IBM Comments to the following EPA documents, released prior to the 9/25/2009 Stakeholders meeting:

ENERGY STAR® Tier 2 Computer Servers Stakeholder Meeting Discussion Guide
ENERGY STAR Tier 2 Computer Server Preliminary Specification

IBM appreciates the opportunity to comment on the "Tier 2 Computer Server Preliminary Specification". IBM believes that an effort to develop a power performance metric for the ENERGY STAR Computer Server Specification is the next logical step in driving energy efficiency improvements for servers and data centers, especially if the idle measurement is integrated into the active mode metric. Server energy efficiency, or more appropriately energy utilization, is the function of two factors: maximizing the amount of time that the server is doing work (virtualization of workloads) and minimizing the power required to maintain the server at idle. It is important that the power management criteria consider and integrate both factors to insure that a given server is optimizing the workload of energy delivered per unit of energy consumed.

IBM also supports the implementation of the blade system criteria through an update of the Tier 1 specification. We have provided proposed modifications to the Tier 1 specification to incorporate blades systems into the specification. IBM encourages EPA to set a power management criteria based on enablement of processor level power management to simplify the implementation of the specification, recognize the complexity of blade systems given the various ways by which overhead is distributed between products and manufacturers and the number of processors installed in a chassis, and allow companies to focus on development of the Tier 2 specification.

We have provided other comments with regards to proposed changes to the specification or in response to specific EPA questions embedded in the text. Please contact Jay Dietrich (jdietric@us.ibm.com) with any questions or for additional discussions on technical issues.

Comments on the Active Mode Rating Tool Options:

The discussion guide offers four options for the reporting of Active Mode data. IBM agrees with the desired tool characteristics EPA detailed on page 2 of the discussion guide. IBM believes that option C in the discussion guide offers the best means to deliver a rating tool which most fully meets the characteristics outlined on page 1 of the Discussion Guide referenced above, with some suggested modifications.

IBM believes that the current effort by the SPEC Power working group to generate a "Server Efficiency Rating Tool ™" offers the best approach to establishing an active mode rating for computer servers. The SPEC benchmark tool is an industry initiative with broad support and, if successful, has the potential to provide a robust metric tool that addresses both the active and idle mode energy-using characteristics of a server.

However, IBM recommends several modifications to the EPA proposal to enhance the applicability of the rating tool and simplify the overall evaluation of energy use by computer servers:

1. The rating tool should focus on performance and power capabilities of the different key functions which are common to all servers including, but not limited to, processor computation, both floating point and integer computations and disk, I/O, and memory access. A tool which focuses on functional, energy efficiency performance offers the best approach to providing an effective comparison of systems with both lighter and heavier configurations and higher powered processors.
2. Idle power requirements should be included as part of the evaluation under this metric, allowing the server to receive a single, overall score which incorporates both active and idle power performance. It may be desirable to provide a higher weighting to the idle power requirements in calculating a score from the tool, but including idle in the metric provides the opportunity to normalize the overall score for higher powered, more capable processors which are currently penalized in the Tier 1 "idle" only criteria.
3. The tool should be designed to allow manufacturers to demonstrate their hardware and system software based efficiency capabilities. Manufacturers have provided their hypervisors, operating systems, and middleware with capabilities that enhance the energy efficiency of the server operation. These capabilities and functions should be recognized for their ability to provide the most energy efficient solution for a customer.

