

From: Aaron Hay

Sent: Friday, February 15, 2008 11:21 AM

Subject: Feedback for the ENERGY STAR Draft 1 Server Specification

Hello Andrew & Rebecca,

Thank you for sending me the draft specification. I commend you for considering a wide variety of key issues in your analysis and I applaud you for being ambitious with what may be included in the final ENERGY STAR server specification.

In particular, I was impressed with your discussion of lower power consumption at idle (synonymous with sleep, or a lesser step-down in consumption) and the related issue of overall server utilization rates.

I would like to offer some feedback in this area. Server power management technology is progressing so rapidly that I feel a sleep-synonymous qualification should be included in, at least, a 'Tier 2' specification for servers. I have a few reasons for this:

1. Although few vendors offer technology that, at this point, allows servers to return to processing with minimal latency, I feel that building this requirement into the higher-ambition levels of ENERGY STAR server specs (e.g. Tier 2) will give vendors an additional incentive to continue research and development in this area. Tier 2 *should* be aggressive – we're at a tipping point where vendors, driven by the market and energy supply realities, have to greatly increase the energy efficiency of the data center ecosystem, and this would be further impetus to do so.
2. Inclusion of sleep-synonymous energy efficiency provides greater flexibility for IT. Some large data center operators are faced with stringent availability/reliability requirements – to the point where they in fact want to virtualize/consolidate – but mission-critical computing requirements prevent (in whole or in part) these actions (even if such actions only create minimal additional risk – something that is not well understood, and is another issue entirely).

If sleep-synonymous energy efficiency is paired with minimal latency and high reliability in returning to computing, these operators can choose equipment which is stand-alone (ie. not virtualized) but is only 'turned on' when *computing* utilization is high enough. Essentially, they have an energy-efficient '2nd best' alternative to pure virtualization – turning off the resources when they aren't in use will increase *computing* utilization across the board (even if actual *equipment* utilization is forced by IT policy/requirements to remain low). ENERGY STAR, I feel, has a responsibility to encourage not only virtualization, but alternatives for special situations as well.

3. ENERGY STAR will be better prepared to offer even more aggressive future specifications in servers – and, perhaps, whole data center operations. Starting to address the utilization issue now is ambitious – but will prime vendors and purchasers for further action and understanding in not only individual equipment efficiencies, but also holistic data center efficiencies.

I was also highly impressed with your inclusion of requirements for power and temperature measurement capabilities. As you well know, this is a 'must-have' for data center operators looking to optimize energy efficiency and start using the management metrics proposed by Green Grid, Uptime Institute, et al. While server measurement capabilities won't account for the draw from the rest of the data center (primarily HVAC/cooling, the power being used before passing

through conversion at the server level, etc), standardizing reporting at the server level is a good start. And, in my opinion, badly needed.

Thanks again for passing this along, and for being so open in your feedback process. Andrew, by the way – I've listened to you at ITAC in Ottawa and Green Grid in San Francisco – I haven't chatted with you yet but I will make an attempt to connect when our paths cross again.

Respectfully, Aaron

Aaron Hay

Research Consultant | Info-Tech Research Group
43 Front Street East
Toronto, ON M5E 1B3

[Info-Tech Research Group](#)

Practical Research that Drives Measurable Results.