

UNITED STATES ENVIRONMENTAL PROTECTION AGENCY
WASHINGTON, D.C. 20460



OFFICE OF
AIR AND RADIATION

December 2, 2011

Dear Data Center Storage Manufacturer or Other Interested Party:

The U.S. Environmental Protection Agency (EPA) would like to share its view on the potential role of Best Foot Forward (BFF) in the development of the Data Center Storage Version 1.0 specification.

The Best Foot Forward (BFF) approach looks at a storage product holistically. It allows the vendor to select and test one or more specific product/family configurations at operating points determined to be at or near maximum performance values, i.e., the "sweet spots." For more information on BFF, please see Section 3.4 of the *SNIA Emerald™ Power Efficiency Measurement Specification Version 1.0*.

EPA believes that BFF is an important development in the effort to successfully identify energy efficient storage system configurations and applauds the work of industry in creating it. However, the ENERGY STAR program's requirements place additional constraints and pressures on any efficiency specification developed, requiring EPA to create extensions to or modifications of the BFF concept for its own needs. EPA requires a very high level of certainty regarding the energy savings realized through the ENERGY STAR label but, at the same time, must keep the testing burden placed on manufacturers to a reasonable level. Labeled products are subject to the new witnessed/third-party verification process, while the program itself is subject to more traditional government auditing and accountability processes. ENERGY STAR specification development has always been a heavily data-driven process and the need for data or at least a path forward to generate necessary data is stronger than ever. A solution must be found that meets the needs of consumers, manufacturers, and the ENERGY STAR program.

To address these concerns, EPA believes that additional criteria that incorporate much of BFF but go beyond its current structure need to be in place to (1) reduce opportunities for gaming and (2) to assure that measured results are able to provide meaningful insights sufficient to the needs of the ENERGY STAR program and consumers.

Premises

Before describing its concerns with BFF, EPA wishes to highlight a few points:

- 1) EPA recognizes that storage media characteristics have a major influence on system energy efficiency. For example, some drives will perform much better on Transaction-oriented workloads that involve a large number of I/O operations per second with low seek times, while others will excel at Capacity-oriented workloads where the ability to stream large, continuous chunks of data is most important. EPA wishes to recognize this fundamental difference, during both specification development and product qualification.
- 2) EPA recognizes that BFF will provide a peak energy efficiency point for a given configuration, and that knowing this point is very important to build an understanding of a system's overall energy performance.
- 3) EPA wants to ensure that, given a particular storage system, there is an opportunity to identify BFFs for both Transaction and Capacity configurations. We expect these two BFF optimization points will exhibit dramatic differences in both the types of storage devices used and in additional configuration options within a given system.

Concerns with BFF

Concern #1: EPA remains unclear as to how system capacity scaling impacts the active power performance of a given system. Any active power metric requires some measure of units of work performed per units of energy consumed, and it remains unclear how systems scale according to such measures. Understanding how systems scale is a key concern for ENERGY STAR and this knowledge is needed to set active mode energy efficiency levels in future specification revisions.

- Data received to date does not lead to any firm conclusions on this point. In the data provided to EPA by industry, there is a single product with significant scaling of drive numbers to support analysis.

# of HDDs	From 30 to 60	From 60 to 120	
% Change (Work/Watt)			Delta
Active "A" - Random Read	13%	6%	-7%
Active "B" - Random Write	2%	-4%	-6%
Active "C" - Sequential Read	-42%	-46%	-4%
Active "D" - Sequential Write	-36%	-44%	-8%
Active "E" - 70R/30W - 25%	12%	5%	-7%
Active "F" - 70R/30W - 75%	4%	-2%	-6%
Active "G" - 70R/30W - 100%	4%	1%	-3%
Ready Idle	33%	7%	-26%

