

Robert J. Meyers
U.S. Environmental Protection Agency
November 7, 2011

Dear Mr. Meyers,

EMC is pleased to respond to Draft 2 of Version 1 of the ENERGY STAR® Program for Data Center Storage Specification with our comments and recommendations. We commend you on the progress that has been made in producing this draft, and hope that our observations and proposed options help further the ENERGY STAR Program. While the attached document addresses a number of specific areas of the specification, we'd like to highlight two general themes to our comments:

- The span of product capabilities within the currently proposed scope is quite varied with respect to capacity, capacity optimization/energy efficiency capabilities, and RAS capabilities. The specification needs to consider the impact of its requirements on the product content/selling price of these products.
- There are many opportunities for potential innovation being explored throughout the Data Center Storage industry. In crafting the ENERGY STAR for Data Center Storage program, all parties must exercise diligence in avoiding prescribing specific technologies or architectures to achieve desired results.

We look forward to continuing our participation in the development ENERGY STAR specification and procedures. Should you have any questions about anything in this submission, or in any other storage topic, please let me know if we can be of any additional assistance. EMC would welcome a discussion with the ENERGY STAR team on these or related suggestions in order to facilitate timely completion of the specification.

Sincerely,

Rona Newmark
Sr. Vice President
Corporate Strategy
EMC Corporation

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 2 of Version 1 of the ENERGY STAR® Program Data Center Storage Specification, issued October 6, 2011. 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 ENERGY STAR® specification for these storage products will be an important component in achieving our shared goal of improving the energy efficiency of data centers.

Draft 2 represents some significant steps toward a useful and practical first version of ENERGY STAR for Data Center Storage. While much work remains, EMC remains focused on helping to develop a clear and meaningful specification that will recognize and foster innovation and accomplishment in the realm of energy efficiency and improvement to the energy efficiency of storage products.

Data Center Storage is a highly complex area, with many different product types meeting a wide variety of business needs. We recognize that any specification that addresses even a portion of this space will need to address several topics and may also be complex. For ease of discussion, EMC's response will focus separately on the Partner Commitments and Eligibility Criteria portions of the Specification. Specific portions of the document will be referred to by document line numbers to avoid ambiguity.

GENERAL COMMENTS

Throughout the Partner Commitments and Eligibility Criteria documents, reference is made to different specifications and publications of numerous organizations. For clarity, as well as for ease of qualification and audit, this specification should refer to such external documents by revision number and or date. This approach will remove any question as to which version of a specification, standard or other document is intended for reference. In addition, a full listing of those documents and web sites should be included as a bibliography.

In the same vein, whenever particular formulas or values are required for qualification purposes, they should be presented directly in this specification, rather than by reference. Again, this should serve to eliminate any potential confusion about specific values or calculations, allowing clearer communication amongst all parties, including the ENERGY STAR program, its partners, the ABs and CBs, test labs, and purchasers of Data Center Storage equipment.

PARTNER COMMITMENTS

1. Line 35: Since it is possible that not all configurations of a product model may qualify for ENERGY STAR recognition, and product collateral, both printed and online, refers to entire models, the requirement that the ENERGY STAR logo appear on collateral presents challenges, the most significant of which is:
 - a. When a product is recognized after its market introduction, a likely scenario when this specification first becomes effective, how much lead time is a partner allowed to update his collateral?
2. Line 49: While we appreciate that the ENERGY STAR program will make every effort to protect sensitive data from FOIA requests, it is our understanding that this protection may not be possible. Does the program accept other forms of masking, such as submissions through a third party, not in your employ, who can anonymize the results?

