ENERGY STAR® Program Requirements
Product Specification for
Computer Servers

Eligibility Criteria
Draft 3, Version 3.0

Following is the Draft 3, Version 3.0 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

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;

C. is targeted to run user-installed applications typically, but not exclusively, enterprise in nature;

D. provides support for error-correcting code (ECC) and/or buffered memory (including both buffered dual in-line memory modules (DIMMs) and buffered on board (BOB) configurations).

E. is packaged and sold with one or more ac-dc or dc-dc power supplies; and

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

2) 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.

(5) *Quarter-height Blade Server*: A blade server requiring one quarter the height of a standard server bay.

(6) *Multi-Node Blade Server*: A blade server which has multiple nodes. The blade server itself is hot swappable, but the individual nodes are not.

B. <u>Blade Chassis</u>: 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. <u>Blade Storage</u>: 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.

3) <u>Fully Fault Tolerant Server</u>: 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.

4) <u>Resilient Server</u>: A computer server designed with extensive Reliability, Availability, Serviceability (RAS) and scalability features integrated in the micro architecture of the system, CPU and chipset. For purposes of ENERGY STAR certification under this specification, a Resilient Server shall have the following characteristics:

A. <u>Processor RAS</u>: The processor must have capabilities to detect, correct, and contain data errors, as described by all of the following:

(1) Error recovery by means of instruction retry for certain processor faults;

(2) Error detection on L1 caches, directories, and address translation buffers using parity protection; and

(3) Single bit error correction (or better) on caches that can contain modified data. Corrected data is delivered to the recipient as part of the request completion.

B. <u>System Recovery & Resiliency</u>: No fewer than six of the following characteristics shall be present in the server:

(1) Error recovery and containment by means of (a) data poison indication (tagging) and propagation which includes mechanism to notify the OS or hypervisor to contain the error, thereby reducing the need for system reboots and (b) containment of address/command errors by preventing possibly contaminated data from being committed to permanent storage;

(2) The processor technology is designed to provide additional capability and functionality without additional chipsets, enabling the design into systems with four or more processor sockets;

(3) Memory Mirroring: A portion of available memory can be proactively partitioned such that a duplicate set may be utilized upon non-correctable memory errors. This can be implemented at the granularity of DIMMs or logical memory blocks;

(4) Memory Sparing: A portion of available memory may be pre-allocated to a spare function such that data may be migrated to the spare upon a perceived impending failure;

(5) Support for making additional resources available without the need for a system restart. This may be achieved either by processor (cores, memory, I/O) on-lining support, or by dynamic allocation/deallocation or processor cores, memory, and I/O to a partition;

(6) Support of redundant I/O devices (storage controllers, networking controllers);

(7) Has I/O adapters or storage devices that are hot-swappable;

(8) Can identify failing processor-to-processor lane(s) and dynamically reduce the width of the link in order to use only non-failing lanes or provide a spare lane for failover without disruption;

(9) Capability to partition the system such that it enables running instances of the OS or hypervisor in separate partitions. Partition isolation is enforced by the platform and/or hypervisor and each partition is capable of independently booting; and

(10) Uses memory buffers for connection of higher speed processor-memory links to DIMMs attached to lower speed DDR channels. Memory buffer can be a separate, standalone buffer chip which is integrated on the system board or integrated on custom-built memory cards.

C. Power Supply RAS: All power supplies installed or shipped with the server shall be redundant and concurrently maintainable. The redundant and repairable components may also be housed within a single physical power supply, but must be repairable without requiring the system to be powered down. Support must be present to operate the system in a degraded mode.

D. Thermal and Cooling RAS: All active cooling components shall be redundant and concurrently maintainable. The processor complex must have mechanisms to allow it to be throttled under thermal emergencies. Support must be present to operate the system in a degraded mode when thermal emergencies are detected in the system components.

