ENERGY STAR®
Data Center Storage
Version 1.0

First Stakeholder Meeting

The Sainte Claire Hotel
San Jose, CA
20 July 2009

Agenda

11:00-11:15 – Welcome & Introductions

11:15-11:25 – ENERGY STAR Program Overview

11:25-11:40 – ENERGY STAR Data Center & Server Initiatives

11:40-12:00 – V1.0 Data Center Storage Objectives & Status

12:00 – Break / Start of Working Lunch

12:15-12:25 – Major Themes from Stakeholder Comments

12:25-15:30 – Detailed Review of Comments

- Support for ENERGY STAR, Terminology & Definitions, Market Segmentation, Efficiency Approach, Metrics & Benchmarks, Power States, Software & Hardware Strategies, Reporting

14:00 – Break

15:30-16:00 – Next Steps & Closing Comments

16:00 – Adjourn

ENERGY STAR Overview

- What is ENERGY STAR?

- A voluntary public-private partnership program
- A strategic approach to energy management
- Recognized by over 70% of Americans
- An internationally recognized brand

ENERGY STAR Overview

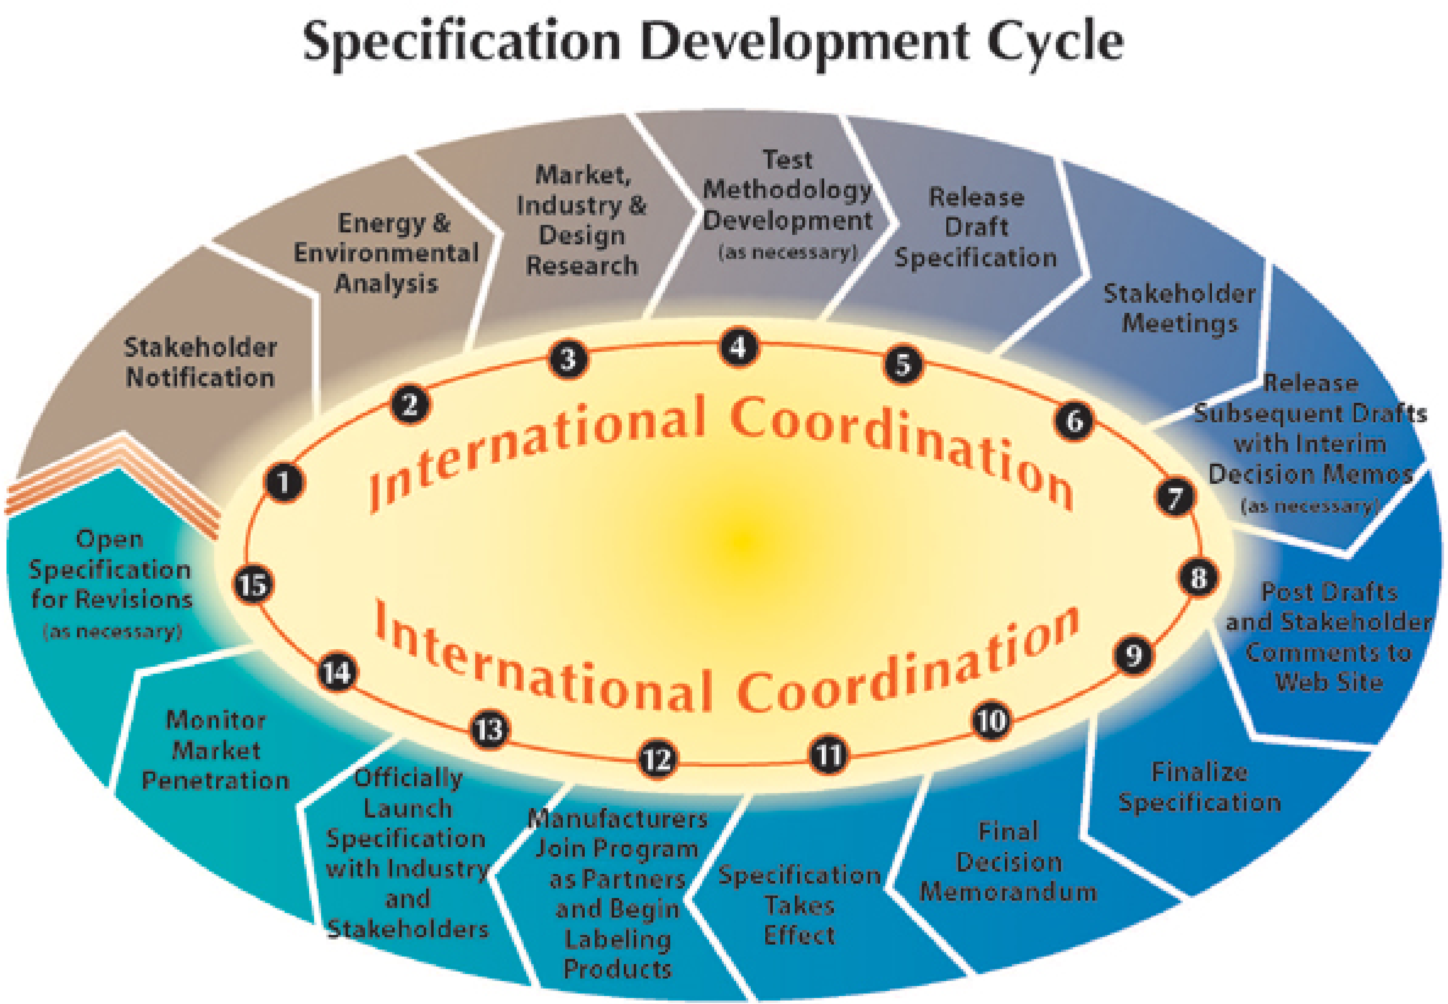

ENERGY STAR Overview

- Guiding Principles of Specification Development
  - Cost-effective efficiency
  - Performance maintained or enhanced
  - Significant energy savings potential
  - Efficiency improvements are achievable with non-proprietary technology
  - Product differentiation and testing are feasible
  - Labeling can be effective in the market

Data Center Building Initiative

- ENERGY STAR for Building Infrastructure
  - How: Build on existing ENERGY STAR platform. Rating on 1-100 scale.
  - What: Building-level assessment of infrastructure (cooling & support systems) performance of both stand-alone data centers & those in an office or other building.
  - Unit of Analysis: Ratio of Total Energy/IT Energy. Ideal metric would be measure of energy use/useful work – this is still under development.
- Final data analysis is underway. Plan to launch as part of the ENERGY STAR Portfolio Manager in January, 2010.

Computer Servers Specification

- Tier 1 requirements effective May 15, 2009
  - Power Supply Efficiency and Power Factor requirements for single- and multi-output AC-DC and DC-DC power supplies.
  - Idle power consumption requirements for 1- and 2-socket servers, with allowances. Requirements scale with performance capability.
  - Power management requirements for 3- and 4-socket servers to reduce power consumption in periods of low utilization.
