

IBM appreciates the opportunity to comment on the ENERGY STAR® Storage Requirements Version 1 Draft 4 Document. The Draft 4 requirements build on the improvements made in Draft 3 while further streamlining the product family testing requirements for gathering the data required to assess the performance power information generated by the SNIA Emerald Test to establish appropriate metrics. IBM has found that the Requirements lack clarity, particularly regarding the ability of manufacturers to test with SSDs and qualify Expanded Maximum and Minimum Configurations, and has made several recommendations to clearly define the testing and reporting requirements. IBM also supports EPA's desire to utilize SNIA Emerald V2.1 as the testing process for Version 1 of the Storage requirements as it provides the best methodology available to assess the energy efficiency of the automated data placement functions that are being incorporated into storage systems.

IBM offers the following specific comments:

Version 1.0 Storage Specification Draft 4

IBM is supportive of the following changes made in Draft 4 of the specification:

1.C.1, Lines 41-43, page 2: IBM supports the revisions to the thin provisioning definition.

1.I.3,4, Lines 213-214, 219-220, page 6: IBM supports EPA's approach to simplifying the drawer rounding approach for the Maximum and Minimum Qualified Configurations in Draft 4 as compared to Draft 3. This gives companies the option to use the drawer loading approach that provides the most efficient system for their products and does not penalize products that require full storage device drawers to operate.

1.I.5,6, Lines 235-50, page 6: IBM supports the increase of the allowable performance per watt range to qualify Expanded Minimum and Maximum systems from 10% to 15%. We believe that the 10% range was unnecessarily restrictive and that the broader range will enable a larger range of configurations representative of the BFF performance to be qualified to the ENERGY STAR specification.

2.2.2.v, Line v, page 10: IBM supports the exclusion of Network Attached Storage devices which cannot perform block function. The unique characteristics of filer systems are not amenable to testing under the SNIA Emerald Specification, making it inappropriate to include these systems in the Version 1 specification.

3.2.1, lines 401-425, pages 10-11: IBM supports EPA's decision to exclude the power supply efficiency and power factor requirements at the 10% load point. IBM has pointed out to EPA that the current power profile of the storage media, where HDDs are typically not idle, likely precludes a storage system from operating below the 20% load point.

3.4.2, lines 452-62, pages 12-13: IBM supports the reduction in the number of available COMS required to qualify a system. The number of COMS proposed as a requirement in

Draft 3 was too high and would disqualify some systems which would otherwise meet the other qualification requirements under the specification.

3.7.3, Lines 674 to 690, page 19: IBM agrees with the EPA approach to match the Storage System reporting requirements to the Computer Server Reporting Requirements, but with reporting of inlet temperature as an optional item for version 1.

IBM has the following recommendations for changes/improvements to Draft 4:

1.J, Lines 298-307, page 8: IBM recommends that EPA maintain the definitions for “Scale Up” and “Scale Out” systems as defined in Draft 3. Scale Up and Scale Out are recognized industry terms and introducing new definitions to represent similar types of behavior will generate confusion. We recommend that EPA allow single instances of “Scale Out” systems to be qualified to the specification, but not require that a “Scale Out” system be qualified with multiple nodes because the test methods have not been proven on a “Scale Out” system where multiple nodes exist independently on the network.

Add an exclusion, 2.2.2.ix: Multi-node scale-out storage products.

Modify included products 2.2.1.iv.c: implement Scale-up or Distributed Controller Storage.

We do recommend that EPA retain the definition for a Distributed Storage Controller: A storage product for which each drawer can contain up to a defined number of storage devices and a dedicated controller and the storage product consists of a minimum to a maximum number of drawers integrated as individual nodes on a single rack.

1.J.6.iv, Lines 334-35, page 8; 2.1.1.iv.a, Lines 370-71, page 9; and 3.4.1.i, Lines 440-41, page 12: IBM supports the use of the “Parity RAID” definition and broader set of technologies as the basis for the requirement that a storage system have an advance error correction and recovery system. The broader definition of the requirement enables companies to utilize the error correction and recovery system they deem most efficient for the performance of the functionality and the energy efficiency of the product. Rather than restricting error correction/detection technologies to RAID, EPA needs to allow other technologies this requirement allows them to be differentiated by the performance/power numbers. IBM strongly supports the TGG proposal for language changes in the specification.

