Thank you for the opportunity to provide comments on the "Energy Star® Program Requirements for Computer Servers Draft 1: Partner Commitments". IBM continues to be committed to working with USEPA and the other stakeholders on developing a specification or specifications for Enterprise Computer Servers and Storage equipment. As has been discussed in previous comments, the complexity and range of capabilities and configurations in the "Enterprise Computer Server" product category necessitate some differentiation between types of server products similar to the differentiation that was done for different classes of computer products when considering power/performance benchmarks. However, IBM believes that a Tier 1 specification can be crafted which utilizes a list of power management capabilities to distinguish server systems that manage power use as workload changes, enable virtualization to deliver more work for each unit of energy used, utilize efficient power supplies, and monitor power use and thermal output to improve power management in the data center. Properly crafted, it may be successfully applied to the majority of server categories listed in the program draft to distinguish systems which maximize the computing power delivered for each watt of power consumed.

1. DEFINITIONS:

Server Definition: We recommend that the server definition provided by EPA be modified as proposed below:

A. Computer Server: A computer that provides various processing, storage, and communication services in response to requests that generally originate and/or are mediated by other client computers and server computer systems. Computer servers generally have the following characteristics:

- Reliability, Availability, Serviceability, and Manageability (RASM) features
- Designed to run Server Operating Systems and/or Hypervisors with Operating Systems (Supervisor code).
- Include some type of network communication capability for inter-node compute and storage communication.
- Designed to operate in a commercial data center environment.
- Dedicated management controller, such as Baseboard Management Controller (BMC) or service processor capable of running on auxiliary power (server itself is powered down).

We recommend that the EPA remove its other characteristics for the following reasons:

- Certified Operating System: This is not a universal requirement for server operating systems. For example, Linux does not certify its operating system for servers.
- Designed for dual processor: Blade servers may have a single processor, so there should not be a lower limit on the number of processors.
- Error Correction/memory buffering and EN?? and EMC requirements: These definitions are too restrictive and may exclude equipment that should be included.

Adjust Blade Server Definition: We recommend that the EPA remove hard drives from the definition. Many blade and non blade servers are available which can boot up from a remote, networked hard drive or storage system and save power in the process. Hard drive should be replaced with "I/O connectivity" as an attribute.

Complex Sub-Types of Servers which should be excluded from the specifications: Clustered Servers, sysplexes, DC-powered servers, super computers, and storage systems should all be excluded from any ENERGY STAR® server specification, as they are specialty servers or networked servers which will not fit within a standard server classification or specification and expose unique complexity that would further delay EPA guidance to the computer industry.

Some adjustments should be made to the Server Categories: For the internal disk item, the line should be changed to "Has the capability to handle up to x internal disks or solid state drives (or 16 or more disks or solid state drives if the server is so configured)." There are servers on the market which do not have any internal storage, instead utilizing networked storage. Also, equipment with solid state drives are being introduced into the market.

> **Server Categories:** Overall, IBM continues to believe that there is value to establishing these three categories of servers, especially as Tier II discussions begin.

Use of workload as an indicator for server type: IBM does not support the use of workload as a means to classify servers. With advances in multi-core processor technologies, blade systems, and virtualization capabilities, IT equipment is increasingly able to run multiple workloads of varying types and at times servers provide critical storage and network functions. As energy efficiency efforts progress to consolidate workloads across the network to maximize the workload on the minimum set of IT equipment and to affect the maximum set of said equipment to enter sleep or hibernate states, workload will become more mobile and less machine dependent. Systems may continue to be optimized for particular types of workload, but they will increasingly be called upon to run varying types of workload to maximize utilization at a minimal energy cost.

Definition for a Storage System: A system designed to provide data storage capacity, data storage networking, and/or data management capabilities such as data compression and de-duplication. While it may contain an embedded processor, this processor is not generally made available to run general-purpose applications but which may execute data specific applications, e.g. backup utilities, data compression, install agents, and other tasks.

