IBM Comments:
ENERGY STAR® Program Requirements for Data Center Storage Draft 2 Version 1.0

IBM appreciates the opportunity to continue to work with the EPA ENERGY STAR® program to develop the version 1 requirements for storage systems. While Draft 2 makes positive steps towards a final specification, IBM is concerned that Draft 2 has not moved boldly enough to address the difficulties inherent in the "bookend" product family proposal regarding cost, testing, and representativeness. Through these comments, IBM hopes to assist EPA in its continued efforts and creative thinking to identify potential changes to the product family process that would reduce the testing burden and identify a limited number of product configurations which would define and assess the energy efficiency of the products actually purchased by data center operators. In these comments, IBM makes two proposals that could serve as a starting point for a workable product family definition. IBM also has concerns with EPA's proposals for power management and Configurable Energy Efficiency Features, also discussed in detail below.

IBM offers the following comments and recommendations with regards to the ENERGY STAR® Program Requirements for Data Center Storage Draft 2 Version 1.0 released by EPA for comment on October 6, 2011.

PARTNER COMMITMENTS:

Section 5.1.1: The EPA proposal to qualify a product based on a maximum and minimum configuration (the so-called "bookend" proposal), representing some subset of the total possible product configurations, creates significant difficulties in using the ENERGY STAR® mark on product literature and websites. If EPA continues with this "bookend" product family definition, they need to work with their partners to identify the way in which the product literature can display the ENERGY STAR® mark and have an appropriate explanation of which configurations qualify for ENERGY STAR® designation, or direction to the appropriate ENERGY STAR PPDS (as required in 5.1.1). IBM's two product family proposals outlined below would eliminate this concern by qualifying the full range of configurations for a machine type.

PROGRAM REQUIREMENTS FOR STORAGE SYSTEMS

DEFINITIONS:

Section I: Product Family (page 5): There are several problems with EPA's Proposal to Establish Product Families based on a combination of "bookend" Configurations and Drive Types. First, it would create a burdensome testing process carrying significant costs for manufacturers. Second, the configuration requirements do not represent systems actually purchased in the marketplace. And third, it creates a qualification taxonomy that does not account for the ease with which a customer can add storage media to the minimum or maximum qualified system configurations.

The bookends product family would require testing of a large number of configurations:
a. Storage systems can utilize 4 to 6 distinct drive types. Even assuming that a company chose a maximum configuration that was 1/3 to ½ populated with storage media (the optimized configuration, or the so-called "sweet spot" configuration), this would

represent 130 to 200 of each type of storage device to fully qualify an OL-4 system to ENERGY STAR® (assuming a maximum capacity of 400 storage devices). Smaller systems would require a smaller number of storage devices for testing. But the enormous capital costs of procuring the storage media and the large amount of set-up and testing time required to qualify the system will make ENERGY STAR® qualification needlessly costly and complicated.

b. For example, IBM tested an OL-3 system with a maximum of 60 storage devices as part of the EPA call for data. Procuring, setting up, and testing the single system with a single type of storage media took over two weeks. Not only does the storage system have to be set up and "tuned", but the servers that drive the workload to the storage device have to be set-up and programmed. Extending this effort to up to 5 additional storage media types and up to 200 storage devices would require a huge effort and one to three months of time. This is a significant time investment, especially when considering that a company must also have arranged for supervisory oversight with its CB.

c. When a system is qualified within the bounds of the minimum and maximum configuration, it is in fact being qualified for the full storage media capacity of the system. A customer who wants, or who is mandated, to buy an ENERGY STAR® storage system, but needs more storage media than the number qualified under a maximum configuration, could simply purchase the maximum configuration and then separately purchase the additional storage media they require, whether or not that unit has been qualified with the additional drive types. Using a more flexible approach, the customer could buy the minimum qualified system type with qualified drives and then populate the remainder of the system with their desired quantity and mix of drives that may or may not have been qualified to the system. In either case, EPA needs to recognize the limitations of the proposed testing structure. The cost burden on the manufacturer to perform testing on a complete range of its systems and the available storage media should be weighed against the ease with which a customer could construct the system they desire by purchasing a minimum configuration of desired, qualified storage media and populating the remainder with their desired mix of unqualified media. EPA should endeavor to design a more rational, affordable qualification system for the manufacturers.

