the green grid™
get connected to efficient IT

THE GREEN GRID RESPONSE TO THE EPA ENERGY STAR FOR STORAGE

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

The Green Grid Association, a consortium of industry leading companies welcomes the opportunity to comment on an early draft of topics under consideration for the ENERGY STAR for Enterprise Storage Specification. Some member companies of The Green Grid Association may in addition have provided additional considerations highlighted by their industry or company's particular perspective. Some members may have also provided their inputs through the Information Technology Industry Council (ITIC), the Storage Networking Industry Association (SNIA), and Climate Savers Computing Initiative (CSCI).

the green grid™

INTRODUCTION

A consortium of information technology providers, consumers and other stakeholders, The Green Grid Association seeks to improve the energy efficiency of data centers around the globe. The organization takes a holistic and comprehensive approach to data center efficiency and understands that addressing this challenge requires a high-level view of the entire data center and cooperation among a wide range of industry principals. Participants in The Green Grid include such diverse companies as major server and storage equipment manufacturers, major software providers, and large end-users / data center owners.

The Green Grid (TGG) believes that the EPA's efforts to create ENERGY STAR for Storage Systems represents a positive step on the path towards making current and future data centers more energy efficient. Many of these efforts are in harmony with TGG's own efforts across the OEM and end-user community in developing standards to measure and mange efficiency across data centers and the larger computing infrastructure. We consider the EPA, and specifically the ENERGY STAR program, as a partner to industry in this effort. TGG supports and acknowledges the EPA and the ENERGY STAR standards for reducing energy consumption across a range of client systems and peripherals. Cooperation and support of ENERGY STAR standards is an essential component in the ongoing effort to the focus on energy efficiency in the data center and in the broader economy.

Storage systems represent a major power consumer in the data center, along with computer servers and network systems. TGG has reviewed the ENERGY STAR® Enterprise Storage Draft Specification Framework released by EPA on June 4, 2009. Overall, the document represents a positive initial first step to develop ENERGY STAR® requirements for storage systems. It identifies and addresses the key components of a future storage specification. However, like computer servers, storage systems and devices provide a diverse range of functions and capabilities and support a variety of workloads, complicating the creation of energy efficiency requirements, particularly in the area of power management. As with the development of the computer server requirements, EPA may choose to work with industry stakeholders and other interested parties and work through several iterations of draft specifications to achieve a final specification which sets an achievable hurdle for the industry and satisfies the ENERGY STAR® mandate to reduce data center energy use and deliver more workload for a single energy input. Specific areas that require additional work on the part of EPA, TGG, SNIA and other interested parties include, but are not limited to, the following:

1. Identify the right subset of storage devices to include in the Tier 1 specification. The selected group(s) will need to be balanced to capture a meaningful percentage of storage equipment energy use while restricting it to one to three groupings for which criteria need to be established and tested.

2. Determine the best means to encourage power management capability in storage systems, be it a maximum ideal power requirement, an individual or specific set of power management function(s), or some other criteria, appropriate to storage systems. As in the server discussion on

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

power management, EPA and the industry need to ensure that a balance is maintained between minimizing power use when no work is present and maximizing the work and/or data storage capacity delivered by a given storage system and the attendant energy input.

3. Properly characterize power supply requirements to represent the unique requirements for storage systems while maximizing the consistency between the server and storage specifications to minimize additional work in this area.

4. Provide a qualification system which enables companies to group multiple configurations of a single product model, where they have a consistent set of energy characteristics, into a product family.

the green grid™

The Green Grid looks forward to working with the USEPA and other interested stakeholders, such as SNIA, in the development of an ENERGY STAR® storage specification which furthers our joint objective of delivering more work for each unit of energy applied to the data center.
For clarity of reference, this response provides feedback on each labelled section of the Framework document issued June 4, 2009.

OVERVIEW

The storage products intended for use in data centers can be categorized by a hierarchical or tiered description that differentiates Storage Products by performance and the criticality of the data they contain. The term "Enterprise Storage," as used in the framework document is typically used in the industry to refer to products in the two highest peforming tiers, which correspond roughly to the Online 4 and Online 5 categories of SNIA's taxonomy. If the EPA is interested in having the specification address the full range of Storage Products present in data centers, the term that more accurately reflects the intended range of products that would be in scope would be "Data Center Storage". This term has the advantage of also meeting the industry definition of these types of products.

