ENERGY STAR® Program Requirements
Product Specification for
Computer Servers

Eligibility Criteria
Version 2.0
Rev. Oct-2013

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

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) 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.

    (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.

4) <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.

5) <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 qualification under this specification, a Resilient Server shall have the characteristics as described in Appendix B of this specification.

6) <u>Multi-node Server</u>: 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. <u>Dual-node Server</u>: A common multi-node server configuration consisting of two server nodes.

7) <u>Server Appliance</u>: 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.

8) <u>High Performance Computing (HPC) System</u>: A computing system which is designed and optimized to execute highly parallel applications. HPC systems feature a large number of clustered homogeneous nodes often featuring high speed inter-processing interconnects as well as large 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 applications;

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

    C. Consist of a number of typically homogeneous computing nodes, clustered primarily to increase computational capability;

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

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

10) <u>Large Server</u>: 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) <u>Product Category</u>: 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) <u>Computer Server Form Factors</u>:

1) <u>Rack-mounted Server</u>: 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) <u>Pedestal Server</u>: 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) <u>Computer Server Components</u>:

1) <u>Power Supply Unit (PSU)</u>: 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. <u>Ac-Dc Power Supply</u>: 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. <u>Dc-Dc Power Supply</u>: 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. <u>Single-output Power Supply</u>: 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. <u>Multi-output Power Supply</u>: 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>Hard Drive (HDD)</u>: The primary computer storage device which reads and writes to one or more rotating magnetic disk platters.

7) <u>Solid State Drive (SSD)</u>: A storage device that uses memory chips instead of rotating magnetic platters for data storage.

E) <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>: 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.

F) <u>Operational Modes and Power States</u>:

1) <u>Idle State</u>: 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 unit under test (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.

6) Auxiliary Processing Accelerators (APAs): Computing expansion add-in cards installed in general-purpose add-in expansion slots (e.g., GPGPUs installed in a PCI slot).

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.

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. 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.

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 lower-price or lower-performance computing platform within the Product Family.

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

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-, 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 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. High Performance Computing Systems;

    iv. Large Servers;

    v.   Storage Products including Blade Storage; and

    vi. 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). Power Supply data generated using Rev. 6.4.2 (as required in Version 1.1), 6.4.3, or 6.5 are acceptable provided the test was conducted prior to the effective date of Version 2.0 of this specification.

    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 and Multi-node Servers: To qualify 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	85%	88%	85%
Single-output (Ac-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. 6.6* (available at www.efficientpowersupplies.org). Power Supply data generated using Rev. 6.4.2 (as required in Version 1.1), 6.4.3, or 6.5 are acceptable provided the test was conducted prior to the effective date of Version 2.0.

    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 or Multi-node Servers: To qualify 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	0.90	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 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 <u>Supervisor Power Management</u>: 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 <u>Power Management Reporting</u>: To qualify 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 <u>Blade and Multi-Node Thermal Management and Monitoring</u>: To qualify 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 <u>Blade and Multi-Node Server Shipping Documentation</u>: To qualify 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 qualifying 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 publically 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 <u>Active State Efficiency Reporting</u>: To qualify for ENERGY STAR, a Computer Server or Computer Server Product Family must be submitted for qualification 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 (both 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 (both 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 <u>Incomplete Reporting</u>: 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.

3.6 Idle State Efficiency Criteria – One-Socket (1S) and Two-Socket (2S) Servers (neither Blade nor Multi-Node)

3.6.1 <u>Idle State Data Reporting</u>: Idle State power (P_{IDLE}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.6.2 <u>Idle State Efficiency</u>: 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.

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 one and two socket systems only.

  ii. Use Section 6.1 of the ENERGY STAR Computer Servers Test Method to determine the Idle State power for qualification.

  iii. The Resilient category in Table 3 applies only to two socket systems that meet the definition of Resilient Server as set forth in Appendix B.

  iv. All quantities (with the exception of installed processors) 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 memory, not supported memory; etc.)

  v. The Additional Power Supply allowance may be applied for each redundant power supply used in the configuration.

  vi. For the purposes of determining Idle power allowances, all memory capacities shall be rounded to the nearest GB[1]

  vii. 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 greater than or equal to 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. This allowance may be applied for each of the following types of I/O functionality: Ethernet, SAS, SATA, Fibre Channel and Infiniband.

  viii. 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.

  ix. 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	Maximum Possible Number of Installed Processors (# P)	Managed Server	Base Idle State Power Allowance, P_{BASE} (watts)
A	1	No	47.0
B	1	Yes	57.0
C	2	No	92.0
D	2	Yes	142.0
Resilient	2	Yes	205.0

[1] GB defined as 1024^3 or 2^{30} bytes.

