ENERGY STAR
Data Center Storage Meeting
Draft 3 Version 1.0 Specification

July 11, 2012

Agenda

Time (all EST)	Topic
1:00 PM	Meeting Introduction
1:10 – 2:10	Product Family Structure
2:10 – 2:45	Product Scope Revisions
2:45 – 3:00	Power Supply Requirements
3:00 - 3:30	Energy Efficient Features
3:30 – 3:45	Brief Break
3:45 – 4:45	Information Reporting Requirements
4:45 – 5:00	Performance Data Measurement and Output Requirements
5:00 – 5:30	Testing / Test Method
5:30 – 6:00PM	Remaining topics, meeting summary and closing

Goals and notes

- High-level review and discussion of key topics presented in Draft 3 (as time allows)
- Opportunity for further detail in advance of stakeholders formulating written feedback
- Note: All slides will be posted to ENERGY STAR Data Center Storage website

Introductions

RJ Meyers
EPA, ENERGY STAR
(202) 343-9923
Meyers.Robert@epa.gov

John Clinger
ICF International
(202) 572-9432
John.Clinger@icfi.com

Al Thomason
TBWC, LLC
(503) 708-7881
thomasonw@gmail.com

Call-in Information

- Audio provided via conference call in:

- **Call in:** +1.877.423.6338 (inside U.S.)
  +1.571.281.2578 (outside the U.S.)
- **Code:** 436598

- Phone lines will remain on mute during presentations, opened during discussion (*please keep phone lines on mute unless speaking*)
- Please refer to the agenda for approximated discussion timing

