGY STAR technical criteria. It is also highly important that server manufacturers be enabled to do in-house certification testing and that there is no interruption of lab certifications when version 2.0 is introduced. The following commentary provides the means to help meet these goals and to foster future improvements in server energy efficiency.

1. **Qualifying Products Eligibility Criteria**

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

1.1. Section 1, Definitions

1.1.1. Line 98. “high performance computing” and “high performance computer” are not interchangeable terms. “High performance computing” is defined by the application and not by the computer hardware. This line would be more appropriately written “Marketed and sold as a computer optimized for high performance computing applications.” This contrasts to computers that are designed for enterprise applications (and may also be considered “high performance” on those applications).

1.1.2. Line 102. The abbreviation “IPC” is not defined. Presumably this means inter-processor communications, but this probably isn’t the correct term. HPC systems have interconnects that range from proprietary inter-processor communication busses (which may be coherent or non-coherent) to non-coherent interconnects like Ethernet, InfiniBand, Servernet, Myrinet, et al. that have a wide range of costs, latencies and speeds. The Top500



September 28, 2012

supercomputer list has solutions that use many different interconnects, but the top two most frequently used interconnects on the Top500 list are Ethernet and InfiniBand.

- 1.1.3. Lines 171-172. An additional attribute that defines solid state drives (SSDs), in the current state of the art, is that SSD storage devices generally have a shorter lifetime than rotating media (Hard Drives) when they are accessed with similar read/write workloads. So an SSD shouldn't be considered a low-power equivalent of a Hard Drive and future specifications should not expect all Hard Drives to be replaced by SSDs.

1.2. Section 2, Scope

HP supports the inclusion of Resilient Servers with similar qualification criteria to 3S/4S Computer Servers. It is not clear in this section if Resilient Servers are eligible or excluded.

1.3. Section 3, Qualification Criteria

- 1.3.1. Line 411 should read “i. SERT main report results; and”
- 1.3.2. Line 412 should read “ii. SERT detailed report results over the entire test run.”
- 1.3.3. Line 415. “workload module” should be “SERT worklet”.
- 1.3.4. HP supports the new allowance to be able to qualify single-processor SKUs in dual-socket server product families by qualifying them with 2 processors installed and using the same idle power limits as dual-processor SKUs.
- 1.3.5. Line 500-515. This section is written for a stand-alone server with an auxiliary processing accelerator (APA), and needs additional description and clarification for testing in blades and multi-node servers that have APAs/GPUs.
- 1.3.6. Line 515. An Idle power budget of 46 Watts is significantly too low for currently available GPUs (APAs). Any particular wattage chosen would not adequately account for variations in GPU peak performance and wrongly favors GPUs with low performance (and low idle power). No idle power below 125W would be reasonable for any GPU pass/fail criteria. The preferable approach would be to qualify a server without a GPU and then allow it to keep its Energy Star certification when APAs are added.

1.4. Section 4, Standard Information Reporting Requirements

- 1.4.1. HP supports the ability to have a broader selection of processors and other system attributes grouped into a single “Product Family”, so that fewer product families are needed to cover each server model. However, we assert



September 28, 2012

that four test configurations are adequate to describe and certify a product family, instead of five. The “typical” configuration for the fifth test will be arbitrarily chosen by each vendor doing the test, whereas the four corner tests will be similar from one vendor to the next, so the “typical” configuration data will not enable valid comparisons and won’t be worth the cost of doing the extra testing.

- 1.4.2. Lines 550-551. HP Agrees with this approach. However, the word “benchmark” needs to change to “rating tool” after the term SERT (and there are two periods at the end of that sentence).

1.5. Section 5, Standard Performance Data Measurement and Output Requirements

No issues at this time.

1.6. Section 6, Testing

- 1.6.1. HP requests that 3-phase AC-DC power supplies be comprehended in the test method. 3-phase power supplies can have superior energy efficiency and phase balance provides additional benefits for distribution power efficiency and reliability.
- 1.6.2. HP requests that DC-DC power supplies be included that have input voltages that span the entire range of what is defined as “Low Voltage” by the National Electrical Code (e.g. <600V). At the very least, power distribution in the 360VDC-400VDC range is becoming more common, can provide additional data center-level energy efficiency, and needs to be comprehended.

