IBM appreciates the opportunity to continue to work with the EPA ENERGY STAR® program to develop the Version 2 requirements for the computer server requirements. EPA has made significant progress on refining the Version 2 requirements and IBM believes that the Version 2 requirements, as proposed in Draft 2, are near final form.

EPA's proposal to continue to require power management enablement as the power criteria for 4 socket processor systems and the choice to initiate blade systems into the ENERGY STAR program, while collecting data on SERT power/performance metrics to inform an active energy metric for Version 3, is an appropriate and workable next step for the ENERGY STAR V2.0 requirements. The complexity and range of configurations for 4 processor socket and blade servers make it difficult to set simplified requirements for these systems. It is important that EPA also refine the Resilient Server definition and provide additional adders for these resilient and scalable systems in the 1 and 2 socket systems to address the impact of the additional infrastructure and redundant components required to achieve resiliency and scalability in one and two socket systems. There are significant differences in the power profile between a managed and a resilient and scalable server and it is appropriate to address those differences in Version 2 of the requirements, both to address current inequities in the evaluation of 1 and 2 socket systems and to set the stage for appropriate evaluation of the SERT metrics across the range of server configurations.

IBM offers the following specific comments and recommendations with regard to the ENERGY STAR® Program Requirements Product Specification for Computer Servers: Eligibility Criteria Draft 2 Version 2.0.

DEFINITIONS:

EPA proposal to remove the ECC/Buffered Memory requirements for systems with more than 50 nodes (lines 23 to 33): ECC/Buffered memory is an integral aspect of an enterprise level computer server and the ECC/Buffered memory function provides important functional capability and affects the power profile of an enterprise server. EPA should not remove the ECC/Buffered memory requirement from the server definition and should not recognize server systems without ECC/Buffered memory

Resilient Server Definition (Lines 70-87): IBM has done an extensive amount of work both internally and in conjunction with the Green Grid ENERGY STAR server working group to refine the resilient server (hereafter referred as a resilient and scalable server) definition and establish a clear distinction between server systems with lower power/functionality processors and limited resiliency and scalability and server systems that are highly resilient and scalable.

IBM has analyzed its server products and determined that it is important to include the resilient and scalable server category, and thereby avoid excluding of higher power servers which provide resiliency and scalability characteristics which are valued for mission critical computing activities, including highly virtualized computing environments and computationally intensive computing applications. The data collected

by EPA for 2 and 4 socket systems (see data in Appendix A) indicates that servers designed for higher Reliability, Availability, and Serviceability (RAS) and to support higher levels of virtualization and connectivity also carry higher infrastructure costs that increase system power use. These infrastructure costs are sufficiently different between high volume managed servers and the proposed resilient and scalable server category to materially differentiate the power profile of the two server types and necessitate separate, distinct power criteria to assess the energy efficiency of the two server categories. The IBM proposal for a resilient and scalable server definition is detailed in the accompanying file "IBM Resilient and Scalable Server Definition Proposal V2D2 Comments.xls". The proposal sets out 6 major criteria, Processor RAS and Scalability, Memory RAS and Scalability, Power Supply RAS, Thermal Cooling RAS, System Resiliency, and System Scalability. With the exception of the Power Supply RAS and Thermal Cooling RAS, each criterion has a set of characteristics of which a server must have some minimum number to qualify as a Resilient and scalable and Scalable Server. The criteria and the corresponding qualification requirements are detailed in the attached file. IBM has assessed these criteria against systems with and without Resiliency and Scalability characteristics and validated that the criteria distinguishes between the two types of servers. See columns G and H of the spreadsheet file for the characteristics exhibited by each system type. Lower power, non-scalable/resilient systems will have some RAS and scalability features and it is expected that the resilient/scalable characteristics will need to be updated with each new version of the requirements as server system technologies and functionality evolve and advance.

IBM proposes that EPA add the resilient and scalable server category to table 3 (page 10 of draft 2) and provide a category specific Resilient and Scalable Server Base Idle Allowance to account for the additional infrastructure and higher processors socket power of this server type and provide additional adders for resilient and scalable components such as Raid 5 controllers and memory buffers. IBM is working to provide available test and vendor data to EPA by July 13, 2012 to provide a basis for setting a Base Idle Allowance and additional component adders for resilient and scalable systems.

