July 3, 2009

TO: Stephen Pantano
ICF International
1725 Eye Street, NW, Suite 1000
Washington, D.C. 20006
spantano@icfi.com

CC: Andrew Fanara
EPA
fanara.andrew@epa.gov

FROM: Chris Hankin
Sun Microsystems, Inc.
chris.hankin@sun.com
202-326-7522

Re: Comments by Sun Microsystems, Inc. on the Energy Star Enterprise Storage Draft Specification Framework of June 4, 2009

Dear Mr. Pantano:

Thank you for the opportunity to provide comments on the EPA's Energy Star Enterprise Storage Draft Specification Framework of June 4, 2009. Sun Microsystems looks forward to continuing to help achieve a successful new specification.

We commend the EPA on their careful consideration of the input provided by Sun Microsystems and by the industry on previous drafts of previous Energy Star specifications.

The comments that follow reflect Sun's deep expertise and technological market leadership in the storage area, and have been made keeping in mind the objective of achieving an energy efficiency specification for storage systems that are useful and valuable to the purchasing decisions of enterprise storage customers. These comments also reflect our mutual learning from the recently completed Energy Star for Servers Tier 1 Specification.

We look forward to working with you to achieve a useful and productive specification for the industry.

Sincerely,

Chris Hankin
Sun Microsystems, Inc.
chris.hankin@sun.com
202-326-7522

COMMENTS BY SUN MICROSYSTEMS, INC.
ON THE EPA ENERGY STAR ENTERPRISE STORAGE DRAFT SPECIFICATION FRAMEWORK

1. Introduction

Enterprise storage is a highly complex area, and products sold in this category have significantly greater diversity and complexity than either enterprise servers or enterprise networking. Any attempt by the EPA to specify an ENERGY STAR framework for enterprise storage must therefore necessarily have well defined scope and limited reach in its initial version (Tier 1).

Energy efficiency of enterprise storage is also significantly more complex and challenging than energy efficiency of enterprise servers. Enterprise requirements for storage systems include, besides energy efficiency, characteristics such as: availability, reliability, access latency, expandability, etc.; in addition, features such as backup, deduplication, and other maintenance activities are also important. Any metric for energy efficiency therefore must also account for these enterprise features.

In addition, it is important that any energy metric for storage recognize and reward innovation that reduces overall energy consumption at the data center level. Experience with the enterprise server specification has shown that picking the wrong basis for taxonomy for server categorization leads to inefficient servers being rewarded with the ENERGY STAR designation while advanced, innovative, energy saving servers are penalized. The ENERGY STAR spec for storage must ensure that the taxonomy basis, and the energy efficiency metrics chosen reward architectures that enable data center level savings and do not penalize them.

2. Building Block #1: Definitions

2.1 Product Categorization

Rather than recreate its own categorization for enterprise storage, Sun recommends that the EPA leverage the excellent work in creating storage taxonomies already done by SNIA. The SNIA enterprise storage taxonomy groups the universe of enterprise storage products into the following major categories: Online, Near Online, Removable Media, Virtual Media Library, Appliance and Interconnect. These are further subcategorized based on Access Pattern, Max Time to Data (TTD) and User Accessible Data. In total, 22 subcategories cover the entire universe of storage products.

We recommend that the EPA pick a subset of categories from the existing SNIA taxonomy to define the scope of applicability for Energy Star for Storage.

2.2 Operational States

The proposed operational states of Idle, Active, Maximum, and Full Load are not meaningful for enterprise storage. A user centric definition of idle might imply that no data I/O is being performed by the storage system in the service of any user application. However, in most enterprise storage systems, when no user I/O of stored data is being performed, many background activities still continue. These background activities might include: backup, deduplication, disk defragmentation, and other maintenance activities. As such, even an "idle" storage system is never really idle. Hence, any semantic that seeks to define operational states as Idle, Active, or Maximum needs to carefully take into account these features of enterprise storage systems.

