IBM RESPONSE TO:
ENERGY STAR® Enterprise Storage Draft Specification Framework: June 4, 2009

IBM appreciates the opportunity to respond to the EPA ENERGY STAR® Enterprise Storage Draft Specification Framework published by EPA on June 4, 2009. IBM worked extensively with EPA and industry group stakeholders on the development of the Computer Server requirements and intends to be actively involved in the development of the Storage requirements. IBM will participate with SNIA and The Green Grid work groups to work to develop an industry consensus on the most effective ENERGY STAR criteria to encourage the deployment of more energy efficient storage devices. As it did on the Computer Server requirements, IBM also intends to provide company specific comments to the EPA and provide access, as needed, to the IBM technical team to assist EPA in understanding the technical nuances of storage systems and the technical justification for IBM or industry positions on various proposals.

Storage systems have similarities and differences with Computer Servers, which will enable the use of portions of the Computer Server requirements as a starting point for the Storage specification in areas such as power supply efficiency and Information and Management requirements while requiring new approaches to addressing other criteria, such as power management requirements. A storage system consists primarily of a controller (a small server), I/O devices, and storage devices. The energy use characteristics of the storage media will have a direct bearing on the ability of the system to manage power use, potentially requiring different criteria for different media types and use scenarios. As the storage criteria are developed, it will be important to build on current industry efforts to improve storage system energy efficiency and set criteria which encourage and enable further innovation in this area. Similar to Computer Servers, there are two components to storage energy efficiency: the power used by the system to do work and the way that data is managed, through de-duplication, compression, prioritization, and other techniques, to minimize access time and the number of storage devices, and hence the energy use, required to manage the data

Following are specific IBM comments to the EPA Framework Draft. IBM also participated in the preparation of The Green Grid and SNIA comments and supports the proposals presented in those comment documents.

BUILDING BLOCK #1:

Definitions: EPA should utilize applicable definitions from the Computer Server Requirements document to assure consistency across the IT equipment programs within ENERGY STAR®. It is recommended that the SNIA Storage Industry Dictionary be used for Storage specific definitions whenever possible.

Computer Server Definition: The definition of "Computer Server" provided in the framework is a combination of the definition of "Computer" and the definition of "Computer Server." The two definitions should be separated into the Computer Server Definition from the "ENERGY STAR® for Computer Server Program Requirements, page 5, section 1.A and the definition for a Computer from the "ENERGY STAR® Program Requirements for Computers Version 5; page 5, section 1.A.

Questions 2 & 3: Storage devices have two types of workloads: client generated I/O load and maintenance and data management activities that enable functions which reduce the quantity of data stored and maintain its quality: compression, de-duplication, defragmentation and other activities. As such, there are three power using states on the system: power use when client generated I/O workload is occurring, power use when maintenance and data management functions are being performed and power use when no work of any type is present. As part of the specification work, it is important to carefully consider the relative values of measuring power use at a "true" idle state or measuring power use over a period of time in a way that includes the power used in the maintenance and data management tasks when setting reporting requirements or ENERGY STAR® criteria. Accounting for the maintenance and data management tasks is important, as these tasks allow more data to be stored on less storage media space, reducing the net energy use in the data center. It is important that EPA not create criterion which penalizes systems which perform data management and maintenance in a way that allows the data to be managed more efficiently on less equipment.

A Storage Product is made up of multiple units of Storage Media. During normal operation, some units will be active while others will not be processing client generated I/O workload. In addition, some systems cache highly active data separate from the storage media to reduce response time and improve data availability. These operating considerations introduce some complications in defining a maximum energy use, as it is highly unlikely that all units of storage media will be fully active at the same time. It is likely that some type of active power measurement (rather than maximum power) will best define the power use of an operating Storage Product.

Question 4: IBM supports the use of the SNIA storage taxonomy detailed in the "SNIA Green Storage Power Measurement Technical Specification."

