Server Energy Use Evaluation – Discussion Document (Update) June 2011

The following discussion reiterates EPA's goals in assembling a dataset for Computer Servers, provides more details on the type of data of interest, and reflects stakeholder feedback received through the month of May. This document does not reflect any changes in the scope or intent of the dataset assembly process but is provided for further clarification to industry stakeholders.

Testing Overview

Goals for Computer Server Dataset

As noted in a similar memorandum in late February, EPA's scope for this effort includes 1-2 socket servers, blade servers and 3-4 socket servers. EPA intends to use the dataset under the following goals:

- **Evaluate the newly proposed structure for product families**. The five-point structure proposed in March is discussed further in this document, along with comments received from stakeholders.
- **Evaluate manufacturer power calculators as a data source**. Stakeholders raised the possibility of leveraging the existing data models manufacturers use to map expected power consumption of end configurations. At the meeting on March 11, EPA asked that stakeholders take the first step of comparing actual laboratory test data ("physical test data") to output from power calculators. Some stakeholders expressed their intent to do so and EPA encourages further related efforts.
- **Evaluate adders**. Data from single variable testing deviations, where one feature or technique is altered between tests on an otherwise consistent server configuration, should allow for investigation of the allowances provided for above-base functions.

As is consistent across the ENERGY STAR program, the dataset will be assembled with voluntary input from interested stakeholders, masked to remove identifying information for the system manufacturer, and then analyzed, with opportunities for stakeholder comment.

Five-point Server Product Family Structure

In a memorandum distributed on February 23, 2011, EPA proposed a five-point family structure (Figure 1) for Server product families. Based on the general support for the proposal received in stakeholder feedback, EPA intends to proceed with this structure as the foundation for the dataset assembly.

Stakeholders suggested small modifications to the definitions for product families: (1) "Machine Type" was a suggested addition to the Common Product Family Attributes; (2) Power Supply Rated Output was recommended as a variable that could change between the representative tests; (3) wording to clarify commonalities across a family's processor options were suggested; and (4) processor socket power was suggested as an additional defining element for the representative test scenarios. These suggestions are reflected in the updated definitions provided below and identified with tracked revisions.

- Common Product Family Attributes: A set of features common to all models/configurations within a product family that constitute a common basic design. All models/configurations within a product family must share the following:
  - Be from the same model line or machine type;
  - Share the same form factor (i.e., rack-mounted, blade, pedestal);
  - Share processors either from a single defined processor series or processors that plug into a common socket.
  - share PSUs ~~with both the same rated maximum output and~~ greater ~~than~~ or equal ~~to the rated~~ efficiencies~~y of the tested configurations at~~ all required load points specified in Section 3.2 (i.e., 10%, 20%, 50%, and 100% for single-output; 20%, 50%, and 100% for multi-output).

- Product Family Tested Product Configurations: This product family proposal requires five tests to define a product family. The variations listed here are explained within the context of an example product family shown in Figure 1, below.
    - Purchase Consideration Variations:
        - Low-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage (HDD/SDD), and I/O devices that represents the lowest-price computing platform within the Product Family.
        - High-end Performance Configuration: The combination of Processor Socket Power, PSUs, Memory, Storage (HDD/SDD), and I/O devices that represents either the highest-price or highest-performance computing platform within the Product Family.
    - Typical Configuration:
        - Typical Configuration: A product configuration that lies between the Minimum and Maximum Power configurations and is representative of a product with high volume sales.
    - Power Utilization Variations:
        - Minimum Power Configuration: The minimum configuration that is able to boot and execute supported OSs. The Minimum Configuration contains the lowest Processor Socket Power, least number of installed PSUs, Memory, Storage (HDD/SDD), and I/O devices, that is both offered for sale and capable of meeting ENERGY STAR requirements.
        - Maximum Power Configuration: The vendor-selected combination of components that maximize power usage within the Product Family once assembled and operated. The Maximum Configuration contains the highest Processor Socket Power, greatest number of installed PSUs, Memory, Storage (HDD/SDD), and I/O devices that is both offered for sale and capable of meeting ENERGY STAR requirements.

