ENERGY STAR® Program Requirements
Product Specification for Computer Servers

Eligibility Criteria
Draft 2 Version 2.0

Following is the Version 2.0 ENERGY STAR Product Specification for Computer Servers. A product shall meet all of the identified criteria if it is to earn the ENERGY STAR.

1 DEFINITIONS

Note: The definitions in this section largely mirror those distributed with EPA's data assembly documentation. Further revisions are indicated with noteboxes.

A) Product Types:

1) Computer Server: A computer that provides services and manages networked resources for client devices (e.g., desktop computers, notebook computers, thin clients, wireless devices, PDAs, IP telephones, other computer servers, or other network devices). A computer server is sold through enterprise channels for use in data centers and office/corporate environments. A computer server is primarily accessed via network connections, versus directly-connected user input devices such as a keyboard or mouse. For purposes of this specification, a computer server must meet **all** of the following criteria:

   a) is marketed and sold as a computer server;

   b) is designed for and listed as supporting one or more computer server operating systems (OS) and/or hypervisors, and is targeted to run user-installed enterprise applications;

   c) provides support for error-correcting code (ECC) and/or buffered memory (including both buffered DIMMs and buffered on board (BOB) configurations)-systems with greater than 50 nodes sharing the same chassis are exempt from this requirement; and

   d) is packaged and sold with one or more ac-dc or dc-dc power supplies; and

   e) is designed such that all processors have access to shared system memory and are independently visible to a single OS or hypervisor.

Note: EPA had proposed removing the ECC/buffered memory provision from the Computer Server definition in response to stakeholder concerns about excluding products relying on non-ECC memory solutions. The products at issue contain a significantly higher number of computer nodes than present in both mainstream and blade configurations and offer the opportunity to save energy through use of less resilient memory. This was offset in their design by the ability to power up nodes dynamically to scale capacity with power consumption. These products are targeted at the data center server market.

Because EPA recognizes a) the scope challenges possible with a complete removal of the existing memory provision; b) the unique nature of non-ECC memory products and concerns brought to EPA's attention as described above; and c) the information-disclosure goals of this version of the specification, an exemption from the ECC/buffered memory provision is proposed for systems with greater than 50 nodes.

2) Managed Server: A computer server that is designed for a high level of availability in a highly managed environment. For purposes of this specification, a managed server must meet **all** of the following criteria:

   a) is designed to be configured with redundant power supplies; and

b) contains an installed dedicated management controller (e.g., service processor).

3) Blade System: A system comprised of a blade chassis and one or more removable blade servers and/or other units (e.g., blade storage, blade network equipment). Blade systems provide a scalable means for combining multiple blade server or storage units in a single enclosure, and are designed to allow service technicians to easily add or replace (hot-swap) blades in the field.

   a) Blade Server: A computer server that is designed for use in a blade chassis. A blade server is a high-density device that functions as an independent computer server and includes at least one processor and system memory, but is dependent upon shared blade chassis resources (e.g., power supplies, cooling) for operation. A processor or memory module that is intended to scale up a standalone server is not considered a Blade Server.

      (1) *Multi-bay Blade Server*: A blade server requiring more than one bay for installation in a blade chassis.

      (2) *Single-wide Blade Server*: A blade server requiring the width of a standard blade server bay.

      (3) *Double-wide Blade Server*: A blade server requiring twice the width of a standard blade server bay.

      (4) *Half-height Blade Server*: A blade server requiring one half the height of a standard blade server bay.

   b) Blade Chassis: An enclosure that contains shared resources for the operation of blade servers, blade storage, and other blade form-factor devices. Shared resources provided by a chassis may include power supplies, data storage, and hardware for dc power distribution, thermal management, system management, and network services.

   c) Blade Storage: A storage device that is designed for use in a blade chassis. A blade storage device is dependent upon shared blade chassis resources (e.g., power supplies, cooling) for operation.

4) Fully Fault Tolerant Server: A computer server that is designed with complete hardware redundancy, in which every computing component is replicated between two nodes running identical and concurrent workloads (i.e., if one node fails or needs repair, the second node can run the workload alone to avoid downtime). A fully fault tolerant server uses two systems to simultaneously and repetitively run a single workload for continuous availability in a mission critical application.

5) Resilient Server: A computer server that is designed with resiliency, RAS, and self-correction features integrated in the micro-architecture of the CPU and chipset to ensure data resiliency and accuracy. A resilient server is often used for a limited set of workloads that may include business processing, decision support, or handling of virtualized workloads. For purposes of this specification, a resilient server must meet **all** of the following criteria:

   a) contains hot-swappable components (e.g., I/O, hard drives, and ac-dc power supplies);

   b) contains multiple physical banks of memory and I/O busses;

   c) provides machine check architecture (i.e., both Fault Isolation and Resiliency;

   d) provides memory fault detection and system recovery though DRAM chip sparing, extended ECC, and mirrored memory;

   e) provides support for error-correcting code (ECC) and/or buffered memory (including both buffered DIMMs and buffered on board (BOB) configurations);

   f) provides end-to-end bus retry; and

   g) supports on-line expansion/retraction of hardware resources without the need for operating system reboot (“on-demand” features).