Review of ENERGY STAR Goals

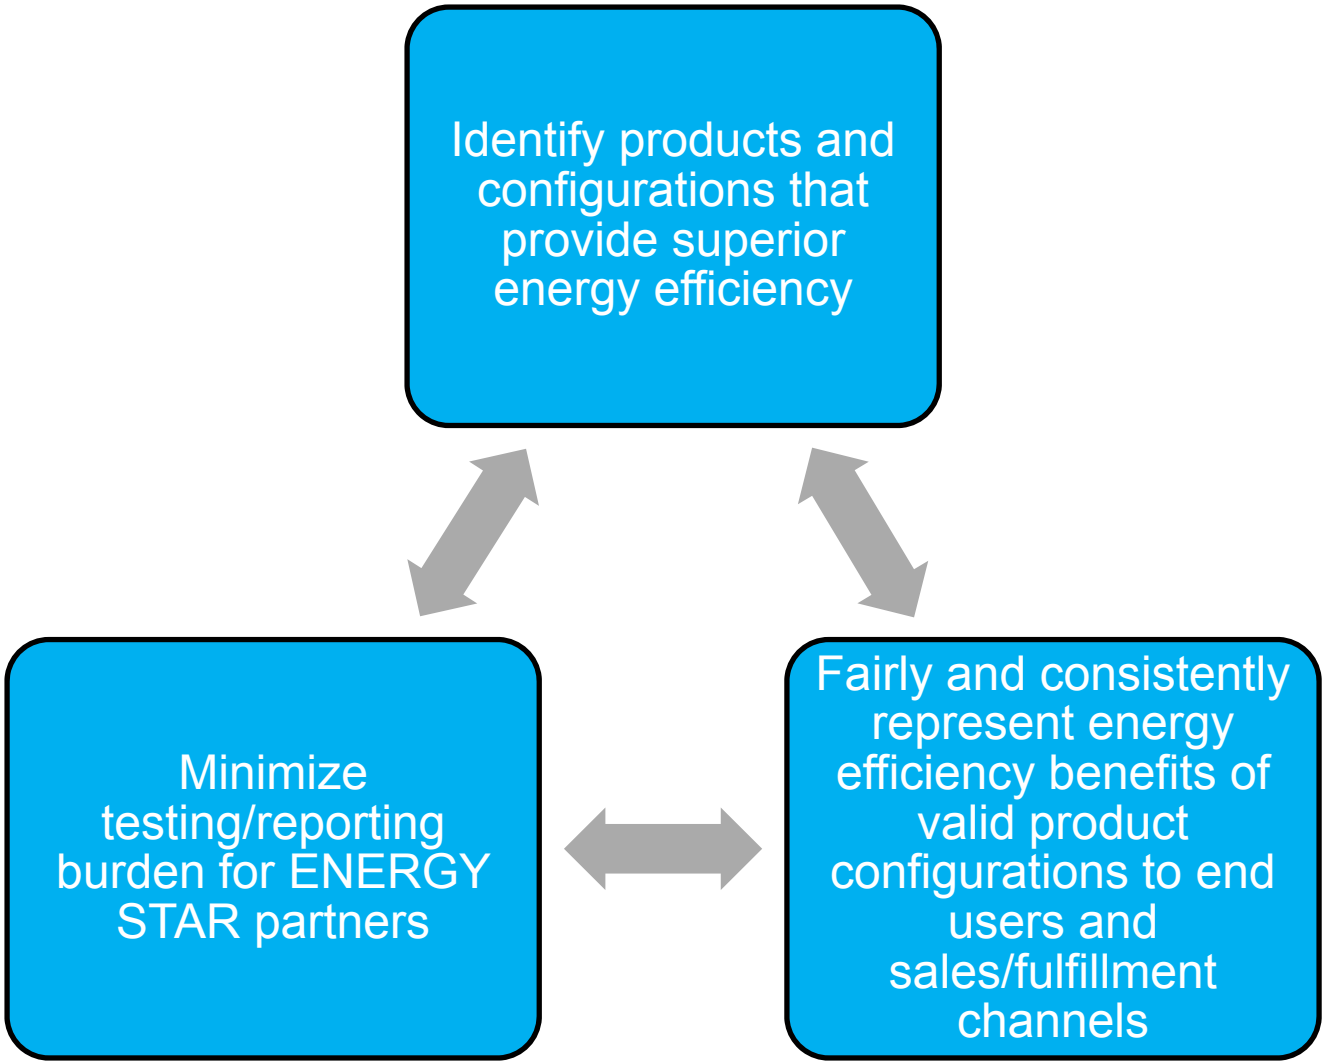

Product Family Structure

- Common Product Family Attributes
- Optimal, Maximum and Minimum Configurations
- Expanded Minimum Configurations
- Combinations of Optimal Configurations

Common Product Family Attributes

- Made by the same manufacturer
- Be from the same model line or machine type
- Utilize the same model of Storage Controller
- Fall under the same taxonomy category
- Controller(s) contain equal or greater amount of cache than the corresponding qualified configuration

- EPA welcomes feedback on additional items for consideration in this list

Optimal, Maximum and Minimum Configurations

- Optimal Configuration
  - Transaction
  - Streaming
  - Capacity
- Maximum Configuration
  - 5% larger in storage device count than optimal configuration
- Minimum Configuration
  - 20% smaller in storage device count than optimal configuration
- Rounding by drawer boundaries for systems with a total device count ≥150

Expanded Minimum Configurations

- Submission of additional data point(s) below the required minimum configuration which:
  - are Transaction or Streaming optimized configurations,
  - are physical data points only, no modeled data permitted for this purpose,
  - are within 10% of the Optimal Configuration performance
    - As defined by the “Workload Weighting Requirements” (Table 6)
- If all conditions are met, the required minimum configuration may be replaced with this new lower bound

Combinations of Optimal Configurations

- Delivered system may consist of combination of approved configurations
  - Applicable with all Optimization Points (Transaction, Streaming, Capacity)
  - And Block-I/O portion of NAS system

- Allows greater flexibility with configuration and delivery of qualified ENERGY STAR systems

Combinations of Optimal Configurations - Example

Example System Optimal points:

- #1 Transaction Optimal Point
  - Dual controllers
  - 8x Drive B (450GB – LFF – 15K)

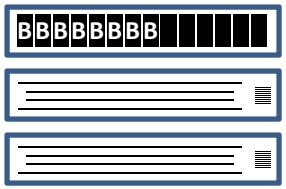

- #2 Sequential Optimal Point
  - Dual controllers
  - 10x Drive E (2TB – LFF – 7.2K)

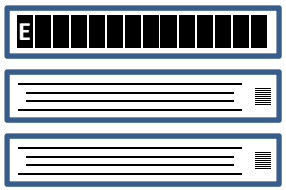

1. Allocate storage media
   - Allocated by % of Optimal test configurations
   - % of allocations must sum to 100%
2. Media Rounding
   - Round UP +5% to nearest whole
   - Round DOWN -20% to nearest whole
     - May use Expanded Minimum Configuration %
3. Drawer Rounding (if applicable)
   - Eliminate –or- fill in partial drawers
   - Keeping overall ratio of drive types the same

