March 25, 2009

TO: Rebecca Duff
ICF International
1725 Eye Street, NW, Suite 1000
Washington, D.C. 20006
rduff@icfi.com

CC: Andrew Fanara
EPA
fanara.andrew@epa.gov

FROM: Chris Hankin
Sun Microsystems, Inc.
chris.hankin@sun.com
202-326-7522

Re: Comments by Sun Microsystems, Inc. on the Draft 4 Energy Star Specification for Computer Servers

Dear Ms. Duff:

Thank you for the opportunity to provide comments on the EPA's Draft 4 Energy Star specification for Computer Servers. Sun Microsystems appreciates the opportunities extended throughout this past year for inclusion in this process, and we look forward to continuing to help achieve a successful new specification.

We commend the EPA on their careful consideration of the input provided by Sun Microsystems and by the industry on previous drafts of this specification. In particular, we note that many of the suggestions provided by Sun and the industry in response to Draft 3 of this specification in December 2008 have been reviewed and included by the EPA in the current draft.

The comments that follow are consistent with our most recent discussions, and are made with the purpose of achieving a specification that better achieves our mutual goals.

We look forward to discussing these points in more detail and to a successful conclusion of Tier 1 of the Energy Star for Servers specification.

Sincerely,

Chris Hankin
Sun Microsystems, Inc.
chris.hankin@sun.com
202-326-7522

Table of Contents

COMMENTS BY SUN MICROSYSTEMS, INC. ON THE EPA ENERGY STAR (DRAFT 4) PROGRAM REQUIREMENTS FOR COMPUTER SERVERS

1. Introduction

Sun commends the EPA on issuing Draft 4 of the Energy Star for Servers specification. This draft represents significant progress towards the goal of managing energy consumption in the data center. Sun applauds the open process that the EPA has followed, including the extensive dialog with the industry and the EPA's willingness to be available for detailed discussions. Sun appreciates the opportunity to meet with the EPA in one-on-one meetings and in industry conference calls, as well as EPA's outreach to the industry at various conferences and symposia.

Sun particularly commends the EPA's transparency throughout the process. The webinar conducted for the industry by the EPA on March 16, 2009 was very informative and educational and went a long way towards explaining the EPA's methodology for arriving at the details of the specification. The EPA's release of the idle power data spreadsheet, as well as the analysis methodology, has helped the industry understand the rationale behind the various aspects of the specification.

While many of Sun's concerns from Draft 3 of the specification have been diligently reviewed, considered, and addressed by the EPA, Sun continues to have concerns about the following aspects of the Draft 4 specification, which are detailed in this document:

1. Idle power for 8 core microprocessors
2. Idle power for depopulated systems
3. Specific concerns around power supplies
4. Concerns about an unrealistic sampling rate
5. New recommendation for bandwidth measurement for input current testing
6. Test equipment accuracy
7. The timing for inclusion of blades in the specification
8. Request for some definitional clarifications
9. Concerns around the effective date

2. Idle Power

A) Idle Power For Systems with 8 Core Microprocessors

Customers size systems for particular workloads that need to be serviced in the data center, such as a total database transactional workload, or a peak web traffic workload, or a maximum HPC computational workload. System sizing is done on the basis of computational requirements (number of cores in the microprocessors), memory requirements (amount of DRAM in the system), storage requirements (the disk capacity) and networking/IO requirements (I/O and network bandwidth).

The data set analyzed by the EPA includes systems built primarily with 2-core and 4-core microprocessors. For systems built with 8-core microprocessors, the EPA has received data only from Sun, as Sun is currently the only vendor shipping with these highly innovative systems.

Vendors innovate for multi-core because of multiple reasons: it not only saves cost and makes the systems cheaper, but it also saves energy, as integrating a greater number of processing cores on a single socket burns less overall system-level energy than distributing those cores over multiple sockets. As such, the energy efficiency (performance per watt) of 8-core systems is greater than that of dual or quad-core systems, because fewer are needed to service a given quantum of workload. Yet, because the EPA grants idle power allowances only on the basis of socket count, not core count, 8-core systems, which are more energy efficient, are penalized, since their lower socket count restricts them to a lower idle power allowance.