Note: EPA understands that stakeholders are actively developing a more cohesive set of criteria describing Resilient Servers. EPA plans to evaluate the resulting proposal once available and welcomes comments on this issue.

6) Multi-node Server: A computer server that is designed with two or more independent server nodes that share a single enclosure and one or more power supplies. In a multi-node server, power is distributed to all nodes through shared power supplies. Server nodes in a multi-node server are not designed to be hot-swappable.

   a) Dual-node Server: A common multi-node server configuration consisting of two server nodes.

7) Server Appliance: A computer server that is bundled with a pre-installed operating system and application software that is used to perform a dedicated function or set of tightly coupled functions. Server appliances deliver services through one or more networks (e.g., IP or SAN), and are typically managed through a web or command line interface. Server appliance hardware and software configurations are customized by the vendor to perform a specific task (e.g., name services, firewall services, authentication services, encryption services, and voice-over-IP (VoIP) services), and are not intended to execute user-supplied software.

8) High Performance Computing (HPC) System: TBD.

Note: A definition for High Performance Computing System was suggested by stakeholders during development of the server dataset. EPA has removed the definition text and intends to engage stakeholders regarding the need to define this product type.

9) Direct Current (Dc) Server: A computer server that is designed solely to operate on a dc power source.

B) Product Category: A second-order classification or sub-type within a product type that is based on product features and installed components. Product categories are used in this specification to determine qualification and test requirements.

C) Computer Server Form Factors:

1) Rack-mounted Server: A computer server that is designed for deployment in a standard 19-inch data center rack as defined by EIA-310, IEC 60297, or DIN 41494. For the purposes of this specification, a blade server is considered under a separate category and excluded from the rack-mounted category.

2) Pedestal Server: A self-contained computer server that is designed with PSUs, cooling, I/O devices, and other resources necessary for stand-alone operation. The frame of a pedestal server is similar to that of a tower client computer.

D) Computer Server Components:

1) Power Supply Unit (PSU): A device that converts ac or dc input power to one or more dc power outputs for the purpose of powering a computer server. A computer server PSU must be self-contained and physically separable from the motherboard and must connect to the system via a removable or hard-wired electrical connection.

   a) Ac-Dc Power Supply: A PSU that converts line-voltage ac input power into one or more dc power outputs for the purpose of powering a computer server.

   b) Dc-Dc Power Supply: A PSU that converts line-voltage dc input power to one or more dc outputs for the purpose of powering a computer server. For purposes of this specification, a dc-dc converter (also known as a voltage regulator) that is internal to a computer server and is used to convert a low voltage dc (e.g., 12 V dc) into other dc power outputs for use by computer server components is not considered a dc-dc power supply.

   c) Single-output Power Supply: A PSU that is designed to deliver the majority of its rated

output power to one primary dc output for the purpose of powering a computer server. Single-output PSUs may offer one or more standby outputs that remain active whenever connected to an input power source. For purposes of this specification, the total rated power output from any additional PSU outputs that are not primary and standby outputs shall be no greater than 20 watts. PSUs that offer multiple outputs at the same voltage as the primary output are considered single-output PSUs unless those outputs (1) are generated from separate converters or have separate output rectification stages, or (2) have independent current limits.

d) Multi-output Power Supply: A PSU that is designed to deliver the majority of its rated output power to more than one primary dc output for the purpose of powering a computer server. Multi-output PSUs may offer one or more standby outputs that remain active whenever connected to an input power source. For purposes of this specification, the total rated power output from any additional PSU outputs that are not primary and standby outputs is greater than or equal to 20 watts.

2) I/O Device: A device which provides data input and output capability between a computer server and other devices. An I/O device may be integral to the computer server motherboard or may be connected to the motherboard via though expansion slots (e.g., PCI, PCIe). Examples of I/O devices include discrete Ethernet devices, InfiniBand devices, RAID/SAS controllers, and Fibre Channel devices.

a) I/O Port: Physical circuitry within an I/O device where an independent I/O session can be established. A port is not the same as a connector receptacle; it is possible that a single connector receptacle can service multiple ports of the same interface.

3) Motherboard: The main circuit board of the server. For purposes of this specification, the motherboard includes connectors for attaching additional boards and typically includes the following components: processor, memory, BIOS, and expansion slots.