High Performance Computing (Lines 101-104): IBM agrees that High Performance Computing (HPC) Systems should be separately defined, and where system characteristics are sufficiently different from a managed or resilient and scalable server, HPC systems should be excluded from the ENERGY STAR requirements. The IBM technical team will evaluate possible HPC definitions and is prepared to work with EPA and industry groups such as Green Grid to establish an appropriate HPC definition for the ENERGY STAR Computer Server Requirements.

Addition of a "Large Server System": A "Large Server System" would be defined as a server product with 4 or less processor sockets where the total server system occupies more than 5 U of rack space and is designed to function as a mainframe server. IBM has a product which has 2 or 4 processor socket configurations where the minimum configuration occupies 19 U, with a Central Electronic Complex (CEC), an I/O drawer with 32 I/O adapters and Power systems and other support equipment. The maximum configuration will occupy 36 U, with two to four I/O drawers, one or more CEC units,

and various communications systems. The size of the system and the extensive, associated peripherals give these server systems a much larger power profile than a typical x86 (4) processor system. The systems have a much smaller sales volume than x86 based systems and are targeted at a specific, defined niche of the enterprise server market. For these reasons, it is not valid to make a comparison between these systems and the more traditional 4 processor systems which occupy 5 U or smaller enclosures.

This additional server type would serve as the basis for an exemption of this system type in Section 2.2: Excluded Products

Product Family (Lines 222-259): IBM is very supportive of the EPA's revised product family definition. By simplifying the definition and allowing a range of processor socket power and core count and PSU output power to be included within a product family, it enables a manufacturer to better group product family data, simplify the communication of ENERGY STAR® qualification to customers, reduce the quantity of testing required to qualify a product while accurately representing the range of power use and performance for a given product model line or machine type.

IBM wishes to clarify EPA's addition of a "Low-end Performance Configuration". Based on IBM's product configurations, IBM would define the system characteristics of each of the 4 corners of the product family as follows:

Minimum Power configuration: Minimum processor power, core count and populated socket count for the machine type.

Low-end Performance configuration: In order to properly bracket the minimum configuration/performance points of the 4 corner product family, our intent had been to test a minimally configured system with the maximum processor power, core count, and minimum populated socket count for the machine type. This may not directly fit the definition of a low performance system. We are also concerned that in some cases the Minimum Power and Low-end Performance Configuration would arguably be the same for a given machine type. We would propose that EPA change the statement "…represents the lowest-price or lower-performance computing platform..." to "…represents a lower-price or low-performance computing platform…". This provides companies more leeway in selecting the configuration to establish one of the two low power corners of the product family.

High-end Performance Configuration: In order to properly bracket the configuration/performance points of the 4 corner product family, our intent had been to configure this system with fully populated sockets with processors with the lowest socket power and core count. As with the Low-End Performance configuration, we request that EPA change the descriptive text for this configuration to "...represents a higher price or higher performance computing platform…". This provides companies more leeway in selecting the configuration to establish one of the two high power corners of the product family.

Maximum Power Configuration: This configuration would be a server with fully populated sockets with processors with the highest power socket power and core count, as well as a component configuration which will draw higher power than all or most other configurations available for the product family.

Given the definitions EPA has provided for the product family, even with the adjustments requested above, it is highly likely that some qualified configurations will exist outside of the power profile envelope defined by the "Product Family Tested Product Configurations". IBM recognizes that it is our responsibility to validate that all products that we market and sell as "ENERGY STAR qualified" meet the applicable requirements, but we also want to verify that EPA intends that qualified products can exist outside of the power profile defined by the 5 tested product configurations.

QUALIFICATION CRITERIA:

Power Supply Efficiency Criteria (Lines 290-294): Because EPA is maintaining the idle criteria for 1 and 2 processor socket systems and the power management criteria for 4 processor socket systems in the proposed Version 2, systems that have been qualified under version 1 with gold or platinum power supplies will continue to be qualified under Version 2. However, the power supplies were tested under Version 1 using Revision 6.6 of the test protocol. IBM asks that EPA explicitly allow gold and platinum power supplies which were qualified under version 6.5 of the testing protocol for server products to continue to be "grandfathered" under Version 2. Power supplies for new server systems qualified after the final publication of the Version 2 requirements would have to be qualified under version 6.6 of the testing protocol.

