

ENERGY STAR Servers v. 2.0 Specification Draft 1 Comments from the European Commission

We provide in the following comments from the European Commission to draft 1 of the Eligibility Criteria of the specification, of the Power and Performance Data Sheet and of the Server Efficiency Rating Tool – Draft Design Document.

We generally appreciate and support the contents of these documents.

The Specification

Line 174 – Volume Servers: We think that adding an additional definition for volume servers will not be very helpful. It would not make much sense to develop a different definition than the one already used by IDC. This may rather cause confusion. Adopting the IDC definition – which however is mainly based on server prices – will not be very relevant for this server specification.

Line 188 – Managed Server: Including the additional criteria – i.e. remote power control, remote reset, and cold start capability – into the definitions for managed servers makes sense from our point of view.

Line 205 and 330 – Definition of Blade Storage and Network Equipment: We believe that it is not necessary to spend very much effort on definitions of hardware types which are outside the scope of these requirements (e.g. blade storage and network equipment). In contrast a clearer definition of some of the relevant standard hardware in the scope of the regulation seems appropriate (see also below).

Line 236–242 – Definition of Fully Fault Tolerant and Resilient Servers: The approach to base terminology for different server types such as fully fault tolerant and resilient servers on already existing standard terminology (e.g. HRG) as far as possible is supported and would add to clarity.

Line 255 – Definition of High Performance Computing Systems: The definition of High Performance computing systems has been improved which is appreciated.

Line 316–327– Definition of Motherboards, Memory, Drives, etc.: The newly introduced definitions for motherboards, memory and hard drives are partly not satisfying. E.g. the motherboard per se does not include the processor or memory. It is proposed to use standard definitions for these components which are used elsewhere in professional literature.

Line 351–365 – Definition of Power States: The power states included should comprise the main power states for which considerable energy is consumed. Therefore we appreciate the intent to include sleep mode and perhaps even an off mode to account for hypervisor–commanded shut down.

Line 356 – Definition of Active State: Inclusion of a definition for the active state is appreciated. However, this definition will have to be revised in parallel to the development of the energy efficiency rating tool (SERT). The current definition does not seem optimised. However, it apparently only makes sense to further develop this definition in co-operation with SPEC and the SERT developments.

Line 367–381 – Controller System and Client Definitions: The approach to use controller system and client definitions from SPEC is supported since these relevant components for the rating tool will have to fit to the overall SERT concept. It should however be considered if the definitions, which are only used for the test procedure, would be better placed in the test procedure annex in order not to confuse the readers of the specification.

Line 395 – Server Processor Utilization: It is understood that the SPEC rating tool will include an assessment of CPU performance and energy efficiency. In this context the levels of utilization also have to be defined. This parameter probably needs to be defined in parallel with the SPEC developments for SERT.

Line 406 – System Configuration: The proposed approach for product family definitions is generally supported, which basically allows disk capacity and speed, memory capacity and processor speed to vary within a product family.

However, the current definition for the different configurations (maximum, minimum configuration etc.) is very unclear. What does “active mode efficiency” mean in this context, what is the definition? Should the headers be “maximum/minimum efficiency configuration”?

The term “maximum configuration” normally would imply maximum hardware configuration: e.g. maximum disk capacity and speed, maximum memory and CPU with highest clock rate etc. If the interpretation rather refers to computing power (which is the more standard approach) “efficiency” should be replaced by “performance” in this context. In any case the concept is unclear at the moment and should be more clearly defined.

We will recommend to use the hardware based configuration terminology.

Line 468 – Multi-node Servers: It appears that multi-node servers are currently gaining more market relevance as a cost-saving technology. We see no reason for excluding multi-node servers from the scope, since testing will not be more complex than for blade servers and therefore support the intent of including these servers in the specification.

A direct comparison of multi-node servers with rack servers or blade servers would require a more complex consideration. However, this is not necessarily to be required in this version of the specification and could be developed in the longer term.

Line 493 – PSU Efficiency Criteria: We also share the view that correct sizing of power supplies should be supported by specific measures. While the net power loss approach seems too complex to achieve in a short term, other measures could be defined supporting the use of well sized power supplies.

For some types of products (e.g. some multi-node servers), already today several power supply models of different size are provided by manufacturers to allow an optimised energy efficient solution for different server configurations. We recommend to consider a possible option for requiring different right sized power supplies for different typical product configurations. This might be an information requirement in the Power and Performance Datasheet.

