Ref. #	Commenter	Topic	Subtopic	Stakeholder Comment	EPA Response
1	Summary	Definitions	Computer Server	A stakeholder noted that it appears the items underneath "Product Type" represent ten different product types and "Computer Server" is one of them. This could be misleading and the distinction is important in the application of the definitions. The commenter recommended that "Computer Server" should replace "Product Type" and the different types of servers be represented in sub-bullets underneath. This stakeholder supported the removal of the error-correcting code (EEC) exemption for systems larger than 50 nodes which share the same chassis.	The current approach, with the Computer Server definition on top, mirrors the approach taken in other ENERGY STAR specifications such as Computers Version 5.2 and UPS Version 1.0.
2	Summary	Definitions	APAs	Several stakeholders requested that Auxiliary Processing Accelerator (APAs) levels not be included in this specification because these are relatively new, unfamiliar products and the technology changes rapidly. They stated that if APAs are included, definitions for APA and GPGDUs need to be refined because there is a significant group of components that can be plugged into server PCI slots to provide additional function. They suggested that sub-categories be defined and offered several examples of APA cards. A commenter recommended that APA cards should not be included in configurations (especially High End and Maximum Cost Configurations) tested for SERT metrics because SERT is not designed to exercise APA cards and the resulting performance and power results could be distorted. They expressed willingness to work with EPA to develop an APA definition and suggested a 46 W limit per card. However, another commenter stated that a 125 W limit would be appropriate because 46 W is too restrictive and could favor lower performance APA cards. A stakeholder questioned how Blade and Multi-node servers that contain APAs/GPUs would be tested. Another stakeholder inquired if the video card used for display output and the card for encryption would be considered an APA.	EPA proposes to change the 46 watt total limit for Auxiliary Processing Accelerator (APA) cards to 46 watts per card. For qualification purposes, all APA cards that can be sold with a qualified product should be tested to provide the consumer with the most data. The results of these tests will be published on the PPDS. EPA has moved the APA defintion to the definitions section of the specification as recommended. However, the definition will remain open and will not contain references to specific APA cards in order to avoid excluding any products/technologies. Blades and Multi-node servers shall follow the same guidelines as all other servers for APA testing. An additional APA test should be conducted on the maximum configurations of the 5 corners, so a sixth test run is necessary. This specification will focus on cards that are exclusively deployed to provide additional computing capability under the supervision of the main CPUs. Cards that have processing capability but that are primarily used for other functions will not be considered APAs (E.g. Video cards, RAID cards, TOE cards, etc.)
3	Summary	Definitions	Computer Server Form Factors	One stakeholder indicated that Blade server is not included in the Rack-mounted category. They recommended that the form factor for Blade server and Multi-node server be specified and should include form factors used by the new class of micro servers (or "scale-out") servers.	Blade systems are often mounted in a rack, but the definition of Rack Server was intended to indicate the type of server rather than to indicate how it is physically mounted. Blade chassis examples have shown that they can be mounted in a rack system or a stand alone chassis.
4	Summary	Definitions	Blade Server	A stakeholder recommended that a fifth type of blade server be added because at least one manufacturer has this product and others are expected: (5) Multi-node Blade Server- A blade server which has multiple nodes. The blade server will be hot swappable, the individual nodes will not.	EPA appreciates stakeholder feedback but will maintain the current approach for Multi-node servers as it does not provide value to the specification to change this definition. It is intended to be general in order to avoid unnecessary constraints on technology.

5	Summary	Definitions	New Blade Server	A commenter requested clarification if Version 2.0 will include a category for New Blade Servers (NBS) because they currently don't fit under the requirements for Blade Servers or Multi-node Servers. NBS do not support real-time chassis temperature monitoring and fan speed management capability and they are designed to be hot-swappable.	EPA continues to support the multi-speed fan requirement and temperature monitoring criteria because of the consumer benefits. If these products can incorporate these devices, they will be able to qualify as a Multi-node server.
6	Summary	Definitions	Multi-bay Blade Server	A stakeholder supported the Blade System definition and agreed that it represents the state of industry today. However, they stated that it might be construed as a constraint in terms of the specific form factors specified. This commenter recommended that these form factors be examples of prevalent blade server form implementations.	EPA has received no indication that the current definition would place constraints on any technologies.