4) Processor: The logic circuitry that responds to and processes the basic instructions that drive a server. For purposes of this specification, the processor is the central processing unit (CPU) of the computer server. A typical CPU is a physical package to be installed on the server motherboard via a socket or direct solder attachment. The CPU package may include one or more processor cores.

5) Memory: For purposes of this specification, memory is a part of a server external to the processor in which information is stored for immediate use by the processor.

6) Hard Drive (HDD): The primary computer storage device which reads and writes to one or more rotating magnetic disk platters.

7) Solid State Drive (SSD): A disk drive that uses memory chips instead of rotating magnetic platters for data storage.

E) Other Datacenter Equipment:

1) Network Equipment: A device whose primary function is to pass data among various network interfaces, providing data connectivity among connected devices (e.g., routers and switches). Data connectivity is achieved via the routing of data packets encapsulated according to Internet Protocol, Fibre Channel, InfiniBand or similar protocol.

2) Storage Equipment: A system composed of integrated storage controllers, storage devices (e.g., hard drives or solid state storage) and software that provides data storage services to one or more computer servers. While storage equipment may contain one or more embedded processors, these processors do not execute user-supplied software applications but may execute data-specific applications (e.g., data replication, backup utilities, data compression, install agents).

3) Uninterruptible Power Supply (UPS): Combination of convertors, switches, and energy storage devices (such as batteries) constituting a power system for maintaining continuity of load power in case of input power failure.

Note: The definition for UPS is updated to align with the Final Draft ENERGY STAR Specification for Uninterruptible Power Supplies.

F) Operational Modes and Power States:

1) Idle State: The operational state in which the OS and other software have completed loading, the computer server is capable of completing workload transactions, but no active workload transactions are requested or pending by the system (i.e., the computer server is operational, but not performing any useful work). For systems where ACPI standards are applicable, Idle State correlates only to ACPI System Level S0.

2) Active State: The operational state in which the computer server is carrying out work in response to prior or concurrent external requests (e.g., instruction over the network). Active state includes **both** (1) active processing and (2) data seeking/retrieval from memory, cache, or internal/external storage while awaiting further input over the network.

G) Other Key Terms:

1) Controller System: A computer or computer server that manages a benchmark evaluation process. The controller system performs the following functions:

a) start and stop each segment (phase) of the performance benchmark;

b) control the workload demands of the performance benchmark;

c) start and stop data collection from the power analyzer so that power and performance data from each phase can be correlated;

d) store log files containing benchmark power and performance information;

e) convert raw data into a suitable format for benchmark reporting, submission and validation; and

f) collect and store environmental data, if automated for the benchmark.

2) Network Client (Testing): A computer or computer server that generates workload traffic for transmission to a UUT connected via a network switch.

3) RAS Features: An acronym for reliability, availability, and serviceability features. RAS is sometimes expanded to RASM, which adds “Manageability” criteria. The three primary components of RAS as related to a computer server are defined as follows:

a) *Reliability Features*: Features that support a server’s ability to perform its intended function without interruption due to component failures (e.g., component selection, temperature and/or voltage de-rating, error detection and correction).

b) *Availability Features*: Features that support a server’s ability to maximize operation at normal capacity for a given duration of downtime (e.g., redundancy [both at micro- and macro-level]).

c) *Serviceability Features*: Features that support a server’s ability to be serviced without interrupting operation of the server (e.g., hot plugging).

4) Server Processor Utilization: The ratio of processor computing activity to full-load processor computing activity at a specified voltage and frequency, measured instantaneously or with a short term average of use over a set of active and/or idle cycles.

5) Hypervisor: A type of hardware virtualization technique that enables multiple guest operating systems to run on a single host system at the same time.

H) Product Family: A high-level description referring to a group of computers sharing one chassis/motherboard combination that often contains hundreds of possible hardware and software configurations.

1) Common Product Family Attributes: A set of features common to all models/configurations

within a product family that constitute a common basic design. All models/configurations within a product family must share the following:

a) Be from the same model line or machine type;

b) Share the same form factor (i.e., rack-mounted, blade, pedestal);

c) Either share processors from a single defined processor series or share processors that plug into a common socket.

d) share PSUs that perform with efficiencies greater than or equal to the efficiencies at all required load points specified in Section 3.2 (i.e., 10%, 20%, 50%, and 100% of maximum rated load for single-output; 20%, 50%, and 100% of maximum rated load for multi-output).

2) Product Family Tested Product Configurations:

a) Purchase Consideration Variations:

(1) Low-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage (HDD/SDD), and I/O devices that represents the lowest-price or lowest-performance computing platform within the Product Family.

Note: The phrase “lowest-performance” has been added to the Low-end Performance definition.