The EPA has noted in Draft 4: "EPA also recognizes the opportunity for savings if future core count multi-core technologies (e.g., greater than four cores per processor) do indeed promise greater energy efficiencies for the same workload. EPA has limited data on multi-core systems and ***may consider alternate proposals for processors with greater than four cores if additional data from multiple manufacturers can be supplied for analysis.*** As part of the comment process, stakeholders are encouraged to provide data on systems with greater than four cores per processor for EPA consideration."

The following data on systems built with 8-core chips will illustrate the energy efficiency advantages of systems built with these microprocessors. We note that while Sun has been the only vendor shipping with these highly innovative systems over the last three years, it is expected that many other vendors will be shipping systems with 8-core chips within the lifetime of the Tier 1 specification.

Recent vendor innovation in multi-core systems has decreased the overall energy consumption in the data center, and ought to be rewarded, not penalized. Multiple cores on fewer chips have been shown to increase efficiencies enormously. A 1 socket 8 core system does just as much work as a 2 socket 4 core system, but consumes significantly less power. As such, a 1 socket 8 core system can be used to replace a 2 socket 8 core system, thereby saving power overall in the data center. Hence, a vendor providing a 1-socket 8 core system ought to be rewarded, not penalized, for doing 8 cores worth of work with only 1 socket, as opposed to other vendors who do 8 cores worth of work with 2 sockets, thereby adding electrical overhead in the system design.

The current EPA specification penalizes, not rewards, systems with 1 socket and 8 cores by granting them a lower idle power allowance than systems with 2 sockets and 4 cores each.