To date, attempts to provide a single workload benchmark that reflects all, or even typical, end-user applications have not proven to reflect the wide variety of workloads applications present to storage products. The wide range of configurations possible for storage products further complicates the prospect of a single metric identifying the most energy efficient solution. Therefore, multiple metrics may be required to cover different parts of the taxonomy. The Green Grid proposes to work with EPA during the development of the Storage requirements to find the best set of criteria to select systems that maximize the work delivered per unit of energy applied.

A further characteristic of most Storage Products found in data centers, given the need for timely access to data, is that they include a high level of redundancy as standard. These systems regularly include n+1 power and redundant fans to allow for component failures. Many also include logic to allow the unit to continue to service requests even if a controller fails. These factors have the potential to alter the baseline energy profiles of Storage Products, and are present because the target customer requires this level of storage availability in some or all of their application environments.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

BUILDING BLOCK #1: DEFINITIONS

Many of the definitions contained in this section do not reflect generally accepted storage industry definitions, while some other terms are not in common use in the industry. Addendum #1 provides the industry definitions or alternative terms that have industry acceptance.
With respect to the questions for discussion:

the green grid™

1. How are Active, Idle, Maximum and Full Load states defined in the industry?
   It is common practice in the industry for background services (maintenance functions, reliability enhancing functions and capacity-enhancing functions) to be scheduled to run when client initiated operations have ceased. This is not actually an idle system – the functions being performed are desired and expected by the customer, who does not want them to interfere with his production use of the Storage Product; in other words, these background services are seen as necessary work by the purchaser, and the system is not seen to be in an inactive state.
   A storage product is considered to be Active when it is servicing client-initiated requests.

   The terms Maximum and Full Load are not used as such in the industry. The characteristics of Storage Products result in different maxima for each type of I/O operations (reads/writes, random/sequential) qualified by some performance or size characteristic, and so there is never a discussion of a single Maximum operating state. Similarly, the concept of Full Load is highly ambiguous, and can refer to many situations where system performance is bottlenecked in some dimension such as throughput, port load or response time, but this is not a specific condition or state.

2. Storage Products are defined to be Idle from a customer demand perspective according to the definition supplied in the framework. This is not a truly idle state as defined in #3 below.

3. Storage Products engage in a number of background services during the times when there is no client-generated I/O load. These services include maintenance functions performed by both the operating software in the product and the individual hardware components in the product, reliability enhancement functions performed by the operating software, and capacity enhancement functions performed by the operating software. These functions are necessary and expected by customers, and they are specifically expected to be performed in such a way that they do not impact the performance of the customers' applications, thus they are performed when the system is "idle". All storage systems do some type of maintenance activities when no I/O based workload is being presented to or executed by the system. The Green Grid proposes to work with SNIA and EPA to develop a meaningful measurement procedure to accurately capture or assess power use in the mode where I/O work has not been presented to the storage system.

5. The SNIA glossary should continue to be a reference for component elements of Storage Product definitions. The Green Grid proposes to work with the EPA to define a simple taxonomy that segregates equipment by its key characteristics.

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

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

- 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 necessitate 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 these types of power supplies are more prevalent across a wider range of Storage Products than was the case for Computer Servers.

- As with server processors or blade systems, customers often 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.

The 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.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

BUILDING BLOCK #2: ELIGIBLE PRODUCT CATEGORIES

The Storage Industry uses multiple approaches to categorizing products. One is by access time, capacity, and availability, similar to the SNIA Taxonomy, and The Green Grid recognizes this as a useful approach. The industry also categorizes storage according to a hierarchy based on the criticality of data to an enterprise and the business purposes to which this data is put. The Green Grid also recognizes the utility of this approach in supporting customer's matching storage products to business needs. A discussion of this hierarchical approach is in Addendum #2. Finally, interface access methods to storage products (such as Fibre Channel, iSCSI, etc.) can have a significant impact on how they are architected, on the internal software that operates them and in their operation in a data center.

the green grid™

These differences can make comparisons across categories of Storage Products difficult, regardless of the categorization method used. In addition, due to the fact that products in one category generally do not substitute for those in another, customers may not find cross-category comparisons will greatly influence their buying decisions. The Green Grid proposes to work with EPA and SNIA to identify the subset of groups of the SNIA storage classification taxonomy which capture a significant percentage of the storage energy use but provides a universe which allows for timely development of a workable set of ENERGY STAR criteria for storage systems.

BUILDING BLOCK #3: ENERGY EFFICIENCY CRITERIA AND TEST PROCEDURES

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 for both single and multi-volt power supplies, as that is the voltage used by the preponderance of storage systems.

