April 14, 2008

To: Rebecca Duff
ICF International

CC: Andrew Fanara
United States Environmental Protection Agency

Re: REVISED Hewlett-Packard Response for the *ENERGY STAR® Program Requirements for Computer Servers - Draft 1*

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

This feedback document is a revision to HP's previous commentary submitted March 14, 2008 and replaces the original HP feedback document in the substantive areas covered. This document may be published.

Hewlett-Packard welcomes this opportunity to provide comments on *ENERGY STAR® Program Requirements for Computer Servers* (hereafter in this feedback document called "Energy Star for Computer Servers") — Draft 1. HP is proud to continue our long standing association with the *Energy Star* program. The comments and issues in the Draft 1 document fall into two categories and are covered in the two major sections of this review:

1. *Energy Star* Partner Requirements
2. *Energy Star for Computer Servers* Product Eligibility Requirements

Several key issues are detailed on the pages below, but an executive summary of the major feedback points can be concisely stated.

1. It is not acceptable to require physical labels on server products.
2. Basing *Energy Star for Computer Servers* version 1.0 pass/fail criteria on power supply thresholds is acceptable, but HP objects to using 10% load efficiency as part of the pass/fail criteria. The ultra-high efficiency, value-added power supplies that HP plans for use in *Energy Star* products are optimized for 20%-100% loads and efficiencies have a higher statistical variance at 10% loads.
3. Aside from the 20%, 50% and 100% load efficiency and power factor thresholds, other thresholds and specific parameters should not be used as pass/fail criteria. HP recommendations for power supply thresholds are detailed in section 2.3.1.2.
4. A server's *Energy Star* rating should be based on the product model and not on the exact delivered SKU configuration.
    - € A single standard information reporting sheet should be able to be used by all *Energy Star* SKUs for a particular model and must not be required to be produced for every possible *Energy Star* server SKU configuration.

5. *Energy Star* threatens to change specifications more rapidly than the multi-year life-spans of servers. As servers get larger and more complex, products change even more slowly. *Energy Star for Computer Servers* should never change more rapidly than biennially and larger or more complex servers demand even less frequent *Energy Star* requirement changes.
6. Do not expect to use the same *Energy Star* requirements for all server types.
   - In section 2.1.1 of this document, and illustrated in Appendix A, HP produces an *Energy Star for Computer Servers* taxonomy to help categorize, type and differentiate the various computer servers that will be in existence for the life of this specification.
   - This taxonomy can help define which server types to include and exclude in Tier I, and may serve as the basis for type-specific differentiation of *Energy Star for Computer Servers* requirements in future revisions.

1. Energy Star Partner Requirements

The first section of the draft specification describes the requirements that HP would have to maintain to become an *Energy Star Partner* for the *Energy Star for Computer Servers* specification. In brief, the draft document describes the following as the commitments that a partner must make:

1.1. Develop and qualify server products that meet the Energy Star for Computer Servers specification.

No HP issues.

1.2. Support possible audits of Energy Star logo'd HP server products by the EPA to see if they actually meet the specification.

No major issues. Need clarification on the auditing process.

1.3. Assure that HP and authorized representatives properly use the Energy Star trademarks.

Possible HP issue with logo size and placement.

- Products like server blades have no available space for an additional 0.375" width logo with 0.125" of clear space all around the logo (*Energy Star Brand Book* https://www.energystar.gov/ia/partners/logos/downloads/BrandBook508r.pdf).

1.4. Provide clear and consistent labeling of Energy Star servers.

HP has a major issue with the concept of requiring physical labels on servers.

- HP servers have no room for an *Energy Star* label on the front of a rack server. Labels on the front or rear can block critical airflow. Consider the detrimental

effect on a 1U or blade server where there is very little space for airflow and even the HP logo is minutely sized.

- HP is opposed to putting *Energy Star* labels on any surface of a server, given that a) some server types have no physical space for a label, b) *Energy Star* may not be possible on all SKUs for a product, and c) the frequency of change for *Energy Star* specifications threatens to be more frequent than the length of server product lifetimes.