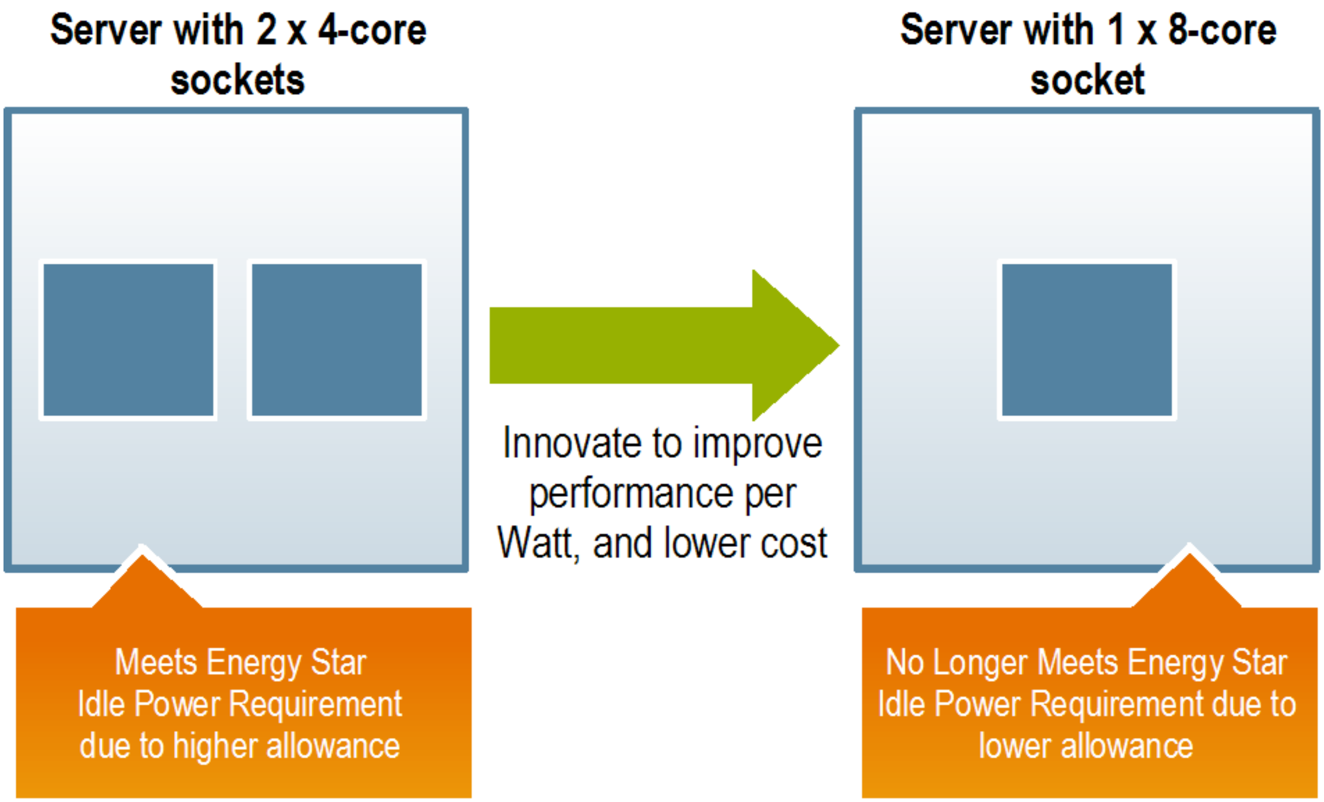

Illustration 1: multi-core innovation to increase energy efficiency

In a later section, we recommend several options as to how the EPA might address these unintended consequences of the current specification, and encourage the industry to move in the direction of energy efficiency innovation.

To prove this point, the following table compares idle power for systems that contain 8 cores: 1 socket 8 core systems verses 2 socket 4 core systems, all containing 16GB of memory. We define this set of servers as the ***equivalency class*** of all 8-core systems with 16GB memory. An ***equivalency class*** is a peer group of systems that all have similar capabilities, and are deployed for similar purposes in the data center.

The data in the table below for this equivalency class is taken from the EPA's published and released idle power data set for Draft 4. Index number on the left indicates the specific EPA data set item number.

Index	Number of Processor Sockets in System	Total Cores	# CPU(s) Installed	Cores/ CPU	Total Installed Memory (GB)	# of PSUs Installed	# of Installed Disk Drives	1 Gbit Ports for Adder	Available External 10Gbit Ports	Fibre or Infiniband	Draft 4 Idle Category	Idle Power	Idle Power per Socket	Idle Power per Core	EPA Allowance
53	2	8	2	4	16	1	1	4	0	0	2P Managed	161.0	80.5	20.1	178
80	2	8	2	4	16	1	1	0	0	0	2P Managed	200.0	100.0	25.0	174
83	2	8	2	4	16	2	1	0	0	0	2P Managed	221.0	110.5	27.6	194
85	2	8	2	4	16	1	1	0	0	0	2P Managed	228.0	114.0	28.5	174
87	2	8	2	4	16	1	1	0	0	0	2P Managed	234.3	117.2	29.3	174
88	2	8	2	4	16	1	1	0	0	0	2P Managed	242.6	121.3	30.3	174
93	2	8	2	4	16	2	2	2	2	3	2P Managed	291.0	145.5	36.4	233
84	2	8	2	4	16	1	4	0	0	0	2P Managed	227.0	113.5	28.4	198
86	2	8	2	4	16	1	6	0	0	0	2P Managed	234.0	117.0	29.3	214
17	1	8	1	8	16	1	1	2	0	1	1P Managed	160.0	160.0	20.0	94

Table 1: Comparing The Equivalency Class of 8 core Systems with 16GB DRAM

The table above compares the equivalency class of systems with 8 cores and 16 GB DRAM from the EPA's data set with the EPA index numbers noted in the left hand column.

