



# ENERGY STAR Computer Servers

## Version 2.0 Draft 1 Stakeholder Web Meeting



Learn more at [energystar.gov](https://energystar.gov)



## Welcome and Introductions

Una Song

US Environmental Protection Agency

[song.una@epa.gov](mailto:song.una@epa.gov)



Learn more at [energystar.gov](http://energystar.gov)

# Draft 1 at a glance

---



- Stated goals for Draft 1 :
  - provide an update on the process to develop a specialized efficiency rating tool
  - establish testing criteria and conditions for blade servers
  - describe proposed modifications to existing Version 1.0 criteria that will continue to form the core of the ENERGY STAR Computer Servers specification
  - refine and present new specification definitions and structure

# Agenda



**Welcome and Introductions**

**5 Min**

**Topic Discussion: *Selected Definitions***

**10 Min**

**Topic Discussion: *Idle State Power***

**10 Min**

**Topic Discussion: *Reporting Requirements***

**5 Min**

**Topic Discussion: *Active Mode Efficiency (SERT)***

**45 Min**

**Topic Discussion: *Blade Servers***

**30 Min**

**Topic Discussion: *Product Families***

**10 Min**

**Next Steps and Closing**



## Topic Discussion: *Selected Definitions*

Evan Haines

ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com)



Learn more at [energystar.gov](http://energystar.gov)

# Product types



- Product definitions
  - Development of cohesive definitions will allow for proper analysis in data collection

217 3) Fully Fault Tolerant Server: A computer server that is designed with complete hardware  
218 redundancy, in which every computing component is replicated between two nodes running  
219 identical and concurrent workloads (i.e., if one node fails or needs repair, the second node can run  
220 the workload alone to avoid downtime). A fully fault tolerant server uses two systems to  
221 simultaneously and repetitively run a single workload for continuous availability in a mission critical  
222 application.

- Selection of stakeholder feedback (Fully Fault Tolerant):
  - *“Full hardware redundancy as currently described, is not required to be considered a resilient server.”*
  - *“The definition for fully fault tolerant server should remain as currently written and not add architecture components or attributes. A system which does not have full redundancy built in, but which also has architecture features to make it more reliable or close to a fully fault tolerant system is included under the resilient system category.”*

# Product types



223 4) Resilient Server: A computer server that is designed with resiliency, RAS, and self-correction  
224 features integrated in the micro-architecture of the CPU and chipset to ensure data resiliency and  
225 accuracy. A resilient server is often used for a limited set of workloads that may include business  
226 processing, decision support, or handling of virtualized workloads, and is often operated at higher  
227 levels of utilization compared to a standard server. For purposes of this specification, a resilient  
228 server must meet **all** of the following criteria:

229 i) designed to accommodate hot-swappable components;

230 ii) designed with multiple physical banks of memory and I/O busses;

231 iii) provides machine check architecture;

232 iv) provides memory fault detection and system recovery (e.g., DRAM chip sparing, extended  
233 ECC, mirrored memory);

234 v) provides end-to-end bus retry; and

235 vi) provides support for on-line modification of hardware resources ("on-demand" features).

- Selection of stakeholder feedback (Resilient Server):
  - Feedback was mixed
    - Direct support of existing definition
    - Concern that the definition was overly-broad
  - The HRG *Availability Environment Classification (AEC)* was raised by EPA as a potential availability metric
    - Support: “*Define the Reliability and Availability aspects of a server for it to be resilient instead of defining the underlying hardware implementation.*”
    - Criticism: “*The ... HRG description of Availability Environment Classification (AEC) system does not contain the quantitative detail required to base an ENERGY STAR category on.*”

# Product types

---



- Other stakeholder feedback
  - HPC
    - Concern with defining the category based on number of memory controllers
    - “Worklet” (modular) approach of SERT noted as not applicable to category
  - Multi-node
    - Recent multi-node servers with hot-swappable motherboard/node capability. Blurring into the blade category...
  - Blades – discussed later in the presentation

# Modes of operation



- Selection of stakeholder feedback (operational modes):
  - EPA requested feedback on the value of adding “sleep” or other low power mode definitions
    - Feedback based on stakeholder expectations of server use (e.g., backup equipment engaged only during primary server failure, low expected use of sleep modes in current datacenters)
  - Suggestion to add ACPI states for Idle State (S0) and other modes as necessary

398	4) <u>Server Processor Utilization</u> : The ratio of instantaneous processor computing activity to full-load
399	processor computing activity at a specified voltage and frequency.

