ENERGY STAR® Program Requirements
Product Specification for
Computer Servers

Eligibility Criteria
Draft 2, Version 3.0

1 Following is the Draft 2, Version 3.0 ENERGY STAR Product Specification for Computer Servers. A
2 product shall meet all of the identified criteria if it is to earn the ENERGY STAR.

3 # 1 DEFINITIONS

4 A) Product Types:

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

12 A. is marketed and sold as a Computer Server;

13 B. is designed for and listed as supporting one or more computer server operating systems
14 (OS) and/or hypervisors;

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

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

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

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

23 **Note**: EPA received feedback in response to Draft 1 that there are no longer any unmanaged servers on
24 the certified product list or in the market. Therefore, the managed server definition is obsolete and has
25 been removed.

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

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

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

37			a blade chassis.

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

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

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

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

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

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

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

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

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

5) <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.

6) <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.

7) <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

86			applications;
87		B.	Designed (or assembled) and optimized to execute highly parallel applications;
88 89		C.	Consist of a number of typically homogeneous computing nodes, clustered primarily to increase computational capability;
90		D.	Includes high speed inter-processing interconnections between nodes.

91 92   8) <u>Direct Current (dc) Server</u>: A computer server that is designed solely to operate on a dc power source.

93 94 95   9) <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.

96   B) <u>Computer Server Form Factors</u>:

97 98 99 100   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.

101 102 103   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.

104   C) <u>Computer Server Components</u>:

105 106 107 108   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.

109 110   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.

111 112 113 114 115   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.

116 117 118 119 120 121 122 123 124   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.

125 126 127 128 129 130   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.

131 132 133   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

| 134 | | | I/O devices include discrete Ethernet devices, InfiniBand devices, RAID/SAS controllers, and Fibre Channel devices. |
| 135 | | | |

134   I/O devices include discrete Ethernet devices, InfiniBand devices, RAID/SAS controllers, and
135   Fibre Channel devices.

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

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

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

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

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

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

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

160   2) <u>Storage Product</u>: A fully-functional storage system that supplies data storage services to
161   clients and devices attached directly or through a network. Components and subsystems that
162   are an integral part of the storage product architecture (e.g., to provide internal
163   communications between controllers and disks) are considered to be part of the storage
164   product. In contrast, components that are normally associated with a storage environment at
165   the data center level (e.g., devices required for operation of an external SAN) are not
166   considered to be part of the storage product. A storage product may be composed of
167   integrated storage controllers, storage devices, embedded network elements, software, and
168   other devices. While storage products may contain one or more embedded processors, these
169   processors do not execute user-supplied software applications but may execute data-specific
170   applications (e.g., data replication, backup utilities, data compression, install agents).

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

174 E) <u>Operational Modes and Power States</u>:

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

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

184   F)   Other Key Terms:

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

187   A.   start and stop each segment (phase) of the performance benchmark;

188   B.   control the workload demands of the performance benchmark;

189   C.   start and stop data collection from the power analyzer so that power and performance
190   data from each phase can be correlated;

191   D.   store log files containing benchmark power and performance information;

192   E.   convert raw data into a suitable format for benchmark reporting, submission and
193   validation; and

194   F.   collect and store environmental data, if automated for the benchmark.

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

197   3)   RAS Features: An acronym for reliability, availability, and serviceability features. The three
198   primary components of RAS as related to a computer server are defined as follows:

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

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

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

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

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

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

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

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

217   **Note**: EPA has revised the APA definition and added two new sub-definitions to better address recent
218   advancements, such as FPGA technology, in server designs where the APA is directly attached to the
219   motherboard or integrated into the CPU package.

220   7)   Buffered DDR Channel: Channel or Memory Port connecting a Memory Controller to a
221   defined number of memory devices (e.g., DIMMs) in a computer server. A typical computer
222   server may contain multiple Memory Controllers, which may in turn support one or more
223   Buffered DDR Channels. As such, each Buffered DDR Channel serves only a fraction of the
224   total addressable memory space in a computer server.

225   G)   Product Family: A high-level description referring to a group of computers sharing one
226   chassis/motherboard combination that often contains hundreds of possible hardware and software
227   configurations. Products within a product family may differ in color.