(2) High-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage (HDD/SDD), and I/O devices that represents either the highest-price or highest-performance computing platform within the Product Family.

b) Typical Configuration:

(1) Typical Configuration: A product configuration that lies between the Minimum and Maximum Power configurations and is representative of a deployed product with high volume sales.

c) Power Utilization Variations:

(1) Minimum Power Configuration: The minimum configuration that is able to boot and execute supported OSs. The Minimum Configuration contains the lowest Processor Socket Power, least number of installed PSUs, Memory, Storage (HDD/SDD), and I/O devices, that is both offered for sale and capable of meeting ENERGY STAR requirements.

(2) Maximum Power Configuration: The vendor-selected combination of components that maximize power usage within the Product Family once assembled and operated. The Maximum Configuration contains the highest Processor Socket Power, greatest number of installed PSUs, Memory, Storage (HDD/SDD), and I/O devices that is both offered for sale and capable of meeting ENERGY STAR requirements.

Note: After examining the results of the data assembly process, EPA is moving forward with the five-point product family approach originally presented in conjunction with dataset development. EPA believes that this structure reduces the product testing burden while providing sufficient information to customers and allows for verification of results.

2 SCOPE

2.1 Included Products

2.1.1 A product must meet the definition of a Computer Server provided in *Section 1* of this document to be eligible for ENERGY STAR qualification under this specification. Eligibility under Version 2.0 is limited to blade-, rack-mounted, or pedestal form factor computer servers with no more than four processor sockets. Products explicitly excluded from Version 2.0 are identified in *Section 2.2*.

2.2 Excluded Products

2.2.1 Products that are covered under other ENERGY STAR product specifications are not eligible for qualification under this specification. The list of specifications currently in effect can be found at www.energystar.gov/products.

2.2.2 The following products are not eligible for qualification under this specification:

i. Fully Fault Tolerant Servers;

ii. Server Appliances;

iii. Storage Equipment including Blade Storage; and

iv. Network Equipment.

3 QUALIFICATION CRITERIA

3.1 Significant Digits and Rounding

3.1.1 All calculations shall be carried out with directly measured (unrounded) values.

3.1.2 Unless otherwise specified, compliance with specification limits shall be evaluated using directly measured or calculated values without any benefit from rounding.

3.1.3 Directly measured or calculated values that are submitted for reporting on the ENERGY STAR website shall be rounded to the nearest significant digit as expressed in the corresponding specification limit.

3.2 Power Supply Requirements

3.2.1 Power supply test data and test reports from testing entities recognized by EPA to perform power supply testing shall be accepted for the purpose of qualifying the ENERGY STAR product.

3.2.2 Power Supply Efficiency Criteria: Power Supplies used in products eligible under this specification must meet the following requirements when tested using the *Generalized Internal Power Supply Efficiency Test Protocol, Rev. 6.6* (available at www.efficientpowersupplies.org).

Note: The test protocol version is updated to the most recent available – Revision 6.6. Power Supply requirements for efficiency and power factor remain consistent with proposals made in Draft 1.

i. Pedestal and Rack-mounted Servers: To qualify for ENERGY STAR, a pedestal or rack-mounted computer server must be configured with **only** PSUs that meet or exceed the applicable efficiency requirements specified in Table 1 **prior to shipment**.

ii. Blade Servers: To qualify for ENERGY STAR, a blade server shipped with a chassis must be configured with **only** PSUs included in the chassis that meet or exceed the applicable efficiency requirements specified in Table 1 **prior to shipment**.

Table 1: Efficiency Requirements for PSUs

Power Supply Type	Rated Output Power	10% Load	20% Load	50% Load	100% Load
Multi-output (Ac-Dc & Dc-Dc)	All Output Levels	N/A	85%	88%	85%
Single-output (Ac-Dc & Dc-Dc)	All Output Levels	80%	88%	92%	88%

3.2.3 Power Supply Power Factor Criteria: Power Supplies used in Computers eligible under this specification must meet the following requirements when tested using the *Generalized Internal Power Supply Efficiency Test Protocol, Rev. 6.5* (available at www.efficientpowersupplies.org).

i. Pedestal and Rack-mounted Servers: To qualify for ENERGY STAR, a pedestal or rack-mounted computer server must be configured with **only** PSUs that meet or exceed the applicable power factor requirements specified in Table 2 **prior to shipment**, under all loading conditions for which output power is greater than or equal to 75 watts. Partners are required to measure and report PSU power factor under loading conditions of less than 75 watts, though no minimum power factor requirements apply.

ii. Blade Servers: To qualify for ENERGY STAR, a blade server shipped with a chassis must be configured with **only** PSUs included in the chassis that meet or exceed the applicable power factor requirements specified in Table 2 **prior to shipment**, under all loading conditions for which output power is greater than or equal to 75 watts. Partners are required to measure and report PSU power factor under loading conditions of less than 75 watts, though no minimum power factor requirements apply.