- Processor Utilization
  - *“The percentage estimate of the server’s compute activities relative to the full operational voltage and frequency of the processor(s)”*

# Power supplies

---



- Ongoing work within Climate Savers and 80plus to refine power supply definitions for better alignment across end use products (servers, storage, client PCs)
- EPA will include updated proposals from these groups in comment summary document
- Proposed modification of definitions will be evaluated as part of the ENERGY STAR development process and based on impact to established efficiency levels, program/market impact



# Next steps - definitions

---

- All comments will be compiled and addressed
- EPA to update in a subsequent document with edits tracked (Draft 2 or interim definition summary)
- Goal is to solidify definitions after Draft 2 to focus work on other areas of the specification



## Topic Discussion: *Idle State Power*

Evan Haines

ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com)



Learn more at [energystar.gov](http://energystar.gov)

# Idle State



- EPA is committed to Idle State power requirements in V2.0
- Draft 1 listed several options for how to incorporate Idle State metrics or measurement into the efficiency rating tool
  - Dedicate a portion of the rating tool operation to automated idle power measurement
  - Fully incorporate Idle State as a factor in the overall system efficiency rating (*note: does would not meet EPA's requirement for Idle State levels*)
  - Scale Idle State allowances based on server function (e.g., efficiency rating, maximum power draw) rather than existing hardware category approach (e.g., processor socket)

# Approaches



- Selection of stakeholder feedback:
  - General support of scaling structure
    - Concern that availability of the SERT tool would not allow scaling approaches to be studied through data collection in a reasonable timeframe
  - SERT tool viewed as an opportunity to integrate measurement and delivery about both active and Idle State power
    - Weigh idle power measurements and measurement at low utilization higher
    - Set required percentage reduction in power use between full power and idle power as measured by the SERT test and incorporate into the SERT metric
    - Evaluate load curve to evaluate slope in different ranges
  - Suggestion to maintain V1.0 structure and tighten limits by 10%
  - Some concerns with V1.0 Idle State method
    - Existing product coverage gaps (2S, 1P servers)
    - Mixed feedback regarding power allowances



# Next steps – idle state

---

- Assess feasibility of approaches based on integration with SERT tool after all feedback received
  - If development cycle will allow for a release of SERT to be used for data collection purposes:
    - Data collection to commence with availability of tool
    - “Categorization” will be possible
  - If this is not possible:
    - EPA will notify stakeholders and immediately begin data collection
    - EPA will work with stakeholders to identify feasibility of coordinating Idle State testing with future retesting with SERT
- Measurement of idle power as part of the SERT tool is planned to automate testing regardless of the final metric



## Topic Discussion: *Reporting Requirements*

Evan Haines

ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com)

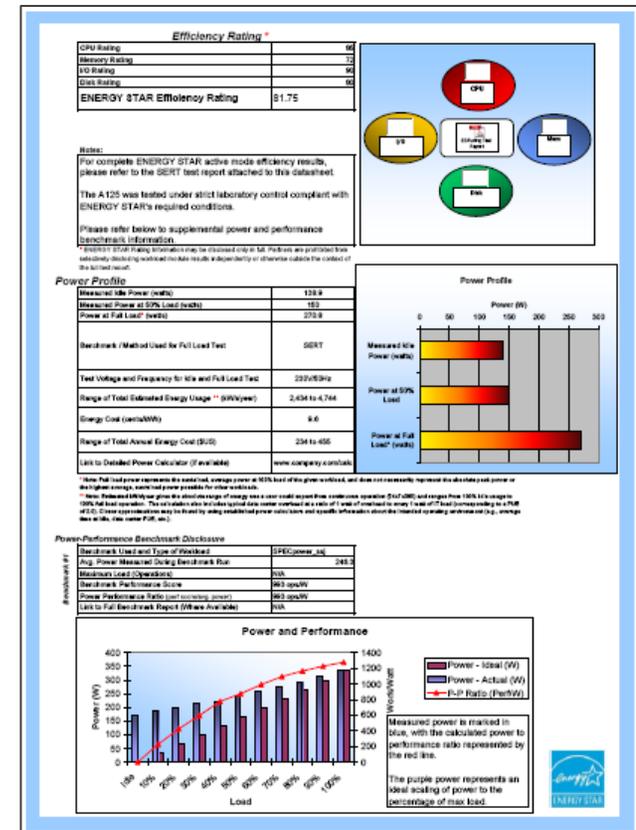


Learn more at [energystar.gov](http://energystar.gov)