7	Summary	Definitions	Resilient Server	One stakeholder appreciated the Resilient server definition and noted that these products have a power signature that is different from a Managed server. The power usage is driven by additional circuitry components that support their functionality. They suggested that a separate idle power limit for 2 processor socket servers is appropriate based on the higher average idle power for currently qualified systems that meet the Resilient definition. This commenter pointed out that there is general industry agreement that Managed servers and Resilient servers have distinct power signatures as a result of different levels of RAS and functionality. Two stakeholders expressed that the functionality of Resilient servers requires additional server infrastructure and more complex firmware capabilities. They demonstrated that the functionality and circuitry differences drives the higher power consumption for resilient servers and offered ranges of the power requirements for some attributes. One commenter gave examples of these additional features that include resiliency to data errors, scalability for large datasets, and fault resolution for big data analysis. Another stakeholder noted that this power consumption is not due to different processors and processor states but rather the redundancy of many other components such as RAM, BUS, etc.	EPA is proposing a separate base idle power allowance for resilient servers in Table 3 of the Final Draft, as well as a Buffered DDR Channel adder in Table 4 of the Final Draft to address the additional power consumption added by Resilient Server functionality. Both of these additions are based on stakeholder provided data. EPA has also made minor revisions in the Resilient Server definition in Appendix B per stakeholder request.
8	Summary	Definitions	Multi-node Server	Two stakeholders requested that the language of the Multi-node server be changed to reflect that server nodes "need not" be hot-swappable because manufacturers offer systems with hot-swappable nodes. Another stakeholder requested that the definition of Multi-node servers be elaborated to include two types in order to have a fair comparison: single processor and multiple processors. Another commenter provided a definition for node that can apply to any form factor. With this definition, Multi-node servers would be two or more independent server nodes that share a single enclosure/blade and one or more power supplies and the power is distributed to all nodes through shared power supplies. A stakeholder requested that Computer Server Form Factors should reference Multi-node servers as a separate form factor.	See Index #4.

9	Summary	Definitions	High Performance Computer (HPC)	Two stakeholder agreed with the HPC (High Performing Computing) definition. They noted that a HPC system can be a standard product or a purpose built system but these products are more than the processor or base server. They are an optimized, highly integrated cluster of server, storage, and interconnect systems that operates as a single data processing system. One commenter offered standard product examples like purpose built systems, which are designed from the processors up to integrated GPUs, memory, storay, and I/O to execute computationally intense workloads efficiently. These systems utilize a dense configuration and thus have to be designed to include highbred water/air or high efficiency air cooling systems. Since these products are large and complex, two stakeholder agreed with their exclusion from ENERGY STAR. Another stakeholder noted that "high performance computing" and "high performance computer" are not interchangeable and so the language would be more appropriately written as "Marketed and sold as a computer optimized for higher performance computing applications." They also suggested that the abbreviation "IPC" be defined. They assumed it represented Inter-Processor Communications and stated that this isn't the correct term because HPC systems have interconnects that range from proprietary inter-processor communication busses to non-coherent interconnects such as Ethernet, InfiniBand Servernet, and Myrinet.	EPA acknowledges and appreciates this additional information on High Performance Computing systems. This language suggestion has been incorporated into the specification to eliminate confusion between "high performance computing" and "high performance computer". EPA believes that IPC is an appropriate differentiator at this time.
10	Summary	Definitions	Large Server	Two stakeholders agreed with the Large server definition. These commenters stated that Large servers are ultimately differentiated from Resilient servers by their increased I/O connectivity. Also, a Large server has a minimum of 32 I/O slots and a resilient server cannot support more than 16 I/O slots. Two stakeholders supported the exclusion of these products because of the different power characteristics and application of these products. Another stakeholder questioned the inclusion of the Large server definition because according to the scope, the requirements are already limited to servers up to 4 sockets, which indicates that mainframes are excluded.	No response required.
11	Summary	Definitions	Storage Equipment	One stakeholder appreciated the consistency with the ENERGY STAR Storage specification because it will minimize confusion and overlap between the product types.	No response required.