  - Standardized reporting of power, temperature & processor utilization.
  - Standardized data measurement and reporting via a Power & Performance Data Sheet.

Computer Servers Specification

- Tier 2 requirements under development
  - Evaluate active performance via adapted benchmark methodologies.
  - Expand Program Scope to include:
    - servers w/ high proc socket count
    - multi-node servers
    - server appliances
  - Enhance requirements for power supply efficiency and sizing.
  - Encourage further application of power management techniques.

Goals for Data Center Storage

- Encourage widespread adoption of energy efficient hardware and software strategies,
- Provide purchasers with the means to identify the most energy efficient enterprise storage solutions for their specific end-use application, and
- Provide tools and information to designers and mangers looking to improve the efficiency of data center operations

Storage Development Status

Framework & Analysis

(May-July)

Draft 1

(August)

Announcement

(April)

Goals for Today

- Review & discuss comments with an eye towards the following:
  - Enable standardized measurement & reporting of key metrics
  - Develop criteria that make ENERGY STAR storage products unique
  - Identify opportunities to leverage servers work
  - Foster competition between OEMs to deliver the best combination of performance and information
  - Drive demand by demonstrating the value of energy efficient storage products
  - Define and encourage best practices

Goals for Today

- Separate the ends from the means
  - Identify key contributors to efficiency versus enabling tactics
- Clearly define the system boundaries
  - Where do you put the label?
  - View the system holistically (software + hardware, etc.)?
- Develop reasonable energy efficiency metrics
  - Proxies may be suitable in the near term
- Identify the similarities in market segments as well as the differences
  - Differences are much easier to see
- Develop a deployment strategy for storage energy efficiency programs
  - Go after low hanging fruit today, lay out groundwork and testing for tomorrow

Litmus Test

- There are numerous product features, functions, and data management strategies that enable energy savings in data center storage.

- There is only one end result that matters: The ability to do more useful work, while consuming fewer resources, in a verifiable and quantifiable manner.