Table 2: Power Factor Requirements for PSUs

Power Supply Type	Rated Output Power	10% Load	20% Load	50% Load	100 % Load
Dc-Dc (All)	All Output Ratings	N/A	N/A	N/A	N/A
Ac-Dc Multi-output	All Output Ratings	N/A	0.80	0.90	0.95
Ac-Dc Single-output	Output Rating ≤ 500 W	N/A	0.80	0.90	0.95
	Output Rating > 500 W and Output Rating ≤ 1,000 W	0.65	0.80	0.90	0.95
	Output Rating > 1,000 watts	0.80	0.90	0.90	0.95

3.3 Power Management Requirements

3.3.1 Server Processor Power Management: To qualify for ENERGY STAR, a server must offer processor power management that is enabled by default in the BIOS and/or through a management controller or service processor. **All** processors must be able to reduce power consumption in times of low utilization by

i. reducing voltage and/or frequency through Dynamic Voltage and Frequency Scaling (DVFS), or

ii. enabling processor or core reduced power states when a core or socket is not in use.

3.3.2 Supervisor Power Management: To qualify for ENERGY STAR, a product which offers a pre-installed supervisor system (e.g., operating system, hypervisor) must offer supervisor system power management that is enabled by default.

3.3.3 Power Management Disclosure: To qualify for ENERGY STAR, all power management techniques that are enabled by default must be itemized on the Power and Performance Data Sheet.

3.4 Blade System Criteria

3.4.1 Blade Chassis Thermal Management: To qualify for ENERGY STAR, a blade chassis that is (1) shipped with an ENERGY STAR qualified blade server, or (2) marketed for use with an ENERGY STAR qualified blade server, must provide real-time chassis temperature monitoring and fan speed management capability that is enabled by default.

3.4.2 Blade Server Shipping Documentation: To qualify for ENERGY STAR, a blade server that is shipped to a customer independent of a blade chassis must be packaged with documentation to inform the customer that the blade server is ENERGY STAR qualified only if it is installed in a blade chassis meeting requirements in *Section 3.4.1* of this document. A list of qualifying blade chassis and ordering information must also be provided as part of product collateral provided with the blade in either a printed format or an alternative format approved by EPA.

Note: References to blade chassis power requirements have been removed. After discussions with stakeholders and review of blade systems currently on the market, EPA has determined that setting levels applicable to the range of chassis solutions available is not possible due to differences in how vendors choose to design shared resources. While a chassis power disclosure option was also considered to allow for investigation in future versions, EPA believes that such information for different blade systems would not be comparable for the same reason.

3.5 Active State Efficiency Criteria

3.5.1 Active Mode Efficiency Disclosure: To qualify for ENERGY STAR, a computer server or computer server family must be submitted for qualification with the following information disclosed in full and in the context of the complete active mode efficiency rating test report:

i. final rating tool results; and

ii. intermediate rating tool results over the entire test run at **all** of the following load levels: [***TBD***].

Public disclosure and formatting requirements are discussed in *Section* 3.7 of this specification.

3.5.2 Incomplete Disclosure: Partners shall not selectively report individual workload module results, or otherwise presenting efficiency rating tool results in any form other than a complete test report, in customer documentation or marketing materials.

Note: Section 3.5 remains to be determined and will be updated as the SERT development process comes to a close.

3.6 Idle Mode and Full Load Efficiency Criteria – One-Socket (1S) and Two-Socket (2S) Servers

3.6.1 Idle Mode Efficiency: Measured Idle State power (P_{IDLE}) shall be less than or equal to the Maximum Idle State Power Requirement (P_{IDLE_MAX}), as calculated per Equation 1.

Note: The Idle State Power requirement structure above is maintained from Version 1 for single and dual socket systems. Guidance for application of adders and allowances also remain consistent with Version 1.

Equation 1: Calculation of Maximum Idle State Power

$$P_{IDLE\_MAX} = P_{BASE} + \sum_{i=1}^{n} P_{ADDL\_i}$$

Where:

- *P_{IDLE_MAX} is the Maximum Idle State Power Requirement,*
- *P_{BASE} is the base idle power allowance, as determined per Table 3,*
- *P_{ADDL_i} is the Idle State power allowance for additional components, as determined per Table 4*

i. These Idle power limits are applicable to single and dual socket systems only, regardless of the number of processors (e.g., a three or four socket system with only one or two processors installed would not be subject to this requirements).

ii. All quantities in Table 3 and Table 4 refer to the number of components installed in the system, not the maximum number of components the system can support (e.g., installed processors, not processor sockets; installed memory, not supported memory; etc.)

iii. Additional Power Supply allowance: TBD.

