

December 5, 2008

**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 3 Energy Star Specification for Computer Servers

Dear Ms. Duff:

Thank you for the opportunity to provide comments on the EPA's Draft 3 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.

The comments that follow are consistent with the dialogue that has occurred to date, and are made with the purpose of achieving a specification that better achieves our mutual goals. You will notice that we are providing substantial feedback on Draft 3's new idle power categorizations, and while we support the categorization approach, we are recommending significant changes.

We look forward to discussing these points in more detail with you next week.

Sincerely,

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

## **Table of Contents**

1.Introduction.....	3
2.Idle Power:.....	3
A)Taxonomy Basis for Idle Power Allowances.....	3
B)Standard vs. High Availability Dichotomy.....	9
C)16GB Memory Dichotomy: .....	9
D)Redundant Cooling and Other Adders.....	9
E)Depopulated Systems.....	10
F)Summary of Idle Power Requests.....	12
3.Power Supply Requirements .....	14
A)1000W versus 1200W threshold limit.....	14
B)Power Supplies that include standby rail power .....	15
C)Small Power Supplies requirements and Table 1 updates .....	15
D)Clarification if load measurements are based on individual or system level loads .....	16
E)Required Accuracy at low load levels:.....	16
4.Effective Date:.....	16
A)Tier 1 Effectivity date .....	16
5.Labeling Requirements:.....	17

# **COMMENTS BY SUN MICROSYSTEMS, INC. ON THE EPA ENERGY STAR (DRAFT 3) PROGRAM REQUIREMENTS FOR COMPUTER SERVERS**

## **1. Introduction**

Sun commends the EPA on issuing Draft 3 of the Energy Star for Servers specification. This draft represents 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. Nevertheless, Sun has significant concerns about the following aspects of the Draft 3 specification, which are detailed in this document:

1. Overt emphasis on low-utilization scenarios, in particular the focus on idle power unqualified by any top-line performance metric
2. The basis of the taxonomy on which idle power allowances are granted
3. The basis of the taxonomy on which power supply efficiency criteria are specified, and the accuracy requirements for power supply reporting
4. The effective date of the specification
5. The labeling requirements of the specification.

## **2. Idle Power:**

### **A) Taxonomy Basis for Idle Power Allowances**

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 energy efficiency (performance per watt) of multicore systems is greater, 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, multicore systems, which are more energy efficient, are penalized, since their lower socket count restricts them to a lower idle power allowance.

Recent vendor innovation in multicore 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 2 socket 4 core system does just as much work as a 4 socket 2 core system, but consumes significantly less power. As such, a 2 socket 4 core system can be used to replace either a single 4 socket 2 core system, or to replace two 2-socket 2 core systems, in either scenario saving power overall in the data center. Hence, a vendor providing a 2-socket 4 core

system ought to be rewarded, not penalized, for doing 8 cores worth of work with only 2 sockets, as other vendors do with 4 sockets.

The current EPA specification penalizes, not rewards, systems with 2 sockets with 4 cores each by granting them a lower idle power allowance than systems with 4 sockets and 2 cores each. We recommend that the EPA reward multicore innovation by vendors by granting an idle power allowance on the basis of *core count* instead of *socket* or *installed processor* count.

To prove this point, the following table compares idle power for systems that contain 8 cores: 2 socket 4 core systems versus 4 socket 2 core systems, all containing 8GB of memory. We define this set of servers as the **equivalency class** of all 8-core systems with 8GB 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 data set. Index number on the left indicates the specific EPA data set item number..

Index	CPU Characteristics	Sockets	Cores/ CPU	Total Cores	Total Installed Memory	Total Installed Hard Drives	Idle Power	Idle Power / Core	Idle Power / Socket	EPA Allowance
63	HPT E5440 2.83G/ 80W/1333	2	4	8	8	1	151	18.88	75.5	151
64	HPT LV L5420 2.33G /50W / 1333	2	4	8	8	1	154	19.25	77	151
65	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	8	1	173	21.66	86.65	151
66	2.67GHz, 2x6MB L2 Cache, 1333MHz system bus	2	4	8	8	1	182	22.75	91	151
67	3.33GHz, 2x6MB L2 Cache, 1333MHz system bus	2	4	8	8	1	186	23.25	93	151
68	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	8	1	186	23.26	93.05	151
69	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	8	1	187	23.38	93.5	151
70	HPT E5440 2.83G/ 80W/1333	2	4	8	8	1	187	23.38	93.5	151
71	2.83GHz, 2x6MB L2 Cache, 1333MHz system bus	2	4	8	8	1	189	23.63	94.5	151
72	3.33GHz, 2x6MB L2 Cache, 1333MHz system bus	2	4	8	8	1	190	23.75	95	151
73	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	8	1	191	23.83	95.3	151
74	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	8	1	203	25.41	101.65	151
107	3.0Ghz 2x4MB L2, 2.93 GHz, 1066 MHz FSB,	4	2	8	8	8	266	33.25	66.5	334