Lunch Break

EPA would like to acknowledge the support of the Storage Networking Industry Association (SNIA) for logistical support and for sponsoring today's food & refreshments

Major Themes from Responses

- Support for ENERGY STAR
- Abundant Opportunities for Efficiency
- Don’t Call it “Enterprise”
- Unique Segments Demand a Unique Approach
- No Consensus on Power States
- Respect RAS Requirements
- Storage PSU ≠ Server PSU
- Definitions Need to be Updated

Support for ENERGY STAR

SELECTED COMMENTS*:

- ENERGY STAR for data center storage makes sense and is applicable to help customers, manufacturers and VARs align the most applicable and effective storage technology to their application.
- Storage is an important area for energy optimization and efficiency improvements. The major power draws for storage systems are:
  - Spinning HDDs and their enclosures (on average, 66%–75%)
  - Controllers and I/O connectivity components (most of the balance)
- The three primary objectives are exactly on target. The second objective, to allow customers to objectively evaluate the energy consumption of storage arrays, is the most challenging.
- The storage industry is working on a three dimensional problem; assessing performance, RAS, and energy efficiency. Often it is easy to configure or design a given storage system for two of the three elements, while degrading the third. The challenge to the storage industry is to embrace all three in an ENERGY STAR specification with no one element being degraded in value.

EPA

Note: All comments in this presentation are abstracted and paraphrased for consistency and clarity. Every effort was made to maintain the author's intent. Original comments will be posted to the ENERGY STAR Web site with permission.

Support for ENERGY STAR

SELECTED COMMENTS:

- Define a standard of improvement that is significantly higher than that which was inferred from the framework.
- Propose that vendors are allowed to meet this standard by choosing the combination of power supply and software that aligns with their specific market focus, customer requirements, and technology portfolios.
- ENERGY STAR would be a good thing assuming that it addresses the various types of storage. The approach needs to factor in both hardware and software functionality without being too overall encompassing so as to be useless.
- Above all, the specification must be of benefit to the buying organizations providing insight into both efficiency while doing work, while storing information and providing a return on investment.

Terminology

SELECTED COMMENTS:

- Recommend using the term Data Center Storage, as there is no exact industry definition for Enterprise Storage. The intent of the EPA specification is to address storage in any data center, with the exception of small/home office (SOHO).
- Concerned with the usage and implication of “enterprise storage systems” as being exclusive to only high-end, expensive storage solutions.
- The draft framework should reference the SNIA Dictionary and its definitions wherever possible.

Market Segmentation

SELECTED COMMENTS:

- There are several approaches to categorizing storage products:
  - Access Time, Capability, Availability;
  - Data Criticality and Business Purpose
  - Interface Methods
- While segmentation is sometimes drawn in terms of communications protocol (NFS, CFS, etc.), these have limited impact on the energy consumption of the underlying product.
- Differences used to exist between block and file (NAS)-based solutions and enterprise (mainframe) solutions, and open systems or modular and monolithic systems, but now the lines are blurred. All of these systems can scale in terms of performance, capacity, availability, physical size, functionality, and connectivity.
- Some storage systems use custom controllers while others leverage general purpose commercial servers including ENERGY STAR compliant models for their controllers.

Market Segmentation

SELECTED COMMENTS:

- Alignment of the most appropriate tier of storage to application needs is an effective technique to address data center bottlenecks without continuing the vicious cycle of sparse storage allocation and later consolidation.
- Storage systems combine various tiered storage media with tiered access and tiered data protection.
- Tiered Access (I/O)
  - High-speed 8-Gb Fibre Channel (8GFC) and 10-Gb Fibre Channel over Ethernet (FCoE) versus 4GFC or low-cost 1-Gb Ethernet (1GbE), or
  - High-performance 10GbE-based iSCSI for shared storage access versus serial attached SCSI (SAS) and Serial ATA (SATA) for direct attached storage (DAS)
- Tiered Data Protection
  - Local and remote mirroring (also known as replication), in different RAID levels, point-in-time (pit) copies or snapshots, and other forms of securing and maintaining data integrity and meet data protection RTO and RPO requirements.

Market Segmentation

SELECTED COMMENTS:

- Tiered Storage Media
  - Tier 1: Online, Highest performance, typically SSD. Data access <80 ms, continuously available, access interruptions less than minutes or seconds annually. Architectures have no single points of failure (SPOFs) and provide for non-disruptive serviceability.
  - Tier 2: Online, very high performance, highest reliability, typically 10K/15K traditional enterprise products (SAS, SCSI). Data access <80 ms, and very highly available, access interruptions less than hours or minutes annually. Architectures usually have no SPOFs and generally provide for non-disruptive serviceability.
  - Tier 3: Near-line, High capacity, high reliability. Data access <80 ms, access interruptions up to a few hours annually, can tolerate the loss of several minutes of data stored. Architectures may have SPOFs or require disruptive service.
  - Tier 4: Offline: Backup medium, removable storage such as tape. Access >80 ms, access interruptions up to a few hours annually, may tolerate “internet-like” times to first access, can tolerate the loss of up to the last 24 hours of data.