d. Systems for ENERGY STAR® testing must be purchased by the testing organization using capital dollars. Testing 3 to 6 maximum configurations to cover all of the media types will cost over $500K for some OL-4 systems. When combined with 2 to 3 months of internal personnel and CB time, ENERGY STAR® qualification becomes unaffordable.

IBM has two proposals for creating a streamlined testing and qualification procedure for storage systems which assures that the systems are properly tested and representative of the types of systems a customer would purchase.

**PROPOSAL 1.**
**STORAGE MEDIUM/STORAGE SYSTEM SWEET SPOT QUALIFICATION:**

Qualify the individual storage media separately and perform a system level qualification on the storage system populated with qualified drive types:

a. Require qualification of storage media by device type: SSD, and high and low capacity and rpm drives (by rpm/capacity combination). As drives represent 70-90% of a storage system's power use, qualifying the drives makes sense. It also removes the need to test each drive type on a system. Each drive can be tested and qualified in a single test in a short amount of time. These tests can be done once and applied many times. While there are not currently established standards for measuring idle and maximum power of storage devices, storage system manufacturers and device manufacturers have established internal procedures to perform this testing. Some storage devices have internal settings for full and idle power which would enable quick testing using a power meter. EPA should be able to work with its industry partners to expeditiously establish storage device testing procedures for maximum and idle power.

b. Because storage media have been tested and qualified, the storage system qualification can be simplified. IBM proposes that a manufacturer test and qualify its optimum configuration of a system (the "sweet spot"), with the manufacturer's chosen combination of ENERGY STAR® qualified drive types. If the sweet spot system qualifies, then the storage system is qualified for ENERGY STAR® with any combination of ENERGY STAR® qualified storage media. A sweet spot configuration is the combination of the controller system and drive types that the manufacturer feels offers the best combination of active and idle performance for the storage system product family.

c. It may be necessary to specify a maximum percentage of SSD drives that can be included in the system, as an idle criterion, is the only proposed power based criteria for Version 1 and a system of all SSD drives will be the top performer on any idle metric.

IBM's Proposal One offers several benefits to the ENERGY STAR program:

a. It removes the complications presented by the almost limitless permutation of drive combinations that could be qualified under an ENERGY STAR® test procedure by qualifying the individual drives. Testing of individual drives, which are used by multiple manufacturers and systems, simplifies the qualification process.

b. By allowing a manufacturer to test an optimized system with the appropriate mix of qualified storage media, the ENERGY STAR® program qualifies the "real world" system that customers purchase and which allows a manufacturer to display the benefits of the full capabilities of its controller system.

c. Qualifying all configurations that can be created with ENERGY STAR® qualified drives, would minimize the uncertainty created by the ability to buy a minimum configuration and populate it with your desired drive types.

**PROPOSAL 2.**
**SWEET SPOT CONFIGURATION WITH TWO REPRESENTATIVE DRIVE TYPES:**

IBM's second proposal, while less desirable than the first proposal due to the need to test an additional system configuration and its qualification with only 3 storage device types, can still offer a significantly reduced testing burden and simplified qualification process. In this proposal, IBM recommends that EPA designate two representative drive types – high capacity and high performance. EPA would need to work with industry to determine the representative form factor (2.5" vs. 3.5"), capacity, and spindle speed for the high capacity and high performance drives. A manufacturer would be required to test their sweet spot configuration with each of the two representative drive types, with an allowance of up to 15% SSD drives (rounded up to the next power of 2) to be representative of how systems are being configured. Again, if the system qualified to the applicable criteria, it would be qualified for all configurations and drive types to simplify management of the ENERGY STAR® brand and avoid the uncertainties for purchasers outlined above. While this option is considered less elegant than the first proposal, it still offers the majority of the benefits detailed for Proposal 1; a simplified testing effort that adequately characterizes the ability of the system to meet customer needs and deliver energy efficient storage performance.