> **Note**: EPA received stakeholder feedback recommending changes to simplify the existing resilient server definition to incorporate the necessary Appendix B language from Draft 2 directly into the specification. The resulting definition above merges the essential characteristics for identifying resilient servers, while shortening the existing definition significantly. EPA welcomes stakeholder feedback on the revised definition.

5) 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.

6) Server Appliance: A computer server that is bundled with a pre-installed OS 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.

7) High Performance Computing (HPC) System: A computing system which is designed and optimized to execute highly parallel applications for high performance, deep learning, and artificial intelligence applications. HPC systems feature clustered nodes often featuring high speed inter-processing interconnects as well as high memory capability and bandwidth. HPC systems may be purposely built, or assembled from more commonly available computer servers. HPC systems must meet ALL the following criteria:

A. Marketed and sold as a Computer Server optimized for higher performance computing, augmented or artificial intelligence, and deep learning applications;

B. Designed (or assembled) and optimized to execute highly parallel applications;

C. Consist of multiple computing nodes, clustered primarily to increase computational capability;

D. Includes high speed inter-processing interconnections between nodes.

> **Note**: EPA received stakeholder feedback suggesting minor updates to the HPC definition to address recent advancements in the HPC market and has proposed revised language to address those changes.

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

9) Large Server: A resilient/scalable server which ships as a pre-integrated/pre-tested system housed in one or more full frames or racks and that includes a high connectivity I/O subsystem with a minimum of 32 dedicated I/O slots.

B) 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/Tower 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.

> **Note**: EPA has clarified that pedestal and tower server form factors are equivalent for the purposes of this specification.

C) 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) <u>I/O Device</u>: 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 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. <u>I/O Port</u>: 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) <u>Motherboard</u>: 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) <u>Processor</u>: 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) <u>Memory</u>: 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) <u>Storage Device</u>: A collective term for disk drives (HDDs), solid state drives (SSDs), tapes cartridges, and any other mechanisms providing non-volatile data storage. This definition is specifically intended to exclude aggregating storage elements such as RAID array subsystems, robotic tape libraries, filers, and file servers. Also excluded are storage devices which are not directly accessible by end-user application programs, and are instead employed as a form of internal cache.

D) <u>Other Datacenter Equipment</u>:

1) <u>Network Equipment</u>: 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) <u>Storage Product</u>: A fully-functional storage system that supplies data storage services to clients and devices attached directly or through a network. Components and subsystems that are an integral part of the storage product architecture (e.g., to provide internal communications between controllers and disks) are considered to be part of the storage product. In contrast, components that are normally associated with a storage environment at the data center level (e.g., devices required for operation of an external SAN) are not considered to be part of the storage product. A storage product may be composed of integrated storage controllers, storage devices, embedded network elements, software, and other devices. While storage products 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) <u>Uninterruptible Power Supply (UPS)</u>[1]: 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.

[1] Input power failure occurs when voltage and frequency are outside rated steady-state and transient tolerance bands or when distortion or interruptions are outside the limits specified for the UPS.

E) 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.

F) 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 unit under test (UUT) connected via a network switch.

3) RAS Features: An acronym for reliability, availability, and serviceability features. 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.

6) Auxiliary Processing Accelerators (APAs): An additional compute device installed in the computer server that handles parallelized workloads in place of the CPU.

   A. Expansion APA: An APA that is an add-in card installed in a general-purpose add-in expansion slot (e.g., GPGPUs installed in a PCI slot).

   B. Integrated APA: An APA that is integrated into the motherboard or CPU package.

7) Buffered DDR Channel: Channel or Memory Port connecting a Memory Controller to a defined number of memory devices (e.g., DIMMs) in a computer server. A typical computer

server may contain multiple Memory Controllers, which may in turn support one or more Buffered DDR Channels. As such, each Buffered DDR Channel serves only a fraction of the total addressable memory space in a computer server.

G) 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. Products within a product family may differ in color.