Market Segmentation

SELECTED COMMENTS:

- Suggest limiting the scope to a subset of the SNIA taxonomy in further developing the specification.
- Need to capture a meaningful percentage of storage equipment energy use while restricting it to one to three groupings.
- Recommendations:
  - Online-2, Online-3, and possibly Online-4.
  - Online-2, Online-3, Online-4, Removable-2, Removable-3
  - Online, Near-online, Removable-media Library, Virtual Media Library
  - Online-2, Online-3, Removable-2, and Removable-3 for 1st release.
  - Online-4, Virtual Media Libraries, Near On-line, and Appliances for 2nd release.
  - Do not recommend Online-1 (SOHO) or Online-5 (small niche).
  - Restrict Tier 1 to "DAS" and "NAS" but only for HDD and SSD.
  - Include one or more from “Near On-line” and “Removable Media Library” to send a clear message to data center operators to tier their storage.

Market Segmentation

SELECTED COMMENTS:

- Concerned about excluding aggregating storage elements such as RAID, tape libraries, filers, virtual tape libraries, etc. Unified storage products that combine blades, external array, block & file, back-up, etc., are becoming increasingly prevalent.
- Products with integrated Information Lifecycle Management (ILM) , should be explicitly included. Moving data from high-performance products to lower performing (and less consumptive) products results in energy savings via the following:
  - Free up capacity on high-performing storage, eliminating the need to bring more of it online, avoiding increasing energy spend
  - Migration to the proper tier allows greater use of energy-saving techniques such as removable media or disk drive spin-down.
- Include “Server Storage Products”: systems composed of a server and one or more forms of storage media (e.g. including hard drive, tape, optical disc, or solid state memory), integrated storage controllers, and software.

Discussion Break

Efficiency Approach

SELECTED COMMENTS:

- Energy avoidance can be accomplished by powering down storage. Energy efficiency can be accomplished by using tiered storage for different needs.
- There is a continued belief that unplugging disks is the best way to reduce the energy bill – classifying data and getting rid of stuff is still the best method to reduce costs, albeit an elusive one.
- For active data, using a high-performance drive to do more work using fewer HDDs is a form of energy efficiency.
- Migrating data to larger capacity drives can reduce power consumption by 5 to 20%.
- Saving power doesn't derive from compressing and de-duplicating bits in a thinly provisioned and tiered infrastructure. Those are tactical approaches that deliver short term gains but no real long term improvements.

Efficiency Approach

SELECTED COMMENTS:

- Concerned that the approach does not accommodate for any high-availability or data integrity features, including redundant power supplies, processors, ports, fans, mirrored memory, "spare" CPU cores, integrated standby power, vault drives, remote replication links, etc.
- Maintaining reliability and integrity of data is never free and must somehow be accounted for in evaluating energy efficiency.
- Reliability and integrity of data are core reasons for purchasing RAID arrays in the first place, so if a purchaser buys a system that uses less energy at the expense of less reliability, they will likely configure greater redundancy into the array and use more energy.
- Concern that a singular pursuit of energy reduction / energy ratings can be orthogonal to performance and data protection, especially when procurement mandates in certain industry sectors may be dictated by the energy ratings alone.

Efficiency Metrics

SELECTED COMMENTS:

- If using storage for active data, consider how much work can be done per unit of power such as IOPS (requests) per watt or bandwidth (bytes transferred) per watt.
- If data is inactive, consider the energy required to support a given capacity density while keeping in mind that some amount of performance will be needed.
  - It is possible to create a storage system that uses less power, but if performance is compromised, multiple copies of the low-power system may be needed to deliver adequate performance.
- Standard throughput ("requests per second" and "bytes transferred per second") and response time ("response time per request") metrics do not necessarily capture the true utilization of a product.
- Maintaining compliance to regulations such as Sarbanes-Oxley and medical, legal, and financial data retention requirements rarely gets recognized as useful work.

Efficiency Metrics

- Measure and compare storage outputs (useful work) versus inputs (resources consumed) across all operating conditions (idle, active, etc.)
- The definition of useful work varies by the type of data being stored:
  - Speed & Bandwidth are most important for Online (Active) Data
  - Capacity is most important Offline (Idle) Data
- All storage products consume the same resources, with consideration given to reliability & availability requirements:
  - Energy to Operate
  - Energy to Cool
  - Physical Footprint
  - Capital Cost

Efficiency Metrics

- Input: Energy
  - The following hardware strategies to improve energy performance have been identified:
    - Power Supply Efficiency
    - Fan Efficiency

- Output: Speed & Bandwidth
  - The following hardware and software strategies to improve speed and bandwidth per unit energy have been identified:
    - High Performance Media
    - High Performance I/O Devices

Efficiency Metrics

- Output: Capacity
  - The following hardware and software strategies to improve capacity per unit energy have been identified:
    - Tiered Storage & Information Lifecycle Mgt.
    - Inactive Media (Tape, Optical, MAID)
    - High Density Media
    - Data De-duplication
    - Data Compression
    - Thin Provisioning
    - Virtualization?
    - Thin Snapshots & Clones?
    - RAID 5/6?