ELIGIBILITY CRITERIA

1. EMC recommends that the ENERGY STAR program consult with SNIA regarding the latest, complete definitions for terms from its dictionary, and that the dictionary version be noted.
2. Line 307: Among the configuration options EMC believes should be considered within family attributes are several items that would be in common with “adders” that are discussed later in this document. Please see the comments regarding Lines 476-482.
3. Lines 310-315: Maximum and minimum configurations are used to different effect in various portions of this draft, sometimes referring to the maximum (or minimum) power draw, and at other times referring to maximum (or minimum) physical configurations. Under some circumstances, these configurations may be the same; under others, they may not. EMC recommends that “maximum configuration” and “minimum configuration” refer to the physical minima and maxima for configurations as defined and sold by the manufacturer. The other usage should be named “maximum/minimum power configuration”.
4. Lines 318-329: EMC believes there are more efficient and cost effective approaches to demonstrating the scalability envelope of contemporary storage systems, particularly those in Online 3 and Online 4 categories. We recommend the approach defined in Appendix 1, which is substantially the same as the approach recommended by The Green Grid and the IT Industry Council.
5. Qualifying Products: Online Storage includes file (NAS) and object semantics as well as the block semantic that has been the focus of SNIA’s Emerald program. Throughout the development of the ENERGY STAR for Data Center Storage program, the focus of data collection has been on metrics defined exclusively for products addressed via block semantics. While both File- and Object-addressed storage can provide the ready-idle capacity/watt metric proposed elsewhere in this draft, there has been little discussion regarding what data would be appropriate to collect for future discussion of active performance metrics for either of these interface types. EMC therefore recommends
 - a. Both file and object access interfaces be excluded from Version 1 of this specification.
 - b. Products that support multiple access interfaces may only be qualified using their block interface and other such configurations must be clearly noted as not being qualified as ENERGY STAR products.
6. Excluded Products: While the specification, as written, does not differentiate between scale up and scale out architectures, very little characterization at either idle or under active workload has been performed on such products in the context of ENERGY STAR. It is therefore our recommendation that scale out architectures be excluded from Version 1.
7. Lines 417-437: EMC recommends that ENERGY STAR use the 80+ power supply rating scheme. It provides further future roadmap possibilities, and is more globally accepted. The EPRI test methodologies will still apply. One major difference will be the elimination of the 10% load point for redundant capable supplies.
8. Lines 439-455: The required use of embedded ENERGY STAR qualified products where categories exist does not always lead to increased energy efficiency in a storage system. Since these servers or other products are running a dedicated custom software load that differs from the ENERGY STAR test load, that testing is not representative of embedded products’ tested performance. Further, the specific configurations required for these embedded functions may not be qualified, leading to either overdesign or underperformance of the storage product. It is EMC’s recommendation that the power supplies of embedded equipment used to perform ancillary functions (i.e. non-core storage functions) should be excluded from the power systems standards defined in this specification. In other words, only the power supplies used for controllers and disk drives (mechanical and solid state) should be held to these standards.

9. Lines 461-474: EMC has several questions about the ENERGY STAR program's intentions and requirements in this section.
 - a. There are multiple potential configurations that could be used to provide the requested data: the configuration used for the Idle data submission, a single "typical" configuration, or configurations optimized for each type of benchmark being collected. We recommend the use optimized configurations for each benchmark, with the configuration clearly defined.
 - b. While we endorse the collection of this data, EMC believes that its public disclosure with manufacturer and model number is premature. As there are not yet any agreed, meaningful definitions of how best to compare systems using these benchmarks, their disclosure does not assist manufacturers in differentiating products. The likely case is that no one system will be superior on all fronts, thus a customer will still need to understand his own needs to find a proper product match.
 - c. Finally, manufacturers will need to understand how this data will be handled in the event of FOIA requests.
10. Lines 475-487: EMC recommends review of the following potential "adder" categories
 - a. Hardware features
 - i. Number and type of front-end connection options
 - ii. Number and type of back-end connection options
 - iii. Cache size and type
 - iv. Hardware RAS features
 - b. Software features
 - i. RAS features
 - ii. Management features
 - iii. Capacity optimization features
 - c. The specific calculation should be included directly in this specification.
11. Lines 488-493: Storage products do not have power management features as defined for computers or computer servers. Section 3.5.1 does not pertain to storage and should be dropped from the specification. As a result, Section 3.5 should be retitled to reduce confusion.
12. Lines 494-496: Since power modeling tools for storage products contain proprietary information, manufacturers should be permitted to make these available as "presales tools" as identified elsewhere in the specification. These tools will be available via the sales channel as part of the purchase selection process, and will not be directly available on a public web site.
13. Lines 499-500: The warranted performance should include a defined accuracy guard band.
14. Energy Efficiency Feature Requirements: As a general comment, the ENERGY STAR program should not be prescribing the use of specific technologies to achieve energy efficiency. This approach fails to recognize that different product classes may find efficiency to be more or less important than data integrity or high availability. This is particularly the case for Online 2 systems.
15. Lines 508-509: This is a data integrity capability not required of Online 2 products.
16. Lines 514-516: The requirement that one or more of the capabilities in Table 4 be shipped with the product and enabled by default to qualify for ENERGY STAR goes beyond defining energy efficiency criteria, and begins to dictate business terms and models. An approach that allows both manufacturers and their customers to define the value they derive from these features while still requiring at least one specific capacity optimization or energy efficiency feature for ENERGY STAR qualification would be to indicate just that: a product will be required to support one or more of these features to be qualified, and a configuration that does not include one of these features will not be able to be designated as ENERGY STAR qualified. Note again that these features are