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. Either share the same form factor (i.e., rack-mounted, blade, pedestal) or share the same mechanical and electrical designs with only superficial mechanical differences to enable a design to support multiple form factors;

C. Either share processors from a single defined processor series or share processors that plug into a common socket type. All configurations shipped as ENERGY STAR within the product family shall contain the same number of populated sockets used during testing.

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).

E. Have all memory channels populated with the same model DIMM. In all cases, the minimum memory capacity is the number of memory channels in the server multiplied by the minimum DIMM size offered in the family.

2) Product Family Tested Configurations:

A. Low-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage Devices, and I/O devices that represents the lowest-performance computing platform within the Product Family. This configuration shall include the lowest processer performance per socket, as represented by the lowest numerical value resulting from the multiplication of the core count by the frequency in GHz, offered for sale and capable of meeting ENERGY STAR requirements.[2] It shall also include a memory capacity at least equal to the number of DIMM slots in the server multiplied by the smallest DIMM size offered in the family.

B. High-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage Devices, and I/O devices that represents the highest-performance computing platform within the Product Family. This configuration shall include the highest processer performance per socket, as represented by the highest numerical value resulting from the multiplication of the core count by the frequency in GHz, offered for sale and capable of meeting ENERGY STAR requirements.[2] It shall also include a memory capacity equal to the value found in Equation 1 below:

Equation 1: Minimum Memory Capacity of High-end Performance Configuration

$$Mem\_Capacity\_High \geq 3 \times (\#\, of Sockets \times \#\, of\ Physical\ Cores \times \#\, Threads\ per\ Core)$$

C. Typical Configuration: A product configuration that lies between the Low-end Performance and High-end Performance configurations and is representative of a

[2] Processor performance per socket = [# of processor cores] x [processor clock speed (GHz)], where # of cores represents the number of physical cores and processor clock speed represents the Max TDP core frequency as reported by SERT for a given processor.

deployed product with high volume sales. It shall also include a memory capacity equal to the value found in Equation 2 below:

Equation 2: Minimum Memory Capacity of Typical Configuration

$$Mem\_Capacity\_Typ \geq 2 \times (\#\, of Sockets \times \#\, of\ Physical\ Cores \times \#\, Threads\ per\ Core)$$

> **Note**: EPA received stakeholder feedback requesting more guidance on memory capacity requirements for the three test product family configurations in the product family definition. These changes are intended to increase the consistency of system configurations within each product testing configuration category. EPA welcomes feedback on these proposed additions to the product family definition.
>
> EPA has also clarified that all configurations shipped as ENERGY STAR must contain the same number of populated sockets as used during testing as defined in Section 6.1.2.

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 certification under this specification. Eligibility under Version 3.0 is limited to Blade-, Multi-node, Rack-mounted, or Pedestal form factor computer servers with no more than four processor sockets in the computer server (or per blade or node in the case of blade or multi-node servers). Products explicitly excluded from Version 3.0 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 certification 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 certification under this specification:

i. Computer Servers shipped with Integrated APAs;

ii. Fully Fault Tolerant Servers;

iii. Server Appliances;

iv. High Performance Computing Systems;

v. Large Servers;

vi. Storage Products including Blade Storage; and

vii. Network Equipment.

3 CERTIFICATION 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 certifying 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. 6.7* (available at www.efficientpowersupplies.org). Power Supply data generated using Rev. 6.4.2 (as required in Version 1.1), 6.4.3, 6.5, or 6.6 are acceptable provided the test was conducted prior to the effective date of Version 3.0 of this specification.

i. Pedestal and Rack-mounted Servers: To certify 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 and Multi-node Servers: To certify for ENERGY STAR, a Blade or Multi-node computer server shipped with a chassis must be configured such that **all** PSUs supplying power to the chassis 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)	All Output Levels	N/A	90%	92%	89%
Single-output (Ac-Dc)	All Output Levels	83%	90%	94%	91%

3.2.3 Power Supply Power Factor Criteria: Power Supplies used in Computers Servers eligible under this specification must meet the following requirements when tested using the *Generalized Internal Power Supply Efficiency Test Protocol, Rev. 6.7* (available at www.efficientpowersupplies.org). Power Supply data generated using Rev. 6.4.2 (as required in Version 1.1), 6.4.3, 6.5 or 6.6 are acceptable provided the test was conducted prior to the effective date of Version 3.0.