	1)	Common Product Family Attributes: A set of features common to all models/configurations within a product family that constitute a common basic design. All models/configurations within a product family must share the following:
228 229 230		
231		A. Be from the same model line or machine type;
232 233 234		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;
235 236		C. Either share processors from a single defined processor series or share processors that plug into a common socket type.
237 238 239 240		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).
241	2)	Product Family Tested Product Configurations:
242 243 244 245 246 247		A. Low-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage Devices, and I/O devices that represents the lowest-performance computing platform within the Product Family. This configuration shall include the lowest processer performance per socket, as represented by the lowest numerical value resulting from the multiplication of the core count by the frequency in GHz, offered for sale and capable of meeting ENERGY STAR requirements.[1]
248 249 250 251 252 253		B. High-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage Devices, and I/O devices that represents the highest-performance computing platform within the Product Family. This configuration shall include the highest processer performance per socket, as represented by the highest numerical value resulting from the multiplication of the core count by the frequency in GHz, offered for sale and capable of meeting ENERGY STAR requirements.[1]
254 255 256		C. Typical Configuration: A product configuration that lies between the Low-end Performance and High-end Performance configurations and is representative of a deployed product with high volume sales.

257 **Note**: EPA is proposing to simplify the product family to three required test configurations that include
258 low-end and high-end performance configurations, as well as the typical configuration. EPA believes that
259 manufacturers will be able to sufficiently capture configurations previously addressed by the minimum
260 power configuration through the modifications to the low-end performance configuration definition which
261 now specifies the lowest performance configuration in the family. This proposed change will simplify
262 testing and the reporting of the product family on the certified product list. EPA welcomes stakeholder
263 feedback on this proposal.

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

2 SCOPE

2.1 Included Products

2.1.1 A product must meet the definition of a Computer Server provided in *Section 1* of this document to be eligible for ENERGY STAR certification under this specification. Eligibility under Version 3.0 is limited to Blade-, Multi-node, Rack-mounted, or Pedestal form factor computer servers with no more than four processor sockets in the computer server (or per blade or node in the case of blade or multi-node servers). Products explicitly excluded from Version 3.0 are identified in *Section 2.2*.

2.2 Excluded Products

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

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

   i. Computer Servers shipped with Integrated APAs;

   ii. Fully Fault Tolerant Servers;

   iii. Server Appliances;

   iv. High Performance Computing Systems;

   v. Large Servers;

   vi. Storage Products including Blade Storage; and

   vii. Network Equipment.

Note: EPA has, to date, not received enough data on integrated APAs to determine if they comprise a meaningful portion of the market, nor the magnitude of adjustment needed for both idle state and active state energy requirements to appropriately address products shipped with them. As a result, EPA is proposing to exclude computer servers shipped with integrated APAs from scope in Version 3.0. EPA welcomes additional data from stakeholders that would allow these products to be covered in the next draft of this specification.

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.

301 3.2.2 <u>Power Supply Efficiency Criteria</u>: Power Supplies used in products eligible under this specification
302   must meet the following requirements when tested using the *Generalized Internal Power Supply*
303   *Efficiency Test Protocol, Rev. 6.7* (available at www.efficientpowersupplies.org). Power Supply
304   data generated using Rev. 6.4.2 (as required in Version 1.1), 6.4.3, 6.5, or 6.6 are acceptable
305   provided the test was conducted prior to the effective date of Version 3.0 of this specification.

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

309   ii. <u>Blade and Multi-node Servers</u>: To certify for ENERGY STAR, a Blade or Multi-node computer
310   server shipped with a chassis must be configured such that **all** PSUs supplying power to the
311   chassis meet or exceed the applicable efficiency requirements specified in Table 1 **prior to**
312   **shipment**.

313 **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%

314 3.2.3 <u>Power Supply Power Factor Criteria</u>: Power Supplies used in Computers Servers eligible under
315   this specification must meet the following requirements when tested using the *Generalized*
316   *Internal Power Supply Efficiency Test Protocol, Rev. 6.6* (available at
317   www.efficientpowersupplies.org). Power Supply data generated using Rev. 6.4.2 (as required in
318   Version 1.1), 6.4.3, or 6.5 are acceptable provided the test was conducted prior to the effective
319   date of Version 3.0.