The last system (Index 17) demonstrates that the 1-socket 8-core system is the most efficient, and has the lowest idle power per core (20W/core) because the vendor innovated to fit 8 cores in a single socket. This system also has the lowest idle power overall in the 8 core category (160 Watts). All the other 8 core systems require 2 sockets of 4 cores each, and are less efficient with idle power per core varying between 20.1 and 36.4 W/core.

In spite of being demonstrably more energy efficient, this system does not qualify for Energy Star because the EPA grants it an idle power allowance of only 94 Watts, because it is classified with all other single socket systems that have a lower core count. **If this system were considered a 2 socket system, it would be granted an idle power allowance of 179 Watts, and would easily qualify for Energy Star.**

The above data demonstrates that the EPA's Draft 4 specification penalizes the most energy efficient and innovative system by granting it a lower power allowance, while other less efficient system designs *in the same equivalency class* are granted generous idle power allowances between 174 and 233 watts.

The following table illustrates the same point for systems for the equivalency class of systems with 8 cores and 64 GB DRAM.

Index	Number of Processor Sockets in System	Total Cores	# CPU(s) Installed	Cores/ CPU	Total Installed Memory (GB)	# of PSUs Installed	# of Installed Disk Drives	1 Gbit Ports for Adder	Available External 10Gbit Ports	Fibre or Infiniband	Draft 4 Idle Category	Idle Power	Idle Power per Socket	Idle Power per Core	EPA Allowance
97	2	8	2	4	64	2	2	2	2	3	2P Managed	345	172.5	43.1	333
95	2	8	2	4	64	2	6	4	0	0	2P Managed	315	157.5	39.4	338
26	1	8	1	8	64	2	1	2	0	0	1P Managed	271	271	33.9	209
25	1	8	1	8	64	2	1	2	0	0	1P Managed	255	255	31.9	209

Table 2: Comparing the Equivalency Class of Systems with 8 cores and 64GB DRAM

The last two rows of this table demonstrate that a 1 socket 8 core system with high memory provides the lowest idle power and idle power per core (31.9W/core) of the group due to the vendors innovation of fitting 8 cores in a single socket. The other 8 core systems in the table require 2 sockets of 4 cores each so have higher idle power and idle power per core that varies between 39.4 and 43.1 W/core.

Yet, the current specification penalizes the most energy efficient and innovative system design in this equivalency class by granting it an idle power allowance of only 209 watts, while less efficient system designs *in the same equivalency class* are granted generous idle power allowances up to 338 Watts.

In spite of being demonstrably more energy efficient, this system does not qualify for Energy Star because it is classified with all other single socket systems that have a lower core count. **If this system were considered a 2 socket system, it would be granted an idle power allowance of 290 Watts, and would easily qualify for Energy Star.**

It is legitimate at this point to ask the question as to whether having the same number of cores is a meaningful criterion to group systems together in an equivalency class. Is it really the case that a system with 8 cores in 1 socket does the same amount of work as a system with 8 cores in 2 sockets, and a system with 8 cores in 4 sockets? Should 1x8 systems, 2x4 systems and 4x2 systems be considered to be in the same equivalency class?

To get a clue to the answer, it is instructive to look at benchmark results. The following table summarizes all published results for the SPECweb2005 benchmark until March 2009, that are above 20,000. This table can be obtained from http://www.spec.org/web2005/results/web2005.html.

