

INTRODUCTION

As the world's leading developer and provider of information infrastructure technologies, services, and solutions that enable people and organizations to transform the way they create value from their information, EMC Corporation appreciates the opportunity to respond to Draft 4 of Version 1.0 of the ENERGY STAR® Program Data Center Storage Specification, issued March 21, 2013. We are committed to acting in a socially and environmentally responsible manner and to being an attentive and thoughtful neighbor in our local and global communities. The development of an effective specification for these storage products will be an important component in achieving our shared goal of improving the energy efficiency of data centers.

Draft 4 represents another step forward in the development of the ENERGY STAR program for Data Center Storage. A number of changes have been made since Draft 3, addressing several concerns and issues raised in that draft. There are some challenges that remain in Draft 4 that we would strongly recommend be addressed prior to finalizing Version 1.0. It is our belief that, given the urgency of bringing this process to a close, the EPA should continue to work closely with the major industry organizations and other key stakeholders to develop the necessary solutions. EMC remains committed to the success of these efforts and is eager to work with you to resolve the open issues.

As this draft of the spec demonstrates, the combinations of technologies and permutations of configurations possible in Data Center Storage lead to some challenges in devising the rule set by which energy efficiency will be reported. For ease of discussion, EMC's response will focus on topic areas by complexity and by similarity. Specific portions of the document will be referred to by document line numbers to avoid ambiguity. While the specification sometimes addresses the same area in different parts of the document, our response will group our comments together for ease of reading.

RECOMMENDATIONS AND DISCUSSION

EMC is in agreement with the following changes:

- The revised definition of Thin Provisioning (line 41)
- The substitution of the term "devices" for "media" (line 119)
- The removal of the 10% power supply loading test point (line 413); we commend your decision to eliminate an unnecessary test point.
- The revised levels of mandatory COMs availability (line 467)
- Reduction of the idle test time to be consistent with the SNIA Emerald Requirements. Test document, lines 74-79.

EMC recommends the use of alternative wording in the following areas:

- The definition of "Optimal Configuration" (line 198) includes the phrase "maximum sellable energy efficiency performance ... for a given workload". What is actually being described is the configuration that represents the maximum energy efficiency performance for a specific device type using a designated workload. Given the fact that many models offer multiple device types which may be configured in different combinations, any data set delivered for a most efficient configuration is really one of a range of potential best choices. Since a customer may choose one of many tested configurations, we recommend that the term "peak performance point" or something else a bit more focused on what it represents to a purchaser. If the Note Box at line 202 is retained, the capacity optimized configuration is expressed in GB/watt, not GB/s/watt.

- The exception listed for object-based storage in 2.1.1.iv.a (line 372) appears to be unnecessary, since object-based products are excluded from the specifications scope in 2.2.2.vi (line 391).
- In the discussion of embedded devices beginning at line 414, the intention of the paragraph would be clarified if the second sentence were to be amended to read “Embedded products that do not power primary components of the storage product are not subject to PSU requirements and are also not required to be individually ENERGY STAR qualified.”
- The adaptive cooling requirement should refer to the “associated cooling needs” to relate the behavior back to the ambient temperature (Line 447).
- Given the expansion of the technologies accepted for error detection/correction now supported in the specification, it would reduce confusion to replace most uses of the term “Parity RAID” or “grid technology” with a more generic term such as “error detection/correction”. The acceptable settings or strength of the error detection/correction schemes used should be appropriately stated in the specification and the actual value used should be recorded along with the data submitted.
- The input power measurement requirements would be more precisely stated as being required to demonstrate a tolerance of less than or equal to 5% of the actual value for measurements greater than 200W (line 651)

The use of Centralized and Distributed Controller Storage (lines 299 and 303) in place of Scale-Up and Scale-Out is acceptable, but seems to be unnecessary. The definitions you provide to offer clarity in to how the ENERGY STAR Program will view these architectures is highly useful, but changing the names does not add value. In a related matter, it seems inappropriate to include Distributed architectures in the scope of Version 1.0 at this late point in the spec development process (line 374). There has been no prior public discussion of this change, and there has been effectively no industry-wide discussion as to the appropriate method of applying the Version 1.0 testing regime to this class of products. Given the range of configuration variations possible within this definition, the minimum test criteria (line 557) may not even be appropriate for all Distributed products. This approach seems to demonstrate an assumption that Distributed architectures are always composed of free-standing modules with fully replicated infrastructure; this assumption is not true of all products in this category. We recommend that Version 1.0 focus solely on Centralized designs, allowing the EPA and the industry to better develop a shared understanding of the range of products that are categorized as Distributed.