12	Summary	Definitions	Product Family - Low/High Configurations	Several stakeholders recommended that the definitions for Lower and Higher Cost and Performance Configurations be revised. They believed that the best method of defining a product family would be to use the lowest socket power, lowest core count processor for the low configuration on one side of the 5 corners and the highest qualified socket power/core count on the other side of the product family. They stated that this approach would minimize the problems with qualified products being outside the 4 corners and fewer product families would be needed to cover each server model. They requested that EPA identify the use of socket power and core count to set the 4 corners as an acceptable approach to defining the product family. A stakeholder stated that the definition needs to fully define the range of the product family. Another supported the socket based definition because it aids in describing the supporting platform features beyond just the additional component and reflects the energy profile of the system. One stakeholder stated that the four test configurations are adequate to describe a product family, instead of five. The "typical" configuration for the fifth test will be chosen by each vendor and will not enable valid comparisons. A commenter noted that because of the limited range of configurations for a one processor socket system, the Minimum Power, Low-end Performance and Maximum Power, High-end Performance are indistinguishable and so there is no benefit to testing five configurations for a one processor socket family.	EPA appreciates the feedback regarding the Product Family configurations. However, the current approach will be retained because the guidance provided in the specification should lead the vender to select what the commenter is describing. For clarifiation, the middle or "typical" configuration enables manufacturers to highlight energy performance of a preferred system, ideally one that they sell in large quantities. The five configuration approach will be continued in the Final Draft. The number of configurations is based on the number of options offered combined with the number of those options that can be supported. One socket machines typically support a lower quantity of options, but the variety should not be dramatically less. EPA remains commited to testing these one socket systems with the five corner approach.
13	Summary	Definitions	Product Family	A stakeholder requested clarification regarding whether several types of processors can be included into the same Family if the processor brand, number of core, or Thermal Design Power (TDP) is different. One stakeholder requested more precise language to describe "Computing platform" because this term is often used to refer to the combination of hardware and operating system.	Different CPUs are expected to be utilized in each of the 5 corners of the product family. EPA understands that software can have a noted impact on performance and energy consumption and the term "Computing platform" is intended to capture the whole system.
14	Summary	Definitions	Product Family - 2 Socket, 1Processor Servers	A commenter noted that many 2 socket systems can accept a CPU that is designed for 1 socket operation because it is typically a lower cost, low to mid-wattage CPU option. They requested that these systems be covered under the 2 socket system product family because they have the same general characteristics and customers choose these products to gain features not found in the low end/lowest cost 1 socket systems. Another commenter supported qualifying single-processor systems in dual-socket server product families by qualifying them with 2 processors installed and using the same idle power limits as dual-processor SKUs.	EPA has reviewed these systems and proposes that two socket servers that can only operate with one processor shall be tested with one processor and will be required to meet the two socket idle state power allowance for Version 2.0. Two socket servers that can operate with two processors must be tested with both sockets populated.

15	Summary	Scope	Multi-node Server	A stakeholder requested that the Included Products should reference Multi-node Servers. Another stakeholder appreciated the inclusion of Multi-node servers as part of the scope. They requested that the language be changed to "…..Multi-node/Blade servers are restricted to have a maximum of 4 processors per node/blade" as opposed to "a maximum of 4 processor sockets". This commenter stated that if this change is not accepted, the restriction would eliminate ENERGY STAR certification for servers based on recent technological developments in the domain of high density, power efficient, "scale-out" servers with each processor/socket having a number of cores with mid-range single thread performance. They requested relaxation in the number of processors/sockets that are required for this class of Multi-node servers because some servers would otherwise meet the requirements.	Multi-node Servers are included in the scope of the specification. EPA appreciates this suggestion on the specification language and has incorporated it into the Final Draft.
16	Summary	Scope	Resilient Server	A stakeholder requested clarification on whether Resilient servers are included or excluded from the scope of the specification.	Two and four socket Resilient Servers are included in the scope of the Version 2.0 specification.
17	Summary	Qualification Criteria	Power Supply Requirements - Power Factor Criteria	A stakeholder questioned if a PSU that has only 12V output and 12V Standby would be considered a Multi-output Power Supply.	If the standby output of a PSU is ≤ 20 W, it is considered single output, otherwise it is a multi-output power supply. This distinction is made in the definitions of single and multi-output power supplies in Version 2.0.