Tester Name	System Name	Cores	Chips	Cores Per Chip	Result
Sun Microsystems, Inc.	Sun SPARC Enterprise T5220	8	1	8	41847
Sun Microsystems, Inc.	Fujitsu SPARC Enterprise T5220	8	1	8	41847
Hewlett-Packard	HP ProLiant DL165 G5	8	2	4	39530
Sun Microsystems, Inc.	Sun Fire X4240	8	2	4	37630
Sun Microsystems Inc.	Fujitsu SPARC Enterprise T5220	8	1	8	37001
Sun Microsystems Inc.	Sun SPARC Enterprise T5220	8	1	8	37001
Hewlett-Packard	HP ProLiant DL165 G5	8	2	4	32687
Sun Microsystems, Inc.	Sun Fire X4240	8	2	4	32288
Hewlett-Packard	HP ProLiant DL385 G5	8	2	4	30007
Hewlett-Packard	HP ProLiant DL380 G5	8	2	4	29591
Fujitsu Siemens	PRIMERGY TX300 S4, Intel Xeon processor X5460	8	2	4	28127
Hewlett-Packard	HP ProLiant DL380 G5	8	2	4	26077
IBM	System x 3650	8	2	4	23716
Powerleader Science & Technology Co.,Ltd	PR4700D	8	2	4	22332
Hewlett-Packard	HP ProLiant DL585 G2	8	4	2	22254
Hewlett-Packard	HP ProLiant BL685c	8	4	2	21470

Table 3: SPECweb2005 Results for 8 Core Systems.
Source: http://www.spec.org/web2005/results/web2005.html

It is clear from the performance results above that 1x8 systems not only are just as good as 2x4 systems, they slightly outperform them. In turn, 2x4 systems are not only as good as 4x2 systems, they slightly outperform them.

It is clear from the above performance table that it is the number of cores, not the number of sockets, that create an equivalency class whose energy characteristics should be judged together.

Although SPECweb does not publish power performance results, it is clear from the data that the SPECweb per Watt number would be the highest for 1 socket 8 core systems, then for 2 socket 4 core systems, and then for 4 socket 2 core systems. This higher performance per Watt at full utilization is compelling evidence of the energy efficiency advantages of massively multi-core systems.

Vendor innovation to reduce the number of sockets required to support a specific number of cores, should be rewarded, not penalized. Idle power allowances should be granted on the number of cores, not on the number of sockets.

Yet, for idle power, the EPA grants 1x8 systems (that lead in enterprise SPECweb performance, as shown in the previous table) the same allowance as 1x2 systems (that could be desktop-derived).

The following two tables conduct a similar analysis from the EPA data set, this time focusing on 16-core systems. The first table analyzes the equivalency class of all 16-core systems with 16 GB

DRAM, while the second table analyzes the equivalency class of all 16-core systems with 64 GB DRAM.

Index	Number of Processor Sockets in System	Total Cores	# CPU(s) Installed	Cores/ CPU	Total Installed Memory (GB)	# of PSUs Installed	# of Installed Disk Drives	Idle Power	Idle Power per Socket	Idle Power per Core	EPA Idle Power Allowance
91	2	16	2	8	16	2	2	289	144.5	18.06	222
110	4	16	4	4	16	2	1	329	82.25	20.56	No idle power requirement; can potentially comply
111	4	16	4	4	16	2	2	337	84.25	21.06	
112	4	16	4	4	16	2	3	342	85.5	21.38	
113	4	16	4	4	16	2	4	369	92.25	23.06	
114	4	16	4	4	16	4	1	371	92.75	23.19	
115	4	16	4	4	16	2	1	373	93.25	23.31	
116	4	16	4	4	16	4	1	384	96	24	
117	4	16	4	4	16	4	5	420	105	26.25	

Table 4: Comparing The Equivalency Class of 16-core Systems with 16GB DRAM (Idle power data for 4 socket systems taken from Draft 3 Idle Power Dataset)

Index	Number of Processor Sockets in System	Total Cores	# CPU(s) Installed	Cores/ CPU	Total Installed Memory (GB)	# of PSUs Installed	# of Installed Disk Drives	Idle Power	Idle Power per Socket	Idle Power per Core	EPA Idle Power Allowance
99	2	16	2	8	64	2	2	353	176.5	22.1	298
121	4	16	4	4	64	2	8	496	124.0	31.0	No idle power requirement
122	4	16	4	4	64	2	4	596	149.0	37.3	
123	4	16	4	4	64	4	5	662	165.5	41.4	

Table 5: Idle Power for 16 Core Systems - Comparing The Equivalency Class of 16 core Systems with 64GB DRAM (Idle power data for 4 socket systems taken from Draft 3 Idle Power Dataset)

The first table above shows that in the equivalency class of all 16-core 16-GB DRAM systems, it is the system with 2 sockets x 8 cores that is the most efficient with the lowest absolute idle power (289 Watts) and idle power per core (18.06 Watts). Yet, it fails to qualify for Energy Star because it is subject to a 2 socket idle power allowance (222 Watts), while the other systems, which also have 16 cores and 16GB of DRAM, can potentially qualify for Energy Star since they do not have an idle power requirement.