Benchmarks

SELECTED COMMENTS:

- There is no single recognized workload benchmark for all ranges of customer workloads or all categories of Storage Products. Existing benchmarks are either reflective of specific workloads or are designed to advantage specific architectures or interface categories.
- To date, attempts to provide a single workload benchmark that reflects typical end-user applications have not proven to reflect the wide variety of workloads present in storage products.
- We welcome the EPA's wish to develop a test procedure to measure energy consumption under load. We encourage use of industry standards where available, but would be in favor of new development if necessary, rather than have no definition.
- We welcome the desire to define utilization of storage equipment and any provision for consideration of features that improve overall utilization. Even small improvements in utilization can have a large overall impact on data center efficiency.

Benchmarks

- SNIA Green Storage Idle Power Metric
  - Metric = Capacity / Power Consumption
    - Capacity = Raw, Unformatted, Uncompressed drive capacity
  - Conditioning for a minimum of 15 minutes. The duration of storage conditioning shall be disclosed.
    - Online & Near-online Conditioning: Transfer size: 8 KiB, Operations Mix: 70% read, 30% write, Transfer Alignment: 8 KiB, Transfer Offset: Randomly distributed throughout the address range.
    - Removable & Virtual Media Conditioning: Transfer size: 128 KiB, Operations Mix: 100% write for the first half of the conditioning phase, followed by a rewind request; 100% read for the second half of the conditioning phase.
  - Idle Test Duration of 24 hours (2 hours for Removable Media)
  - Collect power consumption and inlet temperature data

Benchmarks

SELECTED COMMENTS:

- Other industry benchmarks:
  - SPECsfs2008 for NAS
    - Considering the growth in storage of unstructured data, NAS needs to be considered as a significant piece of data center storage. SPEC based benchmarks are a de-facto standard for file serving measurements.
    - The proposed SPECsfs2008 benchmark would be highly inappropriate as a generic benchmark. It is valid for file system tests only. It would not be appropriate for block-based storage systems, backup tests, or archiving.
    - SPECsfs2008 is highly energy consumptive – multiple high speed drives are kept spinning with highly parallelized data access.
  - SPC is popular for block storage and supported by some vendors, but not endorsed or fully embraced by all, including some industry leaders.
  - Others: Microsoft ESRP for Exchange email application workload, Jet for Microsoft SQL, VMware VMARK, TPC for databases , Various vendor- and application-specific benchmarks.
  - The industry is working to develop new benchmarks for SSD, tape backup, and de-duplication devices.

Benchmarks

SELECTED COMMENTS:

- A benchmark for storage energy efficiency should reward architectures that concentrate data on fewer drives for energy savings, rather than architectures that reduce access latencies.
- There is a need to account for different active workloads, both large and small IOP to address different scenarios for effectiveness.
- A complicating factor is that some storage products are designed to simply hold data, others for I/O performance, and still others are designed for both.
- Storage systems, in order to support Internet and Web 2.0 applications, will need to support variable performance from small random access of meta-data or individual files to larger streaming video sequences.
- Semi-structured email data and unstructured file data has created the biggest data footprint impact. Unstructured data has varying input/output (I/O) characteristics that change over time, as in the case of a video file temporarily becoming popular on a social networking website.

Benchmarks

SELECTED COMMENTS:

- Suggestions for proxy metrics and further development:
  - Until an accepted power-performance benchmark is developed for storage systems, recommend active power be measured as the peak power usage of each product while in an active state. Partners would have to identify the workload used.
  - Recommend energy consumption be measured at Hardware Idle and Peak Power. These are architecture and access interface neutral and allow meaningful comparisons between products.
  - We do not consider power at idle to be a useful benchmark of energy efficiency, unless it is shown to be a useful predictor of power consumption under load.
  - Assuming that a rough representative workload was defined, it would have to be usable for storage products of all different sizes and capabilities. Advise keeping it incredibly simple:
    - 8 kb requests, No writes tagged as write-through, No cache flush requests, Cover the entire space of storage capacity
    - 25% random reads; 25% random writes; 25% sequential reads; 25% sequential writes, with constant number of outstanding requests for each type (workload intensity would depend on this)
    - Run for several hours continuously before capturing power consumption

Discussion Break

Power States

SELECTED COMMENTS:

- Idle Power State:
  - Idle is probably the most important potential power savings area, since individual storage media is idle more often than it is active.
  - The storage industry generally accepts that a backup or archive system spends 60 to 80% of the time in the idle state.
  - The notion of "idle power" seems to be a meaningless data point, as customers tend to buy "Enterprise Storage" when they have 24 x forever data access requirements. It is perhaps by definition that there is never any idle time on an Enterprise Storage Platform.
  - Idle may become a problem in the future, when devices implement multiple low-power states. As more intelligence is being added to storage devices, more activity from otherwise idle systems will arise.
  - For some online storage systems, an idle system consumes as much as 80 to 85% of the power draw of an active storage system.