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

All the systems in the table above do 8 cores worth of work, and are deployed for similar purposes in the data center. The last row in Table 1 (Index 107) indicates that the 4 socket 2 core

system is least efficient and has the highest idle power per core of this peer group. However, the EPA provides this system with the most generous idle power allowance at 334 watts. The other systems *in the same equivalency class*, which also do 8 cores worth of work but do it only with 2 sockets, are penalized by the EPA by only being granted an idle power allowance of 151 watts.

The following table illustrates the same point for the equivalency class of systems with 8 cores and 16 GB DRAM.

Index	CPU Characteristics	Sockets	Cores/ CPU	Total Cores	Total Installed Memory	Total Installed Hard Drives	Idle Power	Idle Power / Core	Idle Power / Socket	EPA Allowance
36		1	8	8	16	1	160	20	160	60
79	HPT E5440 2.83G/ 80W/1333	2	4	8	16	1	161	20.13	80.5	221
80	2x6MB L2 1333MHz	2	4	8	16	2	167	20.88	83.5	236
81	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	16	1	200	25	100	221
82	3.33GHz, 2x6MB L2 Cache, 1333MHz system bus	2	4	8	16	1	221	27.63	110.5	221
83	2x6MB L2 1333MHz	2	4	8	16	4	227	28.38	113.5	252
84	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	16	1	228	28.5	114	221
85	2x6MB L2 1333MHz	2	4	8	16	6	234	29.25	117	268
86	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	16	1	234	29.29	117.17	221
87	3.0GHz, 12MB L2 Cache, 1333MHz system bus	2	4	8	16	1	243	30.33	121.3	221
89	2.66GHz, 8 MB L2 Cache, 1333Mhz system bus	2	4	8	16	2	291	36.38	145.5	236

*Table 2: 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 first system (Index 36) 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. All the other 8 core systems require 2 sockets of 4 cores each, and are less efficient with idle power per core varying between 20.13 and 36.38 W/core.

The above data demonstrates that the EPA's Draft 3 specification penalizes the most energy efficient and innovative system by only granting it an idle power allowance of 60 watts, while other less efficient system designs *in the same equivalency class* are granted generous idle power allowances between 221 and 236 watts.

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

Index	CPU Characteristics	Sockets	Cores/ CPU	Total Cores	Total Installed Memory	Total Installed Hard Drives	Idle Power	Idle Power / Core	Idle Power / Socket	EPA Allowance
39		1	8	8	64	1	255	31.88	255	60
97	E7430	2	4	8	64	6	315	39.38	157.5	268
98	2.66GHz, 8 MB L2 Cache, 1333Mhz system bus	2	4	8	64	2	345	43.13	172.5	236

Table 3: Comparing the Equivalency Class of Systems with 8 cores and 64GB DRAM

The first row of this table demonstrates that a 1 socket 8 core system with high memory provides the lowest idle power and idle power per core (31.88W/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 20.13 and 36.38 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 60 watts, while less efficient system designs *in the same equivalency class* are granted generous idle power allowances between 236 and 268 watts.

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 November 2008, that are above 20,000. This table can be obtained from <http://www.spec.org/web2005/results/web2005.html>.

Vendor	System	Cores	CPUs	Cores/CPU	SPECweb2005 Result
Sun Microsystems, Inc.	Sun SPARC Enterprise T5220	8	1	8	41847
Sun Microsystems Inc.	Sun SPARC Enterprise T5220	8	1	8	37001
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

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.***

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). Further, the EPA grants 2x4 systems (that lead in enterprise SPECweb performance, as shown in the previous table) the same allowance as 2x1 systems (which could also 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	CPU Characteristics	Sockets	Cores/ CPU	Total Cores	Total Installed Memory	Total Installed Hard Drives	Idle Power	Idle Power / Core	Idle Power / Socket	EPA Allowance
88		2	8	16	16	2	289	18.06	144.5	236
110	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	1	329	20.58	82.33	271
111	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	2	337	21.09	84.35	286
112	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	3	342	21.4	85.59	294
113	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	4	369	23.06	92.25	302
114	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	1	371	23.21	92.85	271
115	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	1	373	23.31	93.25	271
116	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	1	384	24	96	271
117	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	16	5	420	26.25	105	310

