

Homogenous vs. Heterogeneous ENERGY STAR Storage Testing

EPA released Draft 3 of the ENERGY STAR Storage v1.0 specification on June 22, 2012. This draft contained a number of new proposals developed in discussions with stakeholders over the previous six months. One of the major changes contained in the document was a proposal to test workload optimized, homogeneous drive systems and then allow them to be combined into mixed drive systems. Some stakeholders have responded by advocating for a shift back to heterogeneous system testing, while others have supported the homogeneous test proposal. To facilitate an open and transparent discussion on this important topic, EPA has developed this document to highlight its concerns with heterogeneous system testing and to continue the existing conversation on the merits of different approaches to system level testing of storage products.

EPA Draft 3 Approach

EPA currently proposes that optimal, maximum and minimum configurations be submitted using only homogenous storage device configurations. This proposal is referenced in Section 1)1)6)viii under “Systems Composed of Combinations of Optimal Configurations”. Lines 380-381 state the following:

“Multiple transaction or streaming optimizations may be submitted to incorporate different storage device technologies and customer usage needs.”

Under this proposal, a manufacturer with multiple types of addressable storage devices that can be used for a specific optimization will have to test a corresponding number of configurations for qualification. For example, if a manufacturer wishes to use both 15k HDDs and SSDs as addressable storage (included in the “Formatted (usable) Capacity” of the system) for transaction workloads, the manufacturer would submit two homogenous transaction configurations, one with the 15k HDDs, and one with the SSDs. Then the provisions in the “Systems Composed of Combinations of Optimal Configurations” can be used to combine the workload optimized systems to create a product family.

EPA would like to emphasize that storage devices which are used for caching purposes (e.g. SSDs) and are non-addressable storage CAN be included in any tested configuration and are not subject to the homogenous qualification criteria that addressable storage devices are held to.

Stakeholder Responses to Draft 3

EPA received stakeholder feedback both for and against this proposal. Feedback from stakeholders with reservations about this approach advocated the testing and qualification of configurations using heterogeneous (mixed) storage device types. The feedback, both for and against, contained three distinct themes:

1. For consistency in data and to allow the “Systems Composed of Combinations of Optimal Configurations” section to function correctly, it may be best to limit optimal configurations to homogenous storage devices.
2. To allow vendors to take full advantage of the Expanded Minimum Configuration, EPA should allow the combination of a single type of SSD and HDD storage device to lower testing burden and optimize data made available to the EPA.

3. Testing systems of homogenous storage devices may not reflect how storage systems are sold. These stakeholders support testing mixed configurations as the optimal point and prefer to “mix and match” storage devices from different configurations to define a product family. Additionally, these stakeholders were concerned that the burden of testing homogenous configurations is significant, especially for stakeholders who choose to use third party laboratories rather than internal laboratories for testing.

EPA Comments on Current SNIA Emerald™ Power Efficiency Measurement Specification Version 1.0

EPA appreciates the development effort from SNIA on the SNIA Emerald Specification for Online storage systems. EPA understands that deployed systems often contain a mixture of storage device types and that such deployments require logic (real time, or as a result of configuration of the entire computer/storage/application environment) to effectively place data on the most appropriate storage devices in the storage product. In order to provide broad coverage of common storage deployments, the EPA is currently looking for an approach that aligns with three major workload types: Transaction, Streaming, and Capacity.

Based on the current SNIA Emerald Specification, and recognizing this limitation is in fact common across many assessment tools, EPA believes that the data produced from heterogeneous storage device configurations would not reflect real world configurations. Real world configurations feature simultaneous servicing of transaction and streaming workloads and direct each type of workload to the appropriate storage devices through partitioning, pre-configuration, or other advanced capabilities.

Although the system configurations used in heterogeneous or “mixed” testing may appear closer to real world systems as they are sold, the EPA is concerned that test results from the current Emerald test method would not produce results reflective of actual deployments of storage products. This, in turn, would not support EPA’s goals of enabling useful comparisons between systems and providing relevant information to end-users. EPA has proposed the combination of optimal configurations to allow stakeholders to address mixed storage device deployments and welcomes additional feedback on the approach of this proposal.

EPA believes that the strength of the current SNIA Emerald Specification, as well as other publically available industry benchmarks, lies in assessing homogenous storage device configurations. EPA believes that addressing heterogeneous storage device configurations would be ideal, but believes that validating the presence and/or effectiveness of requisite intelligent data placement on mixed drive configurations is a necessary condition for allowing heterogeneous storage device configurations to be tested under the ENERGY STAR program requirements.

EPA Comments on SNIA Emerald Specification “Hot Banding” Revision

EPA looks forward to the revision of the current SNIA Emerald Specification which will allow end-users to see the performance benefits of caching in their systems. It is not clear to the EPA that the revision will fully address a way to validate the presence and/or effectiveness of intelligent data placement that we believe is required to address mixed storage device configuration testing as outlined in the sections above. EPA welcomes additional stakeholder feedback on development in this area.

EPA Proposed Path Forward

EPA plans to hold a stakeholder conference call for discussion of the issues raised in this document on September 5th from 3 – 4pm EST. We hope that this will provide all parties an opportunity to further discuss their concerns with the two approaches proposed thus far — homogeneous vs. heterogeneous testing.

EPA will weigh this additional input from stakeholders and move forward with Draft 4 in September, while also monitoring the development of the SNIA Emerald Specification revision for “Hot Banding.” If the revision to the test method effectively highlights performance/watt gains in systems with caching, it could be a welcome addition to the current ENERGY STAR testing requirements. EPA hopes that "Hot Banding" will be ready in time to integrate into Storage Version 1.0 as either a mandatory or optional workload for test.