IBM Comments To EPA ENERGY STAR® Program Requirements for Computer Servers: Draft 3

IBM appreciates the opportunity to comment on Draft 3 of the EPA ENERGY STAR® Program Requirements for Computer Servers. IBM believes that the definitions, power supply efficiency requirements, standard information reporting proposal, and power and temperature measurement and reporting capability requirements are largely complete, requiring only minor modifications to be ready for publication in a final specification. Comments have been provided specific to these items to make improvements on the specification. Some of the comments are made in response to specific requests from EPA for comments on proposals on specific items within the proposed requirements.

IBM is very concerned about the proposals for the idle power criterion and the labeling criterion. IBM understands and supports EPA's desire to reduce server power use when there is no workload is present. However, as IBM has consistently communicated to EPA, maximizing data center energy efficiency and optimizing the delivered workload for each unit of power delivered to the data center is a function of both minimizing the power drawn by the server when no workload is present and maximizing the amount of time that a server is executing workload (maximizing server utilization). Draft 3 sets requirements for idle power, but in so doing, create disincentives for systems which enable virtualization and increase server utilization. IBM details its concerns for the proposed idle criterion on pages 6 to 8 below. In addition, the proposed idle criterion requires that ENERGY STAR® conformance be determined for each model configuration which creates insurmountable process difficulties for manufacturers because systems are sold through multiple channels outside of the manufacturer's direct control, data sheets would need to be created for each model type, and product labeling would have to be done as each model is shipped. This business fulfillment model is not practical and will inhibit conformance with the ENERGY STAR® requirements. IBM has proposed a model based product qualification scheme for the idle criterion on pages 8 and 9 below.

In addition, EPA has set the maximum idle power criteria for four processor systems using a non-representative data set. IBM's specific concerns are detailed on page 4 below. EPA needs to delay setting the idle power criteria for four processor systems until a representative data set is compiled and validated through additional Draft reviews.

IBM respectfully encourages the EPA to modify its proposals on these items to create a workable specification which will not only reduce power use for equipment when it is idled, but which will drive the equally important outcome of maximizing the use of installed equipment. This not only reduces the energy in the use phase, both for operation and from reduced space needs, but also by reducing the energy required to manufacture the equipment where one server does the work of 5, 10 or 100 individual servers.

IBM Comments to Draft 3 ENERGY STAR® Computer Server Requirements

Line 33: Labeling Requirement: EPA's proposal to require an ENERGY STAR label on qualified servers presents several difficulties.
a. IBM's computer server products are distributed through a wide variety of channels including all of the following.

- Direct sales by IBM's sales and account teams
- Direct sales through IBM's Internet web sites. Many of these sales have the form that the customer buys a base model plus options which he or she install or have installed by a third-party contractor hired by the purchaser.
- Sales through value-added retailers (VARs), who typically buy minimally configured servers and then configure them to their customers' requirements.
- Sales by system integrators who purchase and install hardware options in systems which they purchase or which their customers already have ordered through another channel.
- Sales by solution providers who deliver complete packages including properly configured hardware systems and applications to their clients.

The set of hardware options and vendors (CPU , memory, I/O, etc) make not only the set of option types complex, but the raw number of options and the variability of power consumption across those options are daunting. Because IBM often does not complete the final configuration that will be delivered to the customer, another entity other than IBM may control the server at the point at which ENERGY STAR qualification would be made and the label applied. This complicates both the placement (at what point in the delivery process does it get done) and the data contained on the label. This "ownership" of that determination by another entity complicates the labeling process.

b. As proposed, the idle criterion requires that each individual system configuration be assessed for ENERGY STAR requirements. Integrating this determination and the resulting labeling requirement into the fulfillment process will be very difficult and subject to uncertainty. This difficulty is exacerbated by the complications described in 1a above. As proposed, the VARs, system integrators, and other distribution channels would potentially have to register with the ENERGY STAR® program and will have to make ENERGY STAR® qualification assessments.
c. Requiring a label on a server which is often times installed by a technician not associated with the data center operator, in a data center which routinely is not occupied, and in a rack which will obscure the label is not the best method to promote the ENERGY STAR® requirements for servers. It does little to advertise the ENERGY STAR brand and does not justify the additional costs associated with the labeling process.