IBM recognizes there is some risk inherent in the ability of the SPEC working group to meet EPA's timeline of a published Tier 2 specification in October of 2010. IBM urges EPA to consider extending the timeline by up to six months if the SPEC effort demonstrates sufficient progress and technical potential during the first half of 2010. IBM believes this is prudent for several reasons, including:

1. Currently, there are a limited number of computer server products qualified to the ENERGY STAR Computer Server requirements, indicating that the current Tier 1 specification is differentiating the energy efficiency capabilities of the current servers available in the marketplace. Recognizing the server development times are approximately 18 to 24 months, this suggests that manufacturers will not catch up with the Tier 1 requirements until midyear or year end of 2011.
2. While Option D may be attractive, it will result in the ENERGY STAR computer server specification going through three very different criteria sets in the space of two to three years. This creates a moving set of standards which impact design and testing requirements, which is unfair to computer server manufacturers. Given the observation in item 1 above, IBM believes it is better to set an active mode criterion, preferably with an integrated idle power criterion, which will have at least a two year life after it is published. Option D (or B) also suffers from the difficulty of selecting a representative set of benchmarks with both power and performance report aspects. On the one hand, the set of suitable, released benchmarks is still small. On the other hand,

the area is in flux as a number of benchmarking organizations begin to consider the importance of measuring energy-efficiency. More, any benchmark suite constructed under Option D will lack a uniform infrastructure and a consistent set of rule runs. The industry will find it difficult to execute the number of tests likely to be required.

3. There is a significant amount of resource required to test equipment currently available in the marketplace to a new or revised server criteria. The proposed active mode testing, whether done through the rating tool developed by the SPEC Power working group, or using existing workload benchmarks, currently is not done on a routine basis for all product families. It will be difficult to generate benchmark performance data for one iteration of the Tier 2 specification, and nearly impossible to do it for two iterations within 12 months.

If it is determined that the SPEC Power "Server Efficiency Rating Tool" cannot meet a reasonable schedule or exhibit the required characteristic, IBM favors the evaluation of an approach to utilize a select number of current benchmarks to evaluate server Active Mode power. IBM recommends that EPA consider including SPECpower_ssj2008; SPECweb_power2009; SPECvirt_sc2009; and TPC-C and TPC-E benchmarks in the ENERGY STAR criteria. IBM recommends that discussions be initiated to determine which workload benchmarks, or combination of benchmarks, would be most appropriate as an Active Mode Rating Criteria. These discussions should also consider how to provide a uniform infrastructure and testing requirements for power measurement. IBM believes there would be two possible approaches to using these existing benchmarks, with the first choice being the preferred.

1. Develop a rating system for each selected benchmark, requiring that a company report the benchmark rating and the power use for the selected metric. There are several work items to complete:
    a. How would the configuration for the benchmark test be selected? By the manufacturer? Specified by the criteria for each metric? Some other approach?
    b. Would the product family definition continue to be relevant? It will be impossible to complete a benchmark test for every possible configuration in a product family.
    c. How would the idle and active criteria be combined? Would it be done as a single idle/active criterion or as two data points, both of which must be satisfied?
    d. How do you get an adequate sample of benchmark results on product families in the marketplace to assure there is a sufficient sample size to draw meaningful criteria?
2. The second approach would be to select all or some subset of the 5 recommended benchmarks and create a "normalized" metric which would generate a single "performance/power indicator". This would require the completion of all the work items in Item 1 above, as well as the following, additional considerations:
    a. What algorithm do you use to blend and normalize the selected power performance scores?

b. Do you construct a power performance curve, with each point on the curve representing one benchmark and then analyze the curve similar to the way the power supply efficiency curve is analyzed to set that criterion?
c. For a blended benchmark, do you require the use of a single product family configuration for each benchmark or do you allow the manufacturer to select a preferred configuration for each test? How is that configuration determined?
d. How does one create sufficient commonality in infrastructure, run rules and reporting to ensure reasonable results and feasibility? Each benchmark has its own way of doing things.

As these items indicate, the development of an active power criteria based on currently available benchmarks will require a great deal of work and testing. It is not immediately clear that the benchmark selection, validation, and harmonization will require noticeably less time than developing the rating tool. We highlight this point to encourage EPA to establish a Tier 2 specification with an expected longevity of at least two years. While we understand the interest in option D, we are very concerned that the amount of work required to support the development of two concurrent sets of criteria is not achievable, will have many unintended consequences, and result in a less than optimal solution for Tier 2.