Inclusion of Storage and Network Systems in the Tier 1 specifications: IBM supports the development of an ENERGY STAR® specification for Storage Systems, but recommends that it not be included with this specification. Rather, we suggest that EPA consider beginning discussions on a Tier 1 specification for storage 1 to 2 quarters after the release of the Tier 1 server specification. This specification could build on the work

done on the Server specification, using it as a model for the criteria for an ENERGY STAR storage specification. EPA should continue its efforts to engage with Storage and Networking Industry Association on the development of a Tier 1 specification for storage.

Computer Server Power Supply: IBM suggests that EPA remove "lower voltage" from the first sentence of the "Computer Server Power Supply" definition so that it reads: "A component designed to convert AC voltage to DC voltage(s) for the purpose of powering the server." This is appropriate, as some high end servers convert to higher voltage DC, before beginning the step down process (the higher voltage makes for more efficient, lower loss internal power distribution). The focus should be on the inefficiencies of any conversions, to minimize those conversions and not focus on a particular conversion.

Idle State: We suggest the removal of the phrase "and other software" from the definition. Other software is covered by the statement "...and activity is limited to those basic applications that the system starts by default".

2. QUALIFYING PRODUCTS

Products which meet the definition of a blade server or a computer server as defined in Section 1 for ENERGY STAR. Only those computer types which are specifically excluded in Section 1 should be ineligible for ENERGY STAR Tier 1 Server Specification.

POWER SUPPLY EFFICIENCY REQUIREMENTS

3.A Power Supply Efficiency Requirements:

IBM believes that the EPA should look to both Climate Savers and the 80+ power supply efficiency program for guidance on the potential power supply efficiency requirements. As currently proposed, the Climate Savers specification is limited by its focus on 1U/2U single and dual socket or high volume servers. While these servers represent a significant portion of the number of servers sold annual, their characteristics and development times differ from those of blade, medium and large servers as categorized on page 4 of the "ENERGY STAR Program Requirements for Computer Servers: Draft 1".

IBM would propose that EPA use the 80 PLUS Servers Bronze designation[1] (without inclusion of the 10% operating point, see comment below) as the power supply efficiency requirement upon implementation of the specification and the 80 PLUS Servers Silver designation with an effective date 2 years after Tier 1 implementation date. The 80

[1] "Server Research Project, Efficient Power Supplies for Data Center and Enterprise Servers", February 2008, ECOS and EPRI, p.12.

PLUS specifications and testing protocols were developed considering the requirements and characteristics of medium and large servers (up to 7U, with the exception of provision for fan power, see comment below). As EPA considers the available power supply efficiency standards, it is important that any standard referenced by EPA be open to input from all relevant stakeholders. Specifically, IBM would support the following efficiency requirements[2]:

Implementation	Requirement	20% Load	50% Load	100% Load
At release	Efficiency	83%	87%	83%
	Power Factor	.80	.90	.90
2 yrs after release	Efficiency	85%	89%	85%
	Power Factor	.85	.90	.90

The purpose of providing a two year period before increasing the power supply efficiency requirement is appropriate for blade, medium and large servers, as they have development cycles of 24 to 36 months. They require longer lead times to implement the necessary improvements in power supply configurations to achieve higher efficiencies.

For a redundant, multi-volt output power supply, the efficiency requirements at the operating points should be 1% less than for a single volt output supply due to the presence of multiple O-ring devices.

IBM does not support the use of the 80 Plus Gold standard efficiency standards at this time, unless the requirement to include the fan power in the efficiency calculations is removed (see comment on the proposed testing protocol) because it is not clear that the technology exists to meet the 92% efficiency standard at 50% loading where fan power is included in the efficiency calculation.

Removal of the 10% Efficiency Power Supply Efficiency Point: The 10% power supply efficiency point should be removed from the specification. Even with redundant power supplies, the base power needed in blade, medium or large servers to maintain the base management controller or service processor and maintain the various components such as processors, memory, I/O systems and other in sleep mode – where those modes are available; is such that the system power draw is unlikely to drop below the 20% load point on the power supply.

