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. Continued efforts to improve the requirements through changes to the product family process to reduce the testing burden while providing the appropriate range of data on a selected, defined number of server configurations and incorporate requirements for blade servers will strengthen the requirements and maintain ENERGY STAR as the premier brand for energy efficient products. IBM offers the following comments and recommendations with regards to the Energy Server Use Evaluation – Resources and Discussion Documents released by EPA for comment on February 23, 2011.

DEFINITIONS:

1.A: Computer Server: EPA should retain the provision for ECC memory as one required attribute of a computer server. In evaluating the criteria for Enterprise Level Servers, we found that all such servers employed ECC memory, that ECC memory was a distinguishing characteristic of an Enterprise Server, and that the ECC memory contributed significantly to the power profile of the Enterprise server system. Those computer servers that do not provide ECC memory would be classified as "small scale servers" and covered by the ENERGY STAR® Computer Requirements. In addition, the ECC memory referenced in the "Resilient Server" definition (1.B.4) is specifically extended ECC memory, which is incorporate mirrored memory and DRAM chip sparing.

Section 1.B; Add an additional Computer Server Type (1.B.9): The "Large Server System": This 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, 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

Computer Idle definition: (p4 lines 191/92): The inclusion of language specifying the S0 state as the only acceptable idle state for an APCI system is too restrictive and will unnecessarily innovation in power management. Work is underway to use other APCI states into enable idle savings which take into account availability and responsiveness concerns which would limit or preclude a customer's use of the power management settings. The objective of the ENERGY STAR should be to encourage power

management functions which are effective even for high demand or high availability systems.

(page 5 line 220-221) Server Processor Utilization: We recommend two modifications to this definition.
1. The definition currently states "the ratio of instantaneous processor computing activity". We propose that the term "instantaneous" be replaced with "short term average of use over a set of active and/or idle cycles". The use of instantaneous does not accurately depict how the utilization data is collected and reported by the processor systems.
2. Change the final qualifier to "…full load processor computing activity at ***the maximum*** voltage and frequency for that processor." Comparing the current activity level to the maximum activity level is important, as a data center operator is interested in the available capacity on a given system over time in order to understand opportunities for virtualization or improved productivity in the total data center.

This proposed revisions result in the following proposed definition:

Server Processor Utilization: The ratio of the short term average of processor use over a set of active and/or idle cycles

IBM reiterates its position that the utilization requirement continue to be qualitative in nature. We provided these comments to the final draft of version 1 and continue them to be relevant to:

1. Each processor, system and operating system or hypervisor use slightly different technical techniques to estimate or quantify utilization, making it difficult to propose a definition that would suggest that the reporting should be provided on a quantitative basis.
2. 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, so it is important to enable the manufacturers to provide the data over a period of cycles, rather than on an "instantaneous" basis.
3. For the purpose of the Tier 1 ENERGY STAR® computer server requirements, the measurement only needs to be sufficiently accurate for the purposes of enabling decisions around the consolidation or reprovisioning of workloads or the migration of virtual machines. Many data centers today use the current non-instantaneous methodologies to dynamically manage data center power consumption by reprovisioning workloads to minimize the underutilization of machines.

1.H.3: Product Family Definition:

1. Common Product Family Attribute: Products should be the same model line ***or machine type***. For some manufacturers the model line is a sub-group of the machine type.

2. For PSU selections, the product family definition should allow for several PSU outputs. Manufacturers are offering several output choices on a given model line or machine type to enable customers to best match the power supply capacity to their chosen configuration and intended use of the server system. The lowest output power supply should be designated to be used for the two minimum configurations of each model line or machine type and the maximum output power supply should be designated for use with the two maximum configurations of the system.

It appears from the discussion on lines 218 to 236 of the "Server Energy Use Evaluation – Resources" document, that EPA intended to allow a product family to offer several power supplies with different outputs to enable the customer to "right size" the power supply for their chosen configuration and intended use, as the 5 testing configurations offer the "combination of PSUs, Memory, Storage (HDD/SDD) and I/O devices…

4. Add a third Common Product Family Attribute: The processors must be from a single, defined processor series or plug into a common socket.

5. Lines 218-236: Add "processor socket power" to the stated combination of server attributes that should make up the configuration. An example would be:

Low-end Performance Configuration: The combination of ***processor socket power***, PSUs, Memory, Storage….

This addition clearly dictates that the manufacturer should consider the processor socket power in evaluating its choice of testing configurations.

Modifications to the Product Family Definition (Server Dataset Discussion Document):

IBM wishes to state its strong support for the changes which EPA has proposed to the product family definition. By simplifying the definition and allowing a range of processor socket power and PSU output power (see comment above) 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.

