ENERGY STAR® Program Requirements Product Specification for Computer Servers

Eligibility Criteria Draft 1 Version 3.0

Following is the 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) 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 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).

Note: In reviewing the current data set, EPA has not found any examples of unmanaged computer servers certified as ENERGY STAR. EPA welcomes stakeholder feedback on whether unmanaged servers are still sold outside of the ENERGY STAR program, or whether all current computer servers on the market within scope meet the Managed Server definition, making the definition above obsolete.

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

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

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

5) Resilient Server: A computer server 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 characteristics as described in Appendix B of this specification.

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

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

7) Server Appliance: A computer server that is bundled with a pre-installed 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) High Performance Computing (HPC) System: 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) Direct Current (dc) Server: A computer server that is designed solely to operate on a dc power source.

10) 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.

Note: EPA has removed the "Product Category" definition found in Version 2.1 as it is considered redundant and is not used within the specification.

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 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.

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

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

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

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

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

6) Storage Device: 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.

Note: EPA is proposing to consolidate the previous HDD and SSD definitions from Version 2.1 by harmonizing with the Storage Device definition found in the ENERGY STAR Data Center Storage Version 1.0 specification. This definition captures newer non-volatile technologies not covered by the previous HDD and SSD definitions. EPA welcomes stakeholder feedback on this change.

D) Other Datacenter Equipment:

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

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

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).

Note: EPA has removed the sentence "RAS is sometimes expanded to RASM, which adds manageability criteria" found in Version 2.1 based on the note box above questioning the continued need to address manageability as a unique functionality, rather than a core functionality of all current computer servers within scope.

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.

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.

Note: EPA has clarified that differences in product color are acceptable within a product family, given the removal of the previous section titled "Qualifying Families of Products" section in Version 2.1.

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 Devices, and I/O devices that represents the lower-price or lower-performance computing platform within the Product Family.

(2) High-end Performance Configuration: This configuration shall include the highest processer performance per socket, as represented by the highest numerical value resulting from the multiplication of the processor count by the frequency in GHz, offered for sale and capable of meeting ENERGY STAR requirements.[1]

B. Typical Configuration:

Typical Configuration: A product configuration that lies between the Minimum Power and High-end Performance configurations and is representative of a deployed product with high volume sales.

C. Power Utilization Variations:

[1] Highest 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, not the turbo boost frequency.

(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 Devices, and I/O devices, that is both offered for sale and capable of meeting ENERGY STAR requirements.

Note: EPA received stakeholder feedback that the Agency may be able to meet its goal of creating a product family envelope that has sufficient flexibility to allow manufacturers to sell compliant non-tested configurations, while ensuring that all members of the family meet the ENERGY STAR requirements through fewer test points. Stakeholders indicated that this change would reduce testing costs significantly while maintaining the integrity of the program. EPA is concerned that partners will not be able to represent their product families with three points. Further, EPA's dataset shows meaningful differences between minimum power and low end performance configurations, supporting the maintenance of the current test points. EPA seeks stakeholder feedback on the ability of three configurations per product family to fully represent that family. EPA also seeks data that makes clear the redundancy of the current test points.

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. 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 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%

Note: EPA is proposing power supply efficiency requirements equivalent to 80Plus platinum in Version 3.0. Several stakeholders have shared recently that most server products offer a platinum power supply option, and 63% of the configurations tested in Version 2.0 already used platinum level power supplies for certification purposes. EPA welcomes stakeholder feedback on this proposal.

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. 6.6* (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 3.0.

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	0.90	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 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 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 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.

Note: Per request from several stakeholders, EPA has added the requirement to include the xml files generated by the SERT tool in submissions along with the html and text formats. EPA welcomes stakeholder feedback on whether collecting the xml format eliminates the need to collect the html formatted files.

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: **TBD in next draft.**

Note: EPA is engaged in discussions with many stakeholders, including TGG and SPEC, to determine potential options for Active State efficiency metrics to make use of the SERT data generated in Version 2. EPA will continue to work closely with these groups and plans to release a proposed path forward on Active State metrics in Draft 2. EPA strongly encourages manufacturers and other interested stakeholders to come to consensus on a path forward in time for the release of the next draft, but if this is not accomplished, EPA will develop and propose an approach in Draft 2.

EPA welcomes additional feedback on the following areas regarding Active State metrics and efficiency requirements:

1) Should Active State and Idle State criteria remain separated as is currently proposed in Draft 1, or are there technical merits to combining them into a single metric?

2) What guidance can industry provide end-users to better correlate the SERT worklet scores shown on the ENERGY STAR computer servers qualified product list with customer's real life workloads and applications? EPA would like to work with industry to develop and/or disseminate guidance for purchasers as part of the Version 3.0 process.

3) Are the two storage worklets scores in SERT sufficient to differentiate computer server configurations from storage products? EPA welcomes feedback on how to more clearly handle computer server products which are performing heavy storage duties, particularly in light of the Version 1.1 revision of the ENERGY STAR Data Center Storage specification planned for later this year that is expected to expand the scope to include network attached storage products.

3.6 Idle State Efficiency Criteria – All One-Socket (1S) and Two-Socket (2S) Servers

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. 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 3.8 for details on how to calculate P_{BLADE} and $P_{TOT\_BLADE\_SYS}$, and Section 3.9 for details on how to calculate P_{NODE} and $P_{TOT\_NODE\_SYS}$.

Note: EPA has consolidated Idle State efficiency criteria for **all** one and two socket servers into Section 3.6. Additional information on how to calculate per blade idle power (P_{BLADE}) and per node idle power (P_{NODE} is found in Sections 3.8 and 3.9 respectively.