See supplement for additional examples, including Drawer Rounding and NAS

Combinations of Optimal Configurations - Example

Combination	Allocation	Rounding
#1 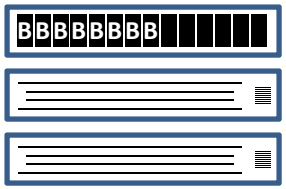 8x Drive B	100% →	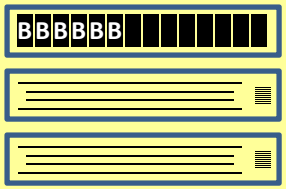 6x Drive B 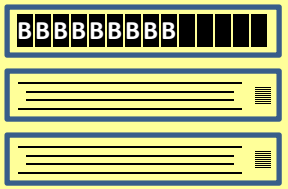 9x Drive B
#1 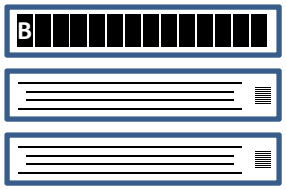 8x Drive B + #2  10x Drive E	40% / 60% →	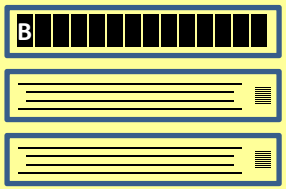 2x Drive B 4x Drive E 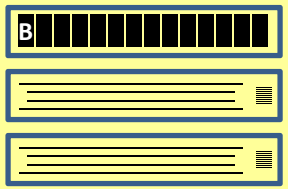 4x Drive B 7x Drive E
	20% / 80% →	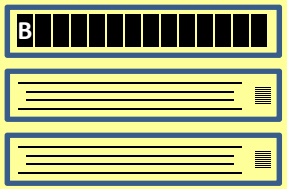 1x Drive B 6x Drive E 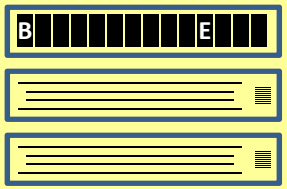 2x Drive B 9x Drive E

EPA

Product Scope Revisions

- Scale-up vs. Scale-out
- NAS
- Additional exclusions

Scale up vs. Scale out

- Scale-up architecture is currently covered in the scope of Version 1.0.
- Scale-out architecture is **not** covered in the scope of Version 1.0 due to:
  - Lack of definition of a node in Scale-out systems.
  - Lack of data to show common behavior of Scale-out storage products as additional nodes are added.

Scale up vs. Scale out

- EPA welcomes feedback on the proposed Scale-up and Scale-out definitions as well as supporting data that can lead to the addition of Scale-out to Version 1.0 prior to Draft 4

NAS

- Addition of Block I/O capable NAS systems to the scope of Version 1.0
  - Tested as Block I/O systems only
  - Only storage devices used for Block I/O are counted in qualification
  - EPA welcomes development of a test method to assess the active performance of NAS file based storage systems with the inclusion of energy efficiency metrics
- NAS systems that cannot perform Block I/O are not included in the scope of Version 1.0

Additional Exclusions from Scope

- NAS-only File Storage Products
- Blade Storage Products
- JBODs
- Object Storage Products

Power Supply Requirements

- Retaining levels consistent with CSCI Silver
  - EPA proposes removing the 10% load point for redundant capable PSUs in exchange for an increase in efficiency requirements of 2% at the 20%, 50%, and 100% load points.
    - Stakeholders have previously suggested removing the 10% load point and focusing on enhanced efficiency at the other load points.
- EPA welcomes additional data on typical PSU load profiles of storage products covered in Version 1.0

Power Supply Requirements

- EPA proposes that PSU requirements only apply to PSUs that power primary equipment
  - Controllers and Drawers
- EPA encourages the use of ENERGY STAR qualified products that can be used in conjunction with qualified storage products when possible

Energy Efficient Features

- Parity RAID
- Revised COM approach

Parity RAID

- EPA has defined Parity RAID as any form of RAID that achieves better efficiency than RAID 1 (mirroring)
- EPA is requiring that Parity RAID is made available on all systems qualified as ENERGY STAR storage products

Recognized COM Features

- The following COMs are currently recognized by EPA
  - Thin Provisioning
  - Data Deduplication
  - Compression
  - Delta Snapshots
- The COMS listed above will be verified using the verification procedures found in the SNIA Emerald™ Power Efficiency Measurement Specification Version 1.0 (8/23/11)

COM Requirements

- EPA proposes to require a certain number of recognized COMs be made available for purchase, with this number determined by the size of the system

Table 4: COM Requirements for Online 2, 3, and 4 Systems

Storage Product Category	Minimum number of COMs required to be made available
Online 2	0
Online 3	2
Online 4	3

Agenda

Time (all EST)	Topic
1:00 PM	Meeting Introduction
1:10 – 2:10	Product Family Structure
2:10 – 2:45	Product Scope Revisions
2:45 – 3:00	Power Supply Requirements
3:00 - 3:30	Energy Efficient Features
3:30 – 3:45	**Brief Break**
3:45 – 4:45	Information Reporting Requirements
4:45 – 5:00	Performance Data Measurement and Output Requirements
5:00 – 5:30	Testing / Test Method
5:30 – 6:00PM	Remaining topics, meeting summary and closing