IBM recommends that EPA consider these alternative options:
a. An ENERGY STAR placard which could be presented to the customer and affixed at the entry to the data center or the server room to designate the support of the ENERGY STAR program.

b. Allowance for the ENERGY STAR label to be affixed to the rack door of a rack which contains ENERGY STAR machines. The label would be installed by the technician in the field.
c. Some form of an electronic "stamp" that gets loaded into the read-only memory of the machine which contains the vital product data and is accessible by standard programming techniques.
d. Provide a .gif file that gets shown when the system bios boots up.

In order to make labeling of ENERGY STAR products possible, both at the product and the data sheet level, EPA must qualify at the level of a service model (see page 9 below).

Line 60: Updated, annual list of ENERGY STAR qualified models. The dependency of the ENERGY STAR qualification on configuration type makes the proposed reporting by model impossible. As proposed, the idle power criterion is proposed, the report would have to be for models which could be ENERGY STAR qualified depending on the configuration. EPA should consider the IBM proposal to qualify Computer Servers by model type to allow for more a more practical approach to this requirement (page 9 of the comments).

Line 159: Delete the phrase "through enterprise channels". The definition of an "enterprise channel" is not clear and there are several methodologies by which computer servers are sold into the marketplace as discussed previously.

Line 228: Definition of a High Availability Server. If this definition is used, the criteria for these systems should include redundant cooling systems and redundant regulator systems. Both of these additional criteria are critical parts of providing a high availability system and they increase the power budget of the server. It is also worth pointing out the term "high availability" has a particular meaning to a wide segment of the industry, identifying fault-tolerant systems and environments with high levels of redundancy. EPA's use of the term is potentially quite confusing.

If the EPA's true intent is to differentiate between single and multiple power supply systems, then the category names should be changed to "Redundant Power Supply Servers" and "Single Power Supply Servers" as we recommended in our comments to Draft 2. This more accurately defines the distinction EPA seeks to make with regard to the presence of redundant power supplies, their impact on system idle power, and avoids the complications of having to consider other system redundancies.

Line 262: Computer Server Power Supply: It is recommended that EPA remove the last sentence of this definition, with the exception of the statement that "The power supply should be separate from and not integrated into, the system motherboard." There is no reason to define the power supply within a specific form factor and connector configuration.

Line 315 and 716: Designation of SPECPower_ssj2008 as the required methodology for measuring idle power. IBM has concerns with the proposed idle power methodology:

1. The SPEC license has specific requirements with regards to the performance of the test and the publishing and reporting of the data which are potentially at odds with the EPA reporting requirements under ENERGY STAR. EPA needs to reach a Memorandum of Understanding with the SPEC board of directors (BOD) detailing how the data can be managed within the SPEC requirements and reported under the ENERGY STAR requirements.
2. Running a complete SPECPower_2008ssj test procedure takes 75 minutes. Under the proposed idle power criteria, manufacturers will need to run many more SPECPower tests for non-idle conditions than they would otherwise normally run on their products to gather sufficient data to determine product qualification under the ENERGY STAR® requirements. It is recommended that the EPA's discussions with the SPEC Board include working toward a way to shorten the measurement requirements and an agreement of how measurements that are completed for ENERGY STAR® reporting will be assessed or qualified for validity under the SPEC license requirements.