3.3.1, lines 427-32, page 11: The phrase “...will be made available to manufacturer qualified purchasers of the product.” is not clear. Based on a reading of the section, it appears to be more appropriate to state “...will be made available to purchasers of an ENERGY STAR qualified storage product.” IBM recommends that this phrase be modified as necessary to clarify the requirement.

3.5, Lines 463-603, pages 13-17: The data required to be reported by the Certification Body to EPA and required to be reported by the manufacturer on the PPDS are not easily determined in this section. Of particular concern, the clarifications on configuration

requirements provided in the document “Proposed Path to Final Data Center Storage Version 1.0.pdf” dated November 5, 2012 are not clearly spelled out in this section.

1. The requirements do not clearly delineate if or how a company can combine a SSD storage device with a HDD storage device to create the 3 or more physical configurations for testing.

- The description in 1.I.7 suggests that you can only have a “single device”.
- 3.5.4.i.b suggests that you can mix SSDs with one HDD drive type for the 5 configuration types.

Given that manufacturers are increasingly marketing systems which combine SSDs with HDDs to increase system performance and the performance/power characteristics of the storage equipment and the availability of the Hot Band test as part of the SNIA Emerald testing procedure which can test multi-device systems with data placement capability, IBM recommends that EPA allow testing of mixed device configurations with a consisting of a single SSD drive type and a single HDD drive type. The manufacturer should be allowed to select the number of SSDs in each configuration type that will permit the best perf/power scores for the system and workload type under test.

In order to incorporate these changes, EPA should make the following changes in the specification.

Section 3.5.3 and 3.5.4: Move (g) to (h) and create a new item “(g) The manufacturer may combine SSD with HDD at a ratio which delivers the optimum performance/power metric for the given controller, configuration and workload types.”

Section 1.I.1: Add an item “vi) contain an equal or greater number of SSD storage devices as the corresponding qualified configuration.”

2. Sections 3.5.3 and 3.5.4 do not indicate that where a company wishes to test multiple types of storage media where those storage media cannot be qualified as “Replacement Media”, the manufacturer just needs to test the Optimal Configuration for additional drive types as proposed in “document “Proposed Path to Final Data Center Storage Version 1.0.pdf” dated November 5, 2012. IBM believes allowing additional drives to be tested at the Optimal Configuration but qualifying to the broader Minimum and Maximum or Expanded Minimum and Maximum configuration(s) makes sense both from a standpoint of minimizing the testing burden and collecting adequate data to assess the Workload Tests. A company will still need to generate Workload Test data for a minimum of three configurations for each system type, Transaction, Streaming, and Capacity to qualify a given product family for all its potential uses which provides a broad data set for each product family. IBM strongly encourages EPA to clearly include the requirement for testing only the Optimal Configuration for additional drive types by adding 3.5.3.i.h and 3.5.4.i.h which states: “The manufacturer determined optimal configuration point for each additional drive type tested beyond the drive type tested in 3.5.3.i.a-c (or 3.5.4.i.a-c).”

3. The relationship between Table 5 and sections 3.5.3.i, 3.5.4.i, and 3.5.7 are not clear. Table 5 should be moved down to section 3.5.7 and sections 3.5.3.i and 3.5.4.i should clearly state that the 6 data points shall be submitted to EPA on the “test report” or whatever the document will be that the Certification Body provides to EPA to demonstrate and document the qualification of a given storage equipment configuration or product family.