Information Reporting Requirements

- Active and Idle State Efficiency Disclosure Requirements
- Workload weighting requirements
- Testing Data Requirements
- Introduction of new PPDS and important product characteristics under consideration
- Storage device replacement requirements

Active and Idle State Efficiency Disclosure Requirements

- EPA proposes to make all required fields public for Version 1.0

Table 5: Public Disclosure Requirements for Active and Idle State Efficiency

Workload Test	Transaction Optimization	Streaming Optimization	Capacity Optimization
Mixed Workload 1	Required	Optional	Optional
Mixed Workload 2	Required	Optional	Optional
Random Read	Required	Optional	Optional
Random Write	Required	Optional	Optional
Sequential Read	Optional	Required	Optional
Sequential Write	Optional	Required	Optional
Ready Idle	Required	Required	Required

Workload Weight Requirements

- EPA proposes workload weighting requirements to provide a common way to define performance when comparing Storage Products within a Product Family
  - Needed to define allowable change in both the expand minimum configuration and storage device replacement proposals

Table 6: Workload Weighting Requirements

Workload Test	Transaction Optimization	Streaming Optimization	Capacity Optimization
Mixed Workload 1	70%	0%	10%
Mixed Workload 2	0%	0%	0%
Random Read	10%	0%	0%
Random Write	10%	0%	0%
Sequential Read	0%	50%	10%
Sequential Write	0%	20%	10%
Ready Idle	10%	30%	70%

Workload Weight Requirements

- Established to assure a level of commonality in different manufacturers' approaches to defining Optimal Configuration.
- Recognize three primary values
  - Ability to receive data for storage
  - Ability to retain received data over a period of time
  - Ability to retrieve received data as requested

Testing Data Requirements

- For Online 2 and Online 3 systems:
  - Physical data for all measurements listed in Table 6 including for the following system sizes
    - Optimal configuration
    - At least 40% smaller in device count than optimal configuration
    - At least 15% larger in device count than optimal configuration
  - Modeled data for all measurements made above can be optionally submitted in addition to physical data (and will not be publicly disclosed).

Testing Data Requirements

- For Online 4 systems
  - Physical data for all measurements listed in Table 6 at the optimal configuration size
  - Modeled data for all measurements listed in Table 6 for the following system sizes:
    - At least 6 points smaller in device count compared to the optimal configuration size with the smallest point being at least 40% smaller
    - At least 6 points larger in device count compared to the optimal configuration size with the largest point being at least 15% larger

Testing Data Requirements

- Additional Online 4 option
  - If accurate modeled data cannot be provided for the Online 4 system, the full physical testing requirements for Online 2 and Online 3 systems may be submitted in place of modeled data.
- EPA proposes that only physical data from optimal configuration points will be made public as appropriate in Version 1.0

Storage PPDS

- EPA will develop a Storage PPDS widget to provide easily accessible data that can be used to directly compare qualified systems
  - Widget can be used on ENERGY STAR website or embedded in stakeholder websites
- PPDS information will be pulled from CB certification submissions

Storage PPDS

- List of important characteristics:
  - System configuration;
  - Controller details;
  - Software configuration;
  - Controller power supply information;
  - Storage media drawer power supply information;
  - Storage media used per optimization points
  - Input power and environmental characteristics during testing;
  - System power optimization capabilities;
  - Inlet air temperature and Power Consumption reporting capabilities.

Storage Device Replacement Requirements

- Goals of replacement requirements:
  - Reduce testing burden
  - Maintain consistency on the qualified product list
- Manufacturers must submit specification sheets from the Storage Device vendor for the original and replacement devices to validate the parameters on the following slides

Storage Device Replacement Requirements

- Proposed HDD Requirements

i. No change in **any** of the following categories:
   (a) Form factor
   (b) Interface type, quantity, and transfer speed
   (c) Cache Size
   (d) Data capabilities (e.g. Self-encryption)
   (e) Power Management related features and capabilities (e.g. Power Down modes)
   (f) Rotational Speed

ii. Performance within +/- 5% in **all** of the following categories:
   (a) Average Seek Time
   (b) Sustained Transfer Rate
   (c) Average Latency
   (d) Reported average power consumption in like modes of operation

iii Capacity equal or greater than the storage device being replaced.