1.5. Annually provide to the EPA a list of Energy Star qualifying computer server models.

No major issues, but the Draft 1 specification doesn't clearly state how many *Energy Star* products HP must make nor if it is acceptable for only certain SKUs of products to be compliant.

1.6. Annually provide to the EPA a confidential accounting of how many units were shipped with an Energy Star logo, or provide some "other market indicator to assist in determining the market penetration of Energy Star."

No major issues. Need clarification regarding the necessary non-disclosure agreement for the data submitted, acceptable reporting processes and acceptable uses by the EPA of the data submitted.

1.7. Notify the EPA regarding the HP point of contact who will be managing the EPA/HP partner program.

No major issues, although the term "partner" may have unintended legal implications.

1.8. Additional recognition can be attained for special distinctions.

No HP issues.

2. Energy Star for Computer Servers Product Eligibility Requirements

The requirements for a server to meet the *Energy Star for Computer Servers* specification are much more sparsely defined in Draft 1 than expected. This may reflect how difficult it has been for the EPA and their contractors during the last six months to sift through the mountain of feedback received after the publication of the framework document. The bulk of this section in Draft 1 is a series of notes regarding what has yet to be described in any draft of the specification. The following is an attempt by Hewlett-Packard to provide feedback and suggestions regarding what is described and alluded to in Draft 1.

2.1. Section 1: Definitions

2.1.1. Section 1A: Computer Server Definition

- This section of the draft attempts to define the term "Computer Server" and also attempts to define "Blade Server" and "Blade Chassis" (see issues below).
  - In section 1A, the definition of Blade Server should not assume that a "hard drive" is included. This is true for two reasons:
    - 1) solid state media may replace some rotating media "hard drives", and
    - 2) some blade servers may have all of their storage located somewhere other than on the same blade and may boot across the network.
  - In section 1A, the definition of Blade Chassis should not assume that the chassis has shared storage resources. Shared storage is not a required blade chassis feature today.
- Notes on this section indicate that there are no finalized definitions of the types of servers that should be separated into different categories. HP has provided some guidance in the past on this subject, and we provide more clarity to that guidance below. *Appendix A* illustrates how HP server product families map onto the proposed Computer Server taxonomy.
  - This specification covers business servers that are of more complexity than the "desktop-derived servers" defined in the *Energy Star 4.0 for Computers* specification. Many types of server will need to be described in enough detail to determine whether they will be included or excluded from participation in the Tier I requirements.
  - The focus of the draft 1 specification appears to be narrowly defined and does not cover the entire span of business and enterprise servers. Several types and subtypes of business and enterprise servers need to be defined, so that the ranges of their features can be comprehended in current or future *Energy Star* specifications. Even if version 1.0 of this *Energy Star for Computer Servers* specification excludes many types of servers, those server types need to be described in this section.
  - ***Energy Star* Server Taxonomy**:
    - The natural world has the taxonomy hierarchy of **Kingdom**, **Phylum**, **Class**, **Order**, **Family**, **Genus**, and **Species**. The

Energy Star program needs to borrow that taxonomy to describe the program's areas of coverage.

- *Energy Star Kingdoms* might be Homes, Appliances, Commercial Buildings, Office Equipment, Home Electronics, etc.
    - Within the Office Equipment Kingdom, the *Phylums* could be defined as Computers, Monitors, Printers, Storage, etc.

§ The Computer Phylum has already been divided into the two Classes that have (or soon will have) *Energy Star* specifications: 1) Computer Servers, and 2) Client Computers (desktop, notebook, workstation, et al. client computer systems). Following this taxonomy, the Class of Computer Servers should have several Orders of server types defined in the *Energy Star for Computer Servers* specification.

§ Computer server Orders tend to be defined by their varying general use environments and computer room densities. Although not a complete list, some of the computer server Orders could generally be classified as:

- Industry Standard Pedestal Servers
- Industry Standard Rack Servers
- Server Blades
- Blade Enclosures
- High-reliability Servers

§ Each one of these computer server Orders has the need for more than one Family in each Order that describe substantive differences in capacities and size. The Families may be divided into Genuses and Species that have differences in their capacity and/or scalability. Section 2 of the specification (Qualifying Products) can select inclusion and exclusion at any level of Order, Family, Genus or Species.