In reference to the Questions for Discussion:

1. Given the earlier discussions of Idle (Building Block #1, point 3) and Maximum (Building Block #1, point 1), The Green Grid recommends that the EPA, SNIA and The Green Grid work to find an acceptable methodology for measuring power use, when no I/O based workload is present on the sytem, and Peak or Active Power as defined earlier in this section as the operational state in which Storage Product energy consumption should be measured. These measurement points should be architecture and access interface neutral and allow meaningful comparisons between products.

4. There is no single recognized workload benchmark for all ranges of customer workloads or all categories of Storage Products. As a result, the existing benchmarks are either reflective of specific workloads or are designed to advantage specific architectures or interface categories. To avoid these biases, and until such time as an accepted power-performance benchmark is developed for storage systems, The Green Grid would recommend that active power measurement be conducted by measuring the peak power usage of each product while in an active state. Partners would have to identify the workload used.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

the green grid™

The proposed SPECsfs2008 benchmark would be highly inappropriate as a generic benchmark for all storage devices. SPECsfs2008 is highly energy consumptive. Systems that perform well on SPECsfs2008 are typically systems with a large number of very high speed drives with sparsely populated data on each drive. Because of this architecture, these systems deliver high SPECsfs2008 numbers because multiple high speed drives are kept spinning with access to their data being highly parallelized. This results in significantly higher energy consumption in terms of GB/Watt and leads to higher overall energy consumption in the data center. A benchmark that is truly reflective of storage energy efficiency would reward storage architectures that concentrate data on fewer drives for energy savings rather than architectures that distribute data sparsely over multiple high speed drives to reduce access latencies. The Green Grid therefore feels that the use of the SPECsfs2008 benchmark in any ENERGY STAR specification for storage would be directly counterproductive to the EPA's objective of saving energy in the data center.

6. A significant fraction of Storage Arrays are made up of many assemblies which utilize their own redundant power supplies. As a result, a method is needed for combining these assemblies to present an aggregate power measurement back to the data center user. In addition, some products integrate Network Switches, Hubs, "off the shelf" UPS modules, or other low power assemblies into the system. Since there are no ENERGY STAR requirements in place today for these assemblies, it is recommended that they should be excluded from the power supply efficiency and thermal monitoring requirements identified in this framework if their total contribution to the system load is less than 10% of the total, and the total power consumption of the system exceeds 1 kW.

7. Given the need for highly available Storage Products, power supplies for these systems are designed to be run in redundant configurations. As a result, they typically run with loads ranging from 20% to 45%, unless one of the supplies in a system has failed.

8. The Green Grid has provided comments against the adoption of the "Net Power Loss" methodology throughout the development of the Tier 1 requirements for computer servers. EPA should reference these comments for details on The Green Grid's concerns with the Net Power Loss methodology.

9. The VAR sales channel has a significant role in the Data Center Storage market. Without an agreed definition of product families, it is difficult to determine what impact the difficulty and expense represented by the requirements above would have on third-party sales channels. 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 of allowing 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 customer's 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

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

value to the ENERGY STAR® brand or to customer understanding of ENERGY STAR®. The more efficient, understandable approach would appear to be 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.

BUILDING BLOCK #4: INFORMATION AND MANAGEMENT REQUIREMENTS RESPONSES TO QUESTIONS FOR DISCUSSION:

1. Storage Products that can be found on the market today typically do not include the ability to measure or report the input power or inlet temperature data identified in the framework. Measuring and reporting capacity utilization is typically available on demand via the management interface(s) to a Storage Product. It is not, however, typical to provide performance utilization information as part of default product capabilities, and it may not be a supported feature on some products. Inclusion of these unsupported or optional features as part of the required feature set for ENERGY STAR labelled products will require changes to both product content and business practices, which should be considered in the decision-making process.

2. A key element of energy consumption in Storage Products is driven by the type (Solid State Drives, Fibre Channel, SATA, e.g), number and spindle speed (7k, 10k, 15k rpm) of disks or the number of SSD or flash drives included in any given configuration. This means that it is not sufficient to specify the capacity at which measurements were taken, but the number and type of each drive class used to achieve the measurement must be delineated.

3. Capacity utilization is routinely discussed in the industry as a fraction of either Raw Capacity or Consumable Capacity as defined in Addendum #1. Performance capacity is measured in a variety of ways depending on the category of Storage Product and its architecture and configuration. There is relatively little commonality to the way this class of data is reported. Purchasing decisions, either of capacity expansion or additional Storage Product systems, are driven as a result of specific understanding of the application environment supported by the storage, the expected capacity growth that will be experienced when compared to the actual consumption, the performance each application is experiencing (as measured by the application, not the storage system), and the archiving and data reduction policies of the data center. Within any given data center the criteria typically vary between application environments, and require information beyond that which the Storage Product can provide.

