EPA Stakeholder Meeting to Discuss Proposed ENERGY STAR®
Server Specification Requirements
October 31, 2007

This document summarizes discussions from the ENERGY STAR Server Working Group held in Santa Fe, New Mexico on October 31, 2007. Additional details can be found in the ENERGY STAR Server Specification Discussion Guide and the Power Point presentation for the meeting. All of these documents are available on the ENERGY STAR Web site: http://www.energystar.gov/index.cfm?c=new_specs.enterprise_servers.

Definitions

Definitions are needed to set the scope of the specification (i.e., distinguish between "server" class and "desktop: class computers). EPA presented attendees with a list of potential characteristics that might be used to define a server:

- Redundancy (including "hot-swappable" components)
- Processor scaling (# of processors/cores)
- Memory capacity and type (e.g., # of Dimms, ECC memory)
- Capable of utilizing external storage (NAS/SAN)
- Dedicated management controllers
- Other components (internal disks, network adapters, option slots)
- Rack Space (e.g., 1u, 2u, 3u, etc.)
- "Server class" operating systems

Question: What are some of the key components and/or characteristics that define a server?

- EPA should review current IEEE and EU Rohs standards, which may provide a server definition/classification.
- It may be possible to create a definition based on shared hardware and software characteristics but EPA would have to determine the cutoff points (e.g., two processor sockets, four DIMMs, etc.).
- Basing the definition on physical components such as sockets could be tricky because this doesn't take into consideration the intended purpose of the server. A definition should consider intended use/workload as well as hardware/software elements.

Product Coverage

EPA intends to cover as all server types and storage equipment, if possible. Some stakeholder comments suggest excluding larger server equipment (i.e., mid-range and high end models) and suggest focusing on volume servers. The bulk of savings will come from these volume and also blade servers so EPA is leaning toward initially focusing on these products. Will it be necessary to further delineate classes of servers for Tier 1?

- It's important that blade servers are addressed by the specification but they would most likely need a separate definition. Each manufacturer has a different idea of what characterizes a blade server.

Power Supply Efficiency

Proposed criteria could be based on (1) efficiency at 10%, 20%, 50%, and 100% loads; (2) no-load power consumption; and (3) minimum power factor (at all load points). Another idea could be to address net power consumption, which would be AC input – DC output during server operation in standardized modes, such as idle and max consumption.

Testing conducted by EPRI show that for multi-voltage power supplies, there is modest variation in efficiency across products tested. Single voltage power supplies show a wider differentiation in efficiency at the various loads but in general reach higher efficiencies. Power factor seems to follow a similar load curve to energy efficiency (i.e., dropping significantly at low loads). Efficiency seems to drop rapidly at 20% load and this (or a condition lower than this) is where many power supplies currently operate during times of low server usage, especially in redundant configurations. This is why the 10% was included in the test procedure.

Questions:

Multi-voltage vs. single-voltage server power supplies?
How should the effects of cooling fans be accounted for?
Is testing at only 208 VAC or 230 VAC acceptable, and does it properly characterize the efficiency of the PS in operation?
How could a specification account for products with shared power supplies (e.g., blade systems or dc power at the rack level)?

- What about those power supplies that can shut down phases to avoid running at lower loads, will they be rewarded? These phase shedding power supplies can intelligently change their characteristics to "act" like a smaller supply. If these product types can avoid running at 10% will they be held to that load point in the specification?
- Standards such as PMBus may help operations run more efficiently but may also come at a small expense to efficiency. How would a specification deal with these products?
- What about DC powered equipment that do not have an internal power supply? How will those be treated in the specification?
- Suggestion: Perhaps those product types that cannot easily be covered can be tabled for Tier 2. However, how will this impact federal procurement?
- Many chassis designs exist that swap out AC to DC and directly compete in the marketplace, so DC powered configurations would be at a disadvantage if not allowed to qualify for ENERGY STAR.
- Measuring the motherboard efficiency and the various losses associated with dc-dc conversion and power distribution at low voltage is very difficult. The Climate Savers Initiative has created an efficiency workgroup that is working on this very issue.
- Since servers are a part of a larger system, at both a rack level and data center level, the question becomes "where is the power conversion occurring"? Ac-ac can happen at the data center, rack, or server level. Similarly, dc-dc conversion can happen at the server level (power supply) or the board level. How can a specification based on power supply efficiency fairly compare products with a multi-voltage supply (ac-ac and dc-dc occurring in the supply), single-voltage supply (ac-dc occurring at the data center or rack level, dc-dc occurring in the supply or the board)?

- We need to be able to represent newer power supply technologies in the test procedure (e.g., digital power supplies).
- It's important to point out that when EPA talks about testing at various load points we are talking about the load on the power supply not the server.
- The power supply load curve is not quite lined up with the actual load on the server.
- Manufacturers are already motivated to increase efficiencies on the power supply for other reasons. EPA should track how efficiencies have increased over time and we may find that power supply efficiency doesn't need to be addressed because that movement is already happening.
- Measuring efficiencies at lower loads addresses the fact that power supplies typically run at these loads because of their redundant configuration.
- The efficiency at 100% load is not as important since power supplies rarely perform there but it is important to test to ensure that the unit does not fail at that load point.
- With 2N redundancy a power supply will never run above 50% load and usually run as low as 10-20% load while in an N+1 situation you will find a power supply able to take on the full load while the other runs at no load. Therefore, it makes sense to have conditions at low and high loads to cover both scenarios.
- In regards to power factor, industry-wide typical PF is 91%. This is a real improvement from 70% a few years ago. Is there enough differentiation in the marketplace today to set a specification for PF?
- When EPRI conducted testing, a minimum load was put on the power supply to test for no load efficiency so it stayed awake. The question is should EPA address no load energy use in the specification?
- How does power supply efficiency vary based on the window of voltage inputs on site? When developing the test procedure EPRI had to choose a reference.