- frequently not offered in Online 2 and Online 3 products, limiting the pool of candidate products in this space.
17. Lines 517 (Table 4): Innovations in efficiency are possible beyond the items identified in this section. The approach of this section appears to limit recognition to only those efficiencies that are common to multiple manufacturers, a negative incentive to innovation. This entire section needs to be more open-ended if the current record of efficiency innovation is to continue. One possible approach would be to accept the deployment of other, unnamed, related or similar technologies. Note that the use of spin down to improve energy efficiency, while a true savings, is inconsistent with the definition of Online storage. It is unclear what is intended to be supplied in the “Implementation” column.
 18. Lines 518-520: The shipment information being requested is both proprietary and complex to derive. The sensitive nature makes disclosure to the ENERGY STAR program, with the attendant publication of that data, a problem for businesses. In addition, not all licenses for these capabilities are sold with ENERGY STAR qualified systems. As a result, it is difficult to interpret what data is expected to be reported. It is also unclear what is meant by the phrase “impacts on manufacturers associated with providing these features”.
 19. Line 540-542: The guarantee of accuracy levels needs clarification – this should be a guarantee that the reported figures are within the accuracy guard band defined elsewhere in the specification, not a statement of absolute accuracy.
 20. Line 543: As noted elsewhere, this capability is not meaningful in the context of storage products, and this subsection should be removed.
 21. Line 546: Which specific data is expected to be reported?
 22. Section 3.8: In general, EMC believes the sampling rate is still too high. More data is going to be produced than can be handled by Data Center Information Managements systems, when looked at in the context of hundreds of pieces of data being produced every 10 seconds by the entire population of a data center. We continue to recommend that this rate be set less aggressively, and be revisited once there is more experience with the use of this data by third parties.
 23. Lines 558-559: The infrastructure required to report this information is typically not included in Online 2 (and potentially Online 3) products due to cost considerations. The approach of using external equipment such as instrumented racks and PDUs is also not economically reasonable in these categories, since that instrumentation may require greater expense than the cost of the storage product in question. EMC recommends that the ENERGY STAR program reconsider this requirement for these classes of products.
 24. Lines 573-574: We agree that it is vital to encourage the development of standards in this area. The wording of this section is very vague, and can be interpreted to indicate that the evaluation standard applied for qualification against this specification will change if a standard is declared. EMC suggests that this wording be modified for clarity.
 25. Lines 575-581: There is a difference between the internal sampling rate supported by a product and frequency with which a monitoring application will request this data. All that the manufacturers can control is the sample rate. Some clarifications are needed for the following:
 - a. Processor utilization (Line 576) is inappropriate for inclusion.
 - b. As has been discussed in the past, a typical Online 3 or Online 4 system contains many power supplies and enclosures. Is the ENERGY STAR program anticipating this data is going to be reported for each individual power supply and enclosure?
 - c. The rationale for a rolling average of input voltage is unclear, and an explanation is requested. Data centers are understandably concerned about the quality and consistency of their power feeds – these are best monitored external to each individual enclosure of a multi-enclosure/multi-rack system.