The discussion of testing Online 4 products beginning at line 526 (Section 3.5.4) should focus solely on the procedures to be used if modeled data is being provided. Direct measurement of Online 4 systems is identical to the methodology and requirements for Online 2 and 3, and should be included in that discussion (Section 3.5.3). This would simplify understanding for all testers of Online 4 equipment. Further, the accuracy requirements for modeling tools must be discussed in more detail. Given the 15% guard-band around acceptable performance/watt above and below the “optimal” configuration, a 10-15% accuracy level in a model could allow a modeled system far more latitude in acceptable configuration than its directly measured counterparts would receive. This disparity could put multiple vendors at a disadvantage with regard to qualifying systems and is at odds with the stringent limits being used for measured data and for device substitutions.

In November 2012, the EPA released a proposed approach to testing storage systems for Version 1.0. This document, while still requiring direct testing of only homogeneous systems, provided a simplified approach to collecting data and to the inclusion of SSD devices when constructing results for heterogeneous systems. Somehow, this proposal did not form the basis for the content of Draft 4, which, instead, maintained much of the burdensome approach taken in Draft 3, and left the question of

how to reflect the use of SSDs in heterogeneous configuration calculations unanswered. This unexpected outcome continues to impose an excessive burden on qualification testing. It is our strong recommendation to return to the proposal in the 11/2012 document, with the following modifications:

- Provide 3 data points for the most popular drive for an optimization type, with those points being
 - peak performance
 - a device count of the vendor's choosing above that point
 - a device count of 80% below peak
- All other device types would be tested only at their peak/Optimal point. Systems would be qualified for device counts within the bounds of the points submitted.
- Up to 10% of the device count could be supplied by SSDs configured as user-addressable storage without direct testing of an optimal SSD drive count without direct measurement or modeling

If the above approach is not acceptable, we offer an alternative. Arguably, the single biggest challenge the Program faces is a lack of data on the behavior of a varied population of storage systems (both types and configurations) under typical load configurations. We believe that in the interest of collecting a broad set of data, and in the spirit of inclusion evident elsewhere in Version 1.0, an appropriate approach would be:

- In keeping with the EPA's desire to obtain direct measurements from a wide a range of devices as possible, the peak, -40% device count, and +15% device count points would be collected for each device as currently outlined; SSDs would be treated as in the option above.
- The qualified range of products would be from minimum to maximum configurations, since in the absence of a rich data set, it is difficult to know what constitutes poor performance in the marketplace.

The reformatting of Section 3.6, Storage Device Replacement Requirements has simplified the spec and the understanding of the requirements. Several items do require further discussion and clarification.

- While the high-level specifications of capacity, rotational speed and bus interface/transfer speeds are well known for device choices, the specific details on particular specification sheets are actually confidential between disk vendors and systems vendors, and may not be shared or posted publicly. This information may be required to be shared with the system vendor's CB, but it should not be required to be submitted to the ENERGY STAR Program, so that confidentiality can be maintained.
- The need to set a limit on the degree to which performance/watt can be degraded by a substitution is understandable. The need to limit the degree to which a substitution can improve performance/watt is less clear. Many times allowable capacity increases occur through technology improvements that may deliver side effects such as increased transfer speed from the device itself. These changes may not translate directly into performance changes due to limitations of the buses to which device are connected, but they will be visible on a spec sheet. An absolute ban on improvements in such as this could limit ordinary capacity increases.
- We recommend that the requirement that replacement drives have the same interface type should be removed from this item (line 613). There are some interfaces, such as SAS and Fibre Channel, which will have the same speed and may appear in the same quantity. 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.

- The upper bound on performance improvements produced by a device substitution without a data resubmission (line 630) will limit vendors' ability to respond to changing parts availability in a timely manner. Since the substitution can occur well after a product is first tested, it may not be practical to perform a complete retest and resubmission, effectively removing a previously qualified product from the QPL because it has gotten more efficient. There is no harm to the customer in this scenario, so this appears to be an arbitrary limit that penalizes a vendor for under-reporting performance/watt, or exceeding the customers' expectations.