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

326   ii. <u>Blade or Multi-node Servers</u>: To certify for ENERGY STAR, a Blade or Multi-node computer
327   server shipped with a chassis must be configured such that **all** PSUs supplying power to the
328   chassis meet or exceed the applicable power factor requirements specified in Table 2 **prior**
329   **to shipment**, under all loading conditions for which output power is greater than or equal to
330   75 watts. Partners are required to measure and report PSU power factor under loading
331   conditions of less than 75 watts, though no minimum power factor requirements apply.

332

Table 2: Power Factor Requirements for PSUs

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

333

Note: In Draft 1, EPA proposed updating the efficiency requirements to 80Plus platinum equivalent levels, which stakeholders supported. EPA also received stakeholder feedback requesting tighter power factor requirements to align with 80Plus platinum equivalent power factor requirements. EPA has harmonized with 80Plus platinum equivalent power factor requirements at the 50% load, which is the only requirement where ENERGY STAR was out of alignment.

3.3 Power Management Requirements

3.3.1 <u>Server Processor Power Management</u>: 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 <u>Supervisor Power Management</u>: 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 <u>Power Management Reporting</u>: 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 <u>Blade and Multi-Node Thermal Management and Monitoring</u>: 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.

359 3.4.2 <u>Blade and Multi-Node Server Shipping Documentation</u>: To certify for ENERGY STAR, a blade or
360 multi-node server that is shipped to a customer independent of the chassis must be accompanied
361 with documentation to inform the customer that the blade or multi-node server is ENERGY STAR
362 qualified only if it is installed in a chassis meeting requirements in *Section 3.4.1* of this document.
363 A list of certified chassis and ordering information must also be provided as part of product
364 collateral provided with the blade or multi-node server. These requirements may be met via either
365 printed materials, electronic documentation provided with the blade or multi-node server, or
366 information publically available on the Partner's website where information about the blade or
367 multi-node server is found.

3.5 Active State Efficiency Criteria

369 3.5.1 <u>Active State Efficiency Reporting</u>: To certify for ENERGY STAR, a Computer Server or Computer
370 Server Product Family must be submitted for certification with the following information disclosed
371 in full and in the context of the complete Active State efficiency rating test report:

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

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

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

377 3.5.2 <u>Incomplete Reporting</u>: Partners shall not selectively report individual workload module results, or
378 otherwise present efficiency rating tool results in any form other than a complete test report, in
379 customer documentation or marketing materials.

380 3.5.3 <u>Active State Efficiency Requirements</u>: Calculated Active State efficiency score (Eff_{ACTIVE}) shall be
381 greater than or equal to the minimum Active State efficiency thresholds listed in Table 3.

382 **Equation 1: Calculation Eff_{ACTIVE}**

383 $$Eff_{ACTIVE} = EXP(0.65 * \ln(Eff_{CPU}) + 0.30 * \ln(Eff_{MEMORY}) + 0.05 * \ln(Eff_{STORAGE}))$$

384
385 *Where:*
386 ▪ *Eff_{ACTIVE} is comprised of Eff_{CPU}, Eff_{MEMORY} and $Eff_{STORAGE}$*
387 *which are defined in equations 2 through 4 below:*

388 **Equation 2: Calculation Eff_{CPU}**

389 $$Eff_{CPU} = Geomean(Eff_{COMPRESS}, Eff_{LU}, Eff_{SOR}, Eff_{CRYPTO}, Eff_{SORT}, Eff_{SHA256}, Eff_{HYBRIDSSJ})$$

390 *Where:*
391 ▪ *$Eff_{COMPRESS}$ is the measured Compression worklet score*
392 ▪ *Eff_{LU} is the measured LU worklet score*
393 ▪ *Eff_{SOR} is the measured SOR worklet score*
394 ▪ *Eff_{CRYPTO} is the measured Crypto worklet score*
395 ▪ *Eff_{SORT} is the measured Sort worklet score*
396 ▪ *Eff_{SHA256} is the measured SHA256 worklet score*
397 ▪ *$Eff_{HYBRIDSSJ}$ is the measured Hybrid SSJ worklet score*
398

399 **Equation 3: Calculation Eff_{MEMORY}**

400 $$Eff_{MEMORY} = Geomean(Eff_{FLOOD2}, Eff_{CAPACITY2})$$

401 *Where:*
402 ▪ *Eff_{FLOOD2} is the measured Flood2 worklet score*
403 ▪ *$Eff_{CAPACITY2}$ is the measured Capacity2 worklet score*
404

| 405 | **Equation 4: Calculation $Eff_{STORAGE}$** |

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

407	*Where:*
408	▪ $Eff_{SEQUENTIAL}$ *is the measured Sequential worklet score*
409	▪ Eff_{RANDOM} *is the measured Random worklet score*
410	
411	