Comments to the Tier 2 Preliminary Specification Document.

Definitions:

Line 261: K. High Performance Computing System: The proposed definition as written is flawed. The definition is not meaningful in that HPC systems are not the only large-scale systems. Their defining characteristic is that they are used to run very large jobs each of which has a very large computational requirement. Their implementation technology varies from very specialized such as the machines from D. E. Shaw Research to ones assumed from off-the-shelf parts such as the cluster systems at the Texas Advanced Computing Center. The traditional view that such systems disable power management has changed over the past two years: power management is increasingly linked to achieving the maximum computational performance. IBM is willing to provide a proposed definition of HPC systems for EPA's consideration.

HPC systems should be excluded from the Energy Star Tier 2 specification due to their diverse requirements and implementations.

Line 268: L. Resilient Server: IBM believes that it is important to include a definition for Resilient Server in the specification. Resilient Servers are engineered with additional and more complex components in their underlying infrastructure in support of the resiliency features, which in turn require more energy to operate, distinguishing them from a computer server without RAS feature. These additional energy requirements will affect the performance of resilient servers on a performance power benchmark, potentially impacting the ability of the servers to compete effectively against servers

without these features. While the intent of the development of the Active Mode Rating Tool is to provide a tool which will normalize for these features, it is possible that resilient servers will demonstrate a unique performance power signature that is fundamentally different from servers without resilient features. In this case it would be appropriate to establish specific criteria for Resilient Servers.

The following features should be used to designate a Resilient Server: Advanced error detection, isolation and recovery (also referred to as machine check architecture), mirrored memory, bus retry, hot swappable components, error correction codes (ECC), ability to perform on-line expansion/retraction of hardware components, and multiple physical banks of memory and I/O adaptors.

Definitions of the features of a Resilient Server are detailed below. A Resilient Server should have four or more of these features.

Mirrored Memory

Mirrored memory means that all or parts of memory have a second, normally hidden copy, for error checking and to insure accessibility. This is an expensive feature, but it is used by IBM and perhaps others. It is likely that a machine with mirrored memory will have many other resiliency features as well and should be classified as resilient, but there are other resilient machines that do not have mirrored memory.

Bus Retry

The data transmission paths in a machine are called *busses*. There are busses between the processors and memory and busses to and from the I/O adapters and devices. On occasion there are errors on a bus. For example, when the processor requests bytes from memory and the bytes do not arrive within a specified time. When that happens, the requester can either give up or reissue the request. Automatically reissuing the request is called *bus retry*. A resilient server, especially with today's technology, is likely to have bus retry.

Hot swappable components

Hot swappable components are system parts which can be replaced without shutting down the system. A common example is a power supply in a multi-power supply machine.

Error Code Correction

Extra bits are stored in memory or in a cache that allows the machine to detect some of the possible errors - such as flips of a single bit - and correct them.

Ability to perform on-line expansion and retraction of hardware resources - also referred to as "on-demand" features.

On-line expansion and retraction of hardware resources is a method of adding processors, memory or devices to a running system or removing them without shutting down the system or stopping the work that it is executing. Generally, support is specific to a

particular category of resource. For processors and memory, support is a relatively new feature.

Multiple physical banks of memory and I/O adapters
Many resilient systems have multiple physical banks of memory with separate paths from the memory controllers to each. Similarly, they have multiple I/O busses.

Line 346: X. Server Utilization: The definition of processor utilization should be changed to read: "The estimation of the server's compute activities relative to their activity at full voltage and frequency of the processor(s)."

It should also be pointed out that this is a definition of processor and not server utilization. There are other server resources such as memory and I/O busses that must be considered when determining how busy a server is.

This definition supports the proposed language for processor utilization measurements.