Note: EPA has updated the latest reference of the Generalized Internal Power Supply Efficiency Test Protocol from Rev. 6.6 to 6.7 in this section to align with Section 3.2.2 above.

i. Pedestal and Rack-mounted Servers: To certify 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 or Multi-node Servers: To certify for ENERGY STAR, a Blade or Multi-node computer server shipped with a chassis must be configured such that **all** PSUs supplying power to the chassis 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
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	0.95	0.95
	Output Rating > 500 W and Output Rating ≤ 1,000 W	0.65	0.80	0.95	0.95
	Output Rating > 1,000 watts	0.80	0.90	0.95	0.95

3.3 Power Management Requirements

3.3.1 Server Processor Power Management: To certify for ENERGY STAR, a Computer Server must offer processor power management that is enabled by default in the BIOS and/or through a management controller, service processor, and/or the operating system shipped with the computer server. **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 certify 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 Reporting: To certify for ENERGY STAR, all power management techniques that are enabled by default must be itemized on the Power and Performance Data Sheet. This requirement applies to power management features in the BIOS, operating system, or any other origin that can be configured by the end-user.

3.4 Blade and Multi-Node System Criteria

3.4.1 Blade and Multi-Node Thermal Management and Monitoring: To certify for ENERGY STAR, a blade or multi-node server must provide real-time chassis or blade/node inlet temperature monitoring and fan speed management capability that is enabled by default.

3.4.2 Blade and Multi-Node Server Shipping Documentation: To certify for ENERGY STAR, a blade or multi-node server that is shipped to a customer independent of the chassis must be accompanied with documentation to inform the customer that the blade or multi-node server is ENERGY STAR qualified only if it is installed in a chassis meeting requirements in *Section 3.4.1* of this document. A list of certified chassis and ordering information must also be provided as part of product collateral provided with the blade or multi-node server. These requirements may be met via either printed materials, electronic documentation provided with the blade or multi-node server, or information publicly available on the Partner's website where information about the blade or multi-node server is found.

3.5 Active State Efficiency Criteria

3.5.1 Active State Efficiency Reporting: To certify for ENERGY STAR, a Computer Server or Computer Server Product Family must be submitted for certification with the following information disclosed in full and in the context of the complete Active State efficiency rating test report:

i. Final SERT rating tool results, which include the results files (in xml, html, and text format) and all results-chart png files; and

ii. Intermediate SERT rating tool results over the entire test run, which include the results-details files (in xml, html, and text format) and all results-details-chart png files.

Data reporting and formatting requirements are discussed in Section 4.1 of this specification.

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

3.5.3 Active State Efficiency Requirements: Calculated Active State efficiency score (Eff_{ACTIVE}) shall be greater than or equal to the minimum Active State efficiency thresholds listed in Table 3.

Equation 3: Calculation Eff_{ACTIVE}

$$Eff_{ACTIVE} = EXP(0.65 * \ln(Eff_{CPU}) + 0.30 * \ln(Eff_{MEMORY}) + 0.05 * \ln(Eff_{STORAGE}))$$

Where:

- Eff_{ACTIVE} *is comprised of* Eff_{CPU}, Eff_{MEMORY} *and* $Eff_{STORAGE}$ *which are defined in equations 4 through 6 below:*

Equation 4: Calculation Eff_{CPU}

$$Eff_{CPU} = Geomean(Eff_{COMPRESS}, Eff_{LU}, Eff_{SOR}, Eff_{CRYPTO}, Eff_{SORT}, Eff_{SHA256}, Eff_{HYBRIDSSJ})$$