Furthermore, we would encourage hardware manufacturers to support the selection of right sized power supplies in server configurators.

Line 537 – Blade System Criteria: It is expected that the development of the active power rating tool for blades will still require significant time. If not possible to include the rating tool in this version of the server specification, it may be a first step to implement idle testing in combination with the blade chassis requirements.

Regarding the overall configuration of the chassis, we understand that testing of fully populated chassis is regarded as to demanding by manufacturers. While higher configurations would be more relevant compared to typical configurations, a 50 % configuration testing seems acceptable though. We suggest though to perform a pre-evaluation of 50 % and higher configuration testing (75 %–100 %) for some samples as part of the requirement development to gain some further insight regarding this aspect.

Although direct comparability of blades with rack and pedestal servers would be nice to have, we believe this should not be a primary goal, because it could delay the development of rating tool.

Line 567 – Blade Chassis Power Consumption: We support setting requirements to blade chassis in idle and full load states in addition to active or idle state requirements for the full blade system. The definitions, test procedures etc. still need to be further clarified.

It is assumed that for blade systems, the redundant PSU configuration is rather standard and not an add-on for additional power allowance.

The approach to refer to suitable chassis types for Energy Star compliance in the product information is supported. However, this information should not only be available in the product documentation attached to the servers, but also made available on typical information sources for buyers like web sites and configurators in particular. The buyer should be informed at the time of product selection and configuration.

Line 602 – Active Mode Efficiency Criteria: We support the approach that this version only considers to collect and disclose active power information, while the quantitative requirements are reserved for further versions.

Line 728 – Power and Energy Data: We support that the idle power together with the measured power at 50 % and full load in Watt will be included in the Power and Performance Datasheet together with the estimated idle and full load energy consumption in kWh/year.

Line 735 – Selection of Benchmarks: We support the requirement of requesting testing based on a limited number of proposed benchmarks, since it facilitates comparison and use of benchmark results.

Line 758–762 – Data Measurement: We agree in the minimum data that the system must provide automatically. However, we should consider for future revisions – with increasing server power density – to additionally measure and report temperature increase and air-flow over the server mainly to enable more energy efficient continuous control of cooling and/or ventilation.

Line 806 – Sampling Requirements: We agree that it is important to include a rolling average because true energy measurements where power is integrated, is always preferable. Else varying power can result in very different measurement when a periodical sampling is used if for example the sampler all the time hits all the high power peaks.

Line 855 – Effective Date: It is not clear if goal of finalizing the structure of the specification by fall 2010 means the complete specification. In any case, we believe that the SERT development is quite difficult to plan and practical problems may occur.

We therefore recommend to have a fallback solution, which could be use of existing benchmarks combined with simultaneous power measurements and disclose requirements.

The Power and Performance Data Sheet

Page 4 – Thermal Information: We recommend to include reporting of minimum and maximum values together with typical values which is very important for energy efficient CRAC design.

Page 4 – Thermal Information: Inlet Air Temp vs. Fan Power should be explained in the data sheet.

The Server Efficiency Rating Tool – Draft Design Document

Basic Concept Description

It appears that the raw concept has been developed but the detailed developments still have to be done, while many challenges are still to be solved.

We propose to develop a clear road map with several milestones segmenting the overall process into packages. We believe it is important to focus on selected hardware and software platforms for the first stage to get a first version of the tool ready soon and to be able to test the capability of the approach.

Workloads vs Worklets

We believe that the approach of using worklets will be suitable for the purposes of this specification.

Information Policy

It is indicated that the SERT concept only allows general Energy Star compliance/non-compliance to be used for product marketing purposes. However, also detailed data on SERT testing except for log-files are indi-

cated to be publicly available. We recommend that these public data to be used for neutral product information data bases like the Energy Star data base

Consideration of Virtualised Servers

The SERT concept is currently not addressing virtualised servers. It is clear that this complex aspect cannot be addressed right from the beginning of the developments. However, due to the high relevance of this technology we would strongly recommend this for a later stage on the roadmap. According to the current concept, it seems completely out of scope for the moment.

Air Flow and Temperatures

We believe that it will be useful if SERT also could be a measuring tool for air flow and outlet temperature (line 128) to be disclosed in the Power and Performance Data Sheet.

We do not agree in deleting requirement for thermal flow rates (line 618).

License Price

We recommend to evaluate further if this would become a barrier to small and medium scale enterprises.