The second table above shows that in the equivalency class of all 16-core 64-GB DRAM systems, it is the system with 2 sockets x 8 cores that is the most efficient, having the lowest absolute idle power (353 Watts) and idle power per core (22.1 Watts). Yet, it fails to qualify for Energy Star because it is subject to a 2 socket idle power allowance (298 Watts), while the other systems, which also have 16 cores and 64GB of DRAM, can potentially qualify for Energy Star since they do not have an idle power requirement.

Industry innovation has already reached the point where single microprocessor servers with 64 GB DRAM, run large terabyte data warehouses, lead in benchmarks for web, application, and database serving; and deliver enterprise-grade performance, reliability, and security. In the examples above, the 1 socket 8 core server indicated in the tables above also has 8 vertical hardware threads on each core, making it look to the operating system like a 64-way server. Grouping such a server in the same category as a 1-socket 2-core pedestal server designed for departmental print serving, and granting it the same idle power allowance, is a gross misclassification.

For a given enterprise workload, EPA's goal should be to encourage customers to use a smaller number of quad-core or higher servers, rather than a larger number of single or dual core servers, thereby saving overall power in the data center. **EPA's data set shows that for a given number of cores, servers with fewer sockets and higher core count per socket are more energy efficient.** EPA should encourage energy efficiency through multi-core innovation.

Request:

Sun proposes the following options to the EPA for recognizing the energy efficiency innovations inherent in 8 core system design, and to encourage the rest of the industry to move towards multi-core innovation to save overall power in the data center:

- ***Option 1:*** *For systems with 8 core micro-processors, create a separate category ("Categories E and F") for idle power allowances*
- ***Option 2:*** *Provide an adder for every core > 4 (20W/core)*
- ***Option 3:*** *Treat systems with 8 core innovation the same as 4 socket systems:*
    - *Exempt from any specific idle power requirement in Tier 1*
    - *Wait for an industry category to emerge after other vendors start shipping, and impose an idle power requirement in Tier 2*
    - *In Tier 1, permit systems to qualify for Energy Star if they meet all the other requirements*

B) Idle Power for Depopulated Systems

Draft 4 is ambiguous as currently worded with respect to 4 socket systems that are sold in depopulated configurations with only 2 processors installed.

Lines 629 – 631 in Draft 4 state, "...all attributes refer to the amount installed in the system, and not the amount the system is capable of supporting (e.g. installed processors, not processor sockets..."

This suggests that 4 socket systems in a depopulated 2 socket configuration must meet the idle power allowance for Category C (if unmanaged) or Category D (if managed)

On the other hand, lines 649 – 650 in Draft 4 say, "...EPA is proposing to hold off on requiring idle power for Computer Servers with greater than two sockets..."

This suggests that 4 socket systems in a depopulated 2 socket configuration are exempt from idle power requirements.

Request:

Sun suggests that the EPA clarify the language on lines 629 – 631 to state that the idle power requirements in Table 3 apply only to systems with 1 and 2 sockets.

3. Power Supply Requirements

A) Power Factors at Light Loads (line 575)

For low power single and multi-output PSUs, there is a lower threshold for power savings at light loads due to fixed losses to operate basic PSU functions. Fixed losses have significant impact on light-load performance targets when compared to output loading as a percentage

Request:

Eliminate power factor and efficiency requirements for loads below 75W output.