Table 5: Comparing The Equivalency Class of 16-core Systems with 16GB DRAM

Index	CPU Characteristics	Sockets	Cores/ CPU	Total Cores	Total Installed Memory	Total Installed Hard Drives	Idle Power	Idle Power / Core	Idle Power / Socket	EPA Allowance
99		2	8	16	64	2	353	22.06	176.5	236
121	2x4MB L2 1066MHz	4	4	16	64	8	496	31	124	334
122	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	64	4	596	37.25	149	302
123	2.93GHz, 8MB L2 Cache, 1066MHz system bus	4	4	16	64	5	662	41.38	165.5	310

Table 6: Idle Power for 16 Core Systems - Comparing The Equivalency Class of 16 core Systems with 64GB DRAM

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 both the lowest absolute idle power (269 Watts) and idle power per core (18.06 Watts). Yet, it is penalized by being granted the lowest idle power allowance in its equivalency class (236 vs. 270 plus for the others).

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 with having both the lowest absolute idle power (353 Watts) and idle power per core (22.06 Watts). Yet, it is penalized by being granted the lowest idle power allowance in its equivalency class (236 vs. 302 plus for the others).

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 mis-classification. Below is the summary of the request from Sun on this issue:

**Discussion Summary:**

- *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 multicore innovation*

**Request:**

*Grant idle power allowance on the basis of core count, not processor or socket count. It would be a much more meaningful basis of comparison if equivalency classes could be created on the basis of core count instead of processor or socket count.*

**Alternative:**

*If the EPA is committed to defining server equivalency classes on the basis of processor or socket count, then it could consider granting an idle power allowance on the basis of processor or socket count plus an adder for every core > 2 (e.g. 50W/core) to encourage multicore innovation.*

**B) Standard vs. High Availability Dichotomy**

The applicability of the high vs. standard availability dichotomy should be extended to any system that has high availability regardless of socket count.

For 2 socket systems the EPA makes a distinction between standard availability and high availability. For 1 and 4 socket systems, there is no such distinction. This penalizes 1 socket systems that are designed for high availability and have enterprise features (because they are multicore) and are restricted to 60 Watts idle power.

**Request:**

*The standard vs. high availability dichotomy should be extended to all system classes, regardless of socket count.*

### C) 16GB Memory Dichotomy:

For 2 socket systems, the EPA makes a distinction between systems with < 16GB DRAM and >= 16GB DRAM. For 1 and 4 socket systems, there is no such distinction; there is only an allowance for > 32GB DRAM. This penalizes 1 socket systems that have between 16 and 32GB DRAM by restricted these systems to 60 Watts idle power since the memory adder does not kick in until after 32 GB of DRAM.

**Request:**

*The Memory idle power specification should be changed to grant an allowance for any system that has >= 16GB DRAM regardless of socket count.*

*The suggested adder for over 16GB DRAM is 3W/GB. This is derived from EPA data, where the EPA grants a base allowance of 169W for a 2-CPU system with 16GB, and a base allowance of 221 W for a 2-CPU system with 32 GB. This indicates that the right allowance for an incremental 16GB of DRAM is 52W, which is approximately 3W/GB.*

### D) Redundant Cooling and Other Adders

The EPA's definition of high-availability does not include redundant cooling, nor is there an adder for redundant cooling. This penalizes systems that have redundant rows of fans. Sun suggests two potential approaches:

**Option 1 (preferred):**

*Change the definition of high availability to include redundant cooling, and increase the baseline allowance for high availability systems by 10W to account for redundant fans.*

**Option 2:**

*Keep the baseline the same and grant an adder of 10W for systems with redundant fans*

Further, the EPA's current spec provides no allowances for features such as high-speed networking, data center fabric I/O, hardware RAID and crypto ASICs. These functions are implemented in hardware and consume power, and hence require additional allowances.