Documentation Requirements; Blade Server Shipping Documentation (lines 338-349) and Products that Do Not Include a pre-installed Operating System (lines 518-521): EPA has provided requirements in the referenced sections to provide documentation in the shipping product packaging for qualified blade servers and for systems which ship without an installed operating system. IBM requests that EPA allow companies to provide the required instructions and blade system documentation on their ENERGY STAR webpage and/or in their on-line product documentation to simplify the process of providing this documentation to our customers. Requiring documentation in the packaging is inappropriate for several reasons:

1. Enterprise server systems are installed in the data center by specialist technicians who are typically working against specific, defined instructions and installation protocols established by the data center operator. As such, they typically do not follow instructions provided with the products, they are not the individuals that are concerned about whether the installation is ENERGY STAR compliant, and oftentimes the software set-up is performed remotely over the network by a different technician. In-package documentation would not have any value.
2. Companies are moving to electronically available documentation for all of their products to reduce resource consumption, improve version management, and simplify availability to the customer. Requiring documentation in the packaging runs counter to this trend.

3. There are many challenges to managing fulfillment systems to insure that ENERGY STAR specific product documentation is added to ENERGY STAR qualified purchases, particularly where a given qualified product family is a subset of a machine type or model. It is more efficient and effective to meet these documentation requirements through an electronic documentation system.

IBM requests that EPA allow companies to satisfy the documentation requirements through their electronic document systems and/or links on their ENERGY STAR web pages.

System Qualification with network attached HDD (Lines 122/123 of the testing procedure): EPA has allowed systems to be qualified where they depend on remote storage for boot-up and storage capabilities. IBM does not object to the ability to qualify this type of systems, but EPA needs to modify Table 3 to include systems which require remote access storage. The Base Idle Power allowance for 1 and 2 socket systems depending on remote storage needs to be reduced by the 8 W HDD power adder.

Power Supply Adder (lines 407-420): IBM is evaluating its power supply data to offer a methodology to EPA to calculate power supply adders and to propose an appropriate power supply adder to EPA. This information will be provided to EPA by July 13, 2012.

Idle Mode and Full Load Efficiency Criteria – Blade Servers (lines 429-433): EPA needs to include language in this section clarifying that the reported idle mode and full mode values should be calculated by dividing the chassis and blade power measurement for the fully and half populated chassis at idle and full load by the number of blades populating the chassis and reporting the lowest value of the full and half chassis measurements. In addition, EPA needs to clarify that the power measurements for the full and half populated chassis at idle and full mode need to be reported under section 4.

STANDARD INFORMATION REPORTING REQUIREMENTS

Power Data for Full and Idle Load (lines 482-483): IBM recommends that the Idle and Full Load be reported as measured by SERT. The SERT system records and reports the idle and maximum power values measured from the full SERT test process.

Power and performance benchmark reporting (lines 484-485): Given that companies are required to test for and provide SERT worklet metrics to EPA under the Version 2 requirements, IBM recommends that the requirement for testing an additional power/performance benchmark be removed from this section. We would propose one of two options for SERT Reporting:

a. As currently proposed, all the SERT worklet metrics for each tested configuration have to be reported. These worklet scores should be allowed to provide performance information for the systems and 4.1.2.vi should be removed from the requirements.

b. Companies are concerned that because SERT worklets are new and companies and data center operators are not familiar with the metrics it will be counter productive to publish the full benchmark set until a broader dataset can be collected and analyzed to determine how best to assess server energy efficiency from the metrics. To this end, we are proposing separately that the full set of SERT metrics not be published. To satisfy the benchmark reporting requirement 4.1.2.vi, IBM proposes that companies publish the test results from the ccsj and flood worklets to provide an indication of the system performance capabilities. Companies should also be required to have the SERT worklet data available to provide to customers during the Request for Quotation or Request for Information process.