Line 320: Qualifying Products

IBM recommends that EPA limit the qualifying products to one and two processor systems, given the proposed idle power criteria and the general lack of data to date about four processor machines. IBM has several concerns with the proposal to include four processor systems:

1. The current dataset for four processor socket systems does not adequately represent the range of system configurations in which four processor socket systems are sold. In fact, it contains very few data points relative to the number of data points for one and two processor machines. About 16.5% of the EPA-collected data is for four processor systems, and most of the machines have memory sizes less than 64GB and only a single disk drive. For the SPECpower_ssj2008 published results, the data sample is even poorer with four processor machines constituting one data point or about 2% of the total sample set. Additional data needs to be developed to gather a complete, representative data set around which an idle power criterion can be set for four processor systems. This cannot be completed on the timeline proposed in Draft 3.
2. Four processor systems are used in many varied applications and are the systems that are being purchased to consolidate applications through the use of virtualization capabilities. With multi-core implementations, a four processor machine may appear to have sixteen processors to systems software and shortly may have up to sixty-four processors visible to systems software. The systems that best support virtualization and consolidation are highly configured systems with higher powered processors (as opposed to low power processors) optimized for very high utilization which are not likely to meet an idle power criterion. EPA needs evaluate approaches to setting an idle power criterion which will enable customers to purchase machines which can consolidate workload and increase their utilization. IBM has provided data to EPA separately which shows the change in system types and system utilizations from 2006 to 2008 at an IBM owned and managed data center. The data shows that the number of multi-processor machines and the server utilization both increased markedly over those three years. It is important that EPA not create a situation where a customers purchase 6 ENERGY STAR servers to do the work that could be performed by four

processor server using 25% of the power of the six servers at maximum and idle. The ratio of workload down to power applied gets maximized when the server is doing work 50% or more of the time and able to reduce its power when it is not doing work. EPA needs to assure that the marketplace utilizes both of these attributes, not just a reduction in idle power, to achieve improved energy utilization.

3. In our comments and data development for Draft 2, we encouraged EPA to use a Power Management Criteria to address achieving idle power reductions. We commented at that time that lower idle power is achieved through the deployment of power management functions, that power management functions were configuration independent, and a criteria that focused on power management functions would give data center operators the latitude to procure ENERGY STAR® servers which met their application needs and would enable virtualization and consolidation of applications to increase server utilization. While we recognize that EPA has dismissed this option, we feel this dismissal is not justified, especially for four processor machines, and that a Power Management Function criterion would eliminate the complications of an idle power plus adders specification while encouraging the purchase of energy efficient machines across the range of models and configurations. We encourage EPA to revisit this position and revisit our Draft 2 comments for details.
4. Pursuant to its comments regarding the use of a power management function criteria, IBM commented that if EPA insisted on a maximum idle power criteria, it should only be established for one and two processor socket systems. IBM commented in its September 19, 2008 comments to EPA: *IBM believes that qualifying Computer Server products for other than Blade Servers should be limited to Computer Servers with up to two processor sockets and a 1U or 2U form factor. As processor capabilities are expanded, with more processor cores available – current designs may include 4 processor cores per processor die – 4 processor systems can deliver 16 or more processor cores of processing power. These systems are capable of managing and utilizing significantly memory, storage and I/O resources, affecting idle power measurements and the utility of SPECPower_ssj2008 as a metric for these larger systems.* These systems are not comparable to one and two processor systems, running significantly more memory, I/O, and storage systems with the capability of operating a virtualized environment to support many independent workloads in a consolidated fashion. As noted in bullet 2, these systems will be the workhorses of the data center as data center operators focus on increasing utilization and the amount of workload done per unit of energy applied and purchase equipment based upon energy efficiency and a performance per watt metric.

IBM recommends that the criteria for four processor systems be completed on the same timeline as the criteria for blade servers, to allow adequate time to collect a representative data set and to explore criteria which recognize the power saving benefits of both reduced idle power and increased utilization.

Line 414: Based on its discussions with EPA, IBM recommends that an additional Power Supply efficiency standard be created for power supplies with a capacity of less than 501 watts. Because of their lower wattage, they typically do not run in the 10 to 20% load

range and the low output at this range creates inefficiencies. IBM recommends that EPA use the efficiency ratings for the ECOs 80+ bronze standard for the 10% and 20% load point power supply efficiency requirements and the efficiency ratings for the ECOs 80+ silver standard for the 50% and 100% load points.