We encourage EPA to issue the revised product family definition as a V1.2 release, rather than wait to the finalization of the V2.0 requirements.

BLADE SYSTEM and TESTING REQUIREMENTS (Lines 571-632)

IBM continues to believe that the most logical way to measure power and performance for blade systems is to test a full blade chassis. This allows for the optimum distribution of the chassis overhead and the most accurate representation of the per blade power use in a chassis. Version 2 should provide manufactures the option to test a fully or half populated chassis to qualify a blade server system to the ENERGY STAR® requirements. The per blade power use should be calculated by taking the total system power use divided by the number of blades in the chassis. If EPA chooses to set idle requirements for blade servers, IBM recommends that the idle power requirements be set based on this "fully burdened" blade power use – the power use measured by dividing the total system power used by the number of blades in the chassis.

We believe that testing a full chassis does not represent any specific financial or procedural hardship if the testing is allowed to cover a product family. A manufacturer should be allowed to choose to test a 50% populated blade chassis to compute the "fully burdened" blade power use, as the result should be conservative due to distribution of the chassis overhead over less than the full number of blades.

We want to reiterate the comment we have made on blade criteria in previous comments to EPA, as we feel the comment continues to be relevant. Because of the complexity of blade systems, IBM continues to recommend that EPA not adopt an idle limit for blade systems. Rather, it is appropriate to implement the requirement for 3/4 processor socket systems for processor level power management enabled on shipped blade products.

Test Data for Blade Systems: IBM is not able to provide any test data for blade systems at this time, as the systems are not available for testing. IBM intends to assess the accuracy of the power calculator estimates for power use compared to available test data for 1,2, and 4 processor systems. If a consistent calculator bias can be determined (blade system power is calculated on the same power calculator as the IBM System x™ rack servers) for these systems, IBM will use the power calculator to provide data on blade system power use.

BLADE CHASSIS REQUIREMENTS:

EPA should not set power requirements for a blade chassis, as each manufacturer distributes overhead differently across its blade chassis. Rather, the blade chassis requirements should be based on the power supply efficiency and functional capabilities of the blade chassis. IBM proposes the following requirements for a blade chassis.

1. PSUs in the blade chassis should meet the computer server efficiency and power factor requirements.
2. The chassis should have variable speed fans.
3. The chassis should be capable of reporting power use and thermal information for the blade system.

This establishes the functional requirements for the chassis, without requiring extensive measurements of power use on the chassis. Power use measurements of a chassis are not

valid for comparison purposes, as each manufacturer configures their blade chassis differently, with different percentages of the "overhead" power for chassis service processor, fans, network, and hard drive components.

In addition, different chassis are configured to support different kinds of workloads. IBM manufactures one blade chassis with dedicated DASD bays where access to storage is important, while another blade chassis has extra I/O slots to support interaction of the blade servers into the network. Each type of chassis will have separate power signature driven by the intended use of the chassis, making comparisons between chassis' difficult.

IBM recommends that EPA require that blade systems be qualified as a combination of a blade server (as designated by a machine type or model line) and a blade chassis. One blade machine type or model line could be qualified with more than one chassis. By qualifying based on a chassis/blade machine type combination, and reporting or setting the idle criteria for a fully burdened blade (as discussed above), EPA will integrate the blade and chassis characteristics into the qualification process. Once a machine type is qualified with a specified chassis, manufacturers should be able to sell ENERGY STAR qualified blade system with any of the qualified blade machine types loaded into a specific, qualified chassis. As an example, Blade Machine Type A and Type B could both be qualified to the ENERGY STAR requirements on Chassis x. Manufacturers should be able to sell the "qualified" blade chassis x populated with blade machine types A and B as an ENERGY STAR qualified blade server system.

Idle and Active Power Test Data for 1, 2, and 4 processor socket systems: The best available data that IBM has for its 1, 2, and 4 processor socket systems has been provided to EPA as part of the ENERGY STAR® qualification process. IBM has provided 54 measurements for 4 processor systems (different configurations with 1,2,3 & 4 sockets populated) and 111 configuration measurements for 2 socket systems (primarily for systems with 2 sockets populated). IBM has not provided data on 1 processor socket systems, as the idle limits are too low to enable IBM to qualify its products to the current requirements. As discussed for blade systems, IBM intends to assess the accuracy of the power calculator estimates for power use compared to available test data for the 2, and 4 processor systems. If the power calculator can be demonstrated to have a consistent bias in its power calculations, IBM will provide power data on additional 1,2, and 4 processor socket machine types to EPA in support of its data request.