EPA has not supplied the acceptable "list of power-performance benchmarks" (line 485) for additional benchmark testing. Without seeing the list, we cannot determine if the request has any validity to it, given our comments above. If EPA decides separately to maintain the SERT worklet scores in a blind database, companies should be allowed to report the ccsj and flood worklets as one of the options in the "EPA list of power-performance benchmarks". IBM strongly submits that companies should not be required to perform multiple instances of power-performance testing under the ENERGY STAR requirements.

Identification of Server Type: Companies should be required to report whether the computer server is unmanaged, managed, or resilient/scalable server. This should be added to the requirements in this section.

Data from required Active State Efficiency Criteria testing (lines 481, Section 4.1.2.iv): IBM encourages EPA to collect the SERT worklet data for qualified products in a blinded, public database, with the exception noted above for reporting of the ccsj and flood worklets as the alternative to or an option within the additional power-performance benchmark data required under 4.1.2.vi. Because the SERT metric is new, the relative and absolute value of the worklets have not been determined, and there is a high degree of interest in the stakeholder community to evaluate and assess the SERT worklets for a variety of purposes, IBM believes that it is best to blind the initial data set to prevent attempts to compare manufacturers systems and/or to establish and assess identified product performance against a single metric before the data set has been adequately vetted and evaluated by EPA, SPEC, and other interested and informed stakeholders.

Input Power, Processor Utilization, and Air Temperature Sampling Requirements (lines 542-547): IBM requests that EPA provide an additional option to the proposed reporting on a 10 second frequency: allow the system to report data up to every 30 seconds where the data is provided with a time stamp. By providing a time stamp on the power reported power data, the data collection/analysis system can match up multiple readings to get a consistent view of the data center power profile while enabling the collection system to poll on a less frequent basis.

Qualifying a Product Family which is a subset of a Machine Type or Model configuration set (lines 578-583): IBM potentially has concerns with the stated requirements for managing qualified configuration(s) which are a subset of all the configurations in a Machine Type or Model group. IBM has recognized this difficulty, both for the server and storage product categories, and has been evaluating options for managing, reporting and identifying subsets of product configurations. Currently, we have not identified an efficient means to manage qualified products in this situation. IBM will provide EPA additional comments on this topic by July 13, 2012.

c. APPENDIX A: ANALYSIS OF ENERGY STAR 2 and 4 PROCESSOR SOCKET SYSTEMS: IDLE AND MAXIMUM POWER AND IDLE TO MAXIMUM POWER RATIO.

An analysis of ENERGY STAR qualified 4 socket systems shows a wide range in the full load and idle power of the systems. The system power profile is affected by the processor type (lower power versus higher power) and the extent to which the processor and the server system support resiliency and scalability capabilities. While there are a wide range of power profiles for the 4 socket systems, Figure 3 shows that power management capabilities can reduce the power use during idle periods as effectively for high power systems as for low power systems.

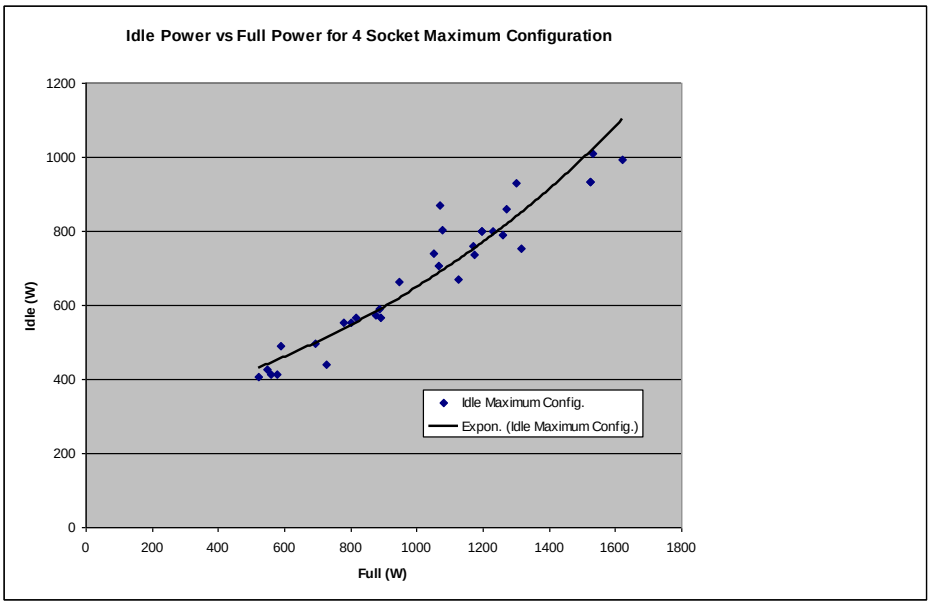

Figure 1: Comparison of Idle and Full Power for Four Processor Socket Systems, Minimum Configuration.