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[(v)]	20 watts per Power Supply
Hard Drives (including solid state drives)	Per installed hard drive	8.0 watts per Hard Drive
Additional Memory	Installed memory greater than 4 GB[(vi)]	0.75 watts per GB[(vi)]
Additional Buffered DDR Channel	Installed buffered DDR Channels greater than 8 channels (Resilient Servers only)	4.0 watts per Buffered DDR Channel
Additional I/O Devices[(vii), (viii), (ix)]	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

3.7 Idle State Efficiency Criteria – Three-Socket (3S) and Four-Socket (4S) Servers (neither Blade nor Multi-Node)

3.7.1 <u>Idle State Data Reporting</u>: Idle State power (P_{IDLE}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.8 Idle State Efficiency Criteria – Blade Servers

3.8.1 <u>Idle State Data Reporting</u>: Idle State power ($P_{TOT\_BLADE\_SYS}$) and (P_{BLADE}) shall be measured and reported, both in qualification materials and as required in Section 4.

3.8.2 The testing of Blade Servers for compliance with Section 3.8.1 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 2.

Equation 2: 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.*

3.9 Idle State Efficiency Criteria – Multi-Node Servers

3.9.1 <u>Idle State Data Reporting</u>: Idle State power ($P_{TOT\_NODE\_SYS}$) and (P_{NODE}) shall be measured and reported, both in qualification materials and as required in Section 4, below.

3.9.2 The testing of Multi-Node Servers for compliance with Section 3.9.1 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 3.

Equation 3: 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.*

3.10 Other Testing Criteria

3.10.1 <u>APA Requirements</u>: For all computer servers sold with APAs, the following criteria and provisions apply:

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

   ii. For Product Families: Idle State testing shall be conducted both with and without the APAs installed in the Maximum Power / High-end Performance Configuration found in 1.H)2). Testing with and without the APAs installed may optionally be conducted and disclosed at the other test points.

   iii. Idle State power measurements taken both with the APAs installed and removed shall be submitted to EPA as part of ENERGY STAR qualification materials. These measurements shall be submitted for each individual APA product that is intended for sale with the qualified configuration.

   iv. Measurements of P_{IDLE} in Sections 3.6 and 3.7, P_{BLADE} in Section 3.8 and P_{NODE} in Section 3.9 shall be performed with APAs removed, even if they are installed as-shipped. These measurements shall then be repeated with each APA installed, one at a time, to evaluate Idle State power consumption of each installed APA.

   v. The Idle State power consumption of each installed APA in qualified configurations shall not exceed 46 watts.

   vi. The Idle State power consumption of each individual APA product sold with a qualified configuration shall be reported.

4 STANDARD INFORMATION REPORTING REQUIREMENTS

4.1 Data Reporting Requirements

4.1.1 All required data fields in the ENERGY STAR Version 2.0 Computer Servers Qualified Product Exchange form 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 qualification 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 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 (unmanaged, managed, scalable, etc.);

    iv. 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);

    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, 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 qualifications, 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 qualification 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 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. Apr-2013)
All	Standard Performance Evaluation Corporation (SPEC) most current[2] Server Efficiency Rating Tool (SERT)

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

    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 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).

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.

---

[2] For the purposes of this document, the most current SERT version is listed in the most recently published Servers 2.0 Clarification Memo, located on the Enterprise Servers Specification Version 2.0 website (https://www.energystar.gov/products/specs/enterprise_servers_specification_version_2_0_pd/)

7 EFFECTIVE DATE

7.1.1 <u>Effective Date</u>: The Version 2.0 ENERGY STAR Computer Servers specification shall take effect on December 16, 2013. 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 on which a unit is considered to be completely assembled.

7.1.2 <u>Future Specification Revisions</u>: 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.

8 CONSIDERATIONS FOR FUTURE REVISIONS

8.1 Active State Efficiency Criteria: EPA intends to set active state efficiency criteria in Version 3.0 for all computer server categories in which it has enough SERT data to adequately differentiate products.

8.2 Right Sizing of Power Supplies: EPA will investigate opportunities for encouraging right-sizing of power supplies in Version 3.0.