**COMMENTS TO SECTIONS OF THE DRAFT REQUIREMENTS**

Section I.1: Common Family Attributes (page 5):

The common family attributes should include:

1. A definition of the appropriate constraints on the controller, such as a family must use the same type of controller.
2. A declaration of the maximum number and types of storage media supported by the storage system.
3. The maximum power the system can draw when fully populated with high performance media (currently 15 K drives).
4. Type of Reliability Management: RAID Parity, mirroring, or other approach.

Section I.2-4: Definition of Maximum, Minimum, and Typical Configuration:

Under the two proposals made above, EPA would need to remove the definitions for a maximum, minimum, and typical configuration from the requirements and replace them with a definition of the "sweet spot" system. This definition could be derived from the SNIA "best foot forward" definition. In general, it should state that a sweet spot configuration is the combination of the controller system and drive types that the

manufacturer feels offers the best combination of active and idle performance for the storage system product family.

If EPA chooses to use the book ending method of qualification, then the maximum and minimum configurations need to be redefined:

Maximum Configuration: The combination of a storage controller system and single storage media type which represents the upper bound of the storage system that the manufacturer wishes to qualify to the ENERGY STAR® requirements. The number of drives assigned to the maximum configuration must be the same for each drive type that the manufacturer chooses to qualify to ENERGY STAR® for a given product family.

Minimum Configuration: The combination of a storage controller system and single storage media type which represents the lower bound of the storage system that the manufacturer wishes to qualify to the ENERGY STAR® requirements. The number of drives assigned to the minimum configuration must be the same for each drive type that the manufacturer chooses to qualify to ENERGY STAR® for a given product family.

The "typical configuration" should be eliminated from the storage system requirements. There is no value to requiring the testing of a third configuration given that manufacturers are already choosing maximum and minimum configurations that represent some subset of the total configurations supported by the product family.

2.1.1 Note (Page 7, Lines 391-394) IBM supports EPA's to defer qualification of Network Attached Storage to a future version of the storage requirements.

2.2 Excluded Products (page 7): EPA needs to specifically exclude blade storage systems. These storage systems do fit within the On-line Storage categories 2-4 but their incorporation into a blade chassis/system result in the blade storage systems having a very different power signature from a conventional, rack mounted storage system. Criteria for blade storage systems should be incorporated into future versions of the storage requirements.

Network Attached Storage (NAS) should be added to the list of excluded products. In addition, EPA needs to add a definition for NAS to the definitions section.

3.2 Power Supply Units (PSUs) Requirements (page 8): IBM recommends that EPA specifically note in section 3.2.1 that PSUs required to meet the efficiency standard are those power supplies which directly power the controller and the storage media devices. On smaller systems, this may be accomplished by a single power supply while on larger systems these supplies may be separate.

3.2.1.i & ii: PSU Efficiency and Power Factor Requirements.
EPA should remove the 10% load point from the PSU requirements. Because of the power load created by the storage devices, it is highly unlikely that the power loads will drop below 20%. Assuming redundant power supplies, the controller uses 10% of the

system power, and that a typical system will have at least 30-40% of the disk drawers populated, the redundant supply will be running at 20% load. Where the disk drives and controllers are powered by separate power supplies, it is still expected that power supplies will operate at 20% of load or higher. Unlike servers, the configuration of storage controllers are relatively fixed making the power use fairly constant over a range of system configurations (dictated by the number of storage media used). Because of the need to be ready to respond to system demands and ongoing maintenance tasks, power management capabilities are typically not enabled on storage controllers (see the power management discussion). So a power supply will be sized closer to the demand of the controller and load is unlikely to be less than the 20% load point. For the disk drives, it is expected that 30-40% will be populated (in order to not underutilize the system investment) placing the minimum load near 20%. IBM is willing to develop and provide EPA examples of system loadings to back up this position.