§ Below are the five server Orders listed above with their generic Family and Genus hierarchies. These structural hierarchies are illustrated with HP product examples in *Appendix A*. Further descriptions of HP server products described in *Appendix A* are available at the following Hewlett-Packard product website: http://welcome.hp.com/country/us/en/prodserv/servers.html .

§ **Order: Industry Standard Pedestal Server**

A pedestal server is capable of being a stand-alone computer without a data center infrastructure. "Industry Standard" refers to the use of x86 instruction set microprocessor(s). Pedestal servers often provide enough capacity and features to handle the compute server and storage needs of the business that it serves. The definitive differences in Families of Pedestal Servers are compute, memory and storage capacities.

- Families based upon RASM feature differentiation
  - € Genuses based upon number of CPU sockets
    - § Species based upon memory and storage capacities

§ **Order: Industry Standard Rack Server**

A rack server is physically mounted in a rack or cabinet with other information technology (IT) equipment and shares power distribution, cooling and communications infrastructure with other IT equipment in a data center. "Industry Standard" refers to the use of x86 instruction set microprocessor(s). Some of its storage may be located in a separate powered enclosure, but traditionally this type of server has the capability to boot off of its own local storage and provide enough storage for applications running on the server.

- Families based upon RASM feature differentiation
  - o Genuses based upon number of CPU sockets
    - § Species based upon memory and storage capacities

§ **Order: Server Blade**

A blade server is differentiated by higher density servers that have a dependence on a surrounding blade enclosure that provides cooling, network connection, management and direct current (DC) power to many blades, with the ability to easily add a blade server to a blade enclosure. Relative to rack servers, the high-density of blade servers causes some limitations on the scalability and capacities of CPU, memory, storage and I/O devices in each blade.

- Families based upon blade interchange compatibility type (e.g. HP c-Class BladeSystem, HP p-Class BladeSystem, AdvancedTCA, etc.)
  - Genuses based upon number of CPU sockets
    - Species based upon how many blade enclosure slots a single blade product occupies

- **Order: Blade Enclosure**

  Blade enclosures house multiple blades. Typically, the blades are all part of the same Family of compatibility (e.g. HP c-Class BladeSystem). A blade enclosure may hold a variety of sizes of both Server Blades and Storage Blades. Blade enclosures provide ACØDC power conversion, DC power distribution, intra-blade connectivity and networking, uplink networking, power management, and management for hot-pluggable blades and network resources.

  - Families based upon blade interchange compatibility type
    - Genuses based upon number of blade slots in the enclosure
      - Species based upon backplane interconnect differences

- **Order: High-reliability Server**

  A high-reliability server has many physical sizes from small rackable servers to large multi-rack servers. A key difference in these types of products is the product lifetimes. Customers demand 3-5 years of longevity for smaller servers and 7-10 years for the larger servers of this Order. Installed systems are highly likely to receive field upgrades of processors, memory, storage and I/O. Sales volumes are currently much smaller than the sales volumes of most industry standard servers.

  Ultra-reliable compute capacity is a prime differentiator, along with larger memory capacities. Some storage may be in the system, but storage is also likely to be in a separate or remote resource, so I/O bandwidth is an important differentiator.

  - Families based upon business usage models
    - Genuses based upon number of CPU cores

- € Species based upon memory size and aggregate I/O capacity to support external storage and networking devices.

- With its focus on power supply efficiency, version 1.0 of *Energy Star for Computer Servers* may have less need for this taxonomy than would a subsequent Energy Star version that is energy performance-based, but it will be very useful to get the definitions written now. With a full taxonomy, Tier I will be able to define specifically which Orders, Families, Genuses and Species of servers are included and excluded, and Tier II will be better prepared for its differentiated development. How *Energy Star for Computer Servers* transitions from version 1.0 to 2.0, etc. should differ from one Order and Family to the next in both requirements and acceptable frequency of change.

2.1.2. Section 1B: Computer Server Power Supply Definition

- € HP supports the exclusion of DCØDC power supplies in Tier I, but section 1B may need some enhancement to prepare for the potential future inclusion of DCØDC power supplies in version 2.0 or later revisions of the *Energy Star for Computer Servers* specification.