4. At this point in time, there are no industry-standard protocols for measuring temperature, utilization, etc.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

ADDENDUM #1: DEFINITIONS

a. Storage Hardware

- Storage Media: For consistency, this should refer to tape media, not tape drives, and optical media, not optical drives.
- Storage Product: The final sentence of this definition is inconsistent with the apparent scope of this specification, and should be stricken. We would further recommend that the definition be amended to read ""A system composed of integrated storage controllers, storage media, optional embedded network elements, and software that provides data storage services to one or more Computer Servers and/or other devices."
- Storage Controller: We would recommend using the SNIA definition of an intelligent controller.
- Storage Product Family: Given the way products are structured and configurations vary with respect to interface speed options as well as capacity options and granularity this definition may require refinement.
- I/O device: We would recommend adding the following to this definition. "I/O devices do not initiate user I/O; they are used for the purpose of transmitting data and/or management information."

b. Storage Characteristics

- Capacity: The definition offered is what is known in the industry as "Raw Capacity". The industry and users distinguish between Raw Capacity and Consumable Capacity, which is the storage capacity available after formatting, system overhead, RAID parity and sparing has been deducted. We would recommend the use of these terms to provide clarity in measurement methodologies.

c. Other Data Center Hardware:

- Network Equipment: We would suggest the first sentence be modified to read ""A product whose primary functionis to provide data connectivity among an arbitrary combination of devices connected to its several ports." to differentiate data center network equipment from embedded switch components.

d. Power Supplies:

- Given the differences between server and storage power supplies, it would be more appropriate to define the secondary ouputs as a percentage of the primary rather than in absolute terms.

e. Operational States:

- Idle: While this does represent a valid state from the perspective of client-initiated workload, Storage Products are generally not in an unproductive state at this time. The only time in which Storage Products are not doing useful work is when all background services have been completed and the system is exclusively in a wait state. Attainment of a state where the storage device is neither executing I/O directed workload or maintenance functions will vary on the machine type and its specific application.

the green grid™

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.

ADDENDUM #2: ELIGIBLE PRODCUT CATEGORIES – HIERARCHICAL STORAGE

Energy efficient storage practices depend on more than energy efficient storage products. The use of the storage products and practices designed for specific tiers in the information hierarchy provide data center operators with the ability to reduce overall energy consumptionby moving information from high-performance (and more energy consumptive) products to lower performing (and less consumptive) products as that information progresses through its lifecycle. This results in energy savings in two ways: first, it frees up capacity on high-performing storage, eliminating the need to bring more of it online, avoiding increasing energy spend, and second, migration to the proper tier allows greater use of energy-saving techniques such as removeable media or disk drive spin-down.

the green grid™

The hierarchy of storage tiers reflects the criticality of information and its availability to a business, and is generally defined as:

- Tiers 1 – Data access is required in <80 ms, and must be continuously available, with access interruptions lasting less than minutes or seconds annually. Business applications using this data demonstrate dynamic workloads with the highest transaction loads of any applications in an enterprise. If a service interruption occurs, the business has a goal of zero or only seconds of data loss for this data. Architectures in these tiers have no single points of failure (SPOFs) and provide for non-disruptive serviceability.
- Tier 2 – Data access is required in <80 ms, and must be very highly available, with access interruptions lasting less than hours or minutes annually. Business applications using this data demonstrate constant, heavy loads with high performance demands. If a service interruption occurs, the business has a goal of seconds to minutes of data loss for this data. Architectures in this tier usually have no SPOFs and generally provide for non-disruptive serviceability.
- Tier 3 – Data access is required in <80 ms, and has an availability target that supports access interuptions lasting up to a few hours annually. Business applications using this data demonstrate primarily read access to information stored in this tier with moderate performance requirements. If a service interruption occurs, the business can tolerate the loss of several minutes of data stored in this tier. Architectures in this tier may have SPOFs or require disruptive service.
- Tier 4 – Data access can take >80 ms, and has an availability target that supports access interruptions lasting up to a few hours annually. Business applications using this data demonstrate primarily read access to information stored in this tier, with moderate performance requirements, and may tolerate “internet-like” times to first access. If a service interruption occurs, the business can tolerate the loss of up to the last 24 hours of data.

©2009 The Green Grid. All rights reserved. No part of this publication may be used, reproduced, photocopied, transmitted or stored in any retrieval system of any nature, without the written permission of the copyright owner.