Where:

- $Eff_{COMPRESS}$ *is the measured Compression worklet score*
- Eff_{LU} *is the measured LU worklet score*
- Eff_{SOR} *is the measured SOR worklet score*
- Eff_{CRYPTO} *is the measured Crypto worklet score*
- Eff_{SORT} *is the measured Sort worklet score*
- Eff_{SHA256} *is the measured SHA256 worklet score*
- $Eff_{HYBRIDSSJ}$ *is the measured Hybrid SSJ worklet score*

Equation 5: Calculation Eff_{MEMORY}

$$Eff_{MEMORY} = Geomean(Eff_{FLOOD2}, Eff_{CAPACITY2})$$

Where:

- Eff_{FLOOD2} *is the measured Flood2 worklet score*
- $Eff_{CAPACITY2}$ *is the measured Capacity2 worklet score*

Equation 6: Calculation $Eff_{STORAGE}$

$$Eff_{STORAGE} = Geomean(Eff_{SEQUENTIAL}, Eff_{RANDOM})$$

Where:

- $Eff_{SEQUENTIAL}$ *is the measured Sequential worklet score*
- Eff_{RANDOM} *is the measured Random worklet score*

Table 3: Active State Efficiency Thresholds for all Computer Servers

Product Type	Minimum Eff_{ACTIVE}
One Installed Processor	
Rack	11.0
Tower	9.4
Blade or Multi-Node	9.0
Resilient	4.8
Two Installed Processors	
Rack	13.0
Tower	12.0
Blade or Multi-Node	14.0
Resilient	5.2
Greater Than Two Installed Processors	
Rack	16.0
Blade or Multi-Node	9.6
Resilient	4.2

Note: After the release of the Draft 2, Version 3.0 specification, stakeholders provided an updated dataset, which included the most recent products released to the market and updated the test data to account for SERT Version 2.0.0, which incorporates a newer, more modern reference server in the benchmark. This dataset represents the best dataset available for server systems and is a good proxy for the market. In addition, EPA participated in lengthy discussions, including two stakeholder meetings, to discuss if SERT Version 2.0.0 sufficiently incorporates idle test results within the overall active state metric. The result of these discussions and the additional test data is incorporated in this proposal, which focuses on active state efficiency requirements with idle state reporting. The proposed levels are based primarily on model data from 2014 and later, though older available data was included for one socket tower systems and four socket resilient systems since there was not sufficient data to set levels using just 2014+ data in those subcategories. For reference, the cutoff date for new data included in the analysis was November 13, 2017. Although the Agency is aware of newer chipsets coming into the market, the current (post 2014) configurations will continue to be sold over the lifetime of the specification and therefore this data was considered. Due to the removal of standalone idle state efficiency requirements, EPA has increased the stringency of the active state requirements from those proposed in Draft 2 (Note: The change to SERT Version 2.0.0 data scaled the efficiency score values down significantly as compared to the Draft 2 values). For the purposes of its analysis, EPA counted a product as passing if at least two of the three applicable, previously measured configurations met the requirements proposed in Table 3. With the proposed levels, EPA found that pass rates in the dataset were at least 24%. This included models between 2014 and November 2017, with the exception of one socket tower products and greater than two socket resilient servers, where there was not enough data in that range to set levels.

With these proposed levels, EPA is recognizing the top quartile of the market in each category where the data support such differentiation, and striving to come as close to this mark as possible where the data do not segment well.

3.6 Idle State Efficiency Criteria

3.6.1 Idle State Data Reporting: Idle State power (P_{IDLE}, P_{BLADE}, or P_{NODE}) shall be measured and reported, both in certification materials and as required in Section 4 for all computer server types. In addition, for blade and multi-node products, $P_{TOT\_BLADE\_SYS}$ and $P_{TOT\_NODE\_SYS}$ shall also be reported respectively. Please see Section 3.7 for details on how to calculate P_{BLADE} and $P_{TOT\_BLADE\_SYS}$, and Section 3.8 for details on how to calculate P_{NODE} and $P_{TOT\_NODE\_SYS}$.

