ENERGY STAR Data Center Storage Discussion Guide
July 2019

Overview

The U.S. Environmental Protection Agency (EPA) is sharing this ENERGY STAR Data Center Storage Program Discussion Guide to invite early stakeholder input on aspects under consideration for the revision of this specification. The topics that EPA feels are of importance for discussion prior to a Draft 1, Version 2.0 release are:

- Simplification of Product Families and Testing Requirements:
  - Configuration structure
  - Certified system ranges
  - Testing data requirements
- Removal of Variation Allowances and Mixed Drive Requirements:
  - Storage product family variation allowances
  - Version 1.1 approach to certifying systems with multiple device types
- Revisions and Additions to specification requirements:
  - Internal power supply requirements
  - Number of required capacity optimizing methods
  - Efficiency requirements (work/watt) for each workload type
- Additional program Information Needs:
  - Status of Vdbench tool, including future availability
  - Identification of any unforeseen impacts of Emerald V4 on ENERGY STAR Version 2.0.

EPA will host a webinar on July 24, 2019 from 1-3 PM Eastern Time to engage with stakeholders on the content included in this discussion guide. Stakeholders are asked to share written feedback with EPA by August 9, 2019. As always, stakeholder engagement is a vital ingredient in the success of the ENERGY STAR program and EPA looks forward to working with all parties to develop the ENERGY STAR Version 2.0 Data Center Storage specification.

Simplification of Product Families and Testing Requirements

Through discussion with stakeholders during the life of Version 1.0/1.1 specification, EPA and stakeholders share a goal to simplify the product family structure and testing requirements to encourage greater participation in the program. To that end, EPA is considering the following changes to its testing approach.

- Require a single optimal (i.e. Best-foot-forward) point test for each workload type (transaction, streaming, composite and/or capacity[1]) certified as ENERGY STAR. This configuration would be determined by the partner and contain only hard drive devices (no SSD etc.) For systems that are *only able to be shipped* with SSD or flash storage, those systems would test with that hardware installed instead.
  - This optimal point test result would be required to exceed the minimum energy efficiency criteria, which will be developed as part of this process. Any other configuration offered in that model line designed for that workload type could be included in the family that is

[1] Note that systems can typically be optimized similar for both streaming and capacity, so that submissions for both are unnecessary as they are redundant.

ENERGY STAR certified, including configurations containing SSD or other flash hardware in a mixed configuration.

- EPA believes that this approach is possible based on the latest round of product data collected, which revealed that the ramp up and drop-off in performance, shown in Figure 1 below, in older products is no longer present in newer products, eliminating the need for capturing additional boundary points in storage device count. EPA believes this development is largely due to increased compute capacity in newer storage controllers.