Power Factor Requirements: IBM supports inclusion of power factor requirements in a power supply specification. At 20% load, the pf should be .85 or greater. Above 20%, the power factor requirements should be set at .9.

[2] ibid, p.12

Additional Power Supply Types: For the most part, power supplies can be classified as single volt or multi-voltage output power supplies. For mission critical servers, there is a third type of power supply in use – the fully redundant, multiple output power supply where all voltage conversions must be redundant". In this case, each regulator has a redundant partner to assure continuity of supply which reduces overall power supply efficiency by 2%. Our proposal would be to set the efficiency requirements for multi-volt power supplies at 1% below the requirements for single volt output supplies and the requirements for a fully redundant multi-volt power supply at 2% below the single volt output supply.

IDLE POWER

IBM understands and appreciates EPA's interest in reducing the power being drawn by a server when it is not doing any work. However, setting a criteria or standard for idle power is not the best means to drive improvements in the power a server draws when it is not doing work. The power drawn by a server when no workload is present is heavily dependent on the configuration: the number of processors, the quantity of memory, the number of active I/O points, etc. The types of server which will have the lowest power level when work is present are those which are least capable of running multiple operating systems and workload to drive higher utilization rates. Unlike a PC, there is no effective power limit or value that can be expected when a system is not doing work or has been put into sleep and hibernate because of the tremendous variability in form and function amongst server configurations. IBM is very concerned that a focus on idle power will drive the wrong behaviors and purchasing decisions. The most efficient server will be the one that maximizes its utilization and the workload that it delivers and has the capability to minimize its power use when no workload is present.

Instead, we recommend that EPA require that an ENERGY STAR server have, at a minimum, some form of processor level power/workload management. X86 processors manufactured by Intel and AMD and Power6 processors manufactured by IBM all have functions which promote power saving when a processor is not doing work although they must be enabled by the equipment manufacturers through firmware enablement or other forms of integration into the system. Below is a list of attributes which EPA should consider:

Processor Sleep;
Power Capping;
Dropping into lower static power modes as workload reduces;
Dynamic Voltage and Frequency Scaling;
Variable Speed Fan Control based on power or thermal readings;
Lower Power Memory States;
Lower Power I/O Interfaces; and
Rack level or processor level liquid cooling
Export power and utilization to System administrators for action.

To achieve an ENERGY STAR rating, EPA can require that a server have a specified number of these capabilities enabled. In order to set the 25% threshold, EPA can require all manufacturers of server equipment to provide a list of their current models and the attributes from the list which are enabled on their current equipment. EPA can then determine how many of these attributes - perhaps two or three - need to be enabled to establish the "25% deployment threshold". It is recommended that EPA work with groups like Green Grid and ITI to collect the data and ensure that the full universe of available products is analyzed to get an accurate representation of the current availability of these functions on the market.

An additional benefit of this approach is that the EPA can resurvey the marketplace two years after the ENERGY STAR server spec takes effect and determine if new attributes are available which can be added to the list, and get the information needed to reset the 25% criteria. This provides EPA a simple and effective way to improve equipment efficiency, drive lower power use when no work is being done, and allow manufacturers to innovate across their product lines. By surveying and utilizing the current list of power management technologies in the ENERGY STAR specification, EPA can encourage all manufacturers to utilize available power management functions.

STANDARDIZED DATA MEASUREMENTS

IBM is supportive of the concept of providing standard information on the capabilities of a server. We have several comments on the information proposal made on page 8 of the draft.

a. The information sheet should be for a specified configuration for a particular model type, not for the specific model configured by the customer. The maximum and minimum power use will be dependent on the chosen configuration, and a system to provide that data for each customer's configuration will be very difficult and costly to develop. Many manufacturers are providing on-line power calculators that customers can use to calculate power use of their specific system – this should be sufficient and manufacturers should be able to provide the web link to the calculator rather than provide specific power numbers. This will allow the customer to calculate the projected power use of their equipment – which arguably they need to do before they purchase the equipment to be sure it can be supported in their data center.
b. The other option would be to require companies to provide the maximum, typical, and minimum power use for that server model for a minimum and maximum configuration. This would allow customers to see the range of power use from that equipment model and provide the impetus needed to use the calculator described in (a) above.
c. SPECpower data should not be required information. Currently, the SPECpower procedure is only applicable to 1U/2U single and dual socket servers, a substantial but very incomplete subset of the total server market. It does not cover blades, medium