> **Note**: As noted previously, EPA believes that the active state efficiency metric adequately incorporates idle power behavior and; therefore, there is no longer a need for standalone idle state efficiency criteria. However, EPA will maintain idle state data reporting requirements and has harmonized that requirement across all computer server products. As such, Sections 3.7 and 3.8 (per Draft 2) have been removed.

3.7 Calculating Idle State Values – Blade Servers

3.7.1 The testing of Blade Servers for compliance with Section 3.6.1 shall be carried out under all of the following conditions:

i. Power values shall be measured and reported using a half-populated Blade Chassis. Blade Servers with multiple power domains, choose the number of power domains that is closest to filling half of the Blade Chassis. In a case where there are two choices that are equally close to half, test with the domain or combination of domains which utilize a higher number of Blade Servers. The number of blades tested during the half-populated Blade Chassis test shall be reported.

ii. Power for a fully-populated blade chassis may be optionally measured and reported, provided that half-populated chassis data is also provided.

iii. All Blade Servers installed in the Blade Chassis shall share the same configuration (homogeneous).

iv. Per-blade power values shall be calculated using Equation 7.

Equation 7: Calculation of Single Blade Power

$$P_{BLADE} = \frac{P_{TOT\_BLADE\_SYS}}{N_{INST\_BLADE\_SRV}}$$

Where:

- *P_{BLADE} is the per-Blade Server Power, $P_{TOT\_BLADE\_SYS}$ is total measured power of the Blade System,*
- *$N_{INST\_BLADE\_SRV}$ is the number of installed Blade Servers in the tested Blade Chassis.*

3.8 Calculating Idle State Values – Multi-Node Servers

3.8.1 The testing of Multi-Node Servers for compliance with Section 3.6.1 shall be carried out under all of the following conditions:

i. Power values shall be measured and reported using a fully-populated Multi-Node Chassis.

ii. All Multi-Node Servers in the Multi-Node Chassis shall share the same configuration (homogeneous).

iii. Per-node power values shall be calculated using Equation 8.

Equation 8: Calculation of Single Node Power

$$P_{NODE} = \frac{P_{TOT\_NODE\_SYS}}{N_{INST\_NODE\_SRV}}$$

Where:

- P_{NODE} *is the per-Node Server Power,* $P_{TOT\_NODE\_SYS}$ *is total measured power of the Multi-Node Server,*
- $N_{INST\_NODE\_SRV}$ *is the number of installed Multi-Node Servers in the tested Multi-Node Chassis.*

3.9 Other Testing Criteria

3.9.1 APA Requirements: For all computer servers sold with expansion APAs, the following criteria and provisions apply:

i. For all configurations, Active State and Idle State testing shall be conducted without any APAs installed which may be offered with the product. Where an APA relies on a separate PCIE switch for communication between the APA and CPU, the separate PCIE card(s) or riser(s) shall be removed for Active State and Idle State testing of all configurations.

ii. Manufacturers shall report the model name and model number as well as idle power consumption of each APA offered as an accessory within an ENERGY STAR product family.

Note: EPA received feedback from several stakeholders showing that there are multiple divergent APA technologies developing which require higher idle power consumption to provide significantly greater processing capability compared to the APA technology considered in Version 2.0. An apple to apple comparison of different high memory bandwidth and performance options (e.g. GPGPUs vs. FPGAs) is not feasible at this time.

Due to the lack of a test method available to assess the performance of APAs, which can no longer be ignored for comparison purposes, EPA is proposing that all computer servers be tested without expansion APAs installed during test, and that manufacturers shall report the APA details and idle power consumption as part of the certification. EPA welcomes stakeholder feedback on this proposed approach.