**Line 628**: *Processor Utilization Measurements: The computer server will provide an estimation of the processor or system utilization that is visible to the operator or user of the computer server through the operating environment (operating system, hypervisor, or other management interface). The number is intended to provide the data center operator a qualitative indication of the amount of load on the system to provide guidance regarding opportunities for virtualization and consolidation of workloads to deliver more work for each unit of energy used. In systems that support active and inactive (powered off) processor states, it is only a measure of the active processors.*

Currently, the industry does not have a single, standardized methodology for measuring processor utilization, particularly in virtualized environments where there are simultaneous threads being managed by the processor. IBM proposes that the utilization requirement continue to be qualitative in nature for the same reasons cited during the review of the same language in the Tier 1 specification as well as some additional reasons which are specific to current technology and the nature of the Tier 2 specification.

1. With the introduction of the power – performance rating or qualification tool, there is now a quantitative measurement of the power efficiency being delivered by the system under test. This reduces the value of reporting other quantitative measurements of system behavior. The tool provides the quantification, and the EPA reports whether or not a particular system meets the current standard for efficiency.

2. There are other components to utilization including memory occupancy and I/O bandwidth consumption. These have to be measured separately. The other components affect the amount of unused capacity on the system. One can easily be misled by traditional CPU utilization alone (which is inaccurate) and believe

that a machine has a large amount of spare capacity when, in fact, very little additional, meaningful work can run on the system. To provide an accurate picture, metrics like memory occupancy and bus utilization would also need to be reported. We do not propose that these functions be included in the ENERGY STAR specification.

3. The language proposed includes measurements "per OS-level CPU, per core, per socket." The natural reading is that all three sets of measurements are required. Most operating systems collect only the first of these, and a few operating systems, hypervisors, or service processors collect the second. Very few, if any, implementations systematically report the third. Thus, the EPA's requirement implies significant coding changes, requiring investment and time to deliver. In the open source case, the community must also accept the new function, which requires months of discussion and cannot be forced by any one company, organization or individual.

4. CPU utilization as defined in the specification appears to be a single number. Unless the system has a single processor socket with a single-core processor in it and is running a single hardware thread, the CPU utilization is per-processor or per-processor-like entity. The system's CPU utilization is a vector.

5. Each processor, system and operating system or hypervisor or service processor uses different technical techniques to estimate or quantify CPU utilization, leading to serious non-uniformities in reporting. This makes the requirement of an absolute or comparative number problematic without a commonly accepted definition of CPU utilization. Further, although the other utilization metrics have fewer measurement techniques, there is no standard way of expressing or normalizing their results.

6. With the introduction of power management functions and hardware multi-threaded systems, the current processor utilization measurement algorithms or functions cannot fully compensate for the presence of more than one thread or reductions in processor frequency. This in turn introduces significant inaccuracies into the measurement which make it unreasonable to establish a quantitative requirement with prescribed levels of accuracy. Each hardware and system software supplier has their own approach to resolving this for their customers and therefore no single formula can represent all implementations.

7. The formula proposed in the specification assumes that the system tracks with reasonable accuracy and records the adjustments made in the frequency. With the latest generation of power management techniques, processors increasingly adjust their own frequencies on very short time scales. Some systems provide a good way to track these changes, but others do not. This means that collecting $F_A$ for the formula in the draft specification may or may not be possible. Operating systems, in particular, are generally unaware of the frequencies actually being used.

8. CPU utilization reported in virtualized environments may be relative to the capacity of the resources awarded to the virtual machine, or it may, alternatively, be relative to the allocation of resources given to the virtual machine. A virtual machine may be allocated 50% of a processor, and it may report 50% utilization, which is either 50% of the physical processor or 25% of it, depending on the scheme used.

9. The most vexing problem with CPU comes with simultaneous multi-thread (SMT) processors. In an SMT environment, the utilizations of the hardware threads are not independent of each other. The load, including the details of the workload, on one SMT thread affects the utilization of all of the other SMT threads of the same processor core. Although it is possible to avoid this problem by mandating that processor cores run in single-threaded mode, many workloads run much more efficiently in SMT mode. The typical default setting is to have SMT active.