# Power and performance datasheet



- Selection of stakeholder feedback:
  - Inlet air temperature vs. fan power chart
    - Stakeholders expressed concern on usefulness
  - Power profile chart
    - One comment regarding “misinterpretation”
  - Power and performance graph (bottom)
    - The particular chart was a sample; in its place would be any graphical output from the vendor-selected supplemental benchmark



# Data measurement and output



- Selection of stakeholder feedback:
  - Pedestal servers
    - Feedback ranging from support of consistent data measurement and reporting requirements with other server types to requests for exemption
    - Concern regarding the ErP Lot 6 requirements in the EU (some pedestal servers can be classified as Class B and fall under these sleep/standby restrictions)
  - Measurement frequency and accuracy
    - Requests for flexibility in averaging method
    - Proposals for alternative sampling frequency

# Supplemental benchmark reporting



738 will open a dialogue with stakeholders to discuss benchmarks for consideration. Following are candidate  
739 benchmarks that have been identified to date: SPECpower\_ssj, SPECweb\_power2009, SPEC's planned  
740 virtualization benchmark, TPC-C, TPC-E, and RPE2. EPA recommends that any candidate benchmark be  
741 evaluated based on (1) presence of an integrated power measurement methodology, (2) applicability to,  
742 and widespread use by the end-user community, and (3) applicability to all servers in the ENERGY STAR  
743 scope, or at least a full subset of servers intended for a specific end-use application.

- Draft 1 presented a candidate list of supplemental benchmarks for reporting on the PPDS
  - Vendors will still be allowed to tailor choices to intended customers
  - The supplemental benchmarks provide more workload/application-specific results than intended by the SERT tool – complimentary in purpose

# Next steps – reporting requirements

---



- Draft 2 will refine data measurement and output requirements
  - EPA encourages further stakeholder review of multiple proposals received upon availability of comment summary
- List of fields and data elements in power and performance datasheet will be updated
  - An updated template will be provided when updates to the SERT section can be made – Draft 2 or 3



## Topic Discussion: *Active Mode Efficiency (SERT)*

Evan Haines

ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com)



Learn more at [energystar.gov](http://energystar.gov)