While storage systems do have component parts in common with servers, there are significant differences between these two categories of ICT equipment. One of the most obvious is that of scale. Servers consist of a single chassis, and have all of their primary power supplies contained in that single chassis. This is true for blade servers as well, although the size of that chassis is larger, and the number of power supplies may also be larger. These differences are not always observed and accommodated when language is transcribed from the server spec into the storage spec, particularly when it comes to the input power sampling requirements (line 675). In a medium to large storage system, the number of power sensors which must be queried and whose data must be collected and massaged can get quite large (10's to 100's). The requirements listed in the spec, of sampling frequency, data manipulation and result management may not scale in systems of this magnitude without the risk of interference with the system's primary task. The detailed behaviors outlined in this spec are, in fact, a hypothetical set of requirements; the actual users of this data do not yet exist. When the management systems being envisioned do come to market, they may dictate a different set of needs that are not met by this specification. We strongly urge the ENERGY STAR program to require only a simple, scalable set of data reporting in Version 1.0, to encourage the creation of an industry-wide spec for this data.

The Test Method document discusses how to configure a NAS system that offers a Block I/O option (lines 95-102). EMC would like to verify that this approach is only needed when the NAS functionality in the product cannot be disabled or deconfigured. Some products offer a NAS option on a Block system; in these systems, the addition of more equipment will only add to the power draw without improving performance, and this equipment is not necessary for the operation Block portions of the system. Since this functionality is optional in these configurations, and is not under test for ENERGY STAR, the hardware required to enable the NAS functionality should not be required to be present.

It is incumbent on the writers of the specification to provide clarity and consistency on the sources and revisions of their reference documents. While this draft does mention the use of SNIA specs, it is inconsistent on the revision numbers involved. That includes ensuring definitions are identical to those in source material, avoiding any potential misinterpretations due to nuanced language. This includes

- Definitions outlined by industry reference materials and industry organizations, with specific annotation as to the source and version number
- Verbatim use of tables defining values, such as Tables 1-3 in the Test Method Specification, also with version annotations
- Stabilizing definitions that might appear in more than 1 product specification by indicating which revision of another product specification was used as reference.

With respect to reporting in general, we strongly recommend that the EPA allow for a period of anonymized reporting, similar to what has been proposed for servers. This is going to be the first large-scale assemblage of data sets from the SNIA test methodology, and analysis of a broad data set based

on a varied population of storage systems and configurations. This data needs to be reviewed by a combination of industry and EPA SMEs to ensure the tests do not have inherent repeatability issues, biases or other anomalies prior to the public posting of data. We strongly recommend that the EPA work with industry groups such as The Green Grid, SNIA and ITI to develop an approach for reviewing and normalizing this data prior to full publication.

Recently, Version 2.0 of ENERGY STAR for Data Center Servers dropped the PPDS in favor of a Data Reporting Template. If the EPA is going to make this same change for Storage, it would be helpful to know this as soon as possible. Given the fact there has not been a data collection exercise in this space since 2010 and the PPDS used then does not reflect the many spec changes that have occurred since then, it is important to all stakeholders if the EPA could publish a proposed PPDS and/or DRT as quickly as possible. Without this document, multiple questions are raised by the items outlined in Section 3.5.7 (line 576):

- Which details about the storage controller are required (line 580)?
- What is specifically meant by “System power optimization capabilities” (line 587)?
- How does Line 587 differ from Line 599 (“A list of power management and other power saving features available and enabled by default”)?
- Which ASRAE Thermal Report is being requested, and what is the “selected data” (line 602)?
- If product members of product families are differentiated only by the need to supply test points for each qualified drive type, i.e. only by device configuration differences, how is the vendor supposed to provide a list of qualified products? They will all have the same model number, and the synthesized heterogeneous configurations (the qualification goal for many products) are the real object of customers’ purchasing interest, and cannot be listed. A template is needed to allow vendors to understand how this information will be reported and communicated.

CONCLUDING NOTES

Although there are still some significant open issues in Draft 4 that must be resolved, EMC applauds the progress the Data Center Storage specification has made since Draft 3. In addition to further discussion of the recommendations included in this document, EMC looks forward to participating in the collaborative effort needed to bring these challenging questions to closure. We also welcome further dialog with the ENERGY STAR Program on the full range of topics that may facilitate the conclusion of development of the final specification.