Question 6: The power supply definitions and requirements from the Tier 1 Computer Server requirements should serve as a good starting point for the power supply requirements for storage. Storage systems have power supply challenges similar to servers.

- Power supplies have fans which may provide part or all of the cooling for both the power supply and the storage system. This is similar to many server power supplies and necessitates that the exclusion of fan power from the efficiency calculations which was granted for server power supplies be extended to storage power supplies. In addition, the fan power exclusion should be extended to multi-volt power supplies.

- As with servers, customers typically do not fully populate the disk slots in a storage device upon their initial purchase of the storage device. Rather, they leave space for future expansion. As such, it is important that power supply efficiencies be measured

by the traditional efficiency at specific load points as they are designed to cover a range of loads for the storage device over its life in the data center.

- Storage systems routinely use redundant power supplies. As was done with the server requirements, it is important to recognize the need for redundant power supply in any power requirements set for storage devices.

- Storage systems typically use 230 V input power to drive improved power system efficiencies. As with the server testing requirements, storage power supplies should be tested at an input voltage of 230 V.

Storage power systems may also have architectures which vary from computer servers due to the power requirements of the storage disks. The Computer Server power supply requirements should be evaluated against the specific technical details/requirements of these power system architectures to verify that they are appropriate for storage systems or to introduce the necessary changes or modifications needed to address the specific technical requirements associated with storage power systems.

Green Grid continues to recommend that EPA does not pursue the Power Loss metric as a power supply criteria for IT equipment. EPA's original 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. As machines are more "heavily" configured, the 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.

BUILDING BLOCK #2:

Question 2: While Solid State Drive (SSD) Storage is included in the Storage Media definition, it is a relatively new product with a power profile that is significantly different from the power profile of a disk system. Careful consideration needs to be given to how SSD media is incorporated into the taxonomy and the power management criteria, as its power profile could potentially make it the only active storage media that could qualify as an ENERGY STAR® product. Given the current product distribution in the industry, this would not be a desirable outcome for the Version 1 requirements.

Question 3: IBM recommends that the EPA focus on the On-line 2, 3, and possibly On-line 4 Storage Products (SNIA taxonomy). These three groups represent a significant portion of the energy use for storage devices in IT equipment use in small and medium

businesses and data centers. On-line classification 4 should be reviewed for inclusion, as it represents a key storage technology and a significant portion of storage energy use in enterprise level (high reliability) data centers.

Consideration should also be given to including one or two classes of removable media systems in the Storage Requirements. Tiered storage systems, where data is stored on devices based on the frequency with which it is expected to be accessed. Data which will be frequently accessed is stored on high response time systems that have a higher energy use while archived or infrequently accessed data is stored on tape media system. The use of tiered storage systems on a network provide the lowest overall energy use for a data management system and the most efficient utilization of hardware resources. Because of the importance of encouraging the implementation of the low energy data center solution, IBM believes that it may be valuable to include one or more classification groups from each of the "near on-line" and "Removable Media Library" categories of the SNIA taxonomy to the ENERGY STAR® Storage Requirements. This will send a clear message to data center operators of tiering their networked based storage systems to optimize the way data is managed and reduce the energy use of the overall system. IBM proposes to work with EPA, SNIA and TGG to identify the subset of groups of the SNIA storage classification taxonomy which will offer the best opportunity to encourage the most efficient use of energy in the data center.

BUILDING BLOCK #3:

Question 1: See the answers to Questions 2 and 3 in the Building Block #1 section for thoughts on this question.

Question 4: SPEC and SPC provide recognized metrics for measuring work output on storage systems. It is important to note the importance of various software algorithms, routines, and strategies which enhance the performance of storage systems (cache of frequently accessed data) and minimize the quantity of storage media requires (compression, de-duplication). These capabilities can significantly enhance overall storage system capability and efficiency while incurring increased energy use on the individual Storage Products in the system. These capabilities will need to be considered within the selected storage product taxonomy and their value recognized as the ENERGY STAR® Storage Requirements are developed.