and large server machines. It is also representative of only a particular type of workload, which may be very different from the customer's workload. Designation of the SPECpower standard by EPA risks setting unreasonable and incorrect expectations on the part of customers. Rather, we recommend that power be reported as described in (b) above and that companies report the SPEC workload benchmark(s) applicable to the particular model of server. This will allow customers to see the typical power use range and workload performance for the configurations available for a given server model or type.

POWER AND TEMPERATURE MEASUREMENT REQUIREMENTS

Real Time Reporting: We believe that it is important that equipment be capable of reporting server power, thermal, and utilization measurements in a standard format and we support EPA's proposal to report AC power consumption, inlet air temperature and processor utilization in order to achieve an ENERGY STAR rating.

We do not support specifying a specific data management protocol, as many IT and equipment manufacturers are marketing or developing systems to manage this data across a full range of OEM equipment. Rather, EPA should require that the data be available in a format such as XML, CSV (comma separated values), or other open standard formats that can be accessed by industry standard data transfer and management protocols. The important thing is that the server level software can collect the data and make it available to higher level system management software for display and system management purposes.

POWER MANAGEMENT AND VIRTUALIZATION REQUIREMENTS

Power Management: See the discussion under Idle Power with regard to the power management requirements.

Virtualization: The ability of a server to support multiple partitions and images (virtualize) is a critical component of equipment energy efficiency. Studies have shown that virtualization can drive server utilization to 50% and higher, maximizing the work delivered per unit energy provided, driving better utilization of the equipment – reducing the number of servers required to do a specified amount of work and also the quantity of material and energy required to manufacture servers; and minimize the size of the physical facility (and hence reducing operating energy as well as the energy required to manufacture the building materials for the larger facility). While there are no currently available benchmarks or metrics to determine the efficiency of a server's virtualization capability, IBM is providing two suggestions for a virtualization metric below. To be useful for an ENERGY STAR specification, the metric details would have to be developed through an industry stakeholder group or appropriate industry association or consortium.

A. The processor is configured such that the hypervisor consumes less than "X"% of the total compute cycles if the machine is running fully utilized with 2 or more virtual images.

Or;

B. Run a single specified workload on a processor, say 250 transactions per minute, and measure the processor utilization. Then, partition the processor into 4 images or instances, run the identical workloads on the 4 images or instances, and record the utilization. The difference between the measured utilization for the 4 images or instances and 4 times the single application utilization would represent the overhead associated with the hypervisor or virtualization manager. A lower percentage utilization would represent a more efficient hypervisor.

For each proposed metric, EPA could work with appropriate industry groups to survey manufacturers to get the capability of their current models and use that sample to establish the metric point which distinguishes the top 25% of the category.

TIER 2 REQUIREMENTS

IBM believes that the Tier 1 spec, if crafted with the approach to power management described above, can effectively recognize systems with superior energy utilization capabilities and drive increased system efficiency over time. We believe that EPA's goal for a suite of power/performance benchmarks that will address all the relevant server categories will take time to achieve. IBM encourages EPA to convene a group of industry stakeholders and SPEC and other benchmarking organizations to develop a workplan to establish the appropriate workloads and power measurement procedures for each server category, prioritized by the order of relative energy use of each category.

TEST CRITERIA

We have one major concern about the 80 PLUS/Climate Savers testing protocol:

Power Supply Cooling Fans: The load from cooling fans drawing power directly from the power supply should be subtracted from the power capacity of the power supply and not included in the efficiency calculation. IBM makes a wide range of server and storage equipment with a diverse range of packaging characteristics. The design teams analyze the system configuration to determine the best way to minimize the number of fans, maximize the fan size and most efficiently cool the server. In some cases on our medium and large systems this includes chilled liquid or refrigeration on the processor book to maximize heat transfer and provide the optimum operating temperature for the processor. As a result, we may have a single fan cooling the power supply and the full server enclosure, a fan cooling the power supply and part of the server enclosure, or a dedicated power supply fan and rack specific fans.