8.3 Inclusion of Dc-Dc Computer Servers: EPA encourages manufacturers to work with SPEC to develop support for dc servers in the SERT, so that dc computer servers may be considered for qualification in Version 3.0.

8.4 Inclusion of Additional System Architectures: EPA encourages manufacturers to work with SPEC to develop support for architectures that are not currently supported by the SERT, but which represent a sizeable portion of the Computer Servers market. EPA will consider any architecture that is supported by the SERT prior to the development of Version 3.0.

8.5 Removal of Adder for Additional Redundant Power Supplies: EPA is aware of technology that allows redundant power supplies to be kept in standby mode and only activated when needed. EPA encourages the adoption of this technology in computer servers, and will investigate whether the current adder for additional redundant power supplies is still necessary in Version 3.0.

8.6 Auxiliary Processing Accelerator (APA) Requirements: EPA intends to revisit and potentially expand APA requirements in Version 3.0, based on APA data collected from Version 2.0 as well as the potential incorporation of APA evaluation in the SERT.

8.7 Thermal Reporting and Testing Requirements: EPA plans to reevaluate current temperature reporting and testing requirements to maximize the value of the data collected for manufacturers as well as data center operators.

APPENDIX A:
Sample Calculations

Idle State Power Requirements

To determine the Maximum Idle State Power Requirement for ENERGY STAR qualification, determine the base idle state level from Table 3, and then add power allowances from Table 4 (provided in Section 3.6 of this Eligibility Criteria). An example is provided below:

EXAMPLE: A standard single processor Computer Server with 8 GB of memory, two hard drives, and two I/O devices (the first with two 1 Gbit ports and the second with six 1 Gbit ports).

1. Base allowance:
    a. Determine base idle allowance from Table 3, provided for reference below.
    b. The example server is evaluated under Category A and could consume no more than 47.0 watts in Idle to qualify for ENERGY STAR.

Category	Number of Installed Processors (# P)	Managed Server	Base Idle Power Allowance (W)
A	1	No	47.0
B	1	Yes	57.0
C	2	No	92.0
D	2	Yes	142.0
Resilient	2	Yes	205.0

2. Additional Idle Power Allowances: Calculate additional idle allowances for extra components from Table 4, provided for reference below.

System Characteristic	Applies To	Additional Idle Power Allowance
Additional Power Supplies	Power supplies installed explicitly for power redundancy	20.0 watts per Power Supply
Hard Drives (including solid state drives)	All installed hard drives	8.0 watts per Hard Drive
Additional Memory	Installed memory greater than 4 GB	0.75 watts per GB
Additional Buffered DDR Channel	Installed buffered DDR Channels greater than 8 channels (Resilient Servers only)	4.0 watts per Buffered DDR Channel
Additional I/O Devices (single connection speed rounded to nearest Gbit)	Installed Devices greater than two ports of 1 Gbit, onboard Ethernet	< 1 Gbit: 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

a. The example server has two hard drives. It therefore is provided with an additional 16.0 watt allowance for each hard drive (2 HDD x 8.0 watts).
b. The example server has 4 GB in excess of the base configuration. It therefore is provided with an additional 3.0 watt allowance for memory (4 extra GB x 0.75 watts/GB).
c. The example server has one I/O card that does not qualify for an adder: the first device has only two Ethernet ports and does not exceed the two-port threshold. Its second device does qualify for an adder: the server is provided with an additional 12.0 watt allowance for the device (six 1Gbit ports x 2.0 watts/active port).

3. Calculate the final idle allowance by adding the base allowance with the additional power allowances. The example system would be expected to consume no more than 78.0 watts at Idle to qualify (47.0 W + 16.0 W + 3.0 W + 12.0 W).

Additional Idle Allowance - Power Supplies

The following examples illustrate the idle power allowances for additional power supplies:

A. If a Computer Server requires two power supplies to operate, and the configuration includes three installed power supplies, the server would receive an additional 20.0 watt idle power allowance.

B. If the same server were instead shipped with four installed power supplies, it would receive an additional idle power allowance of 40.0 watts.

Additional Idle Allowance - Additional Buffered DDR Channel

The following examples illustrate the idle power allowances for additional buffered DDR channels:

A. If a resilient Computer Server is shipped with six installed buffered DDR channels, the server would not receive an additional idle power allowance.

