

From: Peter Hopton [Peter.Hopton@very-pc.co.uk]

Sent: Friday, September 19, 2008 6:07 AM

To: Duff, Rebecca M.; Howard, Arthur

Cc: fanara.andrew@epa.gov

Subject: Energy Star Servers Draft 2 Comments

Hello Arthur and Rebecca,

Here are our comments on Energy Star Enterprise servers draft 2:

1.
 - a. The definition for computer server is acceptable
 - b. Blade servers in some cases included in a fixed chassis (a chassis containing one or more servers) Cooling is not always shared, but power always is.
 - c. In Blade Chassis, the blade can be fixed. Networking hardware, system management or cooling are not always shared, however power always is.
 1. Suggest that Blade Servers and Blade chassis are combined to a category of Blade systems, defined as multiple servers within a chassis with shared power distribution.
 2. Definitions should be as broad as possible to allow for innovation but narrow enough to avoid abuse
 - d. No comments
 - e. No comments
 - f. No comments
 - g. Is this definition impairing iscsi or NAS storage solutions?
 - h. No comments
 - i. No comments
 - j. No comments
 - k. No comments
 - l. No comments
 - m. No comments
 - n. No comments
2. The list of qualifying products looks good, introduction of 8 CPU socket servers should definitely be considered later
3.
 - a.
 1. Surely the efficiency of the PSU is irrelevant? the overall consumption of the unit is what matters.
 2. For efficiency measurements the name plate value of the PSU is irrelevant, the efficiencies should be representative of the actual consumption of the unit, not a name plate value!
 3. Is excluding fan power from this measurement going to promote efficient cooling of supplies?
 4. Measuring power factor at % of name plate value is a poor idea, it should be measured at the servers idle power value, and at the servers 100% load power value.

5. Suggest that 1p, 2p and 4p servers are split out into different power requirements in all categories to represent the fact that they are literally twice as powerful as each other in computing terms!
6. Having 2P and 4P servers in the same category for standard redundancy equipment is going to make it very easy for submission of 2P servers, as 4P servers have twice as much equipment to support at idle load.
7. Having 2P and 1P servers in the same category for high redundancy equipment is going to make it very easy for submission of 1P servers, as 2P servers have twice as much equipment to support at idle load.
8. Splitting out redundant systems out is important, but vendors must measure equipment with both supplies connected, as often vendors cheat by only measuring power numbers with one supply connected
9. Why is SPECpower used for measuring idle? Surely the existing methodology for energy star desktops is suitable?!
10. The use of SPECpower to measure maximum utilisation of servers is flawed, as it only runs the CPU. Using a standard 'burn in' type program should be explored, one which will stress multiple components. There are several such free utilities on the internet, however it would be prudent to select a free open source utility as this would not allow any vendor to obtain an unfair advantage due to exclusive access to source code.
11. Use of SPECpower sets costs and high barriers to entry, use of a compatible 'free' methodology should be considered. If SPECpower is used, a version of the software should be compelled from SPEC to be Free and Open Source!

12. Idle as a percent of maximum power

1. This would discriminate against the use of low voltage/HE CPUs, which would idle at the similar as a not low voltage chip, but load up at a fraction of the power
2. This would actually encourage Big IT to make less efficient kit!

13. Power Saving Checklist

1. I agree with EPA's judgement on this, power management in servers is not as important as absolute watts.

14. Blade Systems

1. Blades idle should be calculated as an average power consumption (mean) of one blade of the whole (fully populated) unit.
2. A requirement on power consumption of underpopulated blade centers (or blades turned off) should be considered – as a % inefficiency of a fully populated unit

b. No comments

- c. These requirements should be optional, not required. The requirement to have energy star servers be able to measure power and air input temperature internally is going to:
 1. Discriminate against smaller vendors who would find it difficult to properly integrate on vendors IMPI with another vendors PSU.
 2. Favour certain vendors with patent protection.
 3. Discriminate against suppliers of cabinet based monitoring systems that also provide a large amount of additional benefit for DC efficiencies
 4. Encourage e-waste by having equipment installed in servers that is also replicated in cabinets/power strips

5. Cause non-standardisation of power management functions across datacentres where these are currently standardised and supported by 3rd parties on a per DC basis
6. Cause variation between equipment of different vendors in measurements and accuracies.
4. Why does SPECpower_ssj2008 (a commercially licensed paid for software where only select vendors have access to source code) need to be used to measure server idling? Surely the Energy Star Desktop 4.0 protocol is sufficient?!
 - a. Load power should also be considered in order to help promote low voltage CPU's that will have a positive impact on power consumption, but not on idling.
5. No Comments
6. No comments

Best Regards

Peter Hopton

Managing Director

VeryPC – PC Pro Environmental Innovator of the Year.

M: 0772 5471 494

T: 0845 6170 081