



ENERGY STAR® Data Center Storage

Test Procedure Workshop

Storage Networking World (SNW)
JW Marriott Desert Ridge Resort
Phoenix, AZ
15 October 2009



Agenda



08:00 – Welcome & Introductions

08:15 – Workshop Goals

08:30 – Industry Presentations / Q&A

(1) Climate Savers Computing Initiative (CSCI) PSU Discussion

(2) Storage Performance Council (SPC)

09:45 – BREAK

10:00 – Industry Presentations / Q&A

(3) Wikibon

(4) The Storage Networking Industry Association (SNIA)

(5) IDC

11:45 – LUNCH

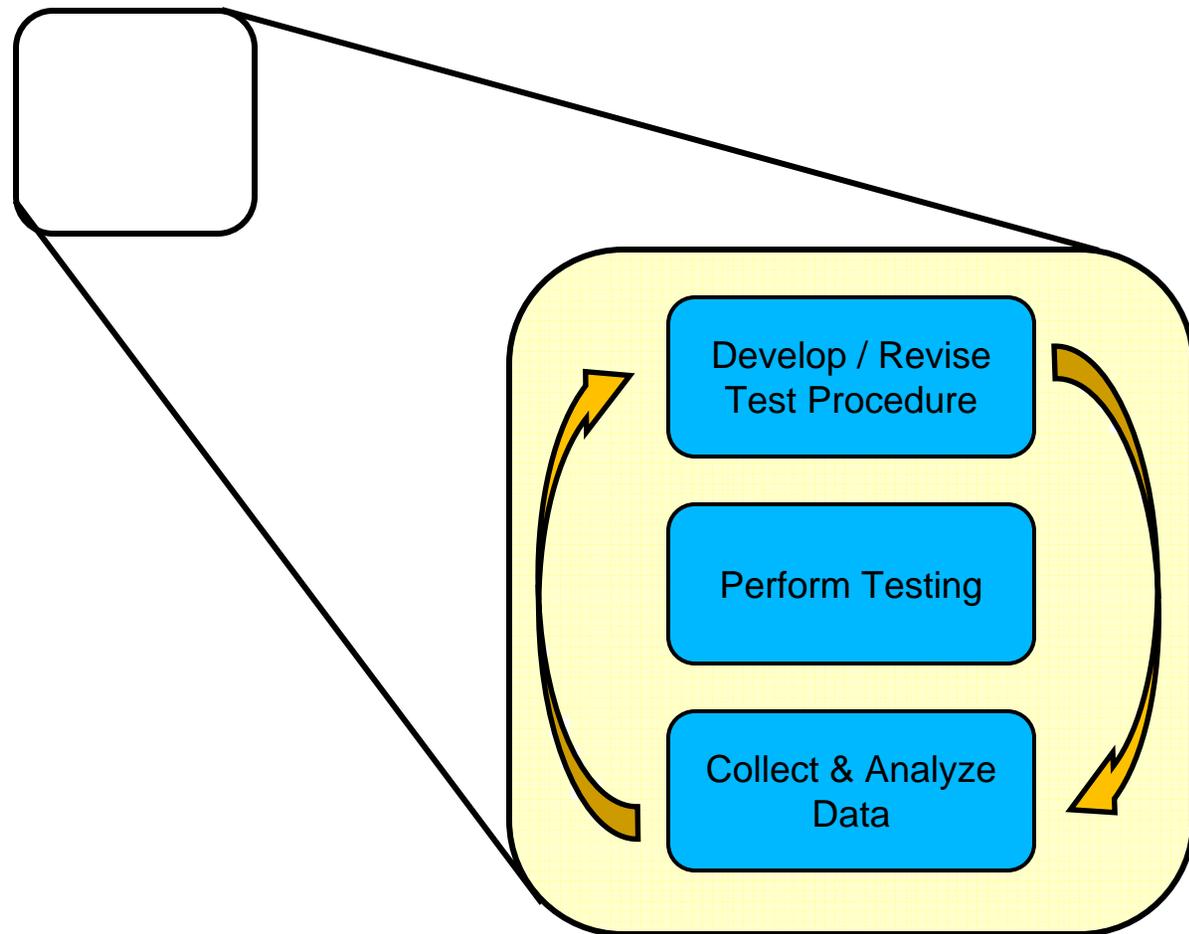
12:15 – ENERGY STAR Storage Workloads & Tools