18	Summary	Qualification Criteria	Power Supply Requirements - Efficiency Criteria	One stakeholder agreed with the decision to accept power supply qualifications performed against the revisions of the Generalized Internal Power Supply Efficiency Test Protocol since there were no material changes in the test procedure that would change the reported data. They also requested that references to Dc-Dc power supplies be removed because it has been removed from the test procedure and SERT does not support these power supplies. They noted that if EPA is interested in qualifying Dc based systems, then a Dc powered server where the Ac powered version of the server has been qualified should be considered.	EPA has removed the Dc-Dc requirements from Table 1 and Table 2 as there is currently no procedure to test Dc servers in the Version 2.0 Computer Servers Test Method. EPA and DOE will revisit including Dc-Dc servers in the Version 3.0 specification revision process.
19	Summary	Qualification Criteria	Power Management - Server Processor	A stakeholder agreed with the recognition of in-band power management options. They requested that language be added to support an electronic only method of distributing documentation as opposed to physical documentation.	The specification currently supports electronic distribution of documentation.
20	Summary	Qualification Criteria	Power Management - Supervisor Power Management	A stakeholder requested clarification regarding whether this section is referring to all techniques configurable in the BIOS or if the EPA is asking a vendor to reveal the internal working of its server's power management.	All methods that impact power management features that are accessible by the end-user must be disclosed - be they BIOS, OS, or other origin. Language has been modified in the specification to clarify this point.
21	Summary	Qualification Criteria	Power Management - Disclosure	A stakeholder requested that it be specified that only the power management techniques enabled by default and listed in the power management section of the Power and Performance Data Sheet be identified. They stated that it is not reasonable to require all power management techniques be disclosed because some are proprietary.	The power management techniques that should be disclosed are those that the end-user will utilize, not proprietary designs embedded in software/hardware.

22	Summary	Qualification Criteria	Blade System Criteria - Thermal Management	A stakeholder requested that this requirement be changed as follows: "..must provide real-time chassis or blade inlet temperature monitoring…" because some chassis systems collect their temperature readings based on the blade inlet temperature to better match fan speed to the temperatures at the blade server. This will not change the intent of the requirement. This commenter appreciated the decision to allow companies to deliver required documentation electronically to the purchaser. They stated that the part of the requirement that specifies that an EPA approved format be used, is confusing and should be removed. They noted that requiring EPA approval will create unnecessary work.	EPA has made the necessary language changes to account for these systems because the point of this requirement is to document the inlet temperature. EPA has removed the language requiring an EPA approved format to be used.
23	Summary	Qualification Criteria	Active State Efficiency Criteria	A stakeholder requested that this section reference "i. SERT™ main report results" and "ii. SERT™ detailed report results over the entire test run". Also, the "workload module" term should be changed to "SERT™ workload". Another stakeholder requested that this section remain TBD depending on the results of the pre-release evaluations of SPEC's SERT™ tool. Since there will be an insufficient experience in the worklets to use the results as a comparative indicator of energy efficiency, this stakeholder requested that the data publication be anonymous until Q2'2014. This method will allow analyses (multiple system, operating system, and architectural) across all SERT™ workloads before identifying indicators that are representative of a product/product family. This stakeholder requested that the following language be added to list limitations of what the values SERT™ represent: "Output values obtained from the SERT™ utility are intended solely for ENERGY STAR qualification purposes, based on a limited, conservative sample set. Actual results may vary. SERT™ output values listed here are intended to represent a precise set of configurations, not necessarily reflective of all available configurations."	EPA has decided that there will be a 9 month period of time after Version 2.0 is published where the data submitted to EPA will be published anonymously. This will provide a period of time for stakeholders to evaluate SERT™ results before presenting them to the consumer. EPA's standard language describing product families and representative testing addresses this concern for the more general case--i.e. not specific to SERT but for all products where families may be qualified through testing of representative units. EPA believes that the existing language in the specification is sufficient but is open to discussions with stakeholders about ways to emphasize this point in other appropriate locations.