Power States

SELECTED COMMENTS:

- "Background" functions include:
  - Disk management: provisioning disks, de-fragmentation, spreading data across multiple disks for performance and protection
  - Data protection services, which includes: encryption, backup, data monitoring, including multiple copies (snap copies, volume copies, remote synchronous/synchronous copies) for different access or locations and snap copy management, and finally data recovery.
  - Storage pool management: data monitoring for lifecycle aging, presentation of different LUN sizes for better mapping of storage resources for optimization of utilization of allocated storage resources..

Power States

SELECTED COMMENTS:

- Active Power State:
  - Active will be difficult for the industry to come to consensus on, since whatever workload is chosen will work well for some products and not as well for others.
  - It is highly unlikely that all units of storage media will be fully active at the same time. Some type of active power measurement will best define the power use of a storage product.
  - Some systems cache highly active data separate from the storage media to reduce response time and improve data availability.
  - Knowing the peak power requirements of a system tells you nothing about the actual power requirements under normal operation. It would be the rare storage array that was continually run at 100% utilization.
  - One could argue that the appropriate operating "sweet spot" is just below that point where any possible component failure would result in a significant reduction of performance.

Power States

ENERGY STAR

SELECTED COMMENTS:

- Operational States are not defined in terms that align with current storage system architectures. Tier 1 storage systems are never truly idle due to management functions these systems provide. Tier 3 storage systems can clearly be classified with idle and active states.
- Alternate Proposals:
  - Idle or Low-power States 1, 10, and 30 – where the number indicates the maximum latency in seconds required to service any 4 kb request that arrives when the storage product is in a low-power state.
  - There are 3 power states: (1) when client-generated I/O is occurring, (2) when maintenance and data management functions are being performed, and (3) when no work of any type is present.
  - Active Idle, Performance Idle, Low Power Idle
  - Host/Initiator Idle/Active, Target Idle/Active, System Idle/Active; as defined by the presence of pending and/or active workload transactions initiated by a host or self-initiated by the storage product.

EPA

Power States

SELECTED COMMENTS:

- Alternate Proposals:
  - Idle: State in which the OS and other software have completed loading, but no active workload transactions are requested or pending by the system. The tiers of disk drives in the system are set into standby mode. No activity occurs on any primary storage controller. In standby mode, the mechanical head assembly is parked within the disk drive, and the spindle is stopped. Fan usage is automatically shut-down in idle state. and the drive logic continues to listen for new commands.
  - Active: A system state where all drives are spinning, the drive logic is issuing disk access and write commands, and maintenance functions are being performed. When active, a storage system's 'background' operations are being performed to support: regulation compliance, data accessibility, security and disaster recovery.
  - Maximum/Full Load: State where all disks are made available to all datasets through a common pool of storage, both performance and capacity utilization are maximized. When disk space is no longer needed by a particular application, it is returned to the free pool and made available to other applications as their storage needs grow.

Discussion Break

Litmus Test

- There are numerous product features, functions, and data management strategies that enable energy savings in data center storage.

- There is only one end result that matters: The ability to do more useful work, while consuming fewer resources, in a verifiable and quantifiable manner.

Software Strategies

SELECTED COMMENTS:

- Some software is very tightly integrated with storage systems and thus should be considered as functionality, while others are add-on and may in fact span server and storage. Software needs to be considered in the big picture, however it also needs to be put into perspective. An emphasis on optional benefit would be desirable, or, create a model where additional software benefits can be realized.
- There are concerns about the ability to demonstrate the efficacy of software innovations in a consistent manner.
- Software to manage data better will save energy, and software that places data on the appropriate tier (including tape) will save energy. However, software that reduces the amount of online data stored will also save energy. The two are complementary and additive.
- It would be useful to quantify customer tolerance for “automatic” software solutions for data center products.

Software Strategies

SELECTED COMMENTS:

- The inclusion of software functionality in the ENERGY STAR program would allow for far more aggressive savings targets. The vendor should be given a choice of how to achieve energy savings versus a base storage array with no software and poor PSUs.
- The potential savings from PSU improvements is limited to about 10-15% of a subset of the storage array market. Software could achieve over 50% savings.
- Software remains the largest opportunity to reduce energy consumption in enterprise storage. Data de-duplication, thin provisioning, thin snapshots and clones, data compression and RAID 5/6 are deserving candidates for adders in the spec.
- Software solutions can be added to existing installations instantaneously. Benefits could be realized by virtually all data centers within a month at zero capital outlay, whereas hardware improvements may take years to propagate through existing stock.