**Request:**

*Additional allowance needs to be provided for redundant fans and Crypto/network ASICS.*

We suggest granting the following allowances as adders:

Feature	Allowance
10GbE Networking	8W/port
PCIe Add-In Cards	10W/card
Hardware RAID	5W
Crypto ASICs	10W

## E) Depopulated Systems

Currently, the EPA only grants idle power allowance on the basis of installed processors, regardless of socket count. This unfairly penalizes higher-end systems which frequently ship in depopulated configurations, as a depopulated configuration of a higher-end system is not the same as a lower-end fully populated system with the same number of microprocessors.

### Question for the EPA:

*Is an 8 socket system eligible for Energy Star compliance when shipped with 4 unpopulated sockets? Does it fall out of scope for Energy Star in the field when a customer adds a fifth CPU to the system?*

Our recommendation is that the EPA should consider designed sockets, not populated sockets, when factoring idle power allowances.

The following table shows that the idle power for a single installed processor is vastly different in 1 socket, 2 socket, and 4 socket systems. Yet, they are granted the same idle power allowance.

Index	CPU Characteristics:	CPU Frequency (MHz):	# of Sockets	# CPU(s) Installed	Cores / CPU	Total Installed Memory Amount (GB):	# of Installed Dimms	Size of Dimms (MB)	Idle Power
EPA Dataset 14	2.4GHz, 2x4MB L2 Cache, 1067MHz system bus	2400	1	1	4	4	2	2048	69
EPA Dataset 15	2.40GHz, 2x4MB L2 Cache, 1067MHz system bus	2400	1	1	4	4	2	2048	78
EPA Dataset 19	3.0GHz, 2x6MB L2 Cache, 1333MHz system bus	3000	1	1	4	8	4	2048	59
EPA Dataset 23	3.0GHz, 2x6MB L2 Cache, 1333MHz system bus	3000	1	1	4	8	4	2048	68
Sun Measured	3.0GHz, 1MB L2 Cache, 1000MHz system bus	2992	2	1	2	8	8	1024	275
Sun Measured	3.0GHz, 1MB L2 Cache, 1000MHz system bus	2992	2	1	2	32	16	4096	346
Sun Measured	3.0GHz, 1MB L2 Cache, 1000MHz system bus	2992	2	1	2	4	4	1024	200
Sun Measured	3.0GHz, 8 MB L2 Cache, 1333Mhz system bus	3000	2	1	2	4	4	1024	200
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	1	8	64	16	4096	553
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	1	8	8	4	2048	386
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	1	8	16	8	2048	423
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	1	8	32	16	2048	486

*Table 7: Idle Power for Systems with a Single Installed Processor - Comparing 1 Socket 1 CPU vs. 2 Socket 1 CPU vs. 4 Socket 1 CPU*

The table above shows that in systems that have only a single installed processor, the idle power is vastly different depending on whether the system is designed with only 1 socket (idle range typically < 100 W) or 2 sockets (idle range between 200 and 350 W) or 4 sockets (idle range between 386 and 553 W).

This suggests that systems designed with different numbers of sockets should be treated differently regardless of number of installed processors.

The table below shows the idle power analysis for systems with only 2 installed processors but with a varying socket count.

Index	CPU Characteristics:	CPU Frequency (MHz):	# of Sockets	# CPU(s) Installed	Cores / CPU	Total Installed Memory Amount (GB):	# of Installed Dimms	Size of Dimms (MB)	Idle Power
EPA Dataset 58	2.6GHz, 2x1MB L2 Cache, 1000 MHz system bus	2613	2	2	2	4	4	1024	186
EPA Dataset 51	3.0GHz, 12MB L2 Cache, 1333MHz system bus	3000	2	2	4	4	2	2048	171
EPA Dataset 44	2.5GHz, Quad Core 2x6MB L2 Cache, 1333MHz s	2500	2	2	4	4	2	2048	111
EPA Dataset 46	3.0GHz, 12MB L2 Cache, 1333MHz system bus	3000	2	2	4	4	2	2048	155
Sun Measured	2.93GHz, 8 MB L2 Cache, 1066Mhz system bus	3000	4	2	2	4	4	1024	262
Sun Measured	3.0Ghz 2x4MB L2, 2.93 GHz, 1066 MHz FSB,	2933	4	2	2	4	4	1024	303
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	2	8	64	8	8192	548
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	2	8	32	8	4096	560
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	2	8	64	16	4096	650
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	2	8	16	8	2048	530
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	2	8	32	16	2048	600
Sun Measured	Model UltraSPARC T2+ 8-core CPUs	1400	4	2	8	64	32	2048	740

Table 8: Idle Power for Systems with Two Installed Processors - Comparing 2 Socket 2 CPU vs. 4 Socket 2 CPU

The table above shows that in systems that have only two installed processors, the idle power is vastly different depending on whether the system is designed with 2 sockets (idle range between 111 and 186 W) or 4 sockets (idle range between 262 and 740 W).