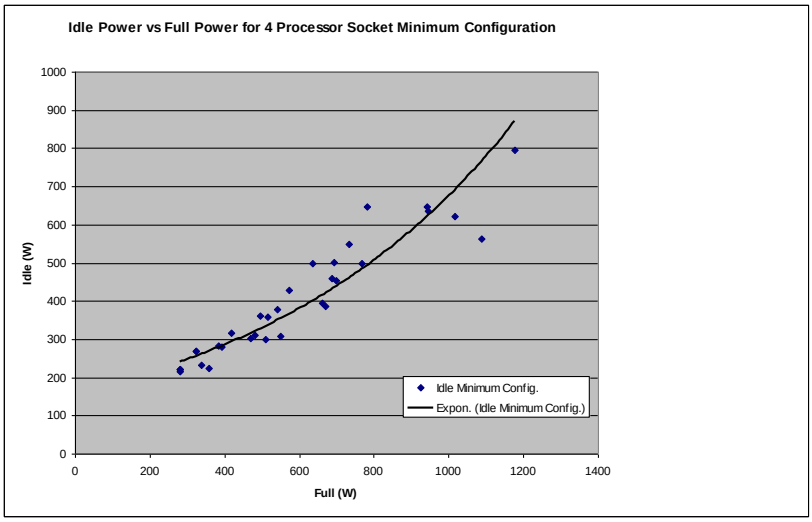

Figure 2: Comparison of Idle and Full Power for Four Processor Socket Systems, Maximum Configuration.

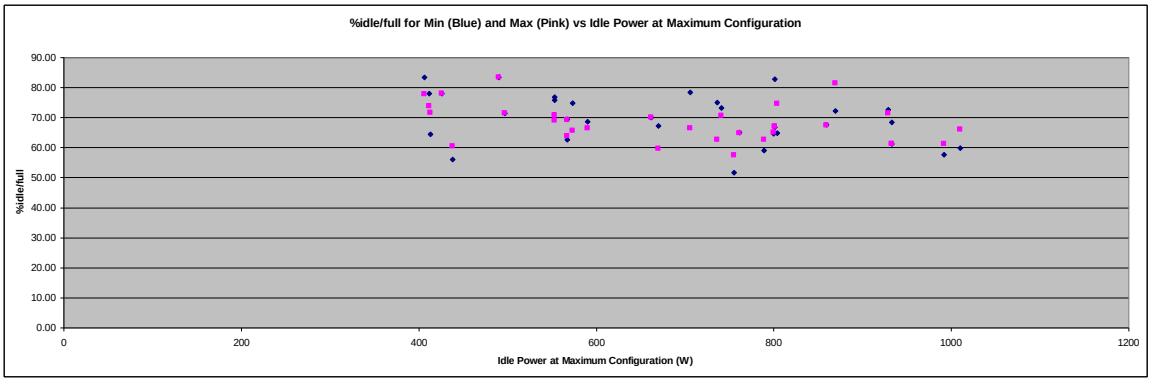

Figure 3: % Idle to Maximum Power for (4) Socket Maximum Configuration systems

2 Processor Socket Systems

Figure 4 below graphs the idle and maximum power of the current ENERGY STAR qualified two socket servers. The data shows that, similar to 4 socket systems, there is a significant range of power use across the ENERGY STAR inventory of qualified 2 processor socket computer server systems. It should also be noted that it there is limited data for resilient and scalable servers in the current ENERGY STAR 1 and 2 processor socket datasets given the difficulty in qualifying these servers to the current idle criteria:

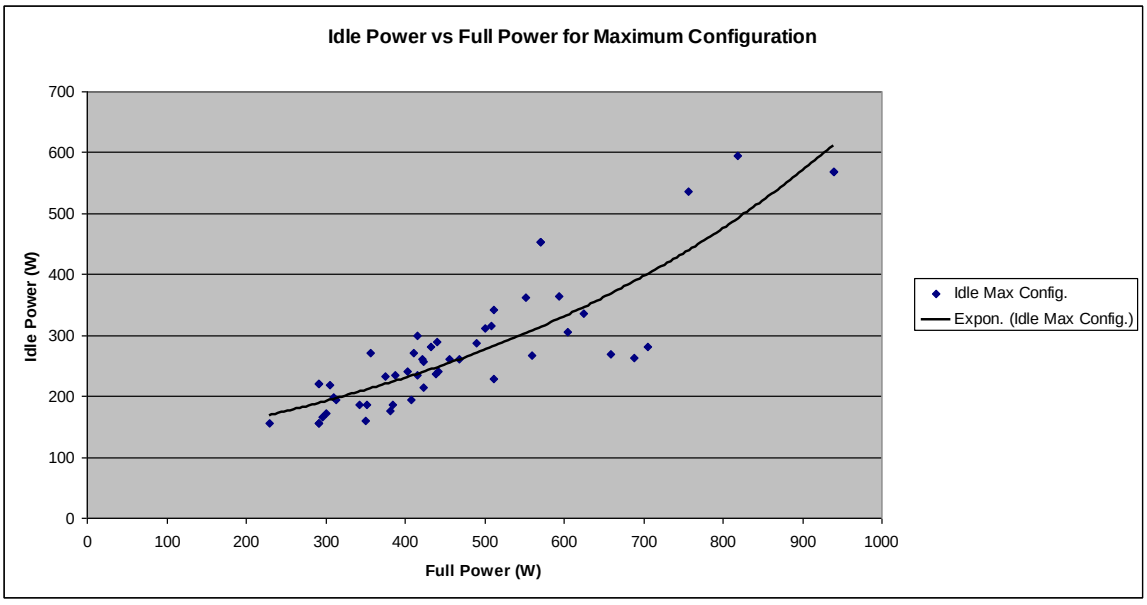

Figure 4: (2) Processor Socket Maximum Configurations: Idle to Full Power

The majority of the 2 processor socket products qualified have minimum configuration full power values of 400 W or less. Higher power systems are difficult to qualify to the idle criteria. However, analysis of the ratio of idle power to the maximum power, the data

shows that higher power products are competitive with lower power products – they can reduce their energy use at idle by equivalent or better percentage as compared to the low power systems. This is depicted in the graphic below:

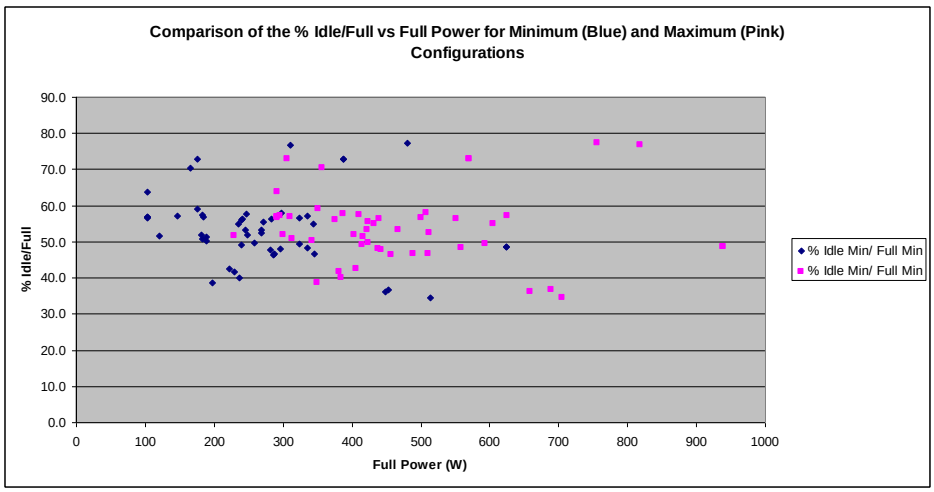

Figure 5: Comparison of Idle/Max Power Ratio for 2 Socket Servers