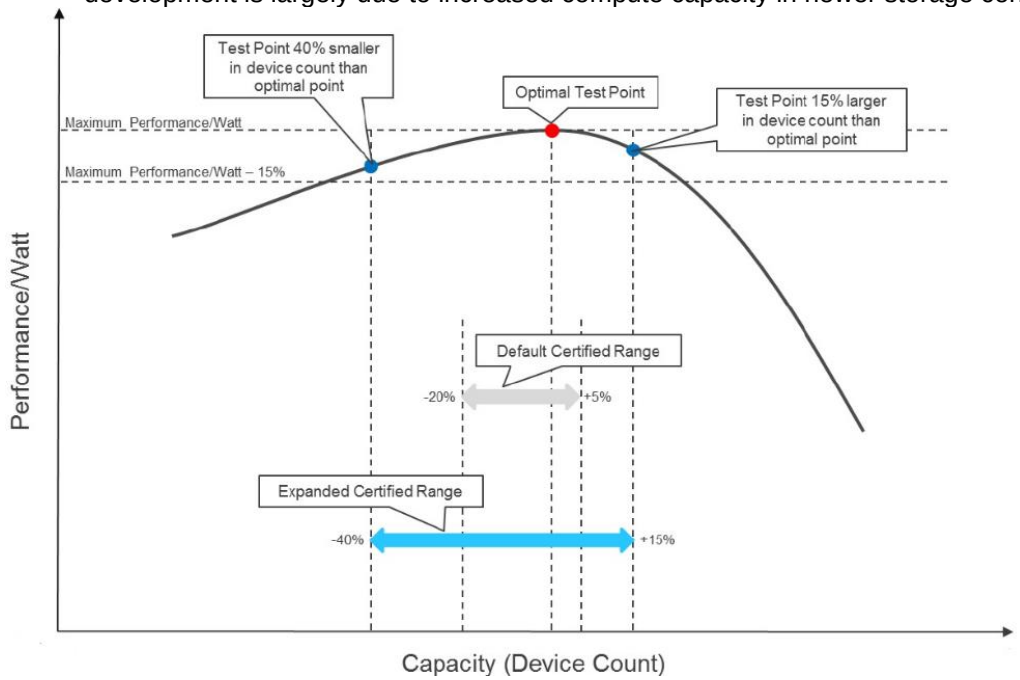

Figure 1: Example of certified ranges used in Version 1.0/1.1

The result of this change is that the total number of testing configurations would be reduced from three per workload type certified to one, for a maximum number of tests of three tests for the three workload types for products with both block I/O and file I/O capability, versus the previous maximum of nine tests.

1. Do stakeholders agree with the rationale behind this new approach, focusing on one optimal point for testing, which emphasizes the efficiency of the controller and associated software?
2. Are there alternative approaches that stakeholders would like EPA to consider when developing the Draft 1 specification?
3. Do stakeholders agree with the simplified testing approach? Are there scenarios that stakeholders could foresee that would not be covered by this testing approach?
4. Is there any reason to suspect that the more recent linear behavior observed in a storage product's performance/watt over a large range of device counts will not continue in the life of Version 2.0?

Removal of Variation Allowances and Mixed Drive Requirements

As part of the development of this document, EPA reviewed analysis and recommendations provided by The Green Grid and noted that updates to storage devices and the firmware used in the storage product over the lifetime of the certification does not impact the efficiency of the product in a meaningful way. The only significant difference observed is typically an increase in the storage capacity. Considering this information and the proposed simplification of the product family, EPA is considering removal of the current Section 3.6 (titled *Storage Product Family Variation Allowances*) and allowing storage device substitution in product families over the lifetime of the certification. This will only be appropriate in cases where the optimal point configuration continues to meet its minimum energy efficiency requirement.

When combined with the changes outlined in the first section, partners will no longer need to account for the inherent testing and reporting complexity of mixed storage device configurations. EPA is proposing to remove the existing guidance on certifying configurations composed of multiple device types in the product family definition as well as device substitution rules that previously allowed mixing multiple test configurations to create a product family. Instead, certification of the product family will occur through the single simplified optimal point configuration using a manufacturer selected composition of either all spinning storage devices, or all solid-state storage devices if the product is not sold with spinning media.

1. EPA requests feedback on the removal of these sections of the specification.
2. Would the removal of these two sections cause any confusion or conflict with other elements of the specification?
3. Does the removal of these sections create unanticipated gaps or loopholes in the requirements?

Revisions and Additions to Specification Requirements

As part of Version 2.0, EPA intends to increase the stringency of certain aspects of the ENERGY STAR criteria and add new criteria to differentiate the top performing models in the market.

Internal Power Supplies (IPS): EPA intends to increase the efficiency requirements for IPSs and is considering two approaches. The first approach would be to increase the requirements at all load points in an incremental step, e.g. by increasing the requirement to 80Plus Gold. Alternatively, EPA could target the tested load points where storage products most frequently operate with more stringent 80Plus Platinum or Titanium requirements and retain the current requirements for the less frequently used load points.