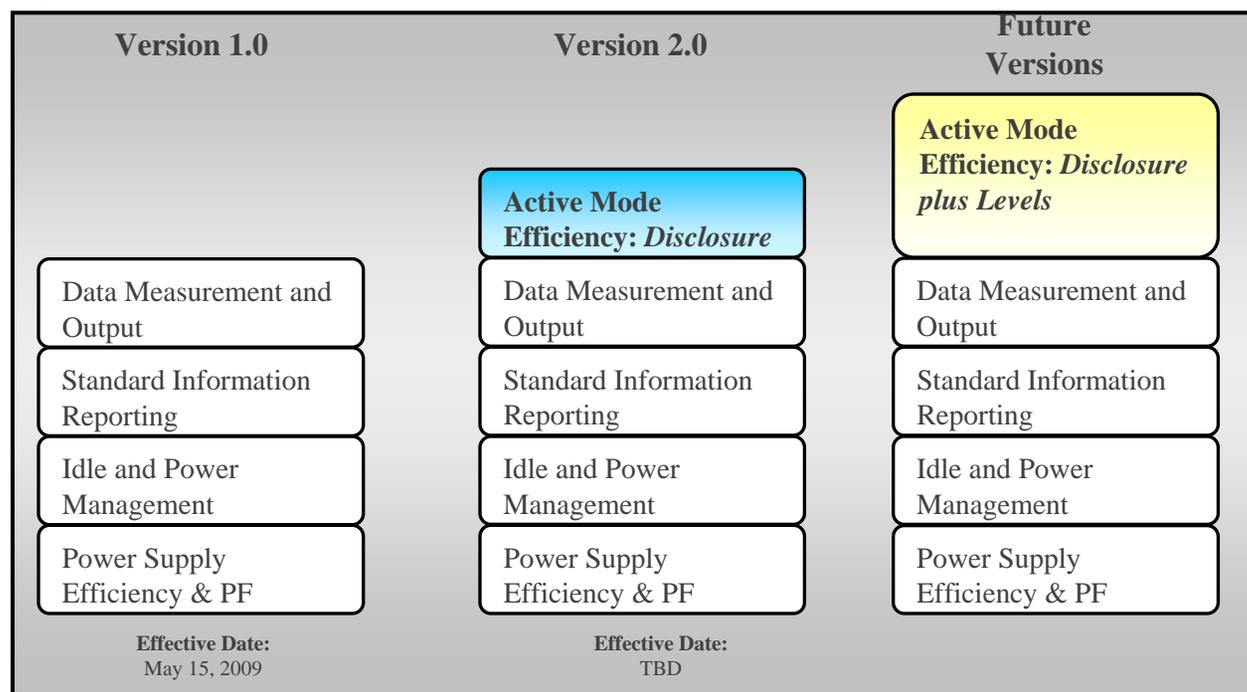
# Active mode efficiency



- Goals
  - Build on the efficiency foundation from existing “rollover” criteria to provide insight into active mode efficiency
  - Collect sufficient data to set future active mode efficiency levels
  - Preserve context to active mode efficiency results
  - Institutionalize efficiency in the purchase decision
- Disclosure model
  - Draft 1 established EPA’s intent to require disclosure of active mode efficiency for all servers in the program
    - Develop a database of active mode efficiency information
    - Other requirements and level structure remains (Idle State, Power Management, Power Supplies, etc.)
  - General support from a variety of stakeholders on this approach for Version 2.0

# Active mode efficiency

- Long-term vision



- Today's presentations: review of early feedback and update from SPEC on SERT development

# Active mode efficiency – selection of stakeholder feedback

---



- Product types
  - Support from end-user community for consistent application of SERT across product categories (blade, rack-mounted, pedestal)
- Versions of the SERT tool
  - Support for EPA's principles of broad architecture and operating system support

# Active mode efficiency - selection of stakeholder feedback

---



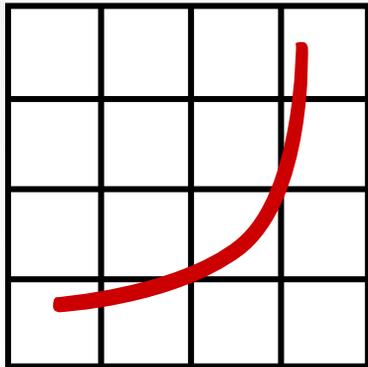
- Structure (*from SERT Design Document*)
  - Support expressed for “Worklets”
- Clarifications
  - Requests for detail about the computation of the SERT metric
  - Suggestion that a list of operating systems and architecture be provided

# Server Efficiency Rating Tool™

May 2010 – Development Update

Klaus-Dieter Lange

Chair, SPECpower Committee, SPEC



spec

# Next steps – active mode efficiency

---



- During development, EPA to provide periodic updates to stakeholders on progress
- Development to continue in parallel with further draft development



## Topic Discussion: *Blade Servers*

Evan Haines

ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com)



Learn more at [energystar.gov](http://energystar.gov)

# Overview



- Draft 1 proposals
  - Consistent metrics applied as for other categories (tailored to the unique considerations of the blade form factor)
  - Combination of two tests:
    - Single blade in chassis – single blade idle/full power
    - Partially-populated, homogeneous blades – active mode efficiency, aggregate idle/full power
    - Derived: chassis power contribution in idle/full
  - Blade chassis approved for use with ENERGY STAR blade servers based on power allowances and feature requirements
    - No direct qualification proposed for the chassis

# Blade server - selection of stakeholder feedback

---



- Test method
  - A number of comments suggesting that blades be compared to other blades rather than across categories (comments in opposition also received)
- Chassis requirements
  - Strong push for features-based chassis requirements rather than power allowances
    - Designs described and difficult to compare from vendor to vendor
  - Suggestions that EPA consider allowing blade chassis to be a qualified product within the program to reduce ambiguity for end-users
- Blade server definition
  - Hot-swappable multi-node servers – possible overlap
  - Processor/memory modules for scale up in standalone servers

# Blade server - selection of stakeholder feedback

---



- Test method alternatives/modifications
  - Augment testing of individual blades with a percentage of shared resource power from a fully or partially loaded chassis test
  - “Power domain”
    - Make use of a complete blade system of minimally configured blades and a full chassis to operate the SERT tool default; allow partially-populated testing based on the number of power supplies installed
    - Power domain = the maximum number of slots supported by a single power supply or a pair of redundant power supplies
    - Derive single blade power by removing one blade and re-measuring