B. If the same resilient server were instead shipped with 16 installed buffered DDR channels, it would receive an additional idle power allowance of 32.0 watts (first 8 channels = no additional allowance, second 8 channels = 4.0 watts x 8 buffered DDR channels)

APPENDIX B:
IDENTIFYING RESILIENT SERVER CLASS

A. *Processor RAS and Scalability* - All of the following shall be supported:

   (1) *Processor RAS*: The processor must have capabilities to detect, correct, and contain data errors, as described by all of the following:

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

   (b) Single bit error correction (or better) using ECC on caches that can contain modified data. Corrected data is delivered to the recipient (i.e., error correction is not used just for background scrubbing);

   (c) Error recovery and containment by means of (1) processor checkpoint retry and recovery, (2) data poison indication (tagging) and propagation, or (3) both. The mechanisms notify the OS or hypervisor to contain the error within a process or partition, thereby reducing the need for system reboots; and

   (d) (1) Capable of autonomous error mitigation actions within processor hardware, such as disabling of the failing portions of a cache, (2) support for predictive failure analysis by notifying the OS, hypervisor, or service processor of the location and/or root cause of errors, or (3) both.

   (2) The processor technology used in resilient and scalable servers is designed to provide additional capability and functionality without additional chipsets, enabling them to be designed into systems with 4 or more processor sockets. The processors have additional infrastructure to support extra, built-in processor busses to support the demand of larger systems.

   (3) The server provides high bandwidth I/O interfaces for connecting to external I/O expansion devices or remote I/O without reducing the number of processor sockets that can be connected together. These may be proprietary interfaces or standard interfaces such as PCIe. The high performance I/O controller to support these slots may be embedded within the main processor socket or on the system board.

B. *Memory RAS and Scalability* - All of the following capabilities and characteristics shall be present:

   (1) Provides memory fault detection and recovery through Extended ECC;

   (2) In x4 DIMMs, recovery from failure of two adjacent chips in the same rank;

   (3) Memory migration: Failing memory can be proactively de-allocated and data migrated to available memory. This can be implemented at the granularity of DIMMs or logical memory blocks. Alternatively, memory can also be mirrored;

   (4) 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. The use of the buffer chip is required for extended DIMM support; they allow larger memory capacity due to support for larger capacity DIMMs, more DIMM slots per memory channel, and higher memory bandwidth per memory channel than direct-attached DIMMs. The memory modules may also be custom-built, with the memory buffers and DRAM chips integrated on the same card;

   (5) Uses resilient links between processors and memory buffers with mechanisms to recover from transient errors on the link; and

   (6) Lane sparing in the processor-memory links. One or more spare lanes are available

for lane failover in the event of permanent error.

C. *Power Supply RAS*: All PSUs 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 degraded mode when power delivery capability is degraded due to failures in the power supplies or input power loss.

D. *Thermal and Cooling RAS*: All active cooling components, such as fans or water-based cooling, 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 degraded mode when thermal emergencies are detected in system components.

E. *System Resiliency* – no fewer than six of the following characteristics shall be present in the server:

(1) Support of redundant storage controllers or redundant path to external storage;

(2) Redundant service processors;

(3) Redundant dc-dc regulator stages after the power supply outputs;

(4) The server hardware supports runtime processor de-allocation;

(5) I/O adapters or hard drives are hot-swappable;

(6) Provides end to end bus error retry on processor to memory or processor to processor interconnects;

(7) Supports on-line expansion/retraction of hardware resources without the need for operating system reboot ("on-demand" features);

(8) Processor Socket migration: With hypervisor and/or OS assistance, tasks executing on a processor socket can be migrated to another processor socket without the need for the system to be restarted;

(9) Memory patrol or background scrubbing is enabled for proactive detection and correction of errors to reduce the likelihood of uncorrectable errors; and

(10) Internal storage resiliency: Resilient systems have some form of RAID hardware in the base configuration, either through support on the system board or a dedicated slot for a RAID controller card for support of the server's internal drives.

F. *System Scalability* – All of the following shall be present in the server:

(1) Higher memory capacity: >=8 DDR3 or DDR4 DIMM Ports per socket, with resilient links between the processor socket and memory buffers; and

(2) Greater I/O expandability: Larger base I/O infrastructure and support a higher number of I/O slots. Provide at least 32 dedicated PCIe Gen 2 lanes or equivalent I/O bandwidth, with at least one x16 slot or other dedicated interface to support external PCIe, proprietary I/O interface or other industry standard I/O interface.