- € HP supports the exclusion of DCØDC converters from the definition of power supplies and their exclusion from being used as criteria for Energy Star version 1.0 for Computer Servers. DCØDC converters are typically used to convert (e.g.) 12VDC that is sourced from the server power supply to the local voltage(s) required by a variety of components inside the computer system enclosure.

2.1.3. Section 1C: Single-Voltage Power Supply

- € This section describes power supplies that convert alternating current (AC) power to a single 12V DC output. AC power inputs on typical distribution systems throughout the world range from 90-240 VAC and commodity power supplies handle that range of AC voltages.

- € In addition to DC distribution, there is also some industry support for >240 VAC local AC power distribution, so support for higher AC voltages might be considered in future drafts.

2.1.4. Section 1D: Multi-Voltage Power Supply

- € This section describes power supplies that convert AC power to multiple low voltage DC outputs. This definition describes a power supply type that is not

currently being used in business and enterprise servers from HP. These are more commonly found in desktop personal computers.

- HP would suggest removing this type of power supply from Section 1D and from Section 3, and if servers ever evolve to need them it could be added in a subsequent specification release.

2.1.5. Section 1E: Idle State

- Idle power should not be used as a pass/fail criterion for Energy Star. A general purpose server changes its idle power with every component that is added to the system, and since components are added to satisfy the performance needs of the server, it makes no sense to choose power thresholds for when (infrequently) a server is doing nothing.
- The definition of idle state could be improved. On small servers the definition should match the active idle definition used in the SPECpower_ssj2008 specification. On larger or more complex systems the idle state definition may need to rely on "idle" definitions in benchmarks that are more appropriate to those solutions.

2.2. Section 2: Qualifying Products

Section 2 has very little detail in Draft 1. The first line refers to server definitions in Section 1, but Section 1 does not yet fully describe the Orders and Families of servers that are possible and Section 1 should describe more Orders and Families than are prudent to promote for Energy Star Tier I. Fleshing out the server definitions in Section 1 and helping to choose the correct qualifying server types in Section 2 are among the most important tasks in this review. Section 2 definitions of "qualifying products" should be a subset of all server Orders and Families defined in Section 1.

2.2.1. Recommendations for inclusion and exclusion of product Orders or Families

- The Order of "high-reliability servers" should be excluded from the Tier I specification. The type of power supplies typically implemented in this type of server is often substantially different in capacity and type from smaller computer systems and the range of power loads do not vary widely, so the Table 1 power supply load variables do not apply. These systems should be re-examined for inclusion in Tier II.
- Non-server blades should be excluded from this specification.
- While the Order of Blade Enclosures is easily covered by the Tier I specification, additional work will need to be done to define how to qualify

DC-powered Blade Servers for *Energy Star*, separately from their Blade Enclosure.

- With the exception of server types noted for exclusion above, HP has no objection to other Orders and Families of computer servers participating in the Tier I specification, although there may need to be different requirements for different Computer Server Orders and Families.
  a. If the Tier I specification chooses to focus on smaller form-factor high-volume server types, then the highest priorities for coverage would be those server Orders defined as 1) the Order of *Industry Standard Rack Servers*, 2) the Order of *Blade Enclosures*, and 3) the Order of *Industry Standard Pedestal Servers.*

2.3. Section 3: Efficiency Requirements

2.3.1. Section 3A: Power Supply Efficiency Requirements

Table 1 in the draft outlines four aspects of an earlier, out-dated ECOS Consulting approach for the "80 PLUS for servers" definitions: 1) Four load levels are suggested by the EPA for testing: 10%, 20%, 50% and 100% of the rated load for the power supply itself (not the server load). 2) Minimum single-voltage power supply efficiencies at those loads. 3) Minimum multi-voltage power supply efficiencies at those load levels. 4) Power Factor thresholds at those load levels.