This suggests that systems designed with different numbers of sockets should be treated differently regardless of the number of installed processors. A 4 socket system with two unpopulated sockets has higher idle power than a 2 socket system because of its enterprise features, and should be granted a higher idle power allowance.

**Option 1 (preferred):**

- *Require Energy Star compliance for models with fully populated configuration only. The Energy Star data sheet should only be required for the models with fully populated configurations.*
- *Require testing for compliance with all CPU sockets and DIMM slots populated, and all PCIe slots unpopulated.*
- *Deem depopulated configs to be automatically compliant if its fully populated config is compliant (because a depopulated config might become a fully populated config in the field).*

**Option 2:**

- *Grant idle power allowances on the basis of **designed** core count (or **designed** socket count with multicore adders), not on the basis of **installed** processors.*
- *Require a “subtractor” instead of an “adder” for depopulated cores (or depopulated sockets) and DIMM slots.*

## F) Summary of Idle Power Requests

The following table presents a summary of our requests on the idle power metric. The first table uses core instead of processor or socket count as the basis for a server taxonomy, and is the most preferred form.

The second table is the less preferred alternative, as it uses processor or socket count as the basis for taxonomy.

The simplicity and elegance of these tables arise from the fact that the notions of standard vs. high availability, and the notion of extended memory, have been extended across all server classes, and not restricted to 2-socket servers only. Both tables have identical “adders” and “subtractor” proposals to account for incremented or decremented hardware features.

# of Cores Designed (not populated)	Base Idle Power Allowance for Standard Availability	Base Idle Power Allowance for High Availability
2	TBD by EPA	TBD by EPA
4	TBD by EPA	TBD by EPA
8	TBD by EPA	TBD by EPA
16	TBD by EPA	TBD by EPA
32	TBD by EPA	TBD by EPA

Feature	Adder
Memory >16GB DRAM	3W/GB
2 <sup>nd</sup> Hard Drive	15W
3 <sup>rd</sup> and Subsequent Hard Drive	8W
10GbE Networking	8W/port
PCIe Add-In Cards	10W/card
Hardware RAID	5W
Crypto ASICs	10W
Redundant Cooling	10W/fan bank

Feature	Subtractor
Unpopulated CPU Socket	35W/core
Memory Config Less than Max	3W/GB

Table 9: A Possible Specification: Alternative One(Preferred)

# of Sockets Designed (not populated)	Base Idle Power Allowance for Standard Availability	Base Idle Power Allowance for High Availability
1	TBD by EPA	TBD by EPA
2	TBD by EPA	TBD by EPA
4	TBD by EPA	TBD by EPA

Feature	Adder
Memory >16GB DRAM	3W/GB
2 <sup>nd</sup> Hard Drive	15W
3 <sup>rd</sup> and Subsequent Hard Drive	8W
10GbE Networking	8W/port
PCIe Add-In Cards	10W/card
Hardware RAID	5W
Crypto ASICs	10W
Redundant Cooling	10W/fan bank
Number of Cores > 2x Number of Designed Sockets	50W/Core

Feature	Subtractor
Unpopulated CPU socket	140W/socket
Memory Config Less than Max	3W/GB

Table 10: A Possible Specification: Alternative Two

### 3. Power Supply Requirements

#### A) 1000W versus 1200W threshold limit

Sun is still concerned about the 1000W versus 1200W PSU cutoff point between large and small power supplies and requests the EPA kindly reassess the 1000W limit. 1200W power supplies have become the power supply of choice for high power applications in 1U/2U system configurations due to the following reasons:

- 1200W is the maximum power that can be provided in low line configurations due to the limitations of power cord supply and infrastructure.
- The majority of the high efficiency, high availability power supplies are in the 500W to 1200W range. The distribution of unit shipments of power supplies for the volume

server market is predominantly clustered in this range, as the IEC connector types C13 and C14 go up to 1200 watts.

- By drawing the dividing line at 1000W, the EPA is effectively creating an artificial boundary right in the middle of this population cluster, instead of drawing the line at the more natural break point (1200W) where the unit volume shipments tail off in the volume server space. The next unit volume population cluster begins at the 1500W range.