4 STANDARD INFORMATION REPORTING REQUIREMENTS

4.1 Data Reporting Requirements

4.1.1 All required data fields in the ENERGY STAR Version 3.0 Computer Servers Qualified Product Exchange form shall be submitted to EPA for each ENERGY STAR certified Computer Server or Computer Server Product Family.

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

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

iii. Whenever possible, Partners must also provide a hyperlink to a 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 The following data will be displayed on the ENERGY STAR Web site through the product finder tool:

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

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

iii. system type (e.g. resilient.);

iv. system configuration(s) (including Low-end Performance Configuration, High-end Performance Configuration, and Typical Configuration for Product Family certification);

v. power consumption and performance data from required Active and Idle State Efficiency Criteria testing including results.xml, results.html, results.txt, all results-chart png files, results-details.html, results-details.txt, results-details.xml, all results-details-chart png files;

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

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

viii. inlet air temperature measurements made prior to the start of testing, at the conclusion of Idle State testing, and at the conclusion of Active State testing;

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

x. for a blade server, a list of compatible blade chassis that meet ENERGY STAR certification criteria.

4.1.3 EPA may periodically revise this list, 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 average 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 and multi-node servers and systems, data may be aggregated at the chassis level.

5.1.2 Computer 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 computer servers in the program must meet the requirement and conditions to report utilization of all logical CPUs.

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. This requirement may be met via either printed materials, electronic documentation provided with the computer server, or information publically available on the Partner’s website where information about the computer server is found.

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.

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*: Average 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 sampled internally to the computer server at a rate of greater than or equal to measurement per contiguous 10 second period. A rolling average, encompassing a period of no more than 30 seconds, must be sampled internally to the computer server 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 sampled internally to the computer server at a rate of greater than or equal to 1 measurement every 10 seconds.

5.4.3 *Time stamping*: Systems that implement time stamping of environmental data shall sample internally to the computer server data at a rate of greater than or equal to 1 measurement every 30 seconds.

5.4.4 *Management Software*: All sampled measurements shall be made available to external management software either via an on-demand pull method, or via a coordinated push method. In either case the system's management software is responsible for establishing the data delivery time scale while the computer server is responsible to assuring data delivered meets the above sampling and currency requirements.

6 TESTING

6.1 Test Methods

6.1.1 When testing Computer Server products, the test methods identified in 6 shall be used to determine ENERGY STAR certification.

Table 6: Test Methods for ENERGY STAR Certification

Product Type or Component	Test Method
All	ENERGY STAR Test Method for Computer Servers (Rev. April-2018)
All	Standard Performance Evaluation Corporation (SPEC) most current[3] Server Efficiency Rating Tool (SERT)

6.1.2 When testing Computer Server products, UUTs must have all Processor Sockets populated during testing.

[3] For the purposes of this document, the most current SERT version will be listed in the most recently published Servers 3.0 Clarification Memo, located on the Enterprise Servers Specification Version 3.0 website (https://www.energystar.gov/products/spec/enterprise_servers_specification_version_3_0_pd)

i. If a Computer Server cannot support populating all Processor Sockets during testing, then the system must be populated to its maximum functionality. These systems will be subject to the base idle state power allowance based on the number of sockets in the system.

6.2 Number of Units Required for Testing

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

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

Note: EPA has revised the number of points required from four to three configurations to match the revised product family definition in Section 1.G)2).

6.2.2 All product configurations within a product family that is submitted for certification must meet ENERGY STAR requirements, including products for which data is not reported.

7 EFFECTIVE DATE

7.1.1 Effective Date: This ENERGY STAR Computer Servers specification shall take effect on **TBD**. To certify 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 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 certification is not automatically granted for the life of a product model.

8 CONSIDERATIONS FOR FUTURE REVISIONS

8.1 TBD

Note: EPA has removed Appendices A and B found in Draft 2, as the idle state efficiency example is no longer applicable and the supplemental resilient server definition language has been condensed and worked directly into the definition section at the beginning of this document (Section 1).