Table 3: Active State Efficiency Thresholds for all Computer Servers

Product Type	Minimum Eff_{ACTIVE}
Rack	50
Tower	50
Blade	55
Multi-Node	55
Resilient	30

Note: The proposed active state efficiency requirements above are a culmination of multiple years of SERT data collection, which included over 600 configurations and over 120 unique product families, as well as consistent collaboration with stakeholders to determine the best way to vet, interpret, and use this data to set appropriate active state requirements across a variety of server types. After performing analysis on the active state efficiency data, EPA is proposing the minimum active state efficiency thresholds presented in Table 3 above. The process of calculating Eff_{ACTIVE} aligns with the metric development work presented to EPA through a collaborative effort between SPEC and The Green Grid.

EPA has balanced the proposed active state and idle requirements so that together the overall energy requirements target the top quartile of the market and represent products that excel in both areas. The overall proposal targets a wide variety of products both by product type and socket count. EPA has received stakeholder feedback that retaining idle state requirements negatively impacts products that have top tier active state results. EPA's review of the available product data does not support this claim, but rather shows that a large majority of the best performing products in active state can also meet their requisite idle state requirements.

EPA welcomes stakeholder feedback on both the proposed metric, as well as the minimum state efficiency thresholds presented in Table 3, including any additional new model product data that may influence either the idle or active state proposed requirements in Version 3.0.

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

3.6.1   <u>Idle State Data Reporting</u>: 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 3.9 for details on how to calculate P_{BLADE} and $P_{TOT\_BLADE\_SYS}$, and Section 3.10 for details on how to calculate P_{NODE} and $P_{TOT\_NODE\_SYS}$.

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

| 441 | | **Equation 5: 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. 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.)

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

v. For the purposes of determining Idle power allowances, all memory capacities shall be rounded to the nearest GB[2]

vi. 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.

vii. 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.

viii. 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 4: Base Idle State Power Allowances for One and Two Socket Non-Resilient Servers

| Category	Number of Sockets	Blade or Multi-Node	Base Idle State Power Allowance, P_{BASE} (watts)
A	1	N/A[3]	37
B	2	No	85
C	2	Yes	105

[2] GB defined as 1024^3 or 2^{30} bytes.
[3] The blade or multi-node differentiation is only applicable for two socket systems. Category A applies to all one socket computer servers.

473 **Note**: Recognizing the changing nature of this segment of the market, the small number of resilient server
474 models in the ENERGY STAR dataset, and lastly the limited differentiation in the idle state performance
475 of these products, EPA is not able to confidently develop idle levels for this product type at this time.
476 EPA is maintaining a reporting requirement for idle state for these products in order to maintain access to
477 this for those interested customers. Further, resilient servers are subject to the new active state power
478 requirements.

479 Due to the removal of the resilient server requirements in the idle state power section, EPA has
480 consolidated the previous Tables 3 and 4 from Draft 1 into a single table covering the socket count and
481 blade / multi-node capability. A figure showing the consolidation of categories and removal of previous
482 resilient categories in Draft 1 is shown below:

483
484

Draft 1   Draft 2

485   A ──────────→ A
486   Ⓑ   ╲→ B
487   C ╲   ╱
488   ╲   ╱→ C
489   D   ╱
490   Ⓔ

491
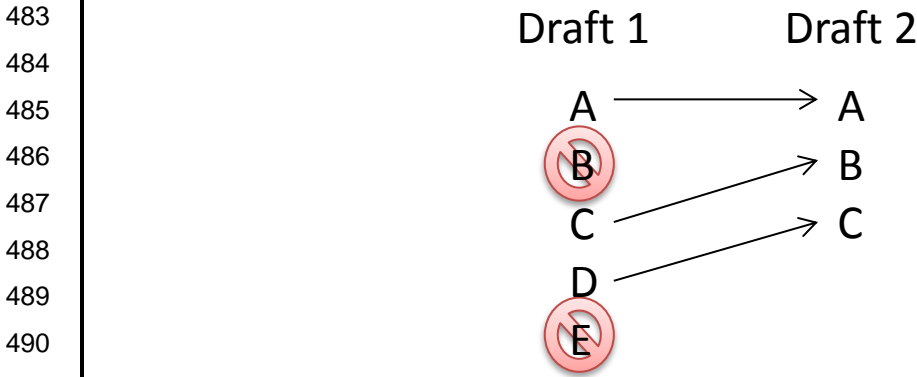
492   **Table 5: Additional Idle Power Allowances for Extra Components**