10. The whole question of how to measure system utilization is a current topic of research. IBM believes the best proposal to date is to measure it using the ratio of the rate of instruction retirement on the system to the maximum rate achievable for the current workload. This means utilization workload-specific, but it also realistically recognizes that the type as well as the intensity of the workload affects utilization.

For the purpose of the Tier 2 ENERGY STAR® computer server requirements, any utilization measurement needs only be sufficiently accurate for the purpose of enabling decisions around the consolidation or reprovisioning of workloads or the movement of virtual machines. Existing CPU utilization measurement algorithms are already sufficiently accurate for this purpose. Many data centers today use CPU utilization numbers as reported by currently shipping operating systems on existing servers. They use this information to dynamically manage data center power consumption by reprovisioning workloads to minimize the underutilization of machines. IBM believes that requiring any particular formula or accuracy for CPU utilization will not provide any particular incremental value to customers and will only increase the cost of the system due to the expensive additional micro-instrumentation required.

As an example of how the current utilization estimate can be used, a data center operations team at a specific facility tracked system utilization. In 2006, over 50% of the 1 and 2 processor servers in the data center were utilized 5% or less of the time and an additional 17% were utilized 5-10% of the time. By 2008, the percentage of 1 and 2 processor machines operating in these utilization ranges were reduced to 13% and 6% respectively[1]. Much of the workload was moved to larger, more heavily utilized servers. Providing a qualitative measure of server utilization provides a data center operator with the information needed to identify these opportunities.

[1] Green Case Study for SCMG.ppt sent to Andrew Fanara and Arthur Howard (USEPA) in email dated November 19, 2008.

The language suggested as an alternative offers a sensible way to require the reporting of computer server utilization recognizing that there is no viable methodology to create a repeatable, quantitative method to measure utilization.

Line 404: Qualifying Products

IBM recommends that EPA add blades to the list of qualifying products as a Tier 1. A proposed approach to modifying the Tier 1 specification to include blade systems is provided later in this document.

For Tier 2, IBM recommends that the specification continue to cover systems with 4 processors or less and blade systems. These systems represent between 80 and 90% of the server market. Maintaining these systems as qualifying products will provide continuity in the specification and an appropriate range of systems to demonstrate the validity of the chosen active mode metric.

IBM believes that the EPA should continue to exclude Server Appliances, Fully Fault Tolerant Servers, and Multi-node Servers from the specification as they represent a small percentage of the market.

Line 418: Blades

In Appendix A, IBM has proposed recommended changes to the Tier 1 Computer Server requirements to include blade servers. The proposed changes would create the following requirements for blade servers:

1. All power supplies used in a qualified blade system will have to meet the power supply requirements, Section 3.A.
2. The blade servers will have to meet the requirement that processor level power management be enabled on all shipped blade servers (section 3.B.2)
3. Standard reporting (power/performance datasheet) will be required for both the blade server (minimum, maximum & typical configuration) and a blade chassis fully populated with minimally configured blade servers. The power on a full blade chassis could be computed either by measuring the chassis fully populated with minimally configured blades or by determining the power use of a partially populated chassis, determining the per blade power draw for that chassis, and then computing the total power for the chassis by adding the power per blade times the number of empty blade slots in the test chassis plus the power use of the partially populated test chassis.
4. A power/performance data reporting sheet will need to be specifically created for blade servers and systems (Section 3.C).
5. The blade system and populated servers will collectively meet the Data Measurement and Output Requirement (Section 3.D).

IBM notes that once a blade chassis and server combination are qualified under ENERGY STAR®, a qualified product would consist of a chassis and blade server from qualified product families. It is proposed that at least 70% of the blades supplied with an ENERGY STAR® qualified system be ENERGY STAR® qualified, allowing a few blades in a qualified chassis to be unqualified blades. This is because customers will mix and match blade servers in a chassis based on their computing needs.