Storage Device Replacement Requirements

- Proposed SSD Requirements

i. No change in **any** of the following:
   (a) Form factor
   (b) Interface type, quantity, and transfer speed
   (c) Data capabilities (e.g. Self-encryption)
   (d) Power Management related features and capabilities (e.g. Power Down modes)

ii Capacity equal or greater than the storage device being replaced.

Storage Device Replacement Requirements

- Performance Improvement Cap
  - Replacement of a storage device in a storage product that leads to a change of greater than 20% of the overall system performance as defined by Workload Weighting Requirements Table (with the exception of the Ready Idle metric) will require testing of a new optimized configuration for inclusion in the product family definition

Performance Data Measurement and Output Requirements

- Air inlet temperature requirement
- Permitted methods for Data Collection in Version 1.0
- Role of iPDUs in Version 1.0

Air Inlet Temperature

- Based on discussions with stakeholders, EPA is proposing to make the air inlet temperature requirements optional for Version 1.0, with the expectation of requiring the measurements in Version 2.0
- The ability to measure air inlet temperature to the level required in Section 3.8.1 will be reported in the PPDS

Permitted Methods for Data Collection

- EPA is requiring the collection of input power using any embedded or add-in solution for Version 1.0
- EPA encourages the collection of air inlet temperature using similar techniques
- EPA intends to require that both of these reporting abilities be conducted by embedded components in Version 2.0

iPDUs

- Performance data measurement and output measurements can be made using iPDUs as long as:
  - iPDUs meet all requirements for accuracy, sampling and data reporting
  - iPDUs are made available for sale and delivery with qualified ENERGY STAR Storage Products

Testing / Test Method

- The Storage Draft Test Method primarily references the SNIA Emerald™ Power Efficiency Measurement Specification V1.0 with following proposed deviations:
  - Online 2 Storage products must include RAID controller
  - 24 hour Ready Idle Test
  - COMs must be disabled during active testing
  - Directions for testing NAS products w/ Block I/O

Testing / Test Method

- Online 2 Storage Products must include RAID capable controller
  - EPA is requiring that all ENERGY STAR Storage Products shall include a RAID capable controller
- The Ready Idle test in the SNIA Emerald specification shall be replaced with a 24 hour Ready Idle Test following the Active Test
  - To ensure all systems are subject to the same Ready Idle testing methodology

Testing / Test Method

- Recognized COMs must be disabled during all active / idle testing
  - To ensure all systems are compared at the same base state, as COMs are optional and not a defining feature of a product family
- NAS testing instructions
  - All usable Storage Devices not needed to enable minimal NAS capability shall be tested as Block I/O
    - Subject to systems maximum block I/O size limitations
  - NAS functionality shall be enabled during all testing

Remaining topics, meeting summary and closing

- Power Modeling Presale Tool
- Development of updated SNIA tool
- Development of NAS file based energy efficiency performance tool
- Remaining timeline for Version 1.0

Power Modeling Presale Tool

- EPA proposes to retain the requirement for stakeholders to make a power modeling presale tool available to prospective customers of Online 4 systems that use modeled data for qualification purposes
  - EPA welcomes feedback on language differentiating between the general public and potential customers
- The warranty requirement from Draft 2 has been removed

Development of updated SNIA tool

- EPA is aware of the current effort to modify the SNIA Emerald specification to better address systems with caching ability
- EPA hopes that development will be completed in the near future in which case the new test method will be considered for use in Version 1.0 if appropriate
  - Adoption of this revision would require minor changes to the current testing and information reporting requirements in Draft 3

Development of NAS File Based Energy Efficiency Performance Tool

- EPA looks to include additional testing for active NAS energy efficiency performance in a future revision
- EPA seeks input on planned development efforts for this type of test method

Remaining Version 1.0 Timeline

- Draft 4 release in August
- Final Draft in late September / early October

Table 8: Specification Effective Date

Effective Date
January 2013

Note: EPA is proceeding with a goal to complete development of the Version 1.0 Data Center Storage program by **October 2012**. As a reminder, the Data Center Storage program would have a slightly delayed effective date to enable CB and lab certification for testing storage products. Therefore, EPA will chose an effective date in January 2013, approximately 3 months following finalization.

References and resources

- ENERGY STAR Data Center Storage specification revision:
  - www.energystar.gov/NewSpecs
  - Select “Data Center Storage”

Reminder

Written comments on Draft 3 due to EPA no later than July 27, 2012. storage@energystar.gov

Thank You!

RJ Meyers
EPA, ENERGY STAR
(202) 343-9923
Meyers.Robert@epa.gov

John Clinger
ICF International
(202) 572-9432
John.Clinger@icfi.com

Al Thomason
TBWC, LLC
(503) 708-7881
thomasonw@gmail.com