24	Summary	Qualification Criteria	Idle Mode and Full Load Efficiency Criteria 1S and 2S - Base Idle	One stakeholder suggested that EPA consider reducing the base idle criterion by 3W and having the manufacturer add the memory adder based on the total quantity of memory in the system to simplify calculations. Another stakeholder noted that the proposed base idle values are the same as Version 1.0 however, improved power management of servers allows for better scaling of power demand dependent on the workload level, so efficient servers should operate at lower idle levels. They requested an analysis of current idle levels for new products as a basis for adjusting the idle levels.	Based on the current data analysis, EPA remains commited to the existing base idle criteria.
25	Summary	Qualification Criteria	Unit Clarification	Another stakeholder requested that the term "GB" be defined because electronic engineers would define it as 10243 power, while others defined it was 109 power.	This specification uses base 2 numbering to define GiB to be consistent with the Data Center Storage specification.
26	Summary	Qualification Criteria	Idle Mode and Full Load Efficiency Criteria 1S and 2S - Additional Power Supplies Idle Power Allowances	A stakeholder suggested that the 20 W adder continues to be a challenging limit despite improvements in PSU efficiency because integration and power densities have increased. However, another stakeholder stated that it was unclear why the 20 W adder is still offered for redundant power supplies since industry has shown products which offer redundant power supplies that can be kept in stand-by mode (power supplies are activated only when needed).	EPA appreciates the feedback regarding the power supply efficiency but will continue to maintain the 20 watt adder as data shows that it is appropriate.

27	Summary	Qualification Criteria	Idle Mode and Full Load Efficiency Criteria 1S and 2S - Additional I/O Devices Idle Power Allowances	A stakeholder questioned if video devices such GPU for display output, USB, External SAS, and FibreChannel are covered under the Additional I/O Device allowance mentioned in the specification. And if so, can they apply the Additional Idle Power Allowances shown in Table 4? A commenter requested clarification regarding whether the additional I/O devices must be pluggable (like PCIe slots), or can they be permanently soldered onto the chassis or integrated into the processor. They suggested that the allowance be the same regardless of the implementation.	EPA has added SAS, SATA, FibreChannel and Inifiband to the Additional Idle Power Allowances for Extra Components list. USB is not eligible for this allowance. Also, to clarify, I/O devices can be soldered onto the chassis, integrated into the processor, or implemented in a pluggable a card.
28	Summary	Qualification Criteria	Idle Mode and Full Load Efficiency Criteria 1S and 2S - Additional Memory Idle Power Allowances	One stakeholder requested that power allowances be added for memory buffers and RAID cards, since these components drive extra power use in Resilient servers. They offered data for two component adders and determined the measured power draw. They proposed a 4 W per DDR port adder for memory buffer.	EPA is proposing a Buffered DDR Channel adder of 4.0 W per installed Buffered DDR Channel greater than 8 channels in Table 4. The power consumption of the first 8 channels has been accounted for in the two socket Resilient Server idle state power allowance in Table 3. RAID cards are not being considered for an allowance as an extra component in Resilient Servers.
29	Summary	Qualification Criteria	Idle Mode and Full Load Efficiency Criteria - Blade Servers	A stakeholder expressed concern with the blade testing procedure and proposal to focus on the use of half chassis data to set qualifications for Version 3.0. The different approaches to distributing system overhead in a chassis will lead to "apples to oranges" comparison between half populated chassis from different manufacturers. The test procedure allows the vendor to decide what the required number of options (I/O, fans, etc.) will be placed in the chassis for the half chassis test which allows for customization to the test rather than the configuration that would typically be used by a customer. They requested that manufacturers test only a full chassis to provide power use and performance data. They offered an example product to point out that this requirement gives advantage to certain products. Their position is that the only test procedure that allows customers to truly compare per blade results from different vendors is to require full chassis testing. The testing burden is considered part of the normal development operation. This commenter requested that if this approach is not accepted, than the EPA should allow companies to provide test data on a full chassis if they choose as it will not risk the integrity of the ENERGY STAR measurement. However, another stakeholder agreed that reporting idle and full load in a half populated chassis is a reasonable compromise and appreciated the considerations for the expense and time testing a fully populated chassis would require. They noted that even half populated chassis testing is very expensive and resource intensive. They recommended considerations for reduced or limited verification for this class of product because of the complexity and expense. One stakeholder requested that more criteria on power management features like power monitoring, power management at blade chassis level, standby for redundant power supplies, etc. be added to account for the limited criteria for Multi-socket and Blade servers.	To address this issue, EPA has changed the requirement to testing after rounding up to the nearest power domain. This will still allow for the reduction in testing burden from testing full chassis. However, full chassis results can still be submitted and shown on the PPDS, if the vendor chooses to test the fully-populated chassis. EPA believes that the criteria set for Blade and Multi-node servers is sufficient based on current data and knowledge of the products.