This approach offers the most effective means to bring blade servers under the ENERGY STAR® computer server requirements given the complexity of the blade systems and the varied approaches to distributing cooling and power distribution overhead between blade servers and blade chassis by different manufacturers. This approach has the following benefits:

1. By testing a full chassis (or a partially populated chassis if the manufacturer prefers) of minimally configured blades, data will be provided to allow calculation of per blade power usage in a fully populated chassis. The minimum blade server configuration was chosen for this test, as it provides the best means to allow meaningful comparisons between manufacturer's systems. It is likely that typical and maximum configurations will vary sufficiently in power use making it difficult to derive meaningful comparisons between systems.
2. It will be possible to estimate the chassis power by subtracting the power use from a minimally configured blade server (times the number of blades in the chassis) tested in a test fixture (see Appendix A, 1. Definitions; Blade Test Fixture (proposed) (p.4,5 of these comments) from the power measured by the fully populated blade system.
3. It allows comparison of the power use of the typical and maximum blade server configuration as tested in a test fixture. As such, a data center operator will be able to compare the relative power use of various systems from the data provided.

We believe that testing the individual servers and a fully populated blade chassis provide the data necessary to allow a data center operator to easily calculate and compare the power use of different configurations and levels of chassis population. While testing a blade chassis fully populated with minimally configured blades will require adjustments in the product testing procedures, it does not represent any specific financial or procedural hardship if the testing is allowed to cover a product family. A manufacturer should be permitted to test a partially populated blade chassis to compute the per blade power use because the per blade power use should be conservative due to distribution of the chassis overhead over less than the full number of blades. The proposed approach provides a workable testing regime for the manufacturers, providing the data needed by the server purchaser to make informed decisions on the blade system power use.

Line 434: Power Supply Efficiency

Net Power Loss: IBM continues to recommend that EPA not pursue the Power Loss metric as a power supply criteria for IT equipment. EPA's earlier analysis of the "power loss" metric based on the published SPEC Power results is skewed towards the "lightly

configured" systems used to generate results for the SPEC Power measurements. The data supplied previously by IBM (July 2008) in the file "Power Config Sort.pdf" shows that idle and maximum power increases significantly by model type, by a factor of 2 to 6, as the system is outfitted from a minimum to a maximum configuration. Thus power supply operating points will be approaching up to 40% on a redundant power supply for a system with a "maximum configuration" and higher where the server system has the capability to actively "sleep" or "switch" off the redundant power supply.

Customers understand the power supply efficiency metric. Setting up two, differing metrics for a single attribute within a standard is not a good practice. We believe that the power supply efficiency requirements are the appropriate metric to use to set a power supply requirement under future ENERGY STAR® Server Specification. Adding Net Power Loss reporting to the Power/Performance data sheet does not add any value to the process or benefit to customers that make use of the information on the data sheet.

Line 440: IBM supports EPA's proposal to increase the power supply efficiency requirements to 80Plus Silver for multi-volt output power supplies and 80Plus Gold for single volt output power supplies.

Line 459: If the active power metric is properly configured to include consideration of idle, then EPA should remove the idle power criteria from the Tier 2 specification. See the discussion on page 1 of this document.

IBM continues to evince that the complexity and range of configurations for three and four processor socket systems makes the development of an idle power criteria for these system difficult and counter-productive. Four processor socket systems are used to virtualize and support multiple workloads, which can be the most efficient approach to computing, but these types of systems will not be identified or recognized by an idle power criterion. Either the active power metric should be constructed in a way that recognizes and evaluates idle power as part of the overall evaluation of a server's power/performance or the requirement for activated processor level power management should be continued in Tier 2.

Line 553: Standard Information Reporting Requirements: IBM has concerns regarding the data required on the power performance data sheet. We are not prepared to provide comments at this time, as we are currently finalizing our first data sheet submittals. IBM intends to provide comments to EPA on the data sheet requirements, and specific clarifications needed, once we have completed our submission of products for ENERGY STAR® qualification. We again believe that a more automated approach to reporting and archiving the reported information is warranted. Entering data directly into a spreadsheet is too error-prone and labor-intensive.