Fan configurations are selected to deliver cooling in the most efficient manor to maximize the energy available to perform work. Penalizing the power supply efficiency measurement because a more efficient, overall system cooling strategy has been chosen is neither reasonable nor fair. It is possible to measure the fan load and subtract it out; this should be allowed in the testing procedure.

Voltage Testing: All testing should occur at 230 V. In addition, the testing procedure should allow testing at 50 Hz, as this is required for the European market. Most systems are tested at 50 Hz, 230 Volts so that we only have to perform a single test for both the EU and North American market.

Comments Regarding EPA Storage ENERGY STAR®

IBM supports extending the ENERGY STAR® program to Storage, and feels some of the currently proposed server criteria and concepts can be extended to Storage Systems. However, IBM recommends that a separate document be developed focusing specifically on Storage Systems due to the form and function differences between storage and server devices.

1) TAXONOMY

In looking at Storage, there is a wide range of offerings and technologies which must be considered. A non-inclusive list of categories would include: Disk based Storage, Tape based Storage, Storage Area Network (SAN) networking equipment, hybrid and value added offerings (combined disk/tape systems, SAN based virtualization, etc.), and Web 2.0/mass scale storage.

Where appropriate, consideration must be taken for scaling within each categories spanning from Entry level to Enterprise offerings though there are likely to be system attributes which are common across the range of systems. IBM believes a simple and effective taxonomy which considers the above issues can be created in a reasonable timeframe.

The EPA must also be aware that different companies can use a variety of technical approaches to architect a storage system to achieve a given storage function. This may lead to implementing key features either close to a storage device, or choosing to offer it away from individual storage devices. A prime example is Storage Virtualization. Some companies have chosen to integrate this capability directly into disk (or tape) based storage offerings, while other have chosen to place this value in the SAN. Each approach has its merits, and an ENERGY STAR® storage specification must be crafted to provide equal recognition under the criteria to any systems configuration.

2) PERFORMANCE MEASUREMENTS

Similar to the comments on the Server specification, IBM does not feel that an idle criterion offers the best methodology to encourage energy efficient storage systems. The best measure of energy efficiency for a storage system will take into consideration consumption, using a mixture of appropriate workloads. It should focus on the ability to optimize power usage with workload, either by reducing power usage as the quantity of work to be done is reduced or by directing the workload to the correct level of storage device, to optimize the work delivered per unit of energy applied.

Like server systems, storage system energy use at no workload is dependent on the system configuration and hardware type and there is no effective way to set a standard for power use at no workload that can effectively incentivize energy efficient system design. Rather, it will be the processing of engineering innovative approaches to power/workload management that will reduce storage power use when no workload is present. We

believe that a useful list of energy efficient attributes for storage, including customer beneficial energy/ workload performance indicators, can be developed in a reasonable timeframe.

3) Potential other features

IBM would propose that a list of potential energy saving features can be created, but they must be mapped against the Taxonomy. For example, it would not make any sense to require a SAN Network equipment to offer spin down disk capability when such devices do not normally contain disks. IBM would caution the EPA from placing feature requirements which might be a mismatch for a given product due not only to its basic nature, but also features which would likely be non-applicable given typical field deployments. Here the Taxonomy and scaling will play a key role. IBM would be interested in working with EPA and other industry stakeholders to develop a list of energy savings features.

In closing, IBM is supportive of extending ENERGY STAR® to Storage, and feels a reasonable separate Tier1 specification can be created. In addition, we feel that it should be done sequentially with the existing Server Tier 1 Specification to allow for better information transfer from the Server Specification. IBM proposes a timeframe of 1-2 quarters after the initial Server Tier 1 specification. Further, IBM strongly encourages the EPA to work with the industry representative SNIA in defining many of these key issues around Storage.