Software Strategies

SELECTED COMMENTS:

- Debates exist about the actual or average storage space capacity utilization for open systems, with estimates ranging from 15% to 30% up to 65–80%.
- Low storage utilization is often the result of several factors, including limiting storage capacity usage to ensure performance; isolate particular applications, data, customers, or users; for the ease of managing a single discrete store system; or for financial and budgeting purposes.
- Over time, underutilized storage capacity can be consumed by application growth and data that needs to be retained for longer periods of time. However, unless the capacity that is to be consumed by growth is for dormant data (idle data), any increase in I/O activity will further compound the I/O performance problem by sparsely allocating storage in the first place.

Software Strategies

SELECTED COMMENTS:

- To increase utilization, it is important to use a balanced approach: Move idle or infrequently accessed data to larger-capacity, lower-performing, cost-effective SATA HDDs, and move active data to higher-performing, energy-efficient enterprise Fibre Channel and SAS HDDs. Using a combination of faster HDDs for active data, larger-capacity storage for idle data, and faster storage system controllers can enable efficiency.
- Multi-tiered storage is constructed of a single-image pool of storage, with integration by segmented storage tiers, all controlled by a unified data storage architecture. Data is promoted or demoted seamlessly. Tiered hierarchy provides structure to safeguard data, make it available for transport, and save it for government regulations and archiving; all with the most energy efficiency.
- Virtualization provides a means to abstract tiers and categories of storage to simplify management and enable the most efficient type of storage to be used for the task at hand.

Software Strategies

SELECTED COMMENTS:

- Data center managers report data-reduction rates from de-duplication as high as 10:1 or even 20:1. Storing less data requires fewer disks and saves energy.
- De-duplication has significant benefit for heavily virtualized environments. Each virtual instance redeploys the operating system files repeatedly. Data de-duplication not only reduces the storage space required by virtualization, but it also increases the likelihood that when data is needed, it's already in the filer's cache.
- MAID storage systems are perfect for archival data, hierarchical storage implementations or backup systems which are seldom if ever accessed. Data centers using MAID storage systems report a 10% to 20% reduction in storage power consumption.
- Unlike consumer or SOHO storage, data center storage has to date not lent itself to being powered down either due to technology or IT customer preferences and risk aversion.
- Other techniques such as boot from SAN are also of interest.

Software Strategies

SELECTED COMMENTS:

- Some de-duplication solutions boast spectacular ratios for data reduction given specific scenarios, such as backup of repetitive and similar files, while providing little value over a broader range of applications.
- Data compression approaches provide lower yet more predictable and consistent data reduction ratios over more types of data and application, including online and primary storage scenarios. In environments where there is little common or repetitive data, de-duplication will have little to no impact, whereas compression generally will yield some benefit across almost all types of data.
- Thin provisioning is a technique that allocates disk blocks to a given application only when the blocks are actually written rather than at initial provisioning and partitioning. By employing thin provisioning along with good storage management software, not only is disk consumption reduced, but the storage management system can accurately project how storage needs based on use history.

Discussion Break

Hardware Strategies

SELECTED COMMENTS:

- Improving the efficiency of PSUs is a useful and productive initiative.
- Possibly need to allow a different set of PSU efficiencies based on a particular categories within the taxonomy.
- Server PSU test procedures serve as a basis for storage systems that use standard servers as underlying controller hardware.
- Some products integrate network switches, hubs, “off the shelf” UPS modules, and other low power assemblies. it is recommended that these items be excluded from the PSU efficiency and thermal monitoring requirements if their total contribution to the system load is less than 10% of the total, and if the total power consumption of the system exceeds 1kW.
- Customers often do not fully populate the disk slots in a storage device upon their initial purchase. It is important that PSU efficiency be measured at specific load points, as these devices are designed to cover a range of loads over the life of the storage device.

Hardware Strategies

SELECTED COMMENTS:

- With respect to measurement accuracy, Energy Star indicates the wish to increase measurement accuracy to "±5% or ±5 W, which ever is greater". This is more rigorous than the current server specification requirements. Recommend maintaining alignment between the two specifications. Suggest delaying the stricter accuracy requirement until the Tier 2 specification, to keep the server and storage specifications aligned and also allow the industry the necessary development time.
- Power supplies should be treated as stand alone hardware as they are modular and can easily be tested separately.
- It will be very difficult to instrument a system to allow measurement of the output loading of both PSUs without interfering with the system cooling. Relying on built-in monitors within the PSU will not give the accuracy required. Bench testing will provide the most accurate data for storage PSU efficiency measurements.