4. IBM is concerned with Table 6, as we continue to believe that it is premature to set workload weightings to assess the three system workload types. However, given that the proposed PPDS reporting requirement enables EPA the majority of the SNIA Emerald data to be masked, IBM believes the current approach is preferable as long as EPA recognizes that it is necessary to collect a sufficient sample size of SNIA Emerald tests to be able to assess which Workload Test or combination of Workload Tests provides a representative metric to assess the performance/power characteristics of a storage product or product family. Given the limited amount of test data that has been generated to date by the SNIA Emerald test method, IBM believes that it is premature to publish the unmasked data, similar to the situation with the SERT data for servers under Version 2 of the Computer Server requirements. Ideally, IBM would prefer to see EPA blind the SNIA Emerald data it collects for 18 – 24 months to give time for EPA and the industry to assess the data and determine representative metrics. Absent this masking, the EPA proposal for publishing limited SNIA Emerald Workload Test data under the PPDS and the EPA on-line system is preferable to publishing all of the Workload Test scores.

IBM does request that only the performance/power scores be reported on the PPDS. The response time data for each the three configuration types should not be provided because it can be misused in comparing systems. SNIA does not require the response time data to be published because of the risk of misuse of the data for inappropriate comparisons of response times for different workload types.

3.6.1.i.b, Lines 613 page 18: The requirement that replacement drives have the same interface type should be removed from this item. There are some interfaces, such as SAS and Fibre Channel, which will have the same quantity and speed. If the transfer rate characteristics (quantity and speed) and the spindle speed and form factor are the same, the storage equipment will demonstrate comparable performance/power characteristics. Differentiating by interface type where transfer rates are the same increases the test burden without providing any differentiation in system performance.

4.1, Lines 701-713, pages 19-20: IBM recommends that EPA reference that all active and idle testing be completed in accordance with the SNIA Emerald Specification Version 2.0 Rev. 1 and eliminate the EPA test method document. The Emerald Specification sets all of the environmental and accuracy parameters and referencing directly to the SNIA specification simplifies the process. Any specific requirements in the test method which are not explicitly detailed in the SNIA Emerald Specification can be included as items under Section 4.

In addition, IBM recommends that SNIA also reference the User Guide in Table 7. The User Guide provides details about administering and executing the Emerald Test Specification to insure the quality and repeatability of the data.

General Discussion of the Testing Matrix: The proposed modifications to the testing of additional drive types and the criteria for defining Replacement Drives, when taken together, are critical to maintaining the workability of the testing matrix. The adjustments of interest are:

1. For a given storage product family, after test and modeling data is generated for one type of HDD storage media, additional storage media should be qualified by completing and reporting the test data only for the Optimal configuration.
2. Enable companies to qualify products with a quantity of SSDs.
3. Removing the common interface requirement for the evaluation of replacement drives enables drive with comparable performance and power profiles to be grouped.

IBM evaluated one of its storage products under our interpretation of the requirements absent the three modifications discussed above. One machine type would require 30 tests. Using the modifications discussed above, we would be able to reduce the number of tests to 16. With each test taking 16 hours, we are still looking at 3-4 weeks of work to qualify a storage product family. However, without the changes we would be looking at 6-7 weeks of work. The current situation is analogous to the difficulties identified in with the server product family definition and we encourage EPA to fix the problem now to make the process more efficient.

In addition to providing the more flexible testing requirements, IBM is pleased with EPA's intent to use the SNIA Emerald Specification V2.01 or V2.10 for the ENERGY STAR storage. IBM believes that EPA has received sufficient data from IBM, HP and Oracle to validate that specification with the Hot Band workload and identify it as the test method for Version 1 of the Storage specification. In addition, incorporation of the Hot Band test is essential given the foreseeable development direction of storage technology – use of the Version 1 Emerald Specification in Version 1 of the Storage Specification risks making the data collection effort irrelevant for the development of metric criteria for the Version 2 ENERGY STAR specification as technology will have passed by the collected metrics data. We urge EPA to use the Emerald Specification V2.xx for the Storage Version 1 Specification.

Version 1.0 Storage Draft 2 Test Method

5.1, Lines 74-79, page 4: IBM supports the reduction of the idle test period from 24 hours to the idle test period specified in the Version 2 Emerald Specification. The period set by SNIA in the Emerald Specification will be sufficient to ensure that a stable idle reading is provided.

Please contact Jay Dietrich at jdietric@us.ibm.com with any questions.