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)]	0.125 watts per GB[(vi)]
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

493 **Note**: EPA is proposing to amend the following adder values from the Draft 1 specification allowances
494 based on feedback from stakeholders and analysis of additional product data made available since the
495 release of the Draft 1 specification.

496 - EPA has removed the Additional Buffered DDR Channel adder since it only applies to resilient servers,
497 which no longer have idle requirements.

498 - After further discussions with stakeholders, along with analyzing an expanded product data set since the
499 development of Draft 1, EPA has determined that the 0.25 watts / GB memory adder was not sufficiently
500 aggressive. EPA is proposing to modify this adder to 0.125 watts / GB above 4GB of installed memory.
501 These requirements along with the new active state requirements highlight the products which perform
502 the best in both idle and active state efficiency earning the ENERGY STAR label.

503 EPA also received feedback regarding the following stakeholder proposed adder changes which require
504 additional data in order to address. They are:

505 - A definition and adder for newer non-volatile memory devices, including those that are flash based. EPA
506 is open to addressing this new type of memory technology, but is unsure whether it should be addressed
507 independently of the existing memory definition and adder which are focused on traditional random
508 access memory. EPA welcomes additional data from stakeholders to determine the best way to address
509 these new memory types including if they merit additional adder allowances, and what value those
510 allowances should be.

511 - Certain stakeholders contend that the current 8 watt per active port adder for Additional I/O Devices
512 greater than 10 Gbit is not sufficient for newer network port technologies. EPA welcomes energy data to
513 support these claims, as well as inform what adder levels would be appropriate for these higher speed
514 ports (e.g. 40Gbit+).

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

517 3.7.1 <u>Idle State Data Reporting</u>: Idle State power (P_{IDLE}) shall be measured and
518 reported, both in certification materials and as required in Section 4.

3.8 Idle State Efficiency Criteria – All Resilient Servers

520 3.8.1 <u>Idle State Data Reporting</u>: Idle State power (P_{IDLE}) shall be measured and
521 reported, both in certification materials and as required in Section 4.

522 **Note**: Although EPA has removed the idle state efficiency requirements for resilient servers, EPA is
523 proposing to maintain a reporting requirement for these products. This information will be posted on the
524 ENERGY STAR website as an additional resource for interested customers.

3.9 Calculating Idle State Values – Blade Servers

526 3.9.1 The testing of Blade Servers for compliance with Section 3.6.1 shall be carried out under all of the
527 following conditions:

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

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

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

538   iv. Per-blade power values shall be calculated using Equation 6.

Equation 6: 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,*

545   ▪ $N_{INST\_BLADE\_SRV}$ *is the number of installed Blade Servers in*
546   *the tested Blade Chassis.*
547

3.10 Calculating Idle State Values – Multi-Node Servers

3.10.1   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 7.

Equation 7: 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.*

3.11 Other Testing Criteria

3.11.1   <u>APA Requirements</u>: For all computer servers sold with expansion 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, 3.7, and 3.8, P_{BLADE} in Section 3.9 and P_{NODE} in Section 3.10 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 clarified that the APA requirements only apply to expansion APAs, as products with integrated APAs are excluded from scope.

4 STANDARD INFORMATION REPORTING REQUIREMENTS

4.1 Data Reporting Requirements

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

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

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

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

4.1.2 The following data will be displayed on the ENERGY STAR Web site through the product finder tool:

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

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

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

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

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

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

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

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

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

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

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

5 STANDARD PERFORMANCE DATA MEASUREMENT AND OUTPUT REQUIREMENTS

5.1 Measurement and Output

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

623 5.1.2 Computer servers classified as Class B equipment as set out in EN 55022:2006 are exempt from
624 the requirements to provide data on input power consumption and inlet air temperature in 5.1.1.
625 Class B refers to household and home office equipment (intended for use in the domestic
626 environment). All computer servers in the program must meet the requirement and conditions to
627 report utilization of all logical CPUs.