Line 442: Idle Power Criteria:

IBM continues to have the concerns with the idle power criteria that were expressed in its September 19, 2008 comments to the Draft 2 document. *The power drawn by a server when no workload is present is heavily dependent on* both the actual and potential *configuration: the number of processors, the quantity of memory, the number of active I/O points, etc., and the availability and enablement of power management functions on the server. The types of server which will have the lowest power level when working are often those which are least capable of running multiple operating systems and workload to drive higher utilization rates. Unlike a PC,* which is both economically and physically constrained to a relatively narrow range of configuration, *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 mode, because of the tremendous variability in form and function among 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 maximize its utilization and the workload that it delivers and have the capability to minimize its power use when no workload is present.*

IBM has multiple concerns with the EPA idle power criteria.

1. Data sets used to establish the maximum idle power criteria: IBM is concerned that the data sets use to set the idle power criteria are not sufficiently representative of the current market to provide a representative idle criteria. Based on a data set analysis by the industry, the one processor criteria appears to be somewhat low. The four processor criterion is not reasonable, because the data set is not representative of the range of products offered in the marketplace today. Specific concerns were discussed previously. The increase in allowable idle power from a two processor to a four processor system is 100 watts, the same adder to move from a processor to a two processor system. This is clearly incorrect since the number of processors is increasing by two rather than just by one. Simple arithmetic gives an adder of 200 watts. Allowing for some economies of scale, the one processor and two processor criteria, the minimum idle power for a four processor system should be somewhere in the range of 375 watts, given that two processors are added. EPA's request for data, sent out by email on 11/25/08 requests additional data for four processor systems. IBM intends to provide IBM configuration level data, generated by testing and its power configurator applications, on three configurations of its four processor models by 12/15/09.
2. Three processor systems should not have the same power budget. That does not make logical sense, given the power overhead associated with each processor.

3. A configuration dependent specification will create difficulties in identifying ENERGY STAR qualified products. IBM markets its System X and low end system P products through Value Added Retailers (VARS), system integrators, solution providers and IBM direct channels. VARs and solution providers often purchase stripped-down version of these products with only processor and limited memory and then build the server to its customer's requirements. The ENERGY STAR determination would be made by the intermediary, not IBM. See the "line 33" discussion above.
4. Setting the processor criteria based on installed processors creates problems in the marketplace, as some customers purchase partially populated machines in order to allow for future expansion or to install specific system components. As the number of processor sockets increase, the system power "overhead" due to increased circuitry and hardware requirements to support more processors and hardware. Such behaviors are much more common in the server marketplace than they are in the PC. This additional overhead consumes more power and makes it unlikely that partially populated multi-socket systems will qualify for ENERGY STAR. IBM will provide data on various one processor and two processor configurations for (2) two processor socket machines by 12/15/08.
5. The proposed maximum idle specification will tend to be biased toward systems with low power processors. Low power systems make trade-offs on performance capability in exchange for reduced power use. This is appropriate for some applications, but not others, and it is important for customers to be able to choose the type of system that they need to meet their performance, reliability and serviceability requirements. In fact, processors that best enable virtualization and consolidation tend to be those with higher power demands to drive the higher performance levels required to support virtualization. It is these types of systems that offer the lowest performance per watt.
6. As proposed, the idle power criteria will exclude the four processor systems best able to virtualize workload and run at higher utilizations. As discussed above, the process of maximizing workload delivered per unit of energy applied in the data center is a combination of lower idle power and higher utilization. By biasing to lower idle power, EPA is inadvertently steering data center operators away from the more powerful servers which will enable high levels of virtualization and reduce the overall power required to deliver work.
7. Line 451: Setting the idle criteria for the number of installed processors and not providing adders for extra components such as more I/O support does not recognize the fact that these additional components have a base energy demand that increases the idle power. Providing for capacity expansion requires additional componentry and potentially a larger power supply, making it exceedingly difficult to design an ENERGY STAR® system with any expansion capability. If it choose to propose a base plus adder idle criteria, EPA should provide adders for expansion capabilities. IBM encourages EPA to establish a model based qualification system. For example, an I/O device that consumes 15w moved the utilization of a CPU complex up by 20% and consumed 5 additional watts while operating. The server is now consuming an additional 20w in this case. Further investigation might show that the alternative facing an operator that needed the additional workload might be to install yet another

1U serve that uses 250w. It would have been better to allow the server to account for 20w within the server and save (250w – 2w) 230w.