1. What is the performance of IPS used in today's storage?
2. What, if any, hindrances are there to moving to 80Plus Platinum or Titanium?
3. How do the savings of requiring greater IPS efficiency for a limited number of IPS load points which represent the most common load condition for storage products compare to those of requiring Platinum or Titanium across all load points? How do other factors, like cost, compare?
4. What load points of those currently tested (10%, 20%, 50%, 100%) best represent a typical operating load for a storage product?

Capacity Optimizing Methods (COMs): EPA currently requires that all ENERGY STAR Online 3 and 4 systems offer at least one COM at the point of sale of the storage product. EPA has become aware that most, if not all, storage products offer thin provisioning, which is listed as one of the applicable COMs to claim in Version 1.1. EPA is considering increasing the number of required COMs that are made available. EPA is seeking feedback on the following:

1. Are there any other COMs that EPA should be considering beyond thin provisioning, data deduplication, compression, and delta snapshots?
2. Are there file based storage COMs that EPA is not covering that should be included in the specification?
3. Would it be clearer to stakeholders if EPA required thin provisioning for all products and maintained the current number of COMs required?
4. Is it appropriate for EPA to consider making all listed COMs available in Draft 1?

Energy Efficiency Requirement: Based on the data available to the Agency, EPA believes that there is room to differentiate more efficient products with the ENERGY STAR without compromising performance. EPA anticipates adding efficiency criteria for all covered storage products that will reflect one requirement for each workload type (a total of four possible requirements).

With the proposed testing changes, this requirement will allow EPA to compare how efficiently the product's controller and software can store data for a given level of performance. The transaction requirement will be in IOPS/watt, the streaming and composite requirements in MiBs/watt, and the capacity requirement in GB/watt. This new requirement will differentiate more efficient products from poor

performers. EPA will consider the distribution of performance data to determine appropriate levels, ensuring sufficient availability of storage products that meet the ENERGY STAR requirements. The ENERGY STAR dataset includes all 210 ENERGY STAR certified products, including both block I/O and file I/O systems. EPA is requesting data for additional, non-certified products.

1. EPA requests any additional data that stakeholders may have on the performance of non-ENERGY STAR certified products to incorporate into its dataset for level setting purposes.
2. EPA has a limited dataset on file-based storage products. Partners that have data on non-ENERGY STAR certified products are encouraged to submit their data to inform level setting.
3. Are there any other considerations EPA should factor in to setting efficiency requirements for storage products in Version 2.0?

Additional Program Informational Needs

There are two industry led areas where EPA is requesting additional feedback in advance of Draft 1:

Vdbench: EPA has been made aware of concerns from stakeholders that the Vdbench tool which is required to run Emerald testing for block I/O storage products may not remain open access going forward. EPA is concerned about the impact this could have on the ability for partners to test their block I/O products in Version 2.0.

1. What is the latest status of Vdbench accessibility?
2. Are SNIA or other industry organizations planning to use alternative benchmarks for block I/O systems if Vdbench becomes unavailable?

Emerald Version 4.0: EPA has been in regular communication with SNIA regarding the Emerald Version 4.0 test methodology update and has plans to reference Emerald Version 4.0 in the ENERGY STAR Data Center Storage Version 2.0 test method. EPA is seeking the following feedback regarding the latest development of Emerald Version 4.0:

1. Are there any concerns regarding implementing the new Emerald taxonomy in Version 2.0?
2. EPA is interested in including Online 5 and 6 in the scope of Version 2.0. Do stakeholders have any feedback on the inclusion of these products within the ENERGY STAR scope?
3. What is the expected timeline for finalization of Emerald Version 4.0?

Version 2.0 Revision Schedule

Following this discussion guide, EPA expects to release a Draft 1 specification by late summer 2019. EPA expects that Draft 2 will come out in the fall of 2019 with the final specification published in early 2020 and an effective date nine months later.

1. Are there any market issues that impact the anticipated timing of this development process that warrant consideration?