- Given the limited data at hand, EPA is unable to conclude that active efficiency scales linearly with number of installed drives and is also unable to characterize the shape of an alternative nonlinear performance curve. While suggestive of a possible performance vs. installed drives trend, a single system is unfortunately not enough to enable EPA to assume a particular curve to use in active power testing.
- BFF represents a single testing point that is the most efficient product in a family, but EPA must be able to either test or mathematically project the point(s) at which a product family comes closest to or crosses below the ENERGY STAR minimum efficiency threshold. Stated another way, ENERGY STAR must know the lower bound of efficiency in a qualifying product family, but BFF is designed to illustrate the upper bound of efficiency.
- It may be possible to test the BFF and mathematically estimate/project to the point(s) that intersect the ENERGY STAR minimum efficiency threshold, thus defining a family based only on a BFF configuration. However, calculating such a projection requires a thorough, data-backed understanding of the performance curve (Work/Watt for varying workloads) of a typical system. EPA understands the great difficulty involved in acquiring this data, but a mathematical projection undertaken in the absence of more data represents a large risk for the ENERGY STAR program and for manufacturers that may face verification testing of their products. EPA's current intention for Storage v1.0—to set idle thresholds and also collect active data—could provide the data necessary for such an arrangement to be constructed for v2.0.

Concern #2: EPA understands that there is no specific guidance or constraint around evaluation of mixed storage media configurations within the Emerald specification. EPA recognizes some mixed media as very appropriate (e.g. SSD + SATA) for a segment of storage deployments, while also believing that other configurations (e.g. FC + SATA *in the same partition*) are likely of no value and not commonly deployed.

- EPA is interested in extensions to the current Emerald Power Measurement Specification to recognize 'hybrid' systems (systems which utilize a mixture of fast access and large capacity

storage media devices, e.g. SSDs and SATA, under the direction of intelligent data placement logic). A pre-described nonlinear data access distribution pattern across the evaluated storage devices may be one potential approach to address this issue.

- EPA is concerned about opportunities for gaming. For example, optimizing for Idle thresholds where 80% deployment is high capacity SATA drives, and hence able to mask any other drive types being used for the remaining 20% of the system which could have a poorer energy footprint.
- EPA is concerned about the inability of uncommon mixed storage device configurations to provide meaningful active insights for end consumers (e.g.: Transaction-oriented and Capacity-oriented drives contained in the same partition).

Summary: EPA believes BFF is an important component to successfully identify efficient configurations, but its use in the ENERGY STAR program will require that additional criteria be put in place to reduce opportunities for gaming and to assure measured results can provide meaningful insights sufficient to the needs of the ENERGY STAR program.

EPAs current proposal is that three data points are to be collected for each system variation submitted:

1. Peak point (Best Foot Forward)
2. Point illustrating a larger scaled system of like configuration
3. Point illustrating a smaller scaled system of like configuration

Points 2 and 3 are needed to provide insight into how a given system scales and how changes are reflected in its energy efficiency under an active metric. In deciding how to determine the scaling points for items 2 and 3 above, EPA sees the following potential options:


- Book-ending: This would represent the largest and smallest configuration the vender wishes to deliver under a qualified ENERY STAR configuration.
- % of peak (BFF) configuration: This would represent configurations at pre-defined scaling points based off the peak configuration, e.g. 50% smaller and 50% larger.

The EPA is interested in discussions around the two options above and is open to any other options that meet the goals of the ENERGY STAR program.

The EPA is also interested in appropriate constraints on the utilization of the Emerald Power Measurement Specification to prevent unhelpful / uncommon configurations, while also allowing for extensions to the tool to allow it to highlight benefits of emerging technology and configuration options (e.g. Hybrid).

Thank you for your continued support of the ENERGY STAR program. Please direct any specific questions to RJ Meyers, EPA, at Meyers.Robert@epa.gov, or 202-343-9923; or John Clinger, ICF International, at jclinger@icfi.com, or 202-572-9432.

Sincerely,



Robert Meyers
U.S. Environmental Protection Agency, Climate Protection Partnerships Division
ENERGY STAR Program