8. A configuration dependent specification, particularly which utilizes a base idle power plus adder, will depend on a product management process with the ability to make an ENERGY STAR® assessment on each configuration of each model. This will require extensive testing and data collection on each component to determine idle power and the ability of the manufacturer's configurator program or a pre-product announcement delineation of ENERGY STAR® qualified configurations within an announced model to assess the conformance of each configuration with the specification. This will require several levels of data quality checks:
   a. Verification of the idle power of each component or component type.
   b. Verification that the support logic to the component allows the component to enter the state in which idle power is consumed.
   c. Verification that the model algorithms correctly calculate the idle power of the configuration and compare it to the ENERGY STAR® standard.
   d. Program to do some Computer Server level power testing to verify that the configurator is correctly reporting the power use of the Computer Server products.
   e. You would need a program to check that the changes made in configurations by resellers and other independent parties do not invalidate the ENERGY STAR® rating given by the vendor. This assumes, of course, that the channel partners do not have to do their own ENERGY STAR® qualification.

Rating each configuration as ENERGY STAR® qualified is not a workable proposal when considering the range of configurations and the customer delivery options discussed above. There would be significant cost to formalize the program, establish the algorithms to determine conformance with the ENERGY STAR® requirements and undertake the necessary testing and verification to satisfy EPA that the configurator results are representative of the performance of the actual Computer Server in the field. In addition, this approach is likely to cause confusion with customers, as different configurations within a specific server model will be qualified or not qualified for ENERGY STAR®.

Specific Comments on the System Characteristics:

EPA should set the memory requirements based on the number of DIMMS, not on the quantity of memory. From our September 19, 2008 comments:
Set the memory criteria for less than 4 DIMM slots (Low Memory Systems) and 4 or more DIMM slots (High Memory Systems). Currently, manufacturers are preparing to release a 16 GB DIMM. Demands for memory will increase in line with processor performance, and the number of cores per processor, *so using the DIMM criteria will better represent actual system purchases and will make the specification more robust over time.*

Proposed Option for an Idle Power Criterion:

IBM proposes that EPA consider an idle power criterion which qualifies specific server models, rather than individual configurations. The proposal is to measure idle power for

the minimum configuration and the maximum configuration on each model in the one processor socket to four processor socket qualified product group (as defined in lines 532 and 526 respectively. The idle power criteria would be set based on the sum of the idle power measurements for that model and would qualify all configurations of that model, including systems which ship with partially good processors and partially populated processor sockets. IBM believes this approach will provide a reasonable indication of the most idle power efficient models while simplifying the identification and administration of ENERGY STAR® qualified products. This approach has the following justifications:

1. While this approach can be subject to criticism, it is no less accurate or meaningful than trying to assign component level numbers for actual power use (by the systems manufacturer) and power adders (by EPA) to a large universe of systems and components and then calculate an ENERGY STAR rating for individual configurations.
2. An analysis of data taken from manufacturer's power configurators and provided to EPA in IBM's July 25, 2008 email ("Power Config Sort.xls") shows that, for the most part, specific models have the lower idle power points across their range of configurations. IBM has also done an analysis on various configurations of a one processor and two processor system. The analysis, done with different components and power management settings, show that each of the models had configurations which would qualify under the ENERGY STAR criteria. A client could purchase an ENERGY STAR® model, where required, and then add the components that they need to meet their performance needs. This suggests that designating specific models, as opposed to configurations, is a sensible way to encourage the purchase of servers with lower idle power while simplifying the process to identify ENERGY STAR® models. This approach still has the concern, detailed above, that the criteria is biased to low power processors and does not necessarily support the implementation of virtualization and consolidation which is best supported by servers with greater capabilities and power consumption.
3. The qualification should be done based on data from systems with all processor sockets populated for the measure of both the minimum and maximum configuration. As discussed above, setting the criteria based on installed processors biases the outcome to systems that have fully populated the available sockets.
4. This approach allows the server manufacturers to qualify their machines within their fulfillment system, eliminating the complexity introduced by the many channels by which machines are configured and delivered to customers.
5. It simplifies identification of ENERGY STAR® servers for data center operators. The client benefits from the fact that the number of models is vastly smaller than the number of configurations, making the purchasing process much simpler.

In order to make this process work, the definitions for Maximum and Minimum configuration need to be modified as detailed below.

Line 526: Maximum Configuration: EPA needs to adjust this definition to read: "The highest performance system with all processor and memory slots fully populated and

enough I/O devices, hard drives, and other devices to drive the server to its maximum wattage rating for the available expansion capabilities." For purposes of this discussion, the I/O slots include hard drives, LAN cards, video cards, etc.

This modification is needed, because if fully populated the expansion capabilities with maximum performance cards, the rack system may draw more power than the power supply can deliver. The definition adjustment makes it clear that the full power capacity of the server needs to be consumed, but prevents any requirement to overload the system.

Line 532: Minimum Configuration: EPA needs to adjust this definition to read: "A base model system that is minimally configured. All processor sockets will be populated, with the lowest memory configuration, a single hard drive, and a single network device."

STANDARD REPORTING REQUIREMENTS

As discussed above, if EPA chooses to use a configuration based approach to product qualification, the process of providing a data sheet indicating that a product is ENERGY STAR® qualified will be extremely difficult. In fact, if taken literally, there will be so many data sheets that matching data sheets and systems will become difficult. Each customer configuration will have to be assessed for ENERGY STAR® conformance and the correct data sheet supplied. This approach would cause customer confusion and significant fulfillment difficulties for the manufacturers and their distribution network. EPA must provide very careful guidelines about how to indicate which configurations qualify, given that most manufacturers will want to provide a single data sheet per computer model. In addition, the EPA should not require that a manufacturer make any statements as to whether a third-party add-on device alters the Energy Star qualification or lack thereof for a configuration of a model. This concern is addressed if EPA adopts the model based qualification process.

Line 528: Insert after "available for the model" the phrase "from the manufacturer". Manufacturers cannot be responsible for third-party items. The request for this edit goes away if the alternate maximum configuration definition supplied above is used.

Line 540: Even if the dominant vendor for a component for a model is not the system manufacturer, Only devices available through the manufacturer can be used in determining the typical configuration. It may be possible to add other components to the system, but the system manufacturer should not be responsible for determining the ENERGY STAR® qualification of a product with third party components.

Also, depending on the channels used, the typical configuration as sold through the manufacturer may be the minimum configurtation. Sales of additional features may be dominated by a third party, such as a Value Added Retailer, a Systems Provider or a Systems integrator.

DATA MEASUREMENT AND OUTPUT REQUIREMENTS

EPA needs to be aware that utilization measurements are dependent on where and by what sub-system they are measured. Windows and some Linux systems measure the utilization through software, while other systems measure it through the hypervisor and the firmware. Further measurement complications are introduced when a system is virtualized and running multiple operating systems and applications. This does not detract from the benefits that a data center operator can glean from tracking utilization (IBM is tracking utilization in many of its data centers), but it does point out the data can not be used for comparative purposes between different platforms and system types. Understanding true utilization for multi-threaded, power-managed systems is a topic for further research.