- d. It is important to efficient data center operation that temperature variation be measured and understood. At the same time, it is impractical to respond to a 10-second fluctuation in temperature via macro-level changes. These changes are best addressed by adaptive cooling, and we believe that this sample rate remains overly aggressive.
26. Line 584: Again, the guarantee is unclear. EMC interprets this requirement to be that we commit to our reporting being within the guard band stated in the requirements in 3.8.3.
27. Section 4.2: The number of units defined here requires clarification.
- a. Since the only data required to be validated for qualification is the Ready-Idle measure, what is the purpose of the 3 data points?
 - b. All active power performance data is for informational purposes and is not subject to CB qualification.
 - c. What is also unclear is the range of data required to be submitted for either the family scaling definition in Appendix 1 or for the Bookending proposal mentioned in this spec. The numbers of test units, and their size, is a clear concern for manufacturers, and will have a significant impact on both the timeliness of product qualifications and the breadth of products submitted for qualification, which should be a concern for the ENERGY STAR program.
28. Section 5: We recognize that the actual date for a final specification is still fluid. This section should also indicate the remaining process steps that are required to be completed for the Data Center Storage Qualified Products List to become a reality. This list would ideally indicate those that can occur in parallel with the remainder of the spec development process, and those which must follow it, so that Federal Procurement officials, as well as other customers, may have a realistic expectation of when qualified products can be expected to be available for sale.

CONCLUDING NOTES

EMC recognizes that while some significant open questions remain in Draft 2, the Data Center Storage specification has made progress since Draft 1. In addition to further discussion of the recommendations included in this document, we 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.

APPENDIX 1

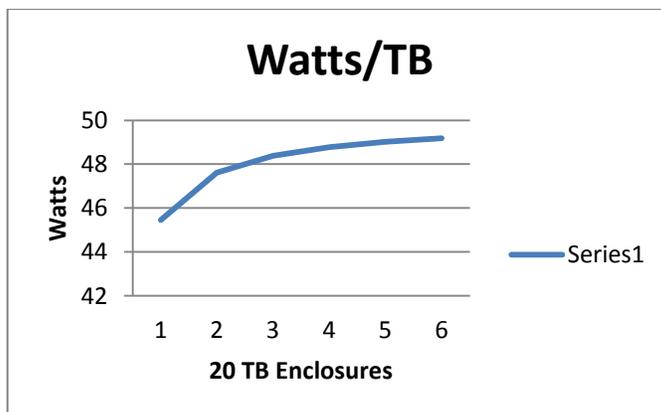
DISK SCALING PROPOSAL TO REPLACE BOOKENDS

Givens:

1. In modular architectures that have separate power supplies for controllers versus disk enclosures (also called shelves and/or drawers), the burden of the controller power draw will be amortized over the drive shelves. This is also true of Online 3 products where some drives are enclosed with the controllers, and capacity expansion occurs with additional shelves.
2. This amortization can be seen to drive to an asymptotically small burden in relatively few drawers for Online 3 and Online 4 systems (in as few as 2-3 disk enclosures for any given drive type).
3. Each disk shelf/drawer contains some overhead logic along with power and infrastructure to support the drives that will be installed. This overhead is additive to the power draw of each drive that is added to the enclosure. Each enclosure draws its maximum power (both at idle and under active I/O load) when all slots are filled. The power draw grows linearly in addition to the fixed overhead.

Example:

1. If each disk shelf/drawer can hold 20 drives, each fully loaded enclosure draws Y watts, where $Y = EO + 20 * DP$ (Enclosure draw = Enclosure Overhead (EO) + Disk Power (DP)*20). If the Controller Overhead (X) = $Y/10$, the controller burden on the first drawer is .1Y (10%); across 3 enclosures, or 3.1Y, the burden drops to <3%.



Recommendation:

Replace bookends by a methodology that recognizes the near-linear scaling of incremental disk enclosures/drawers in modular scale-out storage architectures. Under this method, manufacturers would provide 2 test data points for each type of target storage media. The first test point would be the product's controller(s) as intended to be configured for the idle test datapoint along with a single full enclosure. The second test point would be the same controller configuration with sufficient full drawers to drive the controller burden to <5% per drawer. A method of supporting mixed drive types in a single drawer, or multiple shelves of different drive types would make this form of testing even more productive for both customers and manufacturers.

There remain open questions on how many different drive types must be reported, and on how to most simply use this information to produce an appropriate figure of merit.