

## IBM Comments:

### ENERGY STAR® Program Requirements for Data Center Storage Draft 1 Version 2.0

IBM appreciates the opportunity to provide comments to USEPA on the ENERGY STAR® Program Requirements for Data Center Storage Draft 1 Version 2.0. EPA has made a good start on understanding and evaluating the complexity of storage system technologies, but there is still a significant amount of work to do and technical challenges to address before an initial, workable set of storage system criteria can be established. Storage systems are much more complex and have a larger number of components and configuration permutations to consider and address than Computer Servers. This complexity alone will make it difficult to establish an initial set of requirements that includes an idle or active power metric or criterion. It is also important to limit the initial scope of the specification to products in the Online 2 and 3, and Removable 2 and 3 categories. Storage systems in the online and removable 4 and 5 categories include large number of storage devices and represent a small fraction of the storage market. It is in the interests of EPA to validate the final requirements on storage systems with a more limited number of storage devices to minimize the complexity of the systems and simplify the initial implementation of the requirements. More experience with the on-line and removable 2 and 3 categories will lay necessary groundwork for incorporating the other storage categories.

Following are comments to the Tier 1, Draft 1 document and responses to specific questions or topics raised by EPA with the document. They are organized by topic and citations are provided to the document page and line numbers.

Page 6; L195-201: Definitions: IBM encourages EPA to use the SNIA definitions in the storage requirements. These definitions were created through extensive industry collaboration and represent the industry consensus on storage systems definitions. Utilizing the SNIA definitions across the requirements should simplify their development and reduce the level of discussion required to finalize the definitions.

Page 6; L216-224: IBM concurs with the EPA's conclusion that software based functionality such as data de-duplication, data compression, thin provisioning, storage tiering and other software based capabilities can improve system efficiencies, reduce the size of the system required to deliver a specified workload, and optimize the workload delivered per unit of power supplied to the storage system. However, IBM believes that 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.

Page 10; L370-380 Product Family: Because of the number of possible configurations within each model number or machine type, providing a product family option for storage system will be important to making the product qualification and verification processes affordable and manageable for manufacturers. The most logical approach will be to define a product family by machine type or model number within a given product

## IBM Comments:

### ENERGY STAR® Program Requirements for Data Center Storage Draft 1 Version 2.0

category. Providing data for a product family is complicated by the number and variation of storage devices that can be used within each machine type. IBM is interested in working with appropriate industry groups and EPA to find a workable solution to defining and qualifying product families under the ENERGY STAR® storage requirements.

Page 10; L372-80: The basic definitions for minimum, maximum, and typical configurations make general sense. However, it will be important to consider how these definitions account for the different types of storage devices that can be used to populate the storage product.

The purpose of the base configuration is unclear. Arguably, this defines the minimum configuration for the product and is redundant.

Page 11; L408-9: Reword Availability Features as follows: Features that support a server's ability to continue to perform its intended function at full capacity after the occurrence of one or more component failures.

Page 11; L410-11: Reword Serviceability Features to Read: Those features that affect the duration and skill required to repair a storage system. These features allow a storage system to be serviced in the least amount of time through means such as automatic detection, isolation, and reporting of actual or potential failures, toolless parts removal, color coded touch points, and keyed connectors.

Page 12; L444-48: The qualifying taxonomy categories should be limited to Online 2 and 3, and Removable 2 and 3 SNIA categories.

Page 12; 449-56: EPA should work with SNIA and other interested parties to establish a definition of a Hybrid storage system, as the specific definition is not clear and could address several different product capabilities including different types and levels of system control or the presence of different storage devices in the storage product. The exact definition will enable discussion of whether and how "hybrid" storage products which fit a particular SNIA taxonomy category should be included or excluded from the requirements. IBM is interested in working with appropriate industry groups and EPA to find a workable solution to defining "hybrid" storage systems.

Page 13; Section 3.1 PSU Efficiency Criteria: While there are some differences, PSUs used in storage systems are similar to PSUs used in server systems. The differences exist in the use of cooling fans, UPS systems embedded in or connected to the power supply, and the use of more multi-volt PSUs due to the varied power requirements of different storage devices. Discussions with the IBM power supply engineers indicate that storage system power supply efficiency has not progressed to the degree found in server systems. In order to properly understand the current status of storage power system PSUs configurations and their conversion efficiency and power factors, IBM recommends that EPA perform a survey, in conjunction with EPRI, SNIA, or TGG, to determine the efficiency and power factor characteristics of PSUs on systems currently in the market to

## IBM Comments:

ENERGY STAR® Program Requirements for Data Center Storage Draft 1 Version 2.0

determine what PSU standard should be applied in the storage system requirements. Based on preliminary and incomplete information, it is likely that the ECOs 80+ bronze standard may be the appropriate starting point for the storage requirements.

