

**From:** Liam Newcombe [mailto:[liam@romonet.com](mailto:liam@romonet.com)]  
**Sent:** Wednesday, December 03, 2008 4:35 PM  
**To:** Duff, Rebecca M.  
**Cc:** fanara.andrew@epa.gov; Paolo.Bertoldi@ec.europa.eu  
**Subject:** Re: Draft 3 ENERGY STAR Server Specification

Rebecca, Andrew,

Please find below my comments on the draft specification.

P1 L46 The BCS supports the approach of permanent or temporary labels as discussed. As we previously commented it is not feasible to label the packaging due to the number of component permutations available.

P2 L69, I am sure a vendor has asked this but, could you clarify what is meant by unit data per model? Does this mean chassis model, e.g. Manufacturer Model X or by component permutation e.g. Manufacturer Model X 2x4GB DIMM, 2x73GB SAS HDD, 1xFC Card ?

Page 5 L228, Whilst it is worth collecting data to distinguish between servers with different levels of internal hardware redundancy no allowance should be made for systems with higher levels of redundancy. Resilience is achieved at many levels in the system, the IT equipment is only one of these levels. If the purpose is to group so that once a resilience level decision has been made by a purchaser then this distinction is worthwhile.

Page 9 L414, We entirely support keeping the targets for low load efficiency. High fixed overheads are a major cause of inefficiency in the data centre system and until equipment purchasers understand the energy ROI only this form of approach is likely to correct the issue.

Page 12 L553, Yes, much improved forcing this reporting in a standard form is critical, too many vendors are varying the interface (many from a common chipset), this makes the data centre manager's job harder and software more expensive. I am probably repeating previous comments but we would prefer to see consumed energy than time based averaging of instantaneous readings in the same way that network equipment gives cumulative packet counts instead of instantaneous rates.

Page 15, L679, multi voltage, thanks, this is much more helpful for those of us in the EU.

Page 15, L725, I am sure you already have this from the vendors but, there may well be several 'intended' operating systems, most x86 servers are now supported and validated for at least 3 operating systems. Of course, this makes the idle measurements harder but options would include; specifying the operating system preference order by volume (Windows, Linux, Solaris) and averaging across all approved and supported OS for that machine.

**Power Data:**

Whilst this (draft 3) specification discusses the actual power draw and idle power draw of the server, there is no mention in this specification of the as configured power of the server. The closest approximation to this is the 100% load SPECPower benchmark power in the data sheet. Unfortunately SPECPower only stresses the processor in the system, not the memory, I/O or disk subsystems and as such is not able to give any realistic indication of the maximum power the server could draw in the specified configuration. In the EU Code of Conduct one of the minimum expected practices is to provision power and cooling only to the 'as configured' power of the IT device and not the PSU rating as this is one of the most significant causes of underutilisation and inefficiency in data centres. It would be good to see this data, which the major vendors are able to and most do provide by summing the maximum power draw of the installed components, in this release of the Energy Star specification.

I hope this feedback is of use to you.

Liam Newcombe  
Secretary, [Data Centre Specialist Group](#)  
0774 073 7457