2.3.1.1. HP general feedback:

- HP is strongly opposed to the inclusion of 10% load efficiency pass/fail thresholds. 10% thresholds are not supported by latest test methodology by ECOS Consulting and are not supported by the *Climate Savers Computing Initiative* or *The Green Grid*. The statistical variance of efficiencies at 10% power supply loads across our supply chain makes this load level much less predictable. 10% load efficiencies are not used when specifying power supplies for purchase.
- We support the Climate Savers Computing Initiative test methodology for power supplies.
  - Issue: Are the threshold numbers to be the statistical mean of Energy Star SKU power supplies that HP will ship, or will they be chosen such that (e.g.) the "±2• " (95% confidence) of our power supplies shipped on Energy Star SKUs never fall below that threshold? Testing costs would be less and customer expectations more realistic if the thresholds are based upon the mean value of all power supplies shipped on a particular Energy Star SKU.

- The statistical variance of data samples varies considerably depending upon percent load, so acceptable variances need to be larger at both high and low loads.

€ HP supports the proposal for not including a "no load" efficiency specification.

€ Section 3A defines the test voltage and frequency as 230 VAC and 60Hz. Issue: While 230 VAC makes a lot of sense and yields slightly better efficiency than 208 VAC, the choice of 60Hz with that voltage is curious. HP recommends allowing the option of 50Hz or 60Hz. 50 Hz matches the 230 VAC/50Hz mains that are typical in much of the world outside of the USA (see Wikipedia reference figure below), which may be important as Energy Star is adopted internationally.

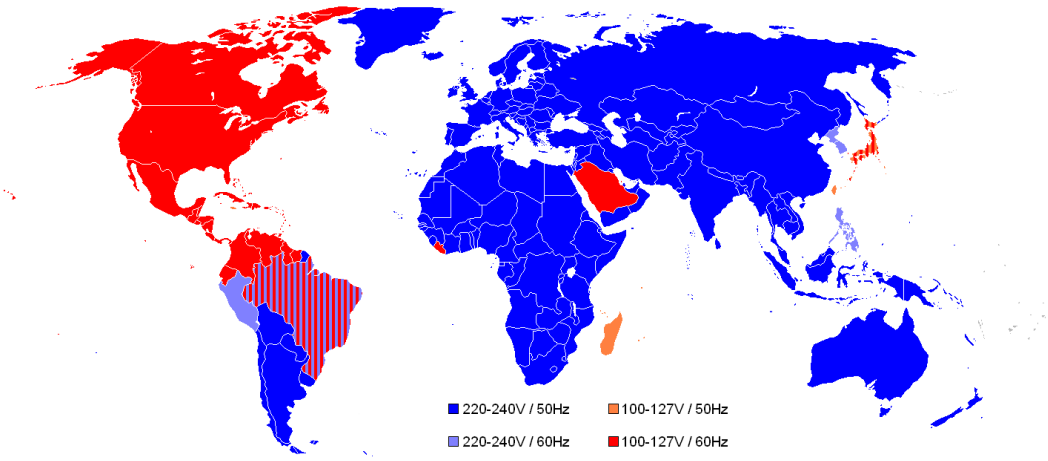

€ Inclusion of multi-voltage power supplies is not supported by HP in version 1.0 of this specification. HP does not fervently oppose its inclusion, but since HP doesn't ship multi-voltage power supplies in servers, we will not provide data to support defining appropriate thresholds and its inclusion dilutes the effort to support more than one single-voltage power supply. A more welcome use for Table 1 would be to add additional single voltage power supply types, so that there are small, medium and large wattage power supply categories and different sets of efficiency thresholds for each category.

€ HP supports the inclusion of power factor thresholds. We provide a suggested set of thresholds later in this feedback document.

€ HP suggests that power efficiency curves for larger power supplies look slightly different than those of smaller power supplies. Instead of the present structure for Table 1, the entries might look like the table below in feedback section 2.3.1.2, to include differing thresholds for small, medium and large output single-voltage ACØDC power supplies. Fan

power inside the power supplies is assumed to be excluded from the conversion efficiency. More data is available in Appendix B.