Line 594: Data Measurement and Output Requirements:

Line 617: Input Power Measurements: IBM believes that the Tier 1 accuracy requirements for measuring input power, 10% or 10 watts, provides sufficient accuracy

for the purposes of managing a data center operation. It is IBM's contention that EPA's proposal to further tighten these requirements add extra cost to the product without any attendant benefit for the data center operator.

Line 646: Sampling Requirements: IBM believes that EPA should continue to allow a reporting of a 30 second average, sampled on 1 second intervals, on 30 second intervals. A data center operator is not likely to be looking at their power usage on a frequency of less than 15 minutes, so use of a rolling average or a 30 second average reported every 30 seconds does not materially impact the data center operator or the use of the information. EPA should continue the Tier 1 Sampling Requirements in Tier 2. Many, perhaps more scalable, high-level data collection tools cannot sample faster than 30 seconds and are often configured to sample more slowly.

Line 796: Effective Date: IBM recommends that EPA be willing to extend the effective date beyond October 2010 if this provides sufficient time to enable the development, testing, and verification of an active mode metric as described in pages 1 to 3 of this document. As discussed on pages 1 and 2, IBM prefers that the Tier 2 active power metric be established for a period of 2 to 3 years after its effective date to allow manufacturers to work to the specification requirements. Frequent adjustments of specifications which affect products that have a 2 to 3 year design cycle are not productive for EPA, customers, or the industry.

Appendix A:

Proposed additions to the ENERGY STAR Computer Server Requirements to include blade servers and systems:

**1. Definitions:**

1. D Modify the blade server definition to include the following sentence: Individual blade servers will be tested in a test fixture to determine the blade server power use (See Appendix A.1.xx) (Section 3.D)
.
Definition of a Maximum Configuration, Minimum Configuration, and Typical Configuration (page 8, Items 1.W, X, & Y): Add a sentence: For a blade server system, the (definition level) configuration definition will be applied to the individual blade servers.

Add the following definitions:

Maximum Configuration; Blade Chassis: A Blade Chassis fully populated with blades with the maximum configuration for that product or product family.

Minimum Configuration; Blade Chassis: A Blade Chassis fully populated with blades with the minimum configuration for that product family.

Typical Configuration; Blade Chassis: A Blade Chassis fully populated with blades at the typical configuration for that product or product family.

Multi-wide or Tall Blade Server: For any blade which is larger than one bay, then the size represented by that unit will be equal to the number of bays occupied. If there are multi-wide (or tall) blades in a chassis, the capacity of the chassis is reduce by the number of bays occupied by the multi-wide/tall blade minus one. (i.e. a 3 wide blade will reduce the chassis capacity by 2 slots).

**2. Qualifying Products**:

Add blade systems and servers to the Qualifying Product list as follows. "The Tier 1 specification coverage is limited to Computer Servers having at most four processor sockets (i.e. Computer Servers with 1-4 individual processor sockets) ***and blade systems***.

Remove the blade system and server bullet from the list of excluded products.

Remove Blade Systems from Tier 2 coverage discussion (page 8).

**3. Efficiency Requirements for Qualifying Products**

Make the follow adjustments to Item 3.B.2

p.11 Item 3.B.2: Computer Servers with Greater than Two Processor Sockets (3S & 4S) ***and Blade Systems***

"All three and four socket ***and Blade System*** Computer Servers must enable…"

**3. C. Standard Information Reporting Requirements**:

p. 12 Modify item(s) and add an item to the list "Each Power and Performance Data Sheet must include the following information:"

Modify: Item 3 "System configuration(s) ***for computer servers with four processors or less and for blade server(s)*** (including maximum…)"

Add an item: For blade systems, report the maximum and idle power for a blade chassis fully populated with the minimum blade server configuration.