IBM recommends that EPA specify the EPRI 80+ Power supply classification for the power supply efficiency requirements. This links the requirements directly to a published, refereed standard and simplifies the compliance process.

3.2.1.iii: Efficiency and Power Factor in Embedded Equipment: For power supplies embedded in other systems, such as network switches included in the storage system rack, IBM recommends that EPA defer action on these power supplies to later versions of the requirements. EPA should consider a data gathering effort to identify embedded systems with separate power supplies, whether the power supplies are AC to DC or DC to DC (powered off of the primary power supply(ies) of the storage systems), the size of those power supplies, and whether the embedded systems (such as switches) are covered by planned future ENERGY STAR product requirements. This will enable EPA to understand the universe of systems that make up embedded systems and formulate the best approach to address the energy efficiency of these systems.

3.3.1: Active State Efficiency Criteria: IBM agrees that it is appropriate to report, but not set criteria for, active state efficiency. Given the cost and time required for testing, allowing the manufacturers to develop and submit active state efficiency data as part of the qualification process simplifies the data collection process and provides EPA the data needed to determine how best to set active state efficiency metrics using the six active energy data points generated for a configuration by the specified test method. It will also allow manufacturers to assess the results and propose workable scoring approaches to EPA.

3.4.1 Ready Idle Efficiency Criteria: Considering the ready idle efficiency criteria in the context of the proposed “book end” method for defining a product family suggests that there will be several complications in setting this criteria. As the GB/watt metric is dependent on the drive type and the number of drives populated in the system (to achieve greater distribution of the overhead power associated with the controller), there are several factors that EPA needs to consider in setting the criteria:

a. A separate criteria needs to be set on groups of drives based on the capacity and spindle speed of individual media. Different media types will have significantly different GB/watt signatures based their capacity and spindle speed and it will be important to compare the results from like media. This complicates the requirements to set ready idle efficiency criteria for systems.
b. EPA needs to consider if there needs to be a range of drive counts that can be tested for the minimum and maximum configurations. The GB/watt metric will be affected by the ability to spread the controller system overhead over the drives. In order to get consistent comparisons of products, it may be necessary to set a range on the number of drives tested.

The two product family proposals put forth earlier deal with both of these problems by allowing a manufacturer to pick the optimum number and combination of drive types that they believe will give the best performance against the metrics and limiting the number of required system level tests.

3.5.1 Power management functions are not readily available on storage systems at this time. While some controllers are based on server systems or server processors, current response time requirements and the scheduling of maintenance workloads preclude the use of power management because systems have to be prepared to respond to network queries or perform the various maintenance functions. Current power management strategies depend more on system management strategies such as tiered storage to provide a better balance of power, capacity, and performance. Rather than requiring specific power management functions, IBM recommends that EPA ask manufacturers to list available power management functionality for their controllers and storage media on the power performance data sheet and that EPA use that data to determine the extent of power management functions available on systems to determine if it is reasonable to set requirements on future versions of the storage system requirements.

3.5.2 There are two issues with the power modeler (hereafter referred to as a power calculator) requirements. IBM agrees that a manufacturer should be required to have power calculators available to their customers, but the manufacturer should be given the option to make that information public or limit the access through the manufacturer's sale staff. IBM currently has power calculators available through its sales staff, but the calculators are not made publicly available. Where a manufacturer has a privately held power calculator, they should be required to demonstrate the capabilities of the model to their CB when qualifying a storage system, performing the power use calculations for the qualified system.