Question 6: The power supply definitions and requirements from the Tier 1 Computer Server requirements should serve as a good starting point for the power supply requirements for storage. See the response to question #6; Building Block #1 for details.

It is likely that the storage PSU requirements can effectively build on the requirements established for servers. There are two critical concerns that need to be extended for the storage power supplies:

1. Fan power should be excluded from all power supplies, both single and multi-volt.

2. Efficiency testing should be performed at 230V, as that is the voltage used by the preponderance of storage systems. It is recognized that some systems will be run at 110 V.

Question 7: Power supplies on IBM Storage Products typically operate in the range of 20-45% loading, as the majority of Storage Products employ redundant power suppliers.

Question 8: IBM continues to recommend that EPA not pursue the "Net Power Loss" metric as a power supply criteria for IT equipment. We provided this comment throughout our responses to the Drafts for the Computer Server Requirements. EPA's original 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. As machines are more "heavily" configured, the 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® Storage Specification.

Question 9: The VAR sales channel has a significant role in the Data Center Storage market. The value proposition of a VAR's offerings is their ability to meet the customer's needs by modifying or adding equipment to meet the customer's application requirements. ENERGY STAR® should consider the efficacy to allow the VAR to purchase the ENERGY STAR® model that satisfies their customers' requirements and then add the necessary peripherals that are required to support the customers' operations. The VAR's ability to deliver this service is restricted if they are required to stay within the defined product family data sheet configuration. The other option proposed by EPA under the computer server requirements, to have each VAR sign on as an ENERGY STAR® partner and qualify its "modified" products, will create a significant data management burden on both the VAR and the EPA. There are thousands of VARs in the marketplace, which will create a multitude of product registrations while bringing no real value to the ENERGY STAR® brand or to customer understanding of ENERGY STAR®. The more efficient, understandable approach is to qualify the base OEM model with key, required components and allow the VAR to add its specialty peripherals and services to that base, ENERGY STAR® qualified product. Without an agreed definition of product families, it is difficult to determine whether the difficulty and expense represented by the requirements above on third-party sales channels. However, there is a benefit to EPA and industry stakeholders to enabling the equipment manufacturer or supplier to determine conformance with the ENERGY STAR® Requirements.

BUILDING BLOCK #4

Question #1: IBM supports the reporting of Storage Product power use and inlet temperature in a format available to third party software systems. IBM intends to make power and temperature data available on its storage system through the same software systems by which it makes its Computer Server systems data available. The trend in the industry is to make this data available to the data center operator to assist in planning and operating of the data center. However at this time, many systems do not report this data.

IBM does not recommend that EPA require the reporting of Storage Product utilization levels at this time. During the development of the Computer Server requirements, there was extensive discussion on the utility and accuracy of utilization data collected from Computer Server processors. It was determined that while there were not universally acceptable methods to providing utilization data and the available information was not highly accurate, it was appropriate to provide the data center user with system information that could help make decisions on consolidation opportunities in the data center. For Storage Products, the validity and utility of utilization data is highly suspect. Because the controller system is performing multiple functions as it manages data entering and leaving the Storage Product and deployment of data to and from various subsystems, utilization on the controller system processors provides limited information on the operation of the system. The validity and accuracy of the utilization data is further limited by the introduction of partitioning within the controller processor. IBM products typically do not report controller utilization because it does not provide meaningful data to the data center operator.

Question #2: IBM recommends that the data requirements for the product data sheet be developed concurrently with the Storage Requirements. Storage systems and their characteristics are sufficiently different from Computer Server systems that material changes to the data sheet are required. Such changes will be most easily developed in conjunction with the development of the Storage Product requirements.

Question #3: See the response to question #1, Building Block #4.

Question #4: Any data reporting protocols will arise at the data center level. There is at least one standard body, DMTF, working on a standard communication protocol for data center equipment.