



ENERGY STAR® Tier 2 Computer Servers

Stakeholder Meeting Discussion Guide

September 25, 2009, Menlo Park, CA

Purpose

This document is intended to provide stakeholders with background information and key discussion questions specific to use of an active mode rating tool in the ENERGY STAR Tier 2 Server Specification. This guide will serve to frame the discussion at the September 25 stakeholder meeting.

Overview and Background

ENERGY STAR computing product specifications are designed to make it easier for purchasers to identify products that provide best-in-class energy efficiency while also meeting their computing performance requirements. The Version 1.0 (Tier 1) specification for computer servers represents EPA's first attempt at providing ENERGY STAR tool to achieve this goal in the enterprise and data center server market. The Tier 1 specification set stringent energy efficiency criteria and created a structure for standardized testing and reporting of energy efficiency performance for qualified products.

Qualified Tier 1 ENERGY STAR computer servers must:

- Utilize power management techniques and/or meet efficiency criteria to minimize power consumption when idle,
- Meet targeted sub-system efficiency requirements to improve the energy performance of various components,
- Supply data center operators with real time temperature and energy consumption data that can be used to improve overall facility energy management, and
- Be tested and characterized according to standard information reporting requirements to promote straightforward comparison of products and provide IT purchasers with a robust energy and computing profile to help inform the decision making process.

Collectively, these criteria recognize those servers that are designed with energy efficiency in mind and provide a procurement tool that, when applied properly, can be one means to reduce energy waste associated with IT.

Tier 2 aims to evolve the program by adding a means to measure the overall efficiency of the server while it is performing actual computing *work*. This "active mode efficiency" criterion was intentionally omitted from the first set of requirements to allow for additional discussion on this complex topic. EPA described this approach during the development of the Tier 1 specification as well as its intent to revisit Active mode soon after Version 1.0 was finalized. EPA expects to include this new requirement as part of the Tier 2 specification, scheduled to take effect toward the end of 2010. Meeting this deadline will require EPA to complete specification development as soon as possible, to provide adequate time for a smooth transition to Tier 2.

EPA is hosting a stakeholder workshop on September 25, 2009, to lay the groundwork for the development of active mode efficiency criteria. During the meeting, EPA will discuss the benefits and drawbacks of various approaches to address active mode efficiency it has considered to date. This discussion will cover the following topics:

- Desirable characteristics of an ENERGY STAR active mode efficiency rating tool,
- Understanding the pros and cons of various approaches to the develop the tool,
- Identify the criteria by which EPA will judge each approach (e.g. effectiveness, time and resources required for implementation), and
- Develop a timeline and identify milestones on the path to completion.

Desired Characteristics of an Efficiency Rating Tool

EPA's intent is to choose a tool that is supported by server manufacturers, server purchasers, and other stakeholders. Whether making use of existing benchmark workloads or developing specialized tools, EPA seeks an approach that limits barriers to implementation and meets the market need for trusted server efficiency and performance data. To meet these needs, EPA will evaluate rating tool and/or workload approaches based on a variety of desirable features and characteristics including the following:

- *Technology-neutral/architecture agnostic* - can be fairly applied to a wide range of platforms and system configurations incorporating today's technologies as well as being adaptable to new technology developments.
- *Limited barriers to implementation* - can be conducted at a reasonable cost, both in capital and labor. Utilizes automation to limit testing burden and requires limited maintenance.
- *Evaluates a variety of end-use scenarios* - e.g. HPC, web services, email services, database management, shared file services, etc., that are reasonably representative of typical types of computing workloads.
- *Comprehensive evaluation* - addresses multiple points on the server's utilization curve.
- *Transparent and standardized* - underlying test conditions and protocol are uniform to avoid test variability and to ensure repeatability of results. Content of the workload is clear enough to allow results to be put in proper context.
- *Available within an acceptable timeframe for use in the specification.*

Charting a Path toward an Active Mode Efficiency Rating Tool

There is no single workload-based rating tool currently available to evaluate server active mode efficiency which meets the needs of all stakeholders. Therefore, EPA feels the development of a hybrid-rating tool incorporating positive elements of different approaches may be necessary. Several different approaches are outlined below, including EPA's preferred approach for a tool to rate server active mode energy efficiency. During the stakeholder meeting, EPA feels a high level review of several approaches would be valuable to the discussion. This discussion would clarify the merits and rationale behind hybrid tool approach as well as identify needed refinements. Subsequent to the stakeholder meeting, EPA will work with stakeholders to identify next steps and milestones to develop this new tool in a reasonable amount of time with available resources.