iv. For the purposes of determining Idle power allowances, all memory capacities shall be rounded to the nearest GB.

v. The Additional I/O Device allowance may be applied for all I/O Devices over the Base Configuration (i.e., Ethernet devices additional to two ports of 1 Gigabit per second (Gbit/s), onboard Ethernet, plus any non-Ethernet I/O devices), including on-board I/O devices and add-in I/O devices installed through expansion slots.

vi. The Additional I/O Device allowance shall be calculated based upon the rated link speed of a single connection, rounded to the nearest Gbit. I/O devices with less than 1 Gbit speed do not qualify for the Additional I/O Device allowance.

vii. The Additional I/O Device allowance shall only be applied for I/O devices that are active/enabled upon shipment, and are capable of functioning when connected to an active switch.

Table 3: Base Idle State Power Allowances for 1S and 2S Servers

Category	Number of Installed Processors (# P)	Managed Server	Base Idle State Power Allowance, P_{BASE} (watts)
A	1	No	55.0
B	1	Yes	65.0
C	2	No	100.0
D	2	Yes	150.0

Note: EPA thanks stakeholders for the information provided during the data assembly process for this version of the Servers specification. Based on the information received as well as data contained in the ENERGY STAR Servers Qualified Product List, EPA has not revised the Base Idle Levels for 1 and 2 socket servers. The above levels, combined with updated allowances for components in Table 6 below, meet EPA's goal of ensuring a sufficient range of labeled products are available to consumers while maintaining the program's focus on recognizing superior energy efficiency.

Table 4: Additional Idle Power Allowances for Extra Components

System Characteristic	Applies To:	Additional Idle Power Allowance
Additional Power Supplies	Power supplies installed explicitly for power redundancy[iii]	[TBD] watts per Power Supply
Additional Hard Drives (including solid state drives)	Installed hard drives greater than one	8.0 watts per Hard Drive
Additional Memory	Installed memory greater than 4 GB[iv]	0.75 watts per GB[iv]
Additional I/O Devices[v], [vi], [vii]	Installed Devices greater than two ports of 1 Gbit, onboard Ethernet	< 1Gbit: No Allowance = 1 Gbit: 2.0 watts / Active Port > 1 Gbit and < 10 Gbit: 4.0 watts / Active Port ≥ 10 Gbit: 8.0 watts / Active Port

Note: Power Adders for Additional Hard Drives and I/O devices are maintained from Version 1 as current data shows that these levels continue to recognize the more energy efficient products.

The adder for Additional Power Supplies is changed to TBD. EPA believes that a revision to both the magnitude and application of this adder is warranted and invites stakeholder comment on the most appropriate approach to determine this allowance.

The Additional Memory Adder is reduced to 0.75 watts per GB to reflect EPA's review of the current generation of memory technology and opportunities to reduce Idle power through memory selection. EPA received a recommendation from a stakeholder suggesting that a reduction to as low as 0.5 W/GB would be feasible based on the capability of the memory itself. In reviewing the ENERGY STAR dataset, EPA noted a) that the Version 1.0 adder of 2.0 W/GB resulted in numerous systems eligible for adders in excess of consumed Idle Power; and b) that there was variation among systems in the dataset when the scaling of system Idle Power was examined in relation to memory installed. The proposed adder is intended to provide an updated allowance for memory without ignoring the possibility of differences in system overhead suggested in the dataset.

3.6.2 Full Load Data Disclosure: Measured Full Load power (P_{FULL_LOAD}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.7 Idle Mode and Full Load Efficiency Criteria – Three-Socket (3S) and Four-Socket (4S) Servers

3.7.1 Idle Mode Data Disclosure: Measured Idle State power (P_{IDLE}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.7.2 Full Load Data Disclosure: Measured Full Load power (P_{FULL_LOAD}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.8 Idle Mode and Full Load Efficiency Criteria – Blade Servers

3.8.1 Idle Mode Data Disclosure: Measured Idle State power (P_{IDLE}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.8.2 Full Load Data Disclosure: Measured Full Load power (P_{FULL_LOAD}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.9 Other Testing Criteria

3.9.1 GPGPUs: For all Computer Servers sold with GPUs installed for general-purpose tasks (GPGPUs), the following additional criteria apply to meet Idle Mode criteria:

i. For single configurations: All Idle Mode testing shall be conducted both with and without the GPGPU installed. Idle Power measurements taken both with the GPGPU installed and removed shall be submitted to EPA as part of ENERGY STAR qualification materials.

ii. For Product Families: All Idle Mode testing shall be conducted both with and without the GPGPU installed in no less than one of the defined test points in 1.H)2).

iii. Idle Power measurements taken both with the GPGPU installed and removed shall be submitted to EPA as part of ENERGY STAR qualification materials.

iv. Idle Power data with the GPGPU removed shall be used as P_{BASE} for the purposes of qualification of the single configuration or Product Family test point.