Figure 1: Product Family Testing

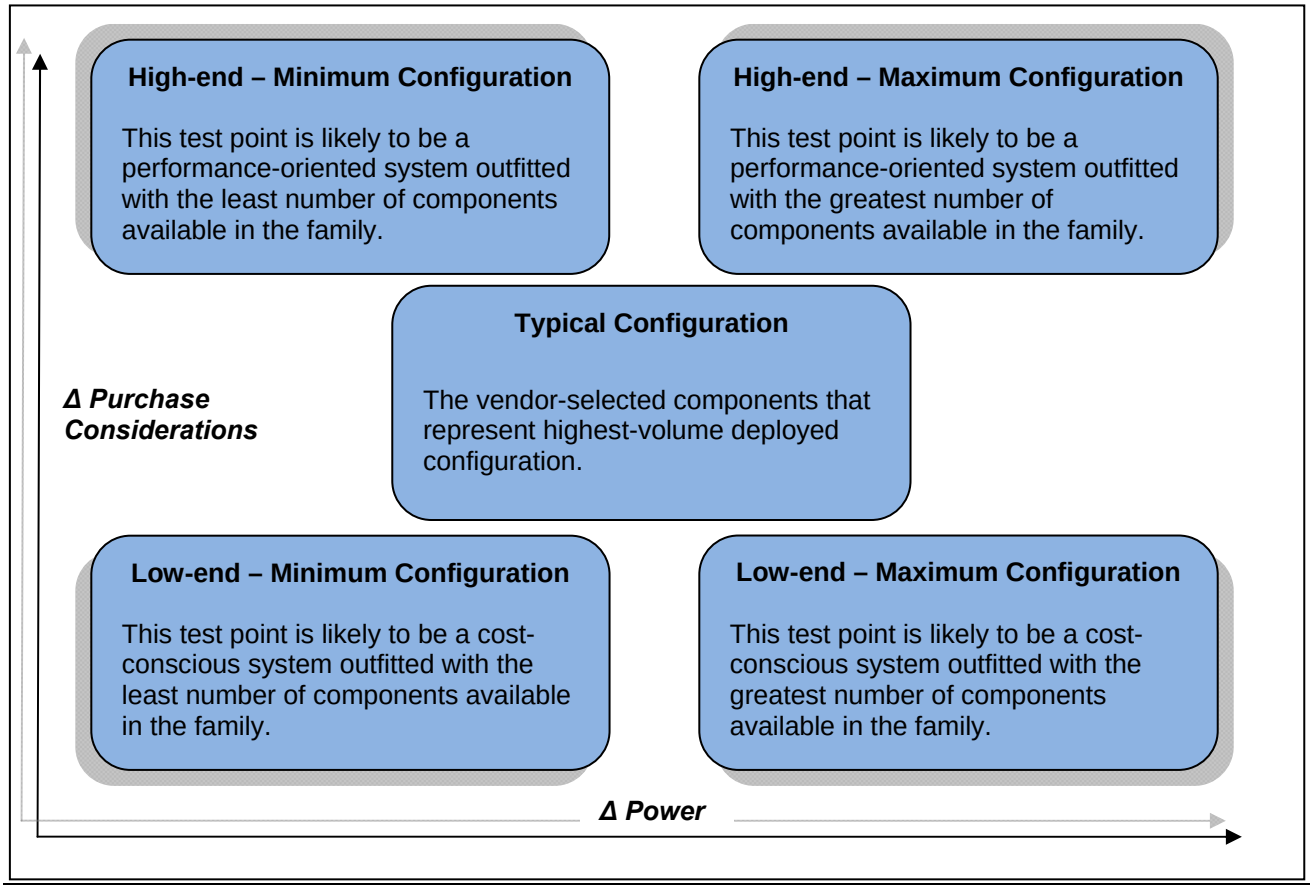

Dataset Contents

To address the goals for the dataset, EPA plans include system-level data, relevant data from power calculators (described further in the next section), and component Idle Power data.

- **System-level test data**. The primary area of interest, Full Load and Idle Power data resulting from actual laboratory testing will be the foundation of the dataset.
- **Power calculator data**. Full load and Idle power derived from power calculators will be used primarily for investigating component idle power. Based on the outcome of comparisons between calculator data and physical test data, it may also be used to supplement five-point physical test data.
- **Component Idle Power data**. To better understand the impact of component choices on system power consumption, EPA will welcome power draw data at Idle and optionally at Full Power for add-in components. The primary focus will be on I/O, which stakeholders have suggested may need review based on experiences with Version 1.0 adder amounts.

Permutations of Interest

As noted in the February memorandum, EPA is interested in evaluating the testing permutations listed below at a minimum.

- **Processor**
  - Choose base configurations with processors for each available core count per socket and rated power that represent, at a minimum, the highest and lowest available rated system speed (FSB). This processor shall be used across a set of typical, maximum, and minimum testing runs.
  - For the typical configuration, test the same configuration with a processor socket unpopulated (if offered as an option to customers).

- **PSU** – Choose power supply options that represent typical options for customers (if applicable, select both PSUs optimized for efficiency and standard models). PSUs should be selected using a manufacturer's standard practices, with a brief, written rationale/explanation of why the PSUs were selected for each configuration accompanying power data (e.g., PSU was tuned to Full Load Power of the configuration, PSU is the only option available across the family).