30	Summary	Qualification Criteria	Idle Mode and Full Load Efficiency Criteria - Multi-Node Servers	A stakeholder questioned whether Section 3.6 (Idle Mode and Full Load Efficiency Criteria 1S and 2S) or 3.8 (Idle Mode and Full Load Efficiency Criteria - Blade Servers) will apply to Multi-node servers.	Section 3.9 has been added to specifically address Multi-Node Server criteria in Version 2.0.
31	Summary	Qualification Criteria	Other Criteria	One stakeholder requested that Other Criteria be expanded to include Multi-node and Blade servers.	Blade and Multi-node Server criteria are covered under Section 3.4 and 3.8 and 3.9 of the specification.
32	Summary	Qualification Criteria	SERT	A stakeholder noted that the SERT™ (Beta-2) Design Document only supports 64b architectures. They suggested a clarification in the SERT™ document to emphasize that SERT™ architecture is agnostic and other architectures or operating systems are not supported due to lack of resource. This additional language will allow power efficient servers based on 32b architectures to qualify. This commenter also noted that the motivation for the restriction of 8 sockets and 64 nodes is unclear. They requested clarification on the relationship between sockets and nodes. This stakeholder also requested that the issue of scalability of the worklets and corresponding measurement of power efficiency be addressed.	EPA supports all architectures and any limitations are a result of the test method. Any questions about the SERT™ tool, should be directed to the SPECpower team.
33	Summary	Standard Information Reporting Requirements	PPDS	A stakeholder requested that the Delta Temperature at Exhaust at Peak Temperature be removed or modified because the power dissipation is reported from the benchmark run where the temperature is 18-27°C. The peak temperature at 35°C is never tested and power use would be somewhat higher due to leakage and fans. Since it is not tested due to additional expense, this value cannot be accurately calculated. They proposed that the requirement be changed to nominal delta temperature, which can be calculated easily with power and nominal airflow. Two stakeholders agreed with the proposal to only consider power and performance data from the SERT™ benchmark but requested that the word "benchmark" be replaced with "rating tool". One stakeholder recommended that EPA conduct trial runs in populating the new Power and Performance Data Sheet (PPDS) in addition to aggregating and posting the information. The SERT™ tool is expected to include hardware detection and reporting tools that could aid in the accuracy and consistency of the data. They also requested that documentation and data entry expectations be reviewed in an ENERGY STAR Servers testing workshop prior to the Version 2.0 effective date. They suggested that annual energy consumption estimates be optional and include condition assumptions/configuration information because this value is dependent on the system configuration, application and supported industry in which the servers are deployed. This stakeholder supported the direction to base idle power compliance on the SERT™ results. They noted that due to resident workloads, the registered idle power may be higher than those observed in the previous non-application loaded test procedures. Testing to idle after running the workloads could thus result in a more challenging specification.	EPA appreciates the input on the Delta Temperature at Exhaust at Peak Temperature value. This will be a priority topic for Version 3.0. Since these are relatively new proposals and will require time to analyze and discuss with stakeholders, they will not be considered for Version 2.0. When referring to SERT™, the language has been changed to "rating tool". Also, EPA will conduct trial runs to validate the performance of the PPDS before it is used to post information publically. EPA confirms data entry expectations for every specification. EPA intends to keep the current approach for estimating energy usage. The assumptions for making this calculation are listed on the PPDS.