# Next steps – blade servers

---



- Presentation of alternative testing approaches in comment summary
  - Work with stakeholders to formalize testing structure
- Investigation of blade chassis feature sets and evaluation approaches
  - Consideration of direct qualification for chassis
- Intent is to collect data for blades in parallel with other server types



## Topic Discussion: *Product Families*

Evan Haines

ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com)



Learn more at [energystar.gov](http://energystar.gov)

# Overview



- ENERGY STAR families
  - The ENERGY STAR structure is intended strictly to ensure that all represented configurations meet ENERGY STAR requirements
- Balance manufacturer testing/reporting burden with data applicability
  - Will the most energy efficiency configuration options be identified and documented?
  - Will end consumer be able to KNOW their configuration is a valid ENERGY STAR server?
  - Will the Channels be able to provide configuration services while maintaining ENERGY STAR compliance?
- P&P datasheet – *how close to purchased configuration?*
- Submittal data – *do representative tests validly apply to the grouped configurations?*

# Product families - selection of stakeholder feedback

---



- Power calculators
  - Stakeholder suggestion received that EPA investigate role for vendors' power model calculators in the ENERGY STAR process
- VARs
  - Limited comments received
- Specific proposals received on Table 1

# Industry proposals



415

Table 1: Product Family Component Requirements

Base Component	Same Part Number Required in All Product Family Configurations	Same Technical & Power Specs Required in All Product Family Configurations	Quantity Required in All Product Family Configurations	Notes
Motherboard	YES	YES	Same across family	
Processor	YES*	YES*	Same across family	* Processors must all be from the same model line. * Processors must have the same core count and power specifications. * Processor speed may vary within a product family.
Power Supply	YES	YES	May vary within the product family	
IO Device	NO	YES	May vary within the product family	
HDD or SSD	NO	NO*	May vary within the product family	* HDD, SSD, and Memory capacity may vary. If so, minimum, typical, and maximum configurations must represent the full range of capacity options.
Memory (DIMM)	NO	NO*	May vary within the product family	

Request for use of worst case components

Support for flexibility added for storage and memory

Request for processor p/n and specification variations

- End result MUST effectively identify those configurations options which provide the best energy efficiency

# Next steps – product families



- Development of product family structure ultimately a data-driven process
- Draft 2 to formalize structure, but subject to change based on trends in collected data
- Near-term goals
  - Formalize evaluation and documentation approach with stakeholders
  - Investigate further use of power calculators
    - Ability to judge quality across vendors is critical for EPA to move forward with investigation



## Next Steps and Closing

Una Song

US Environmental Protection Agency

[song.una@epa.gov](mailto:song.una@epa.gov)



Learn more at [energystar.gov](http://energystar.gov)



# Next steps

---

- Comment summary and proposed actions posted on ENERGY STAR Web site
- Development of Draft 2
  - Changes from Draft 1 will be indicated
- EPA plans to work with SPEC to provide periodic updates via the email list on development

# Schedule Goals



June-July	<ul style="list-style-type: none"><li>• Draft 2 development</li><li>• Metric for Idle State determined</li></ul>
Q3 2010	<ul style="list-style-type: none"><li>• Data collection and analysis – Idle State (1<sup>st</sup> round)</li><li>• Further draft development (rollover requirements)</li></ul>
Q4 2010	<ul style="list-style-type: none"><li>• Data collection and analysis – Idle State (2<sup>nd</sup> round)</li></ul>
Early Q1 2011	<ul style="list-style-type: none"><li>• V2.0 finalized</li></ul>
...	...
Q3-Q4 2011	V2.0 specification effective

# Thank You

---



Thank you for your participation and support of the ENERGY STAR program.

**Please address questions and comments to:**

[servers@energystar.gov](mailto:servers@energystar.gov)

Una Song, US EPA

[song.una@epa.gov](mailto:song.una@epa.gov) • 202.343.9024

Evan Haines, ICF International

[ehaines@icfi.com](mailto:ehaines@icfi.com) • 202.862.1158

Al Thomason, TBWC, LLC

[thomasonw@gmail.com](mailto:thomasonw@gmail.com) • 503.7087881

Materials will be posted to the ENERGY STAR Computer Server Revision page:

[www.energystar.gov/RevisedSpecs](http://www.energystar.gov/RevisedSpecs)