In this regard we have 2 options we would like the EPA to consider:

**Option 1 (preferred):**

*Move the 1000W threshold to 1200W.*

**Option 2:**

*Move the implementation date for >1000W PS to July 2009 in alignment with the CSCI Gold implementation date.*

**B) Power Supplies that include standby rail power**

Power supply vendors produce power supplies at 1000W and 750W but that power does not include the standby rail power. Including standby rail power under the current guidelines -requires compliance to the stringent criteria for larger supplies, placing vendors that provide this high availability feature at a disadvantage in meeting Energy Star compliance.

**Request:**

*Please clarify that the standby rail power is included in the nominal power rating. If standby rail power is included in power supply rating, adjust 1000W and 750W by 50W and 10W respectively (1050W & 760W).*

**C) Small Power Supplies requirements and Table 1 updates**

Sun agrees with the line 425 discussion that "smaller power supplies are at a disadvantage at low loads" and requests the table be updated to include separate requirements for small power supplies ≤ 760W (exempt from 10% efficiency and PFC requirements) along with the other recommended changes to Table 1 below.

**Request:**

*Include the following power supply requirements in Table 1 as noted below:*

*Changes from Draft 3 are noted in RED*

<b>Power Supply Type</b>	Rated Output	10% Load	20% Load	50% Load	100% Load
Multi-Output	All Output Levels	N/A	82.00%	85.00%	82.00%

(AC-DC & DC-DC)					
Single-Output (AC-DC & DC-DC)	≤ 1200 Watts (or 1050W)	75.00%	85.00%	89.00%	85.00%
	> 1200 Watts (or 1050W)	80.00%	88.00%	92.00%	88.00%
	≤ 760 Watts	N/A	85.00%	89.00%	85.00%

#### D) Clarification if load measurements are based on individual or system level loads

In previous EPA correspondence, the EPA indicated that energy star compliance is a system level compliance. Sun would like clarification that all the specification items do apply to a single PSU. Furthermore the EPA was considering allowing exemption from the 10% rule if vendors can demonstrate that the PSU will never operate at a 10% load level in a particular server configuration.

**Request:**

Please clarify if 1) load measurements are based on individual PS loads or system level loads and 2) the requirement to measure at 10% load if it is demonstrated that the PSU will never operate at that load level on a specific server configuration.

#### E) Required Accuracy at low load levels:

The power supply needs to measure the input voltage and input current to provide the input power. The input current measurement typically uses the small value PFC current sense resistor to avoid adding to the power supply losses.

It is not possible to restrict input current measurement accuracy to 10% for low output loads. As the output load reduces the input current, the waveform becomes non-sinusoidal. The current waveform is regular at the full output load and at 50% load but becomes irregular as the load drops through the 30% point, making it difficult to define the RMS value.

**Request:**

Define the Input power measurement +/-10% accuracy to be required over the power supply output load range from 30% to 100%. Remove the accuracy limit requirements below 30% output load. The volume server PS is too dense to permit the extra circuitry required to provide good accuracy at less than 30% load.

### 4. Effective Date:

#### A) Tier 1 Effectivity date

Draft 3 included major changes to the specification. This we welcomed. However, we know



that Sun is not the only stakeholder requesting further substantial changes, and it would seem that a longer process to include at least a Draft 4 and a face-to-face meeting in January is required. A February 1<sup>st</sup> effective date no longer seems reasonable.

Sun recommends that the drafting process be allowed to continue to progress constructively, and yet with appropriate urgency, till a specification has been achieved that truly and reliably differentiates efficient models from less efficient models, a goal not currently met with the current Draft 3 specification. A premature specification put into effect too soon will be to the advantage of no one in the industry, including users. Sun also recommends a full product design cycle between the publication of the specification and its effective date, i.e., 24 months after the issuance of the final specification.

**Request:**

*Sun requests that the effective date be 24 months after the issuance of the final Tier 1 specification.*

## **5. Labeling Requirements:**

- Sun applauds the EPA for removing the packaging labeling requirement and to allow temporary labeling on the product. However, Sun believes that there should be no product labeling requirement for servers for the following reasons:
- It places an undue burden on the vendor to provide special labeling for compliant versus non-compliant versions of the same model.
- The EPA designation will be incorrect if the system is upgraded to a non-compliant level in the field.
- Servers are back-end devices that are commonly operated in a data centers where the label is not visible to the end user.

**Request:**

*Remove the requirement to label the product with an Energy Star mark, except for compliant configurations for whom the Energy Star mark can be expressed on the product's web site.*