It is unreasonable to require manufacturers to warrant that the power modeler in all cases estimates power use equal to or greater than actual power used. A power calculator, is by its nature, an estimating tool based on a combination of a selected sample of actual measurements of system and component power use and associated calculations based on estimated power consumption of the various components of the storage system. There is a great range of variability in component power consumption. For example, different memory DIMMS will have a range of power use based on the manufacturer, the design of

the DIMM, and the manufacturing lot. This range introduces variability in system power use. The model is intended to provide an estimate of the likely power use. As IBM demonstrated with its comparison of the tested power use values and power calculator estimates for ENERGY STAR® qualified server systems, submitted to EPA on July 12, 2011, the power calculator provides a reasonable estimate of the system power use, but overestimates or underestimates the tested power use depending on how the model is set up and how the calculations are performed. EPA should require that the power modeling tool provide a reasonable estimate of the power consumption of the actual system. In addition, legal staff would not warrant the results of the power calculator due to the uncertainties discussed above and the many legal implications of a warranty on the delivered results.

3.6.1 EPA should not require a system to have parity RAID to qualify a storage product to the ENERGY STAR® requirements. While RAID parity may appear to offer power use benefits when compared to a mirroring system, a properly designed mirroring system that is integrated with the hardware system may deliver more energy efficient data reliability than Parity RAID. In addition, requiring RAID parity precludes the introduction of a new, more efficient technology over the time the version of the requirements is in effect. EPA should not be specifying specific system technologies for reliability or other system management tasks. By that logic, the ENERGY STAR® specification should require all ENERGY STAR® systems to have only SSD drives as that offers the most energy efficient storage solution. As we know, this is not the case as other drive types will be used when considering the combination of power use, cost, performance, and capacity required to meet a specific workload. The ENERGY STAR® requirements should be, as much as possible, technology and functionality agnostic. Manufacturers should be able to select and implement their technologies as they see fit and allow the ENERGY STAR® power use criteria to select the best performers.

3.6.2: For Configurable Energy Efficiency Features, EPA should not require that a specified number of features be enabled upon shipment of the system. Many of these features have specific licensing costs and are purchased by the customer when the specific feature benefits their particular workload. Rather, the EPA should set the requirement that the storage system should be able to support x of y of the features included on the list. This assures that features are available to the customer, while allowing the customer to select and pay for only those features that benefit their specific workload.

Unfortunately, metrics are not currently available to measure or quantify the energy efficiency benefits of these functions, as oftentimes the benefits are requirements for less equipment in the data center or a different mix of storage device types on a rack or within the data center. IBM believes that available software functions that optimize energy use and contribute to the delivery of more work per unit of energy applied and/or lower energy use in the overall data center should be listed on the Product Performance datasheet, similar to the listing of power management functions for servers.

Table 4 (line 517) The features list should also include tiered storage.

3.8.1.ii: IBM recommends that EPA remove the requirement for thermal monitoring on storage systems from Version 1 of the storage system requirements. The power use characteristics of storage system controllers and media do not require the level of thermal management that is required for server systems.

3.8.3: EPA needs to increase the minimum sampling interval for power use and thermal reporting (if required, see comment to 3.8.1.ii above) to once every minute. A more frequent sampling rate is not practical or useful in the context of managing a data center with hundreds or thousands of ICT systems. From a practical standpoint, the data traffic and management required to manage reports from every device in the data center made every 10 seconds requires a significant amount of computing and storage capability which is unreasonable given the usefulness of the data. The data center operator and alarm systems will only require updates of a minute or more to track developments in the data center. Individual hotspots tend to develop over time, not instantaneously, and a one minute reporting frequency provides a sufficient minimum time interval to track developments in the data center. If it would be helpful, a call can be arranged with the technical team that develops the IBM monitoring software, Active Energy Manager, to explain the technical difficulties associated with more frequent monitoring requirements.

The IBM team is available to discuss its technical concerns in more detail. Jay Dietrich (jdietric@us.ibm.com) is the IBM interface to the ENERGY STAR® program and would be happy to answer any questions you have or schedule a meeting with our technical team.

Thank you for considering our comments.