B) 1000W Cutoff Limit (line 574)

The 1000W cutoff limit still remains an area of concern. We request the EPA kindly reassess the 1200W limit. We have 2 options we would like the EPA to consider.

Request:

- *Move the threshold to 1200W. 1200W is the natural breakpoint in most power supply capabilities and is the maximum power that can be achieved in a low line situation due to limitations from the cord supply and infrastructure.*

- *If the EPA does not move the band to 1200W we would request that for any PSU above 1000W the increased efficiency targets be moved to align with the CSCI Gold target date (July 2009), as this is the date the industry has been working to.*

C) Power Accuracy (line 886)

The current requirement of +/-10% accuracy is very difficult to meet at loads less then 100W. As the output load reduces the input current, the waveform degrades from being regular at 100% to irregular at 30%, making it difficult to define an RMS value.

When input power is low (below 30% load), the input current through the PFC sense resistor is small and irregular. The accuracy of measuring the voltage across the PFC resistor (as a representation of the current) is influenced by fixed errors and the point on the waveform where the measurement is taken.

Request:

Change the accuracy requirements as indicated in the table below:

Server Input Power Range Per PSU	Maximum Allowed Error
0W ≤ Pin < 100W	+/- 20W
Pin ≥100W	+/- 10%

D) Efficiency requirements for power supplies <500W (line 574)

CSCI and 80Plus.org aligned efficiency targets to CSCI bronze silver and gold. Energy Star server specification Draft 4 deviates from this alignment which will create confusion in the industry.

Request:

Align power supply efficiency targets to the CSCI and 80Plus.org bronze or silver targets for power supplies <500W.

4. Sampling Rates (line 895)

The current Draft 4 specification requires a sampling rate of one sample per second. In a server environment the power does not change that fast. Power increases are seen over a longer period of time. This is also true of temperature data and processor utilization data. Sampling the power draw, temperature and processor utilization once per second generates an extremely large volume of relatively static, unchanging data which conveys little meaningful information to the customer and unnecessarily fills up the customer's monitoring database.

In addition, servers today typically have many sensors to read and take readings from each sensor in turn. Depending on the amount of sensors to read and the workload of the server at the time, a poll of 1/sec can not be achieved. These sensors typically communicate over an industry standard I2C bus, which has very low bandwidth and can only communicate data at low speeds, thereby inhibiting its use as a real-time sampling bus with one second intervals.

With larger systems there are significant numbers of sensors distributed throughout the different modules, not only those in the power supply. Without increasing the speed of the interfaces and the service processors (and likely cost and power consumption) these cannot all be accessed once per second.

Typically, the power supply has a primary side micro controller that process the input sensors before passing the data back over the primary to secondary barrier and often there is some filtering/averaging applied at this stage to convert the samples to an RMS equivalent before the sensor data is accessed by the system.

For conveying meaningful information to a customer, many customer's have indicated that a reporting rate of once every fifteen minutes is sufficient, because the underlying data does not change very fast. Customers are more interested in longer terms patterns of change (e.g. over a 24 hour cycle or a seven day cycle) in the power, temperature, and utilization data than in second-by-second readouts.

Request:

- *Specify that hardware polling rates of embedded sensors must meet a minimum of one sample per minute.*

- *Specify that rolling averages shall be reported at least once every 5 minutes for power and utilization data.*

- *Specify that temperature data shall be reported at least once every 5 minutes. Eliminate the rolling average requirement for temperature data.*

5. Bandwidth Measurement for Input Current Testing

In order to maintain accurate results from test equipment when measuring input current a bandwidth range should be added to either the specification or the power supply test procedure referenced within the energy star specification.

Request:

Specify minimum and maximum required bandwidths of 3kHz and 20kHz respectively.

6. Test Equipment Accuracy (line 1254)

On line 1254, in Appendix A, there is a requirement for the power meter to have an accuracy of 0.01 Watt or better for power measurements of 10W or less.