Review of Key Elements of Different Approaches:

- A. *Refined Tier 1 Criteria plus Individual Benchmark Workload Reporting.* Require manufacturers to publish data in the ENERGY STAR Power and Performance Data Sheet (PPDS) for servers operating existing benchmark workloads which have been extended to include a performance vs. energy consumption metric. Near term examples could include SPEC system-level benchmarks, TPC benchmarks, etc.
 - Pros: Streamlined transition from current requirements; insight into active mode using existing tools.
 - Cons: Appropriate tools would have to be selected and evaluated. Assigning one or more expected end uses for a general-purpose server may be unrealistic; use of benchmark data may focus attention on the performance score and not the *efficiency* result; past vendor hesitance to widely publish benchmark results could be a barrier to participation in the program.
 - Required Resources: Stakeholder input on effective benchmarks for use in the power-performance publishing requirement; continued idle data collection for rackmount, pedestal, and blade servers.
 - Time to Completion: Short, but the timeline would need to account for availability of appropriate benchmark workloads under development by third-party organizations.

- B. Refined Tier 1 Criteria plus Unified Benchmark Workload Reporting. Require manufacturers to publish data in the ENERGY STAR PPDS for servers under a single *blended metric* that combines power and performance results for several existing benchmark workloads. Blended metric score of a certain value or level would serve as means to determine qualification in combination with other criteria.
- o Pros: avoids immediate need to develop specialized rating tools; provides insight into active mode; single score fits well with general ENERGY STAR structure.
 - o Cons: Appropriate workloads would have to be decided. Results would likely be in incompatible formats, making the merge into a single result challenging; single result might obscure good performance in one expected end use if the server performs poorly in an unsuitable benchmark; use of benchmark data may focus attention on the performance score and not the *efficiency* result.
 - o Required Resources: Stakeholder input on effective benchmarks for use in the power-performance publishing requirement; continued idle data collection for rackmount, pedestal, and blade servers; creation of the metric to tie various test results together.
 - o Time to Completion: Medium to Long.
- C. **PREFERRED APPROACH:** Refined Tier 1 Criteria Plus Specialized Efficiency Rating Tool. Require servers to meet refined Tier 1 criteria, with active mode efficiency addressed through development and implementation of a multi-workload rating tool.
- o Pros: specialized tool would be developed with efficiency in mind rather than adapted from tools focused on performance; ENERGY STAR stakeholders would have opportunity to comment on structure of tools; tool development could occur in parallel with other development efforts (idle data collection, blade test procedure development).
 - o Cons: requires development time for the rating tool; the underlying components of the workloads would need to be clear to users to ensure that results were held in the proper context; capturing all servers in the scope of the program would be challenging.
 - o Required Resources: Software development time to build the rating tool; rating tool data collection upon availability of a pre-release version and/or the final version.
 - o Time to Completion: Long.
- D. Blend of Approaches B & C. Implement Approach B as part of Tier 2 for some specified period of at least one year. Use Tier 2 data for development of Tier 3 specification similar to Approach C.

This list is not intended as exhaustive; rather it is intended to outline the merits of several broad approaches including EPA's preferred approach of a rating tool that could provide value to a variety of stakeholders. The plausibility of the preferred approach depends on identifying resource constraints and time to market. Lastly, it is critical that in terms of program implementation, the output of any rating tool must be used in its proper context for the end user, so as to not misrepresent the results.

Questions for Discussion

- Are there unforeseen barriers that preclude the use of an active mode-rating tool?
- Are there developments underway that support EPA's active mode efficiency goals for the Tier 2 specification?
- Are there any characteristics of an active mode efficiency-rating tool that EPA did not identify in this document? If so, what effect do these characteristics have on the viability of the preferred approaches presented in this document?
- Which workloads would be appropriate for inclusion? Which of those workloads would be

- reasonable to combine into blended metric?
- What infrastructure is required for EPA to implement each of these approaches? How might the selected approach effect manufacturers who wish to qualify products to Tier 2?
- What challenges exist to implementation and how must the PPDS be augmented to facilitate the display of active mode efficiency in the proper context?
- What role can other industry organizations play in the development of an active mode efficiency-rating tool? How can stakeholders facilitate the process?

Next Steps

Prior to development of the next draft of Tier 2 requirements, EPA seeks stakeholder comments on the approaches and topics presented in this document. EPA is interested in thoughts on what is achievable given the anticipated Tier 2 development schedule, plus suggestions for development of a future tool to evaluate server active mode energy efficiency. Stakeholders wishing to provide comments may direct their suggestions to Evan Haines, ICF International, at ehaines@icfi.com.