**3.D: Data Measurement and Output Requirements:**

First sentence: modify to: "One and two socket (1S and 2S) Computer Servers which meet the definition of a Managed Server in Section 1.G, ~~and~~ all Computer Servers with greater than three or four processor sockets (3S and 4S)***, and blade servers*** must have the ability…"

Insert a sentence in the first paragraph, after the first sentence and before the second (or last) sentence:

"For blade systems, the inlet air temperature can be reported for the blade chassis."

**4.B Idle and Full Load Power Testing:**

Last sentence of the first paragraph should be modified as follows: "…for all Computer Servers, including three socket (3S)***,*** ~~and~~ four socket (4S)***, blade server, and a blade chassis fully populated with the minimum configuration of the blade server."***

**4.D Qualifying Configurations and Families Under this Specification**:

Add the following to the third bullet: "For a blade chassis, report the idle and maximum power for a blade chassis fully populated with the minimum configuration blade server as defined in 1.xx."

Add a final paragraph to this section:

"For blade systems, a product or product family will consist of a model of blade server and one or more supporting chassis. A blade system purchased by a customer will be qualified based on the following requirements:

a. Has a chassis which supports a qualified product or product family.
b. Is populated with blade servers from qualified products. An ENERGY STAR® qualified blade system could be populated with blade servers from one or more qualified products or product families.
c. At least seventy percent of the installed blade servers in a chassis must be from an ENERGY STAR® qualified product or product family. The remaining purchased blades can be support blades (such as network or storage blades) or specialty blade servers which are not from a qualified product or product family."

**5. Effective Date:**

Provide a bullet for the effective data of the Computer Server Requirements v1.

**APPENDIX A:**

I. Definitions:

Modify the UUT definition by adding the following sentence: For blade servers, a UUT is an individual blade server, which must be tested individually in a test fixture as required in Section xx.

Add a definition for: CUT: "CUT is an acronym for "Chassis Under Test" which in this case refers to the blade chassis fully populated with a minimum configuration (1.??) blade server."

Blade Server Test Fixture: A blade server test fixture may be one of the following:

a. A specially designed fixture which provides power and network connectivity allowing the blade server to be operated independent of the blade chassis. The power demand for the test fixture must be quantified and reported as part of the testing report.
b. A single blade server can be tested in a blade chassis, with the power use reported through the Standardized Data Measurement capability of the server.

II. Testing Requirements:

Required Power Analyzer Attributes:

Add the following statement at the end of the bullets:

"…Calibration must be current and within the past year."

Or

"The power must be measured with the power measurement capability provided as part of the "Data Measurement and Output Requirements" (Section 3.D) In this case, the measurement accuracy requirements for the "Data Measurement and Output Requirements" (Section 3.D) will be the requirement for the test."

Test Configuration

Modify the first sentence under "Test Configuration" as follows"

"Power consumption of the UUT or CUT shall be measured and tested…"

Modify the first sentence of the next paragraph as follows:

"The UUT or CUT must have at lease one port connected to an Ethernet network connection…"

III. Test Procedure for all Computer Server Products

Add a second sentence to the introduction, which states the following: For blade systems, the Test Procedure should be applied to a minimum, maximum and typical blade server configuration tested in a fixture and to a blade chassis fully populated with minimally configured blade servers.

A. UUT Preparation for All Computer Server Products

Modify Item 2 as follows: Ensure that the UUT ***or CUT*** is connected to a live Ethernet connection…

Modify Item 4 as follows: Plug the UUT ***or CUT*** into the measurement power …

Modify Item 4.b as follows:

UUTs ***or CUTs*** with multiple power supplies… and

…included in the measurement of idle power for the UUT ***or the CUT.***

B. Measuring Full Load and Idle Power

1. For all instances of UUT in this section, change it to UUT(s). This is to indicate that for a blade system, all of the UUTs in the blade chassis need to be managed according to the test procedure.