13:15 – ENERGY STAR System Characterization & Datasheet

13:45 – Timeline & Path Forward

14:00 – ADJOURN

Where Are We?



Workshop Goals



- Understand industry experience and perspective on Active and Idle testing.
- Identify best-practices for procedures, instrumentation, measurement, test operations, and data collection / reporting for initial data collection.
- Plan to complete the first round of testing by January 1. Results from Round 1 of testing will inform development of the ENERGY STAR test procedure.

Taxonomy of Interest



 Storage Taxonomy Summary	Online Storage	Near Online Storage	Removable Media Libraries	Virtual Media Libraries	Infrastructure Appliances	Infrastructure Interconnect
<p>Prime storage, able to serve random as well as sequential workloads with minimal delay</p>	<p>Intended as second tier storage behind Online Storage. Able to service Random and Sequential workloads, but perhaps with noticeable delay in time to 1st data access.</p>	<p>Archival storage used in a sequential access mode. A typical example would be Tape based archival, both Stand Alone and Robotically assisted libraries.</p>	<p>Storage which simulates removable Media Libraries. Will typically use non tape based storage and as such are able to respond to data requests more quickly</p>	<p>Devices placed in the storage SAN or network adding value through one or more dedicated Storage enhancements. Examples include: SAN Virtualization, Compression, De-duplication, etc.</p>	<p>Devices which enable a SAN or other Storage Network data switching or routing.</p>	
<p>Maximum Capacity Guidance <small>Note: Maximum Capacity Guidance reflects the maximum capacity a given offering can be purchased with and/or field upgrades. It is intended to be used as a guideline as opposed to an absolute value. There will be cases where a device may have greater or small capabilities, but otherwise is an appropriate match for a given classification due to other criteria, e.g.: redundancy capabilities</small></p>	<p>Max Storage Devices (Up to 80 Ms MTTD)</p>	<p>Max Storage Devices (Over 80Ms MTTD)</p>	<p>Max Tape Drives</p>	<p>Max Storage Devices Supported*</p>	<p>Max Port Count</p>	
<p>Group 1) SoHo & Consumer</p> <p>Storage which is designed primarily for home (consumer) or home / small office usage. <small>-Often Direct Connected (USB, IP, etc.) -No option for redundancy (will contain SPOFs)</small></p>	<p>Up to 4 Devices</p>	<p>MTTD = Max Time to Data Maximum time needed to access any data stored in any place on the storage system</p>	<p>Stand Alone Drive (No Robotics)</p>	<p>Note: * Infrastructure Appliances by definition have no intrinsic storage, other than what is used for local processing and/or local Caching of data.</p>	<p>Storage Devices Support in this case refers to the number of storage devices controllable down stream of the Appliance</p>	
<p>Group 2) Entry, DAS, or JBOD</p> <p>Storage which is dedicated to one or at most a very limited number of servers. Often will not include any integrated controller, but rely on server host for that functionality. <small>-Often Direct Connected (SATA, IP, etc.) -May optionally offer limited number of redundancy features</small></p>	<p>More than 4 Devices</p>	<p>Up to 4 Devices</p>	<p>Up to 4 Drives</p>		<p>Up to 32</p>	
<p>Group 3) Entry / Midrange</p> <p>SAN or NAS connected storage which places a higher emphasis on value than scalability and performance. This is often referred to as 'Entry Level' storage. <small>-Network connected (IP, SAN, etc.) -Has options for redundancy features</small></p>	<p>More than 20 Devices</p>	<p>More than 4 Devices</p>	<p>More than 4 Drives</p>	<p>Up to 100 Devices</p>	<p>Support for up to 20 Devices</p>	
<p>Group 4) Midrange / Enterprise</p> <p>SAN or NAS connected storage which delivers a balance of performance and features. Offers higher level of management as well as scalability and reliability capabilities. <small>-Network connected (IP, SAN, etc.) -Has options for and often delivered with full redundancy (no SPOF)</small></p>	<p>More than 100 Devices</p>	<p>More than 100 Devices</p>	<p>More than 24 Drives</p>	<p>More than 100 Devices</p>	<p>Support for more than 20 Devices</p>	
<p>Group 5) Enterprise / Mainframe</p> <p>Storage which exhibits large scalability and extreme robustness associated with Mainframe deployments, though are not restricted to Mainframe only deployments. <small>-Mainframe connectivity with optional network connection (IP, SAN, etc.) -Always delivered with full redundancy (no SPOF) -Often Capable of non-disruptive serviceability</small></p>	<p>More than 1000 Devices</p>		<p>More than 11 Drives</p>	<p>More than 100 Devices</p>	<p>Support for more than 100 Devices</p>	