- This would give a large wattage power supply, optimized for blade server enclosures that are deployed in data centers, a fairer chance to achieve Energy Star and enable it to improve at a different rate in future specification revisions.
- By also differentiating between small server power supplies and mid-size server power supplies there may be a fairer comparison between similar product types. Mid-size servers can choose power supplies that are more reasonable for their range of loads and the two categories of power supplies can evolve at different rates.
- DC output Wattages dividing the small, medium and large power supplies could be proposed, for example, as 550W and 1750W.
- DCØDC power supplies are coming unto use in both experimental and production data centers trying to improve total data center energy efficiency. HP supports their exclusion in this version, but there should be some consideration for their inclusion in Tier II.

2.3.1.2. Proposed New *Table 1: Efficiency Requirements for Computer Server Power Supplies.* Power supplies tested and reported at 230 VAC; excluding P/S fan power.

(10% load efficiency and power factor must not be part of the pass/fail criteria)

Percentage of Rated Power Output	20%	50%	100%
Minimum Efficiency Requirement Single Voltage <550W	82	90	82
Minimum Efficiency Requirement Single Voltage 550W to 1750W	81	89	81
Minimum Efficiency Requirement Single Voltage >1750W	80	86	80
Minimum Power Factor Requirement	0.8	0.9	0.95

2.3.2. Section 3B: Idle Power

HP strongly urges that no Idle Power threshold and no Idle Power formula be used as a pass/fail criterion for *Energy Star for Computer Servers*. Even the smallest servers have enough memory, I/O and storage options to make the as-shipped product power unpredictable.

- € For small systems where a simple java benchmark is appropriate, the methodology to measure idle power using SPECpower_ssj2008 is reasonable. For mid-size and larger or

more complex systems, the SPECpower_ssj2008 methodology would not be appropriate.

- Two typographic errors were published in this section of the Draft 1. In the final paragraph of page 7, the word "world" should be "would", and the word "patters" should be "patterns".

2.3.3. Section 3C: Standard Information Reporting Requirements

As a concept, an online, available standard information reporting sheet for each product model is supported by HP. Items 1 and 2 in the notes on page 8 of Draft 1 are similar to the information that HP creates for every build-to-order invoice. Item 3 (Power Management Features) is currently available through HP documentation today.

Issues include:

- It is unclear how to describe Item 4 "Virtualization Capability". Since all servers have the ability to virtualize using one of the VM providers, this item is unnecessary. Servers will not typically be sold with software installed, so there can be no requirement that a server is virtualized as-shipped.
- For Item 5:
  - Idle power is not important to most types of servers.
  - SPECpower_ssj2008 cannot be the only benchmark alternative and even the idle power methodology of SPECpower_ssj2008 may not apply to larger or more complex systems.
  - Benchmark results published in this section must not be required to match the as-shipped configuration of the server. Any benchmark results published would be for a similar server model and power supply, but likely with a different actual configuration of memory, storage and I/O devices.
  - Results of benchmarks like SPECpower_ssj2008 must follow publication rules specified by the owners of the benchmark(s). It is economically infeasible for every one of the thousands of Energy Star server "as-shipped" configurations to have certified benchmark results, and what's more, the required benchmark software, application optimizations and Java Virtual Machine to achieve a published benchmark score will not be sold as part of the Energy Star server product SKU.

2.3.4. Section 3D: Power and Temperature Measurement Requirements

HP supports a slightly different set of requirements. HP support supplying *Input Ambient Air Temperature* and *Estimated System Power* readings via a networked interface. Issues:

- At no time must there be a requirement for air temperature measurements on the air output of the server. Output air temperature is misleading without correlating airflow data.
- System power, as measured by an internal server power meter, is not going to be as accurate as an external AC power meter.
- Must not specify DMTF SMASH as the required interface to access this type of data. HP also uses other standard means to achieve this functionality.
- The text says to provide "real time" data on AC power consumption, inlet air temperature and processor utilization.
  - First, "real time" is not possible. Even professional instrumentation provides sampled data, not "real time" data.
  - Second, "processor utilization" is typically provided by an operating system. Operating systems are not a required part of all server systems when shipped, thus there would not be any ability to deliver this feature out-of-the-box on every *Energy Star* product. Every enterprise operating system provides processor utilization information, so remove this requirement.