Note: 3.9.1 is intended to support acquisition of data showing the power impact of GPGPUs installed in Computer Servers. EPA welcomes stakeholder input on any existing industry-recognized benchmarks that might be referenced as part of this requirement.

A stakeholder raised the topic of add-in compute capability via I/O cards as a similar topic to GPGPUs. EPA is interested in discussing this issue further with stakeholders, particularly on the appropriate scope (GPGPUs, add-in I/O cards, etc.) and measurements to take in support of the possible future inclusion of adders for these components.

4 STANDARD INFORMATION REPORTING REQUIREMENTS

4.1 Power and Performance Datasheet (PPDS)

Note: As part of the ENERGY STAR Uninterruptible Power Supply (UPS) specification development effort, EPA has initiated development of an improved approach to the Power and Performance Datasheet process. Under this approach, data will be maintained in a centralized database instead of individual forms, with a consumer-facing “widget” available for embedding on websites. It is anticipated that such as system would also have the benefit of tying data from qualification directly into the online PPDS, reducing the level of effort for ENERGY STAR Partners.

Information about this approach will be provided in conjunction with further drafts of the Version 2.0 Server Specification.

4.1.1 Data for a standardized Power and Performance Data Sheet (PPDS) shall be submitted to EPA for each ENERGY STAR qualified Computer Server or Computer Server Product Family.

i. Partners are encouraged to provide one set of data for each ENERGY STAR qualified product configuration, though EPA will also accept a data set for each qualified product family.

ii. A product family PPDS must include data for all defined test points in 1.H)2), as applicable.

iii. Whenever possible, Partners must also provide a hyperlink to a more detailed power calculator on their Web site that purchasers can use to understand power and performance data for specific configurations within the product family.

4.1.2 Templates for the Power and Performance Data Sheet can be found on the ENERGY STAR Web site at www.energystar.gov/products.

The PPDS contains the following information:

i. model name and number, identifying SKU and/or configuration ID;

ii. system characteristics (form factor, available sockets/slots, power specifications, etc.);

iii. system configuration(s) (including Low-end Performance Configuration, High-end Performance Configuration, Minimum Power Configuration, Maximum Power Configuration, and Typical Configuration for Product Family qualification);

iv. Data from required Active State Efficiency Criteria testing;

v. power data for Idle and full load, estimated kWh/year, link to power calculator (where available);

vi. additional power and performance data for at least one benchmark as chosen by the Partner from the EPA list of power-performance benchmarks;

vii. available and enabled power saving features (e.g., power management);

viii. power consumption and performance data, along with guaranteed accuracy levels for all power and temperature measurements, disclosure of the time period used for data averaging, and a hyperlink to a detailed power calculator, as available;

ix. a list of selected data from the ASHRAE Thermal Report;

x. for product family qualifications, a list of qualified configurations with qualified SKUs or configuration IDs; and

xi. for a blade server, a list of compatible blade chassis that meet ENERGY STAR qualification criteria.

4.1.3 EPA may periodically revise this PPDS, as necessary, and will notify and invite stakeholder engagement in such a revision process.

5 STANDARD PERFORMANCE DATA MEASUREMENT AND OUTPUT REQUIREMENTS

5.1 Measurement and Output

5.1.1 A computer server must provide data on input power consumption (W), inlet air temperature (°C), and utilization of all logical CPUs. Data must be made available in a published or user-accessible format that is readable by third-party, non-proprietary management software over a standard network. For blade servers and systems, data may be aggregated at the chassis level.

5.1.2 Servers classified as Class B equipment as set out in EN 55022:2006 are exempt from the requirements to provide data on input power consumption and inlet air temperature in 5.1.1. Class B refers to household and home office equipment (intended for use in the domestic environment). All servers in the program must meet the requirement and conditions to report utilization of all logical CPUs.

Note: Provision 5.1.2 is added to exclude "value" products intended for non-datacenter use from the scope of the measurement and output requirements. Stakeholders suggested an alternative approach to instead exclude single processor socket pedestal servers. EPA welcomes comments on whether there are products in the single socket pedestal category that might be managed remotely.

5.2 Reporting Implementation

5.2.1 Products may use either embedded components or add-in devices that are packaged with the computer server to make data available to end users (e.g., a service processor, embedded power or thermal meter (or other out-of-band technology), or pre-installed OS);

5.2.2 Products that include a pre-installed OS must include all necessary drivers and software for end users to access standardized data as specified in this document. Products that do not include a pre-installed OS must be packaged with printed documentation of how to access registers that contain relevant sensor information;

5.2.3 When an open and universally available data collection and reporting standard becomes available, manufacturers should incorporate the universal standard into their systems;

5.2.4 Evaluation of the accuracy (5.3) and sampling (5.4) requirements shall be completed through review of data from component product datasheets. If this data is absent, Partner declaration shall be used to evaluate accuracy and sampling.