Page 14, Section 3.3 Active State Efficiency Criteria: Developing an active state efficiency criteria, given the complexity of the storage system and the varied power use profiles associated with different types of storage devices, may be problematic. EPA will need to consider several options to approach the active and idle efficiency criteria.

- a. Require specific mixes of storage devices for qualifying a storage system based on the functional characteristics of each category of storage products. This would enable consistent configurations for qualifications, but may be complicated by system differences (such as the control structure or the number of storage devices that a product can support) between manufacturers.
- b. Create qualification criteria for each type of storage device, storage element, and/or the controller system that makes up the overall storage product. This would introduce some complications, but it may be the best way to address the varied quantity and types of storage devices that can be supported by a given storage product.

These 2 options for establishing active and idle efficiency criteria for storage systems are provided to stimulate thought and demonstrate that there may be several viable options for these criteria. However, the final methodology chosen will depend on the generation of data for the range of storage devices available in the marketplace and analysis of the data to determine which of these ideas, and other ideas not offered here, offers the best means to incorporate consideration of active and idle efficiency metrics into the requirements.

Page 14; L519-24: Idle state energy efficiency would not be a suitable proxy for active energy efficiency. The power differences between the active and idle states are affected by too many variables to allow a meaningful correlation of the two criteria.

Page 14; L525-30: Power Management Requirements: Storage systems do not typically have power management capabilities similar to those found in server systems. Improving power utilization in the system is more typically done using the software functions discussed on page 1, the comments for page 6; L216-224. Listing the available software functions on the product performance datasheet will serve to identify the available capabilities.

Page 14, 15: Product Performance Datasheet (PPD): Because the datacenter storage requirements are largely undefined, it is premature to consider what data should be provided on the PPD. IBM will address the PPD requirements once the storage requirements are better defined.

## IBM Comments:

### ENERGY STAR® Program Requirements for Data Center Storage Draft 1 Version 2.0

Page 15; L577-79: Storage systems do not have the ability to report power use and thermal conditions; the only available option to supply this information is through the use intelligent power distribution units. The frequency of the power and thermal readings provided by the iPDUs do not approach those stated in this draft. The specific requirements identified in the draft will require significant development by the industry, both in modifications to PSUs or other hardware to integrate the components needed to enable this reporting directly from the storage systems, and in the development of software to report this information into the existing data reporting systems utilized by the storage systems. This integration work will take time and should be delayed to the next tier of the storage system requirements.

Page 16; L593-596: Input Power: The requirement for 5% measurement accuracy at the system level is not consistent with a power supply accuracy that allows  $\pm 10\text{W}$  below 100W input. Accuracy will not make a jump from 10% to 5% when the power supply crosses a given power threshold. In addition, systems with multiple power supplies will keep increasing (by 100 W for each power supply added) the minimum error levels. IBM recommends that the measurement accuracy be measured at the power supply level and that it be maintained at 10% accuracy below 100 W to be consistent with the power supply requirements.

Page 16; L601-2: Inlet temperature probes are currently capable of  $\pm 2^{\circ}\text{C}$ , but the validation process creates a measurement/validation accuracy of  $\pm 3^{\circ}\text{C}$ . It is our recommendation that we keep the specification requirement at  $\pm 3^{\circ}\text{C}$ , as that is the consistent with the level of accuracy that can be verified by the validation test.

Page 16; L603-8: IBM objects to the requirement to report a 30 second rolling average ever second. This quantity of data, when considering that a typical data center can have hundred's or thousand's of storage devices, is completely unmanageable and requires a significant amount of computing infrastructure to just collect and file the data. Cumulating and storing all this data will involve several minutes of latency – so the data availability will be delayed by several minutes. The data should be collected at a minimum of 30 seconds for power readings and one to 5 minutes for temperature readings. There is no value to collecting data on a more frequent basis.

The IBM team is available to discuss its technical concerns in more detail. Jay Dietrich ([jdietric@us.ibm.com](mailto: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.