2.3.5. Section 3E: Power Management and Virtualization Requirements

This section is confusing and redundant to the information provided by the standard information reporting sheet. Issues:

- Virtualization software, much like operating systems, is not shipped on most servers. The hardware is all capable of virtualization, so this requirement seems unwise.
- Power Management is described in another section of this document.

2.4. Section 4: Test Criteria

HP supports the ability for companies to self-certify. HP supports the *Climate Savers Computing Initiative* test plan for power supplies.

April 14, 2008

2.5. Section 5: Effective Date

HP supports the plan to announce an *Energy Star for Computer Servers* version 1.0 specification in 2H 2008, with pass/fail criteria based upon power supply efficiencies. The plan for replacement by Tier 2 requirements must be coupled with a means for long product-life server families to either be exempt from *Energy Star* or to be able to stage the changes in *Energy Star* requirements differently on different types of products.

Higher sales-volume server products must not see *Energy Star* requirements change more rapidly than every 2 years. Lower Sales-volume servers included in this *Energy Star* specification must have *Energy Star* requirements that change much less frequently.

2.6. Section 6: Future Specification Revisions

See comments on Section 5. HP is not asking to grandfather products to older *Energy Star* specifications, but rather for *Energy Star for Computer Servers* to structure its revisions such that:

1. *Energy Star* requirements on the highest-volume servers should not change more rapidly than once every 2 years.
2. *Energy Star* requirements should not be changed on longer product-life servers every time that specification requirements change for small servers.

April 14, 2008

3. Appendix A — Energy Star Server Taxonomy Mapped to HP Server Products

This appendix will illustrate how the Hewlett-Packard server product line maps into the taxonomy outlined in section 2.1.1 of the Hewlett-Packard feedback document for *ENERGY STAR® Program Requirements for Computer Servers - Draft 1*.

.

3.1. Order: Pedestal Servers

In the HP ProLiant product line all servers with a prefix of "ML" are in the Order of Industry Standard Pedestal Servers. Sizes are not necessarily equivalent to rack mount servers, although some pedestal servers can optionally be mounted in racks, so the rack server sizing terms of 1U, 2U or 4U height have no meaning when the server is free-standing.

Family: ML100 Series

Genus: CPU socket capacity = 1

Example: ML110

Example: ML115

Genus: CPU socket capacity = 2

Example: ML150

Family: ML300/DL500 Series

Genus: CPU socket capacity = 1

Example: ML310

Genus: CPU socket capacity = 2

Example: ML350

Example: ML370

Genus: CPU socket capacity = 4

Example: ML570

3.2. Order: Industry Standard Rack Servers

In the HP ProLiant product line, all servers with a prefix of "DL" are in the Order of Industry Standard Rack Servers.

Family: DL100 Series

The DL100 series provides dual CPU socket compute capacity with a minimum of server hardware and software features. Low cost is the primary motivator for purchasers of the DL100 series.

Genus: CPU socket capacity = 1

Species: Internal storage drive bays = 2

Example: DL120 (1U)

Genus: CPU socket capacity = 2

Species: Internal storage drive bays = 2

Example: DL140 (1U)

Example: DL145 (1U)

Species: Internal storage drive bays • 4

Example: DL160 (1U)

Species: Internal storage drive bays • 12

Example: DL180 (2U)

Example: DL185 (2U)

Family: DL300/DL500 Series

The DL300 and DL500 series servers include many world-class data center features, like hot-plug hard drives, hot-plug fans, hot-plug power supplies, and value-added management subsystems. Added reliability and availability features, ease-of-use, ease-of-deployment, ease-of-management and ease-of-upgrade are several ProLiant advantages that differentiate DL300 and DL500 series servers from both the DL100 series and competitors' servers.