Line 575 Sampling Requirements: Sampling requirements for the power supply should be set at a minimum 30 second average and the EPA should not require multiple averages be reported (the end user could take the 30 second sample data and average them for longer reporting periods) IBM products typically also report a maximum value measured during the 30 second window. A 15 second polling rate for the power supply is too frequent, not industry standard, and a 30 second sampling period provides comparable information. Providing average and maximum values allows operators to get a sense for system behavior and potential anomalies. Customers are typically reporting up to their monitoring systems on 5 or 10 minute periods, so a minimum 30 second averages for data collection at the device is adequate. A one second minimum hardware polling rate for the power sensors are not required to be specified since the EPA has specified a required accuracy, unless the EPA wishes to include the requirement of reporting peak power (which is the maximum 1 second average value during the 30 second window).

Proposed wording:

Sampling Requirements: Hardware polling rates of the embedded sensors must be sufficient to provide power measurements of a 1 second average value. Data must be collected to report a minimum of a 30 second rolling average and the peak (1 second average value) for the rolling average window reported.

Line 608: Air flow is not typically measured in server systems. Air flow is controlled by system temperature and the temperature can provide an adequate indication of system health. Air flow is not needed.

Line 612: All server systems should have the same reporting requirements. There should not be exclusions for 1p systems. By the way, comments requesting the exclusion of the reporting of data are coming from stakeholders that manufacture 1p systems without BMCs or service processors, which are required to report out the requested data. In turn, the presence of a BMC adds a power overhead which affects the 1p idle power metric.

Line 622: TIER 2 Proposal:

Line 639: IBM continues to recommend that EPA does 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 system efficiency metric, making this the appropriate metric to use to set a power supply requirement under future ENERGY STAR Server Specification.

TEST CRITERIA

EPA should specify the test voltage for DC-DC power supplies from 52 to 54 VDC. 48V input dc/dc power supplies are largely targeted to the Telco market or locations where direct battery backup of the 48V is used. The actual nominal input to the power supplies in every datacenter is in the 52 to 54V dc range. This is the nominal "float" voltage that the power supplies are operating from.

If EPA continues to require the use of SPECPower_ssj2008 for idle measurement, EPA must arrange with SPEC for the free, unlicensed, unrestricted use of SPECpower_ssj2008 by any and all parties for this purpose and/or ensure that using the benchmark in this way does not require a submission of the full benchmark results to SPEC and its acceptance of them for publication. This may require that SPEC make the code freely available for download without charge or restriction and it requires a legally binding agreement between the EPA and SPEC.

IBM prefers that EPA approach SPEC to develop a "stripped down" SPECPower_ssj2008 that would use specified portions of the total SPECPower test but be completed in 15 minutes and require only the publishing of the idle number and work point (say power use at 10%). This proposal would keep the test procedure within the parameters of the SPEC procedure, could be done through SPEC (per comment from the IBM representative to the SPEC organization) but would simplify the process. We believe this approach would achieve the intent of getting a consistent, accurate idle power measurement while shortening the testing time and the required resources.

Lines 725 and 731: The vendor should also indicate what hypervisor, if any, was used during testing. Any such procedure would have to be freely available and usable under the terms of a legally binding agreement between EPA and SPEC.

Line 744: If EPA requires that each configuration be certified as qualifying for ENERGY STAR, this requirement is unreasonable and impossible to fulfill. This would require testing of each server as it left the factory or the IBM distribution partner. Manufacturers must be allowed to test a representative sample of models and configurations and be allowed to make the determination of ENERGY STAR qualification through the use of their power configurators. The real point though, as discussed previously, is that use of an infinite configuration qualification is not sustainable for the manufacturers, third part channel partners, EPA, or the customers. EPA needs to use a model level qualification of ENERGY STAR® products as proposed on pages 8 and 9 above.

EFFECTIVE DATE:

IBM does not believe that an effective date of February 1, 2009 is achievable for the specification, given the concerns about the data set and the need to create a model based qualification procedure for all affected server models. IBM recommends that EPA circulate at least one additional Draft of the specification, given the extensive comments provided above and the need to gather additional idle data for the qualifying products.