1.7. Section 7, Effective Date

Products shipped during the interim period from November 9, 2012 and August 1, 2013 should be allowed to test and claim certification with either ENERGY STAR version 1.1 criteria or version 2.0 criteria.

1.8. Section 8 Considerations for Future Revisions

Using SERT data collected during version 2.0 submissions as the basis for choosing idle and active mode pass-fail criteria for version 3.0 would not accurately portray the energy efficiency profile of the entire server market. ENERGY STAR has a stated goal of choosing the top quartile of market performers. Since version 2.0 submissions only provide data from servers in the current top quartile, then any expectations set by only using that data would skew the reality of the market.



September 28, 2012

2. Draft Test Method

This feedback applies to the Draft Test Method document, dated August 2012.

- 2.1. Dc Server input voltages in the range of 360V-400V are not uncommon and should be added to the list of approved test methods.
- 2.2. There is inconsistency in the capitalization of AC and DC in the specification, test method and Power and Performance data sheet. The preference is to use “DC” and “AC” for all direct current and alternating current abbreviations, but instances currently use a variety of capitalization choices: e.g. Ac, Dc, ac, dc.
- 2.3. Line 19. Standard international voltages that are supported in most benchmarks are 100V, 110V, 200V, 208V, 220V, 230V and 400V $\pm 5\%$.
- 2.4. Line 21. Minimum temperature should be 20 °C, which aligns to SPEC requirements. 18 °C “...could increase hours of chiller operation and increase energy use.” (http://tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf)
- 2.5. Line 22. Should add “The Ambient temperature upper limit should be within the documented operating specification of the UUT.”
- 2.6. Line 23. Should add “The Relative Humidity should be within the documented operating specification of the UUT.”
- 2.7. Line 44. The measurement accuracy should be aligned with SPEC’s guidelines: “Accuracy - Measurements must be reported by the analyzer with an overall uncertainty of 1% or better for the ranges measured during the benchmark run. Overall uncertainty means the sum of all specified analyzer uncertainties for the measurements made during the benchmark run.” [see SERT Design Document (DD) http://www.spec.org/sert/docs/SERT-Design_Doc.pdf] Actually, the whole section D should be replaced with section 4.3 of the SERT DD
- 2.8. Line 61. Recommend adding “The SERT Run and Reporting Rules include specific tuning instructions for supported environments to ensure fair measurement of the loads being tested by the tool.”
- 2.9. This test method should also comprehend 3-phase power supplies.
- 2.10. Lines 127-140. For half (or full) chassis tests, we suggest that all but one of the blade servers be kept at a minimal configuration, and only one blade server be required to be changed to perform the “four corners” (plus typical) configuration testing for product families; which would vastly reduce both the time and parts costs for certifying each blade product family. The focus can then be on the one blade that changes.
- 2.11. Line 156. Change “Install manufacturer specified workload software...” to “Install the SERT software...”



September 28, 2012

- 2.12. Line 164. The 5-15 minute window of time to let a server boot to an idle state is overly prescriptive. It is unknown if 15 minutes is enough time for a large enclosure of blade servers to boot up, since it may need to sequence the server start-up times in order to keep circuit breaker current below its rated limit. The sentence reads like it cannot be longer than 15 minutes, when it should be the undefined length of time needed to let the server(s) boot and all become ready to run applications.
- 2.13. Line 164-179. Recommend replacing with “Between 5 and 15 minutes after the completion of initial boot or login of all UUTs, launch SERT according to the product’s Run and Reporting Rules; document and retain the associated output files. Information on SERT and the associated Run and Reporting Rules can be found at www.spec.org/SERT.”

3. Power and Performance Data Sheet

This feedback applies to the Power and Performance Data Sheet draft.

- 3.1. Line 40. Should clearly state SERT as the test software.
- 3.2. Line 45. Formula in the spreadsheet, but the text is not. Should be “24x365”.
- 3.3. Line 89. Change SPECpower_ssj2008 to “Server Efficiency Rating Tool (SERT)”
- 3.4. Lines 24, 25, 39, 92, 146, 148, 150 and 153. The term “Power Meter” should change to “Power Analyzer”.