© SNIA 2009



The Ideal Test Procedure



- Easy to run & reconfigure
 - Simple & low cost
 - Can be reproduced outside of a test lab environment
 - Allows many software & hardware configurations to be exercised
- Broadly applicable to the taxonomy
 - Compare energy performance across categories
 - Disk & Tape
- Provides accurate, intuitive, actionable results
 - Isolate SW, HW, and RAS impact on energy performance

Industry Perspectives



1. Climate Savers PSU Discussion
2. SPC
3. Wikibon
4. SNIA
5. IDC

Round 1 Test Objective



- Understand energy performance impact of single-variable changes in hardware configuration
 - Single vs. Redundant Controller
 - Capacity vs. Performance Drives
 - RAID Level
 - Total Raw Capacity

Round 1 Proposal



- Test 4 representative system configurations.
 - Transaction-oriented Value System
 - Transaction-oriented High-reliability System
 - Archival/Streaming-oriented Value System
 - Archival/Streaming-oriented High-reliability System
- Test Active & Idle
 - 4 corners?

Round 1 Proposal



- Model the same 4 representative configurations to determine model accuracy
- Model single-variable changes to hardware configuration to understand impact on energy consumption.
 - Increase raw capacity
 - Single vs. redundant controllers
 - Drive type (SAS HDD, SATA HDD, SSD)
 - RAID level

Discussion Questions



- Modes of Operation
 - Block IO and/or File IO
 - Random and/or Sequential (as appropriate)
- Workload Tools
 - Iometer
 - VDBench
 - FileBench
 - Others?

Discussion Questions



- Configuration Boundaries
 - Sample of disk technologies
 - Single vs. Mixed technology (e.g. SSD and HDD)
- Test Equipment
- Workload Boundaries
 - 4 corners (as appropriate)
 - Ramping workload



Discussion Questions

- Allow Calculated / Modeled Results
 - What constraints should be put in place?
 - What assumptions must be documented?
- Data Collection / Reporting
 - Data Collection Format
 - HW & SW configuration details?
- Data Analysis
 - How will the data be used?

Next Steps



- Send follow-up comments from this meeting to storage@energystar.gov
- Draft test procedure & data collection form distributed for review by October 23. Comments due November 6.
- Test procedure released November 13.
- Rolling data collection begins. First round of testing closes December 30.

Contact Information



- Andrew Fanara
fanara.andrew@epa.gov // 206.553.6377
- Steve Pantano
spantano@icfi.com // 202.862.1551
- Al Thomason
thomasonw@gmail.com // 503.708.7881

Email: Storage@energystar.gov

Web: www.energystar.gov/NewSpecs