3. Building Block #2: Eligible Categories

There are multiple approaches to categorizing storage products. Below is a brief summary of these approaches:

- **By Tier**: storage products are organized in tiered hierarchies. Data that needs to be immediately available at low latency is stored in Tier 0 or Tier 1 storage. Data that is required to be online but can tolerate slightly longer access latencies is stored in Tier 2 storage. Data that has lower availability requirements is stored in Tier 3 storage. Archival data that can tolerate access latencies of up to a few hours is stored in Tier 4 storage.

- **By Capacity**: Storage products are sometimes categorized by the total amount of storage capacity (GB/TB/PB).

- **By Technology**: Storage products are sometimes categorized by their technology, e.g. tape storage, disk storage, and flash storage. The industry is innovating rapidly to create hybrid products that combine multiple technologies in order to best serve the storage needs of specific applications.

- **By Availability**: Storage products are categorized by the degree of redundancy engineered within the product. Products with higher redundancy will consume more energy, but are better suited for enterprise applications where data retention is critical.

- **By Network Interfaces**: Storage products are categorized by the access method. Access methods can vary from proprietary direct access storage (DAS) to various networked access mechanisms (FibreChannel, iSCSI, etc.).

Sun recommends that the EPA work with SNIA to identify a subset of the categories in the existing SNIA classification taxonomy which will fall within scope of ENERGY STAR for Storage Tier 1.

4. Building Block #3: Energy Efficiency Criteria and Test Procedures

4.1 Power Supplies

Sun recommends that any power supply specification for ENERGY STAR compliant storage systems remain consistent with the power supply specification already in effect for ENERGY STAR for Servers Tier 1. This is because many vendors manufacture both servers and storage systems and for reasons of supply chain efficiency use common power supplies for both their server and storage chassis.

4.2 Net Power Loss Approach

The net power loss approach for power supply efficiency is untested and unproven. None of the organizations that are domain experts in power supplies (Climate Savers, 80Plus.org) use this approach. The most common method of specifying power supply efficiencies is through specifying the shape of the efficiency vs. load curve. The EPA has successfully used this approach in ENERGY STAR for Servers Tier 1. We recommend remaining consistent with this approach and against the net power loss approach.

4.3 Operational Modes

As specified above, the proposed operational modes of Idle, Active, Full Load, and Maximum are confusing because their semantics remain unclear. They are also not standard terminology in the storage industry. Sun recommends against the use of these modes, and suggests that the energy efficiency of storage systems be determined by a combination of metrics involving throughput, reliability, availability, access latencies, and other features on a per Watt basis.

4.4 Workloads and Benchmarks

No single workload benchmark for storage is reflective of the multitude of ways in which customers use storage systems. All existing benchmarks are reflective of specific workloads. As such, Sun recommends that the ENERGY STAR for Storage specification be required to report peak power usage of each product while being actively used. The storage vendor would be allowed to use any workload for this purpose as long as it was identified.

4.5 Proposed Use of SPECsfs2008 Benchmark

The proposed SPECsfs2008 benchmark would be 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. While this approach may yield good numbers for GB/second throughput, it 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. Sun 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.

5. Building Block #4: Information and Management Requirements

Sun supports the objective of transparency in making as much information about the energy use of storage systems available to customers as possible. For this reason, Sun makes the following recommendations on the information reporting and management requirements for storage systems:

- Many storage systems being shipped today do not have mechanisms to report system level power consumption, inlet air temperature, or utilization. For the ones that do, a variety of different protocols are used. In order to take the first step toward transparency and to encourage the next generation of storage systems to start reporting this information, Sun recommends not standardizing on any particular protocol just yet. For Tier 1, manufacturers should be permitted to report the necessary information in any protocol of their choosing.

- The proposal for measurement accuracy for input power (+/- 5% or +/- 5 Watts, whichever is greater) is too stringent. The reason this proposal is onerous is because instrumented power supplies cannot report measurements of this accuracy, especially toward the low end of the range. This accuracy requirement for input power is also inconsistent with the accuracy requirement articulated for servers in ENERGY STAR for Servers Tier 1. In order to maintain consistency in the power supply industry, and also to permit manufacturers to single source power supplies for both their storage and server products, we recommend that the measurement accuracy for input power be maintained at +/- 10% or +/- 10 Watts.