

IBM Comments to the Final Draft of Version 1 of the ENERGY STAR® Program
Requirements for Data Center Storage

IBM appreciates the opportunity to comment on the Final Draft of Version 1 of the ENERGY STAR® Program Requirements for Data Center Storage. This Final Draft improves on previous drafts by providing more accurate definitions of a product family, updating testing matrices for a reasonable assessment of the power use characteristics of a storage product, and defining an appropriate range of products to include under Version 1. We believe that this Final Draft, with some adjustments and clarifications discussed in the comments below, should be implemented as requirements.

The development of this Final Draft has highlighted the challenges of assessing and differentiating complex storage product categories. IBM looks forward to working with EPA to overcome those challenges. Our comments are below.

Comments

A. IBM supports the provisions identified below and provides the following comments:

1. IBM supports EPA's decision to collect the SNIA Emerald Metrics data for Online 2, 3, and 4 Storage Taxonomy Categories for Block I/O storage functions. However, as the SNIA Emerald test protocol was finalized just before the publication of Draft 4 of the requirements; there are less than 10 published test results available today. We believe that EPA needs to gather substantially more data from a range of storage product manufacturers so that EPA and interested stakeholders have sufficient data to assess how best to use the SNIA metrics to assess and measure storage product energy efficiency. There is substantial work to do to determine the appropriate approach.

2. Page 6, I.4, lines 227 to 259: IBM supports the allowable percentage difference of 15% in the performance per watt metric to qualify Expanded Minimum and Maximum systems. We believe that the 15% allowable range will enable a larger range of configurations representative of the BFF performance to be qualified to the ENERGY STAR specification while still providing a meaningful assessment of the energy efficiency of a storage product.

3. Page 7, 1.I.6.v, lines 273-275; Drawer Rounding: IBM supports EPA's approach to simplifying the drawer rounding approach for the Maximum and Minimum Qualified Configurations. This gives companies the option to select the approach that provides the most efficient system for their products and does not penalize products that require full storage device drawers to operate.

4. Page 9, 1.J.7.iv, lines 371-374: Advanced Data Recovery Capability: Expanding the definition of Data Recovery to include technologies other than RAID technologies is an important addition. Companies are pursuing a variety of data back-up technologies outside of the RAID technology to leverage and exploit specific technical capabilities of storage systems and to address specific client needs. Limiting the requirements to only RAID based back-up technologies would unfairly exclude some

IBM Comments to the Final Draft of Version 1 of the ENERGY STAR® Program
Requirements for Data Center Storage

products from qualifying to the ENERGY STAR requirements. The SNIA Emerald Test should identify those products which provide superior performance.

5. Pages 13-14, Sections 3.5.1 and 3.5.2, lines 489-500; IBM supports the clarification of Tables 5 and 6 in the Final Draft which clearly delineates what metric data needs to be reported to EPA through the CB and what data will be published on the ENERGY STAR web site.

6. Pages 13-14, Sections 3.5.1 and 3.5.2, lines 489-505: IBM supports the addition of the Hot Band workload both to the data that is required to be reported to EPA and published on the website. Use of the Hot Band workload is critical to properly assess storage system performance/power characteristics given current and expected developments in automated storage tiering technologies and will improve the ability to assess the performance/power characteristics of storage systems.

7. Pages 15-16, Section 3.5.4, lines 566-615: IBM supports the inclusion of language which enables manufacturers to use performance and power models to generate qualification data for storage products. This offers manufacturers a simplified route to qualification of storage products.

B. IBM is concerned with and/or recommends changes to the following provisions:

1. Page 6, Section 1.I.3, lines 241 to 259: Flexible Minimum, Maximum and Mixed Qualification Ranges: We encourage EPA to allow manufacturers to set flexible limits at -25% or less and +10% or higher for a configuration. Products that do not meet the 15% performance threshold at -40% or +15% may meet the threshold at -25% or +10%. In such a case, manufacturers should be able to expand the qualification range beyond the -20% and +5% stated in the requirements where they cannot meet the requirement at the specified Fixed Maximum or Fixed Minimum Testing range.

2. Page 8 lines 329-336 section 1.J.3: Automated Storage Tiering Definition: It would be helpful if EPA could provide one or two examples of how the “single device” concept would be used in managing a multi-storage device product. It appears to mean that the relationship of device counts between groups of two, three or four storage devices can be represented by a ratio of 1:x:y:z, where the 1 is set against the smallest number of devices in the multi-device system and x, y, and z represent the ratio of the other devices to the device type with the smallest number of devices. As discussed in the comments to Section 3.5.3, we are concerned that the ratio of different storage devices in products sold to clients will depend on the customer application and will likely not match a single relative ratio of devices.

3. Pages 13-14, Section 3.5.2, lines 496 to 505: We continue to be concerned that since the SNIA Emerald test results are not yet understood, any attempt to modify or weight the sequential read and write scores is premature. At a minimum, we think the

IBM Comments to the Final Draft of Version 1 of the ENERGY STAR® Program
Requirements for Data Center Storage

sequential read and write scores should be reported unweighted and weighted on the ENERGY STAR webpage so that the actual and weighted score results are clear.

4. Page 14, Section 3.5.2, lines 499 to 500: The mechanics of the 70%/30% split on the sequential reads and writes and the methodology to incorporate the weighting into the published score are not clear. We assume that the sequential read score should be multiplied by 0.7 and the sequential write score by 0.3 and the two weighted scores should be added together to get a single weighted score. However, we recognize that we may be mistaken (are the two weighted scores reported individually? Or are the unweighted scores reported to EPA and the weightings are publically reported?). The Final Requirements can benefit from a clarification of the methodology for calculating and reporting the weighted metrics.