3.6.2 Idle State Efficiency: Measured Idle State power (P_{IDLE}, P_{BLADE}, or P_{NODE}) 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 or Table 4,*
- *P_{ADDL_i} is the Idle State power allowance for additional components, as determined per 5.*

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 certification.

iii. The Resilient category in Table 3 and Table 4 applies only to 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, Table 4 and, Table 5 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[2]

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.

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

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 all One Socket Servers

Category	Resilient	Base Idle State Power Allowance, P_{BASE} (watts)
A	No	37.0
B	Yes	130

Table 4: Base Idle State Power Allowances for all Two Socket Servers

Category	Blade or Multi-Node	Resilient	Base Idle State Power Allowance, P_{BASE} (watts)
C	No	No	85.0
D	Yes	No	105
E	No	Yes	297

Note: EPA is proposing the above base idle requirements for one and two socket servers. Per stakeholder suggestion, EPA did investigate separating one socket servers by whether they are blade or not, but the data in this category did not support such a differentiation. As a result, blade/multi-node form factor was not used as a differentiating factor for one socket servers. EPA did observe an increase in energy consumption of blades and multi-nodes in the two socket category, and included it as a differentiating function along with resiliency for that category.

Table 5: 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]	10 watts per Power Supply
Storage Devices	Per installed storage device	4.0 watts per Storage Device
Additional Memory	Installed memory greater than 4 GB[vi]	0.25 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

Note: EPA is proposing the following adder value revisions from the Version 2.0 allowances based on recent improvements in technology.

- Aligning in observed trends in power supply efficiency, EPA changed the Additional Power Supplies adder from 20 watts per power supply to 10 watts.

- Recognizing improvements in storage device performance in a range of other products, EPA is proposing changing the Storage Device adder from 8 watts per Storage Device to 4 watts.

- Based on understood state of the art performance gathered during the development of Version 2.0, EPA proposes changing the Additional Memory adder from 0.75 watts/GB above 4GB of installed memory to 0.25 watts/GB above 4GB of installed memory.

EPA intends to set levels such that the revised Idle State energy efficiency requirements above, when combined with the upcoming Active State energy efficiency requirements in Draft 2, will result in recognizing approximately the top quartile of most efficient computer server products on the market. The proposed Idle State requirements above result in an average reduction in P_{IDLE_MAX} of 40 watts for one socket rack servers and 78 watts for two socket rack servers compared to previous Version 2.1 Idle State efficiency requirements.

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

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

Note: EPA does not have enough data to propose meaningful three or four socket idle state efficiency levels at this time. EPA welcomes additional Idle State data on these products to aid in the development of idle requirements for them.

3.8 Calculating Idle State Values – Blade Servers

3.8.1 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 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 Calculating Idle State Values – Multi-Node Servers

3.9.1 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 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 APA Requirements: 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 certification materials.

ii. For Product Families: Idle State testing shall be conducted both with and without the APAs installed in the High-end Performance Configuration found in 1.G)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 certification materials. These measurements shall be submitted for each individual APA product that is intended for sale with the certified 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 30 watts.

Note: EPA has reviewed publically available test data on the latest two generations of high performance consumer grade GPUs (price range $500-1000) and has observed average idle values ranging from 7-15 watts. The same data shows variability with spikes as high as 30 watts and dips below 5 watts while in idle. For this reason, EPA is proposing that the Idle State power consumption requirement for APAs be lowered from 46 watts to the maximum observed energy consumption spikes of the two latest generation high performance cards (30 watts). As proposed, this requirement balances the desire to encourage the use of the latest generations of GPUs with the best performance per watt with the recognition that use of newer APAs can be a more efficient computing approach than general purpose servers for specific workloads.

EPA welcomes stakeholder feedback on if this level is sufficiently aggressive to remove older (typically less energy efficient) generation APAs from consideration without negatively impacting the use of newer APAs.

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 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., managed or resilient.);

iv. system configuration(s) (including Low-end Performance Configuration, High-end Performance Configuration, Minimum Power 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-2016)
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.

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 four 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).

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.

Note: EPA has removed the section previously titled "Qualifying Families of Products" from Version 2.1 as the contents previously found in that section are sufficiently covered in the product family definition and Section 6.2.2 above.

[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)

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

APPENDIX A:
Sample Calculations

Idle State Power Requirements

To determine the Maximum Idle State Power Requirement for ENERGY STAR certification, determine the base idle state level from Table 3 or Table 4, and then add power allowances from Table 5 (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 storage devices, 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 for one socket servers, provided for reference below.
    b. The example server is evaluated under Category A and could consume no more than 37.0 watts in Idle to certify for ENERGY STAR.

Category	Resilient	Base Idle State Power Allowance, P_{BASE} (watts)
A	No	37.0
B	Yes	130

2. Additional Idle Power Allowances: Calculate additional idle allowances for extra components from Table 5, provided for reference below.

System Characteristic	Applies To	Additional Idle Power Allowance
Additional Power Supplies	Power supplies installed explicitly for power redundancy	10.0 watts per Power Supply
Storage Devices	All installed storage devices	4.0 watts per Storage Device
Additional Memory	Installed memory greater than 4 GB	0.25 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 8.0 watt allowance for each hard drive (2 Storage Devices x 4.0 watts).
    b. The example server has 4 GB in excess of the base configuration. It therefore is provided with an additional 1.0 watt allowance for memory (4 extra GB x 0.25 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 58.0 watts at Idle to qualify (37.0 W + 8.0 W + 1.0 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 10.0 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 20.0 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)

Note: EPA has revised the examples in Appendix A to reflect the new proposed idle state requirements in Version 3.0.

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.