<u>Idle Energy</u>

Measured data shows that servers spend a majority of their time idling. Idle mode involves zero latency and is relatively easy to measure.

Question: Are manufacturers already addressing / measuring idle?

- End users look at application performance versus power when considering new equipment.
- Industry currently playing around with the idea of what to do during server downtime (e.g., sleep, idle, etc.)
- We need to identify what the server is intended for regarding idle energy (i.e., amount of time spent in this mode). Some servers are intended to be active and fully utilized most of the time.
- Virtualization and other technologies that seek to reduce time in idle area already being worked on by manufacturers.
- EPA is potentially rewarding the wrong element of design if it only looks at idle and ignores other opportunities for energy efficiency such as energy reduction under other load conditions.
- EPA's focus should be to ensure that ENERGY STAR servers use less energy while performing its intended use.

- Perhaps idle could simply be a reporting requirement, with no levels of criteria. Could these measurements come from the SPECpower output?

Standardized Labeling

Standardized labeling would provide for consistency in the marketplace and allow customers to compare performance, capability, and other characteristics to inform their purchasing decision. EPA recognizes that right sizing and proper static provisioning are key to achieving energy savings. Some elements that might be included: system characteristics; SPECpower benchmark results; and maximum energy use.

Questions:

What additional information could be provided?
How should information be presented to purchasers – product labels or other means of reporting?

- Labels should be aimed at IT professionals (i.e., not a consumer label).
- Since servers are not purchased in retail spaces a "label" would most likely not be appropriate. Standardized information would be more helpful as online fact sheets (e.g., xml tools).
- Server configurations are highly variable. Very few customers buy a standard configuration so reporting based on this standard configuration will be tricky.
- Customers are savvy enough to make choices and compare servers based on standard configuration – it's only a tool in the decision making process.
- The label should include a number of characteristics such as the benchmark used and score, energy use, configuration, etc. so the customer gets the full picture. It would also be preferable to include several benchmarks for multiple workloads so that IT professionals can compare to what is most appropriate for their needs.

Power Management

There are significant savings that could be had through power management with minimal effect on performance when properly implemented. Many servers aren't operating at full utilization – is there an opportunity to reward machines that scale back power during times if low utilization? What might criteria be based on? EPA could provide a list of features/techniques and require a small number to be installed (e.g., power scaling, power capping, sleep mode).

Questions:

What technologies are available for power management in servers?
What might the criteria look like?

- Servers typically come with power management capabilities but it is up to the end user to enable it on site.
- Power management is based on end user implementation to ensure savings and as such this cannot be the main criteria of the specification.

- By including a requirement, EPA is trying to ensure the means to simplify implementation of data center wide energy saving schemes (i.e., data center operators could be confident that an ENERGY STAR server could plug into their power management schemes.)
- Power management varies in applicability across different server types and workloads (i.e., there isn't one power management solution for all servers).
- Power management need 3 elements to work: (1) when does the policy get imposed; (2) what are the savings; and (3) what is the latency?
- Listing power management requirements might not be needed because Specpower reporting may address it. Power management features could be a reporting requirement under ENERGY STAR.

Access to Operational Data

Requiring access to power and temperature data would allow monitoring of equipment status in real time and be able to manage those products. Outputs could include inlet/outlet temperature, airflow rate, power draw, and processor utilization. Criteria could be based on a standardized Ethernet protocol to collect the information.

Questions:

What data should be measured?
Do a sufficient number of current products on the market have this ability?
Could collaboration with the Distributed Management Task Force (DMTF) or use of the Simple Network Management Protocol (SNMP) produce such a standard?
Could reported data include processor utilization?

- At every point in the data center this is already happening. Nearly all servers now have temperature sensors for inlet temperature.
- Outlet temperature is the least important piece of information unless you also have flow rate information.
- Power draw information might require an advanced power supply specification such as the PMBus standard, which would add extra cost to the power supply.
- What is needed is a uniform protocol for reporting the data. Currently this data is being collected using many different protocols. Green Grid is currently working on this issue.
- EPA should ask end users whether this would be useful information to have.
- End users are moving away from SNMP because of security issues.

Virtualization

A number of stakeholders suggested to EPA that virtualization should be a part of the specification because of the significant savings that could be realized by reducing the number of servers needed and thus increasing server utilization. This is a low cost solution that promises significant savings in the short term. However, it is not the long term solution to addressing the energy used by servers.

Question:

How can EPA specify support for virtualization: support for number of virtual images? Utilization potential of the virtualized server? Others???

- Virtualization is primarily a software function so in essence all servers are capable of being virtualized, though some servers may have special software of hardware installed to make this easier.
- EPA must ensure, at a minimum, that the specification does not penalize virtualization.
- The industry is already moving in this direction so perhaps EPA does not need to include it in the ENERGY STAR specification to encourage adoption.

Potential Tier 2 Requirements

Server Performance Benchmark: For Tier 1, EPA intends to require the submittal of SPECpower output information for server qualification. Reducing

- A specification based on an energy-efficiency benchmark may take too long to develop.
- The SPECpower benchmark is a good starting point but EPA will have to include additional workloads if the specification is to be applicable to all servers.

Energy-Efficiency Ethernet: Reducing network link rates saves power and allows other connected devices to also reduce power (e.g. switches and routers). IEEE 802.3az committee has begun work and anticipates products which meet the standard in the Tier 2 time frame.