Note: 5.2.4 is added to clarify CB responsibilities in evaluating requirements in Section 5.

In conversation, a stakeholder raised the topic of required timestamping of environmental data (power and temperature) as a necessary prerequisite to generation of useful output. EPA seeks stakeholder feedback on this topic and the capability of Computer Servers to meet a possible timestamping requirement for power and temperature.

5.3 Measurement Accuracy

5.3.1 *Input power*: Measurements must be reported with accuracy of at least ±5% of the actual value, with a maximum level of accuracy of ±10W for each installed PSU (i.e., power reporting accuracy for each power supply is never required to be better than ± 10 watts) through the operating range from Idle to full power;

5.3.2 *Processor utilization*: Utilization must be estimated for each logical CPU that is visible to the OS and must be reported to the operator or user of the computer server through the operating environment (OS or hypervisor);

5.3.3 *Inlet air temperature*: Measurements must be reported with an accuracy of at least ±2°C.

5.4 Sampling Requirements

5.4.1 *Input power and processor utilization*: Input power and processor utilization measurements must be collected at a rate of ≥ 1 measurement per contiguous 10 second period. A rolling average, encompassing a period of no more than 30 seconds, must be reported at a frequency of greater than or equal to once per ten seconds.

5.4.2 *Inlet air temperature*: Inlet air temperature measurements must be collected at a rate of ≥ 1 measurement every 10 seconds.

6 TESTING

6.1 Test Methods

6.1.1 When testing Computer Server products, the test methods identified in Table 5 shall be used to determine ENERGY STAR qualification.

Table 5: Test Methods for ENERGY STAR Qualification

Product Type or Component	Test Method
All	ENERGY STAR Test Method for Computer Servers, Rev. TBD

6.2 Number of Units Required for Testing

6.2.1 Representative Models shall be selected for testing per the following requirements:

i. For qualification of an individual product configuration, the unique configuration that is intended to be marketed and labeled as ENERGY STAR is considered the Representative Model.

ii. For qualification of a product family of all product types, one product configuration for each of the five points identified in definitions 1.H)2) within the family are considered Representative Models. All such representative models shall have the same Common Product Family Attributes as defined in 1.H)1).

Note: The section above is added to specify representative testing based on the definitions for Product Family.

6.3 Qualifying Families of Products

6.3.1 Partners are encouraged to test and submit data on individual product configurations for qualification to ENERGY STAR. However, a Partner may qualify multiple product configurations under one Product Family designation if each configuration within the family meets one of the following requirements:

i. Individual products are built on the same platform, are eligible under and meet the same specific requirements in this specification, and are identical in every respect to the tested, representative product configuration except for housing and color; or

ii. Individual products meet the requirements of a product family, as defined in Section H), above. In this case, partners must test and submit data as required in Section 6.2.1ii.

6.3.2 Partners are required to submit a Power and Performance Data Sheet for each product family that is submitted for qualification.

6.3.3 **All** product configurations within a product family that is submitted for qualification must meet ENERGY STAR requirements, including products for which data was not reported.

6.3.4 If a Partner wishes to qualify individual product configurations within a product family that contains non-qualifying products, the Partner must assign a unique identifier to ENERGY STAR qualified product configurations. This identifier must be used consistently in association with qualifying configurations in marketing collateral and on the ENERGY STAR Qualified Product List (e.g., model A1234 for baseline configurations and A1234-ES for ENERGY STAR qualifying configurations).

7 EFFECTIVE DATE

7.1.1 Effective Date: The Version 2.0 ENERGY STAR Computer Server specification shall take effect on the dates specified in Table 6, below. To qualify for ENERGY STAR, a product model shall meet the ENERGY STAR specification in effect on its date of manufacture. The date of manufacture is specific to each unit and is the date (e.g., month and year) on which a unit is considered to be completely assembled.

7.1.2 Future Specification Revisions: EPA reserves the right to change this specification should technological and/or market changes affect its usefulness to consumers, industry, or the environment. In keeping with current policy, revisions to the specification are arrived at through stakeholder discussions. In the event of a specification revision, please note that the ENERGY STAR qualification is not automatically granted for the life of a product model.

595 **Table 6: Specification Effective Dates**

Effective Date	
Version 2.0	TBD

596 # 8 CONSIDERATIONS FOR FUTURE REVISIONS
597

598 ## 8.1 TBD

599

APPENDIX A:
Sample Calculations

Note: This appendix will ultimately include sample calculations for reference in calculating performance levels for products covered in this specification.