- **HDD or SSD** – Include storage elements that represent a range of customer preferences (i.e., high performance HDDs versus high efficiency HDDs, etc.).

- **Memory (DIMM)** – Consider testing several vintages of memory technology such that EPA can evaluate the impact of new memory technologies on overall server efficiency.

- **I/O Device** – Provision systems with minimal I/O add-in cards; for testing, ensure the server offers at least one Ethernet port (using a single add-in card only if no onboard Ethernet support is offered). EPA welcomes device-level power information on I/O devices to build on the existing I/O dataset from Version 1, available at http://www.energystar.gov/ia/partners/prod_development/new_specs/downloads/servers/Final_Draft_Server_IO_Data_Set.xls.

Blade Testing

In the February memorandum, EPA stated its intention to test blades using an N-1 approach: a half-filled chassis of homogeneous blades would first be tested for Idle and Full-load Power, with a single blade removed for a second test in order to calculate the per-blade power value. At the subsequent meeting with stakeholders, response was mixed between support and suggestions for a return to EPA's previous proposal of a separate single-blade test. EPA asked at the meeting that stakeholders follow up with written description of the desired testing model. No written comments on this point were received, and EPA is proceeding with the N-1 testing approach.

The following testing guidelines for Blades reflect other comments received from stakeholders on blade testing after publishing of the February memorandum. In summary, these were: (1) an option should be provided for full chassis testing should a manufacturer choose this over a half-filled chassis; (2) the same blade configuration should be used among all blades in a test.

- **Chassis population**: A fully-populated blade chassis may be optionally tested instead of a half-populated chassis. The half-populated chassis scenario remains the default and the only other accepted option in the dataset.
- **Blade Server configuration for a single test point**: All blades must share the same configuration ("homogeneous" blade selection).
- **Blade Server configurations among the five test points**: The Blade Server configurations should reflect the general guidelines provided for each of the five test points. For example, the *High-end – Maximum Configuration* would involve a half (or fully) populated chassis with all configurations representing a performance-oriented system with the greatest number of components offered in the family.
- **Characteristics to identify when recording test results**:
    - Form factor would be added: e.g., single-wide, double wide, half-height

The testing process is otherwise as described in EPA's February testing memorandum.

Other comments were received on the format of requirements, such as using the qualification model in place for 3S and 4S Servers, considering variations in overhead from different vendors chassis, allowing Blade Chassis and Blade Servers qualified in independent tests to be considered qualified when combined. EPA will take these into account, but does not feel they impact the initial step of developing a dataset for Blade Servers.

Role of Power Calculator Data

ENERGY STAR will accept a combination of both physical test data and calculator data. EPA strongly encourages manufacturers to submit physical test data, but understands that time and resource constraints will make calculator data more readily available in many situations.

EPA anticipates the following uses of calculator data, in order of value:

1. Correlating output of calculators to physical test data to compare results and understand the variation between laboratory testing and power-predicting models.
2. Evaluating the impact of single-variable changes in components by comparing simulated calculator data with base physical test data.
3. Supplementing physical test data once the relationship between power calculator predictions and physical test data is better understood.

Comparing Sources

Before submitting data from calculators, EPA asks that manufacturers take existing test data from ENERGY STAR qualified units, compare it to calculator-generated data for those same units, and analyze and report to EPA the results for calculator biases. EPA also encourages manufacturers to compare calculator data to physical test data from non-ENERGY STAR models where possible.

Single Variable Changes

To examine the effects of single component permutations (i.e., PSU choice, storage type/capacity/spindle speed, vintage/quantity/density of installed memory, I/O technology), EPA believes that calculator data could provide an effective balance between broad insight and testing burden. EPA strongly encourages physical testing where possible for comparison purposes.

Supplementing Physical Test Data

To provide insight into the effect of the 5-point family definition, EPA intends that data in each of the five points be

collected from physical testing. EPA encourages manufacturers to enhance these tests as desired with calculator data for configurations outside the five physical tests.

Manufacturers should contact EPA to discuss their calculators, the results of their calculator bias analyses, and to identify any issues that may arise from the use of their calculators in data generation. If any of the above guidelines conflict with what you understand to be the aim, scope, or methodology of the ENERGY STAR Servers data assembly process, please contact RJ Meyers at Meyers.Robert@epa.gov or (202) 343-9923 or contact Evan Haines at ehaines@icfi.com or (781) 676-4081 with your concerns.

Development Timeline Update

Based on stakeholder feedback, EPA has updated the development timeline.