5. Page 14, Section 3.5.3.vii.b, lines 536-539: While IBM understands EPA's intent that SSDs can be included in a qualified mixed drive configuration built on qualified single device configurations (as described in section 3.5.3.i to vi) without testing of a SSD drive based system, we think that fact should be explicitly stated in this section. We recommend that EPA add a statement at the end of the first sentence of 3.5.3.vii.b stating, "...physical test point results ***in order to include SSDs in a system of mixed storage devices derived from the certification of configurations of individual storage devices.***"

6. Section I.7, Section J.3, and Section 3.5.3: It is not clear how the testing and qualification of mixed device configurations will work. While we understand the qualification of a single configuration, it is not clear whether it is possible to test Optimal Configurations of different ratios of the multiple device types to create a range of the number of storage devices of each type as described in Section 3.5.3. Such a scenario with multiple devices is needed to allow different ratios of devices to be chosen to optimize a mixed device system for a customer's workload.

For example, take a storage system whose optimum, auto-tiered configuration consists of 100 devices total, with three media types: 10K, 15K, and SSD. For Optimal, Fixed Minimum and Fixed Maximum configuration testing for a storage product consisting of 100 devices, assume it uses 40 10 K drives, 40 15 K drives, and 20 SSD drives. For the -40% case, assume it uses 24 10 K drives, 24 15 K drives and 12 SSDs and for the +20% case we would have 48 15 K, 48 10 K and 24 SSD. When tested, we can then qualify a system with the 40%/40%/20% ratio of the three types of drives.

Further, if we broaden the range of percentages for the different drive types to test/qualify an optimal configuration of 20 10 K drives, 70 15 K drives and 10 SSDs, the ratio is 20%/70%/10%. Arguably, under the requirements of 3.5.3, we can "combine" the 4 test results to create qualified auto-tiered systems within the following range of drive types:

10K 20-40 drives; 15K 40-70 drives; SSDs 10-20 drives

IBM Comments to the Final Draft of Version 1 of the ENERGY STAR® Program
Requirements for Data Center Storage

Another optimal configuration with 60 10K drives, 10 15K drives and 30 SSDs should allow us to increase the range of drives for qualified systems to:

10K 10-40 drives; 15K 10-70 drives; SSDs 10-30 drives

The result allows us to create qualified systems with between 80 and 105 drives (or 60 and 115 drives assuming my tests at -40% and +15% had performance/power scores within 15% of the optimal configuration) with any combination of the three drives as long as the number of each of the individual drive types was within the range of drives established for the individual drive types. This approach meets the requirements as established in the three referenced sections. However, it is not clear that it meets EPA's intent in establishing this qualification option. We request clarification from EPA as to whether this testing and qualification approach is allowed or if EPA intended a different set of restrictions on the mixed device with automated storage tiering testing approach.

It should be noted that it is possible to qualify a similar range of configurations by testing Optimal, Fixed Minimum and Fixed Maximum configurations for the largest volume drive type and Optimal Configurations of two other drive types, combine the three drive types in the desired ratio to meet the customer requirements, and include the automated storage tiering software as part of the system. In our view, this approach likely offers the most efficient testing method. We request that EPA confirm that this is a valid approach to qualifying systems which incorporate automated storage tiering software.

7. Page 17, lines 657-658 and page 12, Section 3.3.1, lines 461-467: We request that EPA provide a definition for "energy efficiency performance data." It is not clear if this is considered to be all of the data requested in Table 5, the more limited, published data detailed in Table 6, or some other combination of performance/power, power supply, and other storage product information. The appropriate meaning of the term is confused by the fact that section 3.3.1 separately references the performance per watt data from the energy efficiency performance data.

8. Page 17, Lines 657-60: We request that EPA specify the units for active and idle data in Table 7. We assume the units should be "IOPS per watt" for active measures and "GB/watt" for capacity measure, but we feel a documented clarification through a text discussion would be appropriate.

9. We are unsure how to communicate which storage device types and the number and combination of those devices are ENERGY STAR qualified. The large number of permutations of storage devices and their combinations will make publication of a succinct, easy to understand description of ENERGY STAR qualified configurations problematic. While this is an issue that will have to be resolved as manufacturers generate and submit test results, we request that EPA be flexible in accepting approaches for defining qualified systems and be open to discussion with manufacturers regarding their chosen approach to these communications.

IBM Comments to the Final Draft of Version 1 of the ENERGY STAR® Program
Requirements for Data Center Storage

C. IBM provides the following comments on test methods:

1. Section 5.1.B.5, line 74 page 4: The reference to the COMS test section should be Section 6.5.

2. Section 6.5, lines 99-106, page 5: Add an item, section 6.5.B which states: “Where the SNIA test process for verifying the Delta Snapshot COM does not test the specific Delta Snapshot implementation utilized by the manufacturer, the manufacturer can work with the CB to establish modifications to the test protocol to properly assess the Delta Snapshot implementation. The CB will be responsible for validating the test process and certifying the results.”

Not all implementations of the Delta Snapshot are amenable to testing by the “SNIA Emerald™ Power Efficiency 100 Measurement Specification Version 2.0.0, Rev 1: Section 7.4.5: Online Capacity Optimization Test” Methodology. Modifications are required to properly test the implementation; the Requirements need to allow these modifications to be made.