628 **5.2   Reporting Implementation**

629 5.2.1 Products may use either embedded components or add-in devices that are packaged with the
630 computer server to make data available to end users (e.g., a service processor, embedded power
631 or thermal meter (or other out-of-band technology), or pre-installed OS);

632 5.2.2 Products that include a pre-installed OS must include all necessary drivers and software for end
633 users to access standardized data as specified in this document. Products that do not include a
634 pre-installed OS must be packaged with printed documentation of how to access registers that
635 contain relevant sensor information. This requirement may be met via either printed materials,
636 electronic documentation provided with the computer server, or information publically available on
637 the Partner's website where information about the computer server is found.

638 5.2.3 When an open and universally available data collection and reporting standard becomes
639 available, manufacturers should incorporate the universal standard into their systems;

640 5.2.4 Evaluation of the accuracy (5.3) and sampling (5.4) requirements shall be completed through
641 review of data from component product datasheets. If this data is absent, Partner declaration
642 shall be used to evaluate accuracy and sampling.

643 **5.3   Measurement Accuracy**

644 5.3.1 *Input power*: Measurements must be reported with accuracy of at least ±5% of the actual value,
645 with a maximum level of accuracy of ±10W for each installed PSU (i.e., power reporting accuracy
646 for each power supply is never required to be better than ± 10 watts) through the operating range
647 from Idle to full power;

648 5.3.2 *Processor utilization*: Average utilization must be estimated for each logical CPU that is visible to
649 the OS and must be reported to the operator or user of the computer server through the operating
650 environment (OS or hypervisor);

651 5.3.3 *Inlet air temperature*: Measurements must be reported with an accuracy of at least ±2°C.

652 **5.4   Sampling Requirements**

653 5.4.1 *Input power and processor utilization*: Input power and processor utilization measurements must
654 be sampled internally to the computer server at a rate of greater than or equal to measurement
655 per contiguous 10 second period. A rolling average, encompassing a period of no more than 30
656 seconds, must be sampled internally to the computer server at a frequency of greater than or
657 equal to once per ten seconds.

658 5.4.2 *Inlet air temperature*: Inlet air temperature measurements must be sampled internally to the
659 computer server at a rate of greater than or equal to 1 measurement every 10 seconds.

660 5.4.3 *Time stamping*: Systems that implement time stamping of environmental data shall sample
661 internally to the computer server data at a rate of greater than or equal to 1 measurement every
662 30 seconds.

663 5.4.4 *Management Software*: All sampled measurements shall be made available to external
664 management software either via an on-demand pull method, or via a coordinated push
665 method. In either case the system's management software is responsible for establishing the
666 data delivery time scale while the computer server is responsible to assuring data delivered
667 meets the above sampling and currency requirements.

6 TESTING

6.1 Test Methods

668

669

670 6.1.1   When testing Computer Server products, the test methods identified in 6 shall be used to
671   determine ENERGY STAR certification.

672   **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[4] Server Efficiency Rating Tool (SERT)

673 6.1.2   When testing Computer Server products, UUTs must have all Processor Sockets populated
674   during testing.

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

6.2 Number of Units Required for Testing

679 6.2.1   Representative Models shall be selected for testing per the following requirements:

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

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

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

689 6.2.2   All product configurations within a product family that is submitted for certification must meet
690   ENERGY STAR requirements, including products for which data is not reported.

7 EFFECTIVE DATE

691

692 7.1.1   <u>Effective Date</u>: This ENERGY STAR Computer Servers specification shall take effect on **TBD**. To
693   certify for ENERGY STAR, a product model shall meet the ENERGY STAR specification in effect
694   on its date of manufacture. The date of manufacture is specific to each unit and is the date on
695   which a unit is considered to be completely assembled.

---

[4] 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)

696 7.1.2 <u>Future Specification Revisions</u>: EPA reserves the right to change this specification should
697   technological and/or market changes affect its usefulness to consumers, industry, or the
698   environment. In keeping with current policy, revisions to the specification are arrived at through
699   stakeholder discussions. In the event of a specification revision, please note that the ENERGY
700   STAR certification is not automatically granted for the life of a product model.

701 # 8 CONSIDERATIONS FOR FUTURE REVISIONS

702 **8.1 TBD**

703

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

| 805 | | | 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.