Hardware Strategies

SELECTED COMMENTS:

- Do not recommend NPL measurements for power supplies; more accurate data can be obtained by testing for efficiency.
- EPA's original analysis of the "power loss" metric based on the published SPEC Power results is skewed towards the lightly configured systems used for SPEC Power measurements. As machines are more "heavily" configured, the idle and maximum power increases by a factor of 2 to 6.
- Concerned that the NPL approach would mean that storage systems with redundant power supplies would be unlikely to comply with EPA standards. The only arrays that would be likely to comply would be entry JBOD systems with single power supplies.
- Strongly oppose the NPL approach. Storage and server systems often use the same PSU designs across different platforms. The industry is currently developing PSUs to meet the recently released server specification. Given 9-12 month PSU development cycles, the industry will not have time to modify designs by the expected implementation date of January 2010.

Hardware Strategies

SELECTED COMMENTS:

- Recommend testing all power supplies without fan power, and allowing for slightly lower efficiencies in multi-output power supplies.
- Two types of power supplies are likely to be used in a storage array, though neither approach has a significant inherent advantage in terms of overall efficiency:
  - A fully custom storage PSU will likely have two primary outputs of 5V and 12V and will have fans that are capable of cooling the entire array.
  - A server type power supply will only have a 12V primary output and its fan is not capable of providing the cooling that the array requires. Elsewhere in the array an additional CDC module to provide 5V is needed and more powerful cooling fans will also be present.
- It is entirely appropriate to consider the inefficiency of the internal fan in a server PSU, but a customized storage PSU includes far more powerful fans to cool the array. If fan power is included in calculations, the storage PSU appears less efficient, when in reality the integration of the system cooling fans is actually the more efficient overall design.

Hardware Strategies

SELECTED COMMENTS:

- PSU definitions for both Computer Server and Storage Products should be the same.
- Recommend a third PSU category, to apply to the Tier 1 storage specification and the Tier 2 server, storage, and desktop computer specifications:
  - Single O/P PSU, with the same efficiency and power factor (PF) requirements as the current Energy Star server specification. These calculations would exclude the fan.
  - Multi O/P PSU, used predominantly for low end servers and desktop computers. These would have the same efficiency and PF requirements as the current Energy Star server specification. These calculations would include the fan.
  - Multi O/P PSU for Storage. Given that the storage fan within the power supply is integral to a properly functioning storage system and not just the power supply, efficiency and PF calculations for this power supply category should exclude the fan and be measured at 230V.

Discussion Break

Reporting

SELECTED COMMENTS:

- Recommend that data requirements for the product data sheet be developed concurrently with the requirements.
- Recommend reporting the average request size, and both read-specific and write-specific statistics to inform administrators who understand the specifics of their workload.
- Energy consumption is driven by the type, number, and spindle speed of disks or the number of SSD or flash drives in a given configuration. These details must be delineated.
- Enumerate the background processes performed during idle to provide customers the ability to better understand the probable source of differences between systems.
- Include thermal parameters for each operational state.
- Peak or circuit breaker power in Amps and KVA, typical or nominal energy usage, and expected cooling requirements in BTUs should be reported on a label on the device, in product documentation, and via sense code inquiry status information from a storage system.

Reporting

SELECTED COMMENTS:

- At this point in time, there are no industry-standard protocols for measuring temperature, utilization, etc. Products on the market today typically do not include the ability to measure or report input power or inlet temperature data.
- Measuring and reporting capacity utilization is usually available on demand via a management interface, but not typically by default.
- An initial approach may be to measure / report power at the Product level using an external PDU. More granular power, temperature, and utilization reporting could be considered in the future.
- SNMP is the most widely used reporting technology and is less heavyweight than most others. Prefer not to specify a reporting technology at all, so that market forces may apply.
- The information in the SMI-S standard is helpful for monitoring and managing environmental states in a storage system.
- DMTF is developing a standard communication protocol for data center equipment.

Reporting

SELECTED COMMENTS:

- The validity and utility of storage utilization data is highly suspect. Because the controller system performs many functions, 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.
- Capacity utilization is routinely discussed in the industry as a fraction of either Raw Capacity or Consumable Capacity. There is relatively little commonality to the way this class of data is reported.
- Granularity must be addressed correctly. It is difficult to justify individual power information from hundreds of PSUs in a large array, for example. We desire a moderate approach, with sensors limited such that there is a readily apparent payback on the expense.
- Some systems include COTS products from companies that are not storage system providers. If a PDU is not supplied by the storage vendor, but it would be up to the data center to procure and install the PDU to meet potential data reporting requirements.

Next Steps

- Review this presentation and send additional feedback to Storage@energystar.gov.
- Reach out to your peers and colleagues to bring more participants to the conversation.
- Continue your open dialogue and creative collaboration with the EPA.
- Watch for Draft 1 and other working documents to be distributed in the coming weeks.

Contact Information

- Andrew Fanara

  fanara.andrew@epa.gov // 206.553.6377
- Steve Pantano

  spantano@icfi.com // 202.862.1551

Email: Storage@energystar.gov

Web: www.energystar.gov/NewSpecs