Genus: CPU socket capacity = 1

Species: Internal storage drive bays = 2

Example: DL320 (1U)

Species: Internal storage drive bays • 4

Example: DL320p (1U)

Species: Internal storage drive bays • 14

Example: DL320s (2U)

Genus: CPU socket capacity = 2

Species: Internal storage drive bays = 6

Example: DL360 (1U)

Example: DL365 (1U)

Species: Internal storage drive bays • 8

Example: DL380 (2U)

Example: DL385 (2U)

Genus: CPU socket capacity = 4

Species: Memory capacity = 16 DIMMs

Example: DL580G4 (4U)

Example: DL585G2 (4U)

Species: Memory capacity • 32 DIMMs

Example: DL580G5 (4U)

Family: DL700 Series

The DL700 family supports larger CPU socket counts with larger internal disk arrays, memory capacities and I/O bandwidth.

Genus: CPU Socket capacity = 4

Example: DL750 (7U)

Genus: CPU socket capacity = 8

Example: DL760 (7U)

3.3. Order: Blade Enclosures

Because server, storage and other sorts of blades may be plugged into many different sizes of blade enclosures, the blades and blade enclosures need to be categorized separately. Since *Energy Star for Computer Servers* version 1.0 will be focused on thresholds for power supply efficiency and blade enclosures hold the power supply infrastructure for multiple blades, it makes sense for Blade Enclosures to receive separate Energy Star ratings from the blades that plug into them. Defining a procedure for certifying specific server blades for *Energy Star for Computer Servers* version 1.0 may be difficult, while the certification requirements for blade enclosures is more easily defined.

HP has multiple Genuses of c-Class BladeSystem Family enclosures (e.g. c3000 and c7000) that are optimized for different numbers of blades and different types of data centers. They have dramatically different power supply capacity, availability and reliability requirements. However, the same types of server blades, storage blades, network switches and management subsystems can be plugged into both the c3000 and c7000 enclosures.

Family: HP c-Class BladeSystem

Genus: Blade Enclosure — Blade slots = 8

Example: BladeSystem c3000 (6U)

Genus: Blade Enclosure — Blade slots = 16

Example: BladeSystem c7000 (10U)

Family: HP p-Class BladeSystem

Genus: Blade Enclosure — Blade slots = 8

Genus: Blade Enclosure — Blade slots = 16

3.4. Order: Server Blades

Family: c-Class BladeSystem

Genus: CPU socket capacity = 2

Species: Size = 1 blade slot

Example: BL460c

Example: BL465c

Species: Size = 2 blade slots

Example: BL480c

Example: BL860c

Genus: CPU socket capacity = 4

Species: Size = 2 blade slots

Example: BL680c

Example: BL685c

Species: Size = 4 blade slots

Example: BL870c

3.5. Order: Storage Blades (and other non-server blades)

Family: c-Class BladeSystem

Storage Blades should be excluded from *Energy Star for Computer Servers* v1.0. However, given the rapid growth of storage and the significant energy consumption of storage, storage devices deserve their own *Energy Star* Phylum and need one or more *Energy Star* Class specifications.

3.6. Order: High-reliability Servers

HP recommends that the *Energy Star for Computer Servers* version 1.0 specification exclude the High-reliability Order of servers.

Server products with high-reliability have several defining features:

- Microprocessor families with native 64-bit instruction sets, i.e. other than Intel/AMD x86.
- Large memory capacity relative to the number of CPU sockets.
- Large aggregate I/O bandwidth, measured by how many I/O slots are available per CPU socket and the throughput capacity of those I/O slots
- High levels of security, virtualization, reliability, availability, serviceability and management features.

Family: *Integrity* entry-level and midrange series servers

Genus: CPU core capacity = 4

Species: Memory capacity = 8 DIMMs

Example: rx2660 (2U)

Species: Memory capacity = 24 DIMMs

Example: rx3600 (4U)

Genus: CPU core capacity = 8

Species: Memory capacity = 48 DIMMs

Example: rx6600 (7U)

Genus: CPU core capacity = 8

Species: Memory capacity = 96 DIMMs

Example: rx7640 (10U)

Genus: CPU core capacity = 16

Species: Memory capacity = 48 DIMMs

Example: rx7620 (10U)

Species: Memory capacity = 192 DIMMs

Example: rx8640 (17U)

Genus: CPU core capacity = 32

Species: Memory capacity = 96 DIMMs

Example: rx8620 (17U)

Family: Integrity Superdome

Family: Integrity NonStop