34	Summary	Standard Information Reporting Requirements	Public Disclosure	Another stakeholder requested that the SERT™ data be anonymous for the first 18 months of Version 2.0 because this metric is new and relative/absolute value of the worklets have not been determined. They recommended allowing time for the stakeholders, EPA, and SPEC to evaluate the data set. They suggested a proposal to accomplish this, by using two data sheets. The PPDS without the SERT™ data could be used to complete public information on the ENERGY STAR site. The complete PPDS could be used to compile a blinded datasheet. The manufacturers should be required to have the complete PPDS with the SERT™ data available for customers.	See Index #23.
35	Summary	Standard Performance Data Measurement and Output Requirements	Measurement and Output	A stakeholder requested the following language change: from "…utilization of all logical CPUs…" to "AVERAGE utilization of all logical CPUs". In cases where there are many hardware threads running on many cores on several processors, it would not be practical, nor would it provide value, to report the utilization values for every thread. Another stakeholder suggested including criterion concerning warranty from manufacturers that guarantees normal server operation and lifetime of equipment at inlet temperature up to 27 C. Data center managers are usually skeptical of recommendations on temperature regimes and often keep their data center at 20 C or below but this warranty would support data center manager in choosing appropriate temperatures.	EPA believes that the current language is sufficient and aligns with the common usage of CPU utilization within industry. In terms of warranty criteria, EPA welcomes this idea any may consider it in the future specifications but it is outside of the current scope.
36	Summary	Standard Performance Data Measurement and Output Requirements	Sampling Requirements	One stakeholder requested further information on the difference between the requirements for Input Power and Processor Utilization and Inlet Air Temperature. The requirement for input power states "a rate of >= 1 measurement per contiguous 10 second period" while for inlet temperature the requirement is "a rate of >= 1 measurement every 10 seconds". Is there a different technique that should be used for sampling input power? A stakeholder suggested a 60 second reporting frequency because collecting data from hundreds or more servers on a 10 second frequency will consume a significant portion of the data center network infrastructure with no benefit in clarity of thermal conditions of operational response time. This stakeholder recommended that the reporting interval for time stamped data be set at 10 minute intervals (of 30 second averages and 20 data points) will provide adequate information and generate sufficient response times to identify and react to non-catastrophic thermal excursions in a data center.	EPA has revised the current language to eliminate confusion on this topic. Also, EPA has decided that one minute reporting frequnecy is sufficient to measure the delta temperature as it would not be expected to change significantly within a minute.
37	Summary	Effective Date		A stakeholder requested that products shipped during the interim period from November 9, 2012 and August 1, 2013 should be allowed to test and claim certification to either Version 1.0 or 2.0 criteria. Another stakeholder suggested that systems that qualify under Version 2.0 be listed on the Qualified Product List (QPL) for Version 1.0 until the Version 2.0 QPL is effective without the need for resubmission.	Manufacturers can test to either Version 1.0 or 2.0 during this interim period when the specification is finalized but not yet effective. Also, once the test method is finalized for Version 2.0, products that meet the new criteria will be listed on the Qualified Product List.

38	Summary	Consideration for Future Revisions		A stakeholder stated that using SERT™ data as the basis for choosing idle and active mode criteria will not accurately portray the energy efficiency profile of the entire server market.	EPA has several ways of developing a more accurate picture of the entire server market. In order of reliability: (1) EPA can do a data collection for V3.0. (2) EPA can acquire non-qualifying data from manufacturers that ran the benchmark without the intention to qualify for ENERGY STAR. (3) EPA can look at our market share in V2.0 and correct for that to some extent. (4) EPA can look up non-qualifying products and make assumptions about their energy consumption. There might be some problem for blade servers due to cost, but overall we should be able to correct for this and get reasonably close to a 25% level in V3.0. This is a problem we face in every specification.
39	Summary	Resilient Servers		A stakeholder provided data for EPA to set accurate criteria levels for Resilient servers.	No response required.
40	Summary	General		A stakeholder made the following suggestions for all three documents (Specification, Test Method, and PPDS): • Dc and dc -> DC • Ac and ac -> AC • UUT and SUT is used to describe the same thing; it is preferred that only "SUT" is used. • Lines 24, 25, 39, 92, 146, 148, 150, 153: Power Meter" should be "Power Analyzer"	EPA appreciates this feedback but will maintain the current language in order to be consistent with other ENERGY STAR specifications.