- This requirement greatly increases the cost of the meter, making it more difficult for manufacturers to make the meters available to development staff. Since the smallest power that will be required to be measured under the Draft 4 specification is 55W, there is no reason to require this level of accuracy at values under 10W.

- Given that the spec also indicates that power numbers should be rounded to the first decimal place, an accuracy to 2 decimal places is redundant.

- This requirement seems to be at odds with the language in lines 1237 through 1243 indicating that a requirement to measure very low power levels had been removed.

Request:

Remove the requirement for the power meter to have an accuracy of 0.01 Watt or better for power measurements of 10W or less.

7. Blades

Sun applauds the EPA for the inclusion of blades in the Energy Star for Servers specification. However, Sun remains concerned that the idle power data collection exercise for blade servers will not complete or be analyzed in time for the date that the EPA anticipates for Tier 1. Hence, Sun would like to request that the inclusion of blades in Energy Star be deferred until Tier 2 or a later revision of the Tier 1 specification.

Request:

Defer Blade requirements until Tier 2 or a later revision of the Tier 1 specification.

8. Definitional Clarifications (line 633)

In Table 4, the use of "Port" and "Device" in the right hand column for different I/O "Devices" is unclear. Sun suggests that the EPA clarify the exact semantics of “port,” “device,” and “interface” with precise language. Below are some suggested definitions for each of these terms:

- ***Interface:*** *an abstract term representing a type of I/O or networking capability in the system. An interface may be provided by many devices and may have many ports. For example, a system's Infiniband capability may be provided by an Infiniband interface with multiple chips and have multiple Infiniband ports,while a system's Gigabit Ethernet capability may be provided by a GigE interface which may have multiple chips and multiple Ethernet ports.*

- ***Device:*** *Actual physical logic circuits that are needed to provide an interface of a particular type. A single device may have many ports. A device can be an external add-in card (a printed circuit board with ASIC chips on it), or be justan ASIC chip by itself that's designed into the motherboard, or be just logic circuitry that's designed in as a feature of the microprocessor.*

- ***Port:*** *The place in the system where an actual networking or I/O session can be established. A port is not the same as a connector receptacle, e.g. it's possible that a single receptacle on the system that accepts a single connector actually services multiple ports of the same interface.*

The following examples will help make the above definitions more clear.

- Some interfaces are provided by external add-in cards that are added into PCIe slots, in which case they are considered 1 device with up to N ports

- Some interfaces are provided by ASICs on the motherboard, in which case they should be considered 1 device (if provided by 1 ASIC or by 1 PHY+MAC chipset) which could also have up to N ports. If they are provided by 2 different ASICs (or 2 different PHY+MAC chipsets) then they are 2 devices.

- Some interfaces are provided on the microprocessor in which case they should also be considered 1 device with up to N ports, because even if the MAC layer is built into the microprocessor the PHY layer is usually not, and is a separate device. So if there is 1 separate PHY ASIC on the motherboard, that's 1 device, if there are 2, that's 2 devices, etc.

Request:

Clarify precise definitions of interface, device, and port along the lines suggested above.

9. Effective Date (line 1150):

In the Energy Policy Act of 2005 (Section 131 of PL109.58), which amends the Energy Policy and Conservation Act (42 USC 6294a. Sec. 324), the duties of the EPA administrator with respect to the Energy Star program are specified as follows:

> "The Administrator and the Secretary shall provide appropriate lead time (which shall be 270 days, unless the Agency or Department specifies otherwise) prior to the applicable effective date for a new or a significant revision to a product category, specification, or criterion, taking into account the timing requirements of the manufacturing, product marketing, and distribution process for the specific product addressed."

Since the Energy Star for Servers specification is a new specification, and server vendors need the appropriate lead time for manufacturing, marketing, and distributing products compliant with this new specification, we recommend that the EPA not deviate from the 270 day lead time notification suggested in the legislation above.

Request:

Sun recommends an effective date of 270 days following the publication of the Tier 1 specification for the reasons stated in the Energy Policy Act of 2005.