

ENERGY STAR Program Requirements for Servers 1.0 – Draft 3 – Comments from the ECESB Technical Working Group

We present in the following a summary of comments from the European Community Energy Star Board Technical Working Group (ECESB WG) on ENERGY STAR Program Requirements for Servers 1.0 – Draft 3. The comments are public.

Definitions

Line 33 – Permanent or temporary front side labelling: Front-side labelling should be no problem as there is still the option of non-permanent labelling. Limitation of air flow would only be an issue if the label covers a significant area of the server front.

Line 65 – Note on product qualification: It is highly appreciated to have updated product lists. However, there is no reference in the text, which we will recommend. It should be a general requirement in the partner commitment for all Energy Star products to keep the qualified lists updated.

Line 164 – Computer server definition: We support the definitions, however, with the remark that the requirement in line 169, packaged and sold with power supply(s), should be moved into the Qualifying Products section (line 320). This avoids potential contradiction with the definition for Blade Server and allows subsequent revisions to address computer servers without a power supply.

Line 236 – High availability servers: The terminology has been changed from “redundancy” to “availability”. We are still not sure if this is the best way to distinguish the different server configurations since “high availability” is also used in a different context at system level. One additional problem in the current definition is that there may be servers with a service processor but equipped with only one power supply. Consequently power supply redundancy can not easily be coupled with the service processor criterion.

Line 296–311 – Idle definition: It may make sense to directly state here that idle power is measured with the SPEC procedure and consequently to combine the paragraphs on “operational states” and on “other key terms”. There would be a direct link to the idle definition and “other key terms” could be deleted.

Qualifying Products

Line 334 – Blade servers: We agree to exclude blade servers for the first tier of the requirements since some issues need to be solved before this class of products can be addressed properly. Determining idle consumption for blade systems in general should be feasible but cannot be done properly based on the current SPECpower version. It would be preferable in this context to support testing based on a commonly accepted

standardized procedure. Using SPECpower as a basis will be an advantage for a next stage when the power consumption of servers under workload is addressed because SPECpower supports a multiple workload level concept.

In general, blade servers and chassis can not be tested separately but the complete blade system needs to be tested. Different configurations regarding number of installed blades are possible. Minimum configurations (e.g. 1 blade per chassis) clearly are inefficient. Consequently, a practical approach would be to define specifications for a reasonable minimum configuration (e.g. 60 % blade capacity installed) as well as for the maximum configuration.

Since blade servers are to be tested and Energy Star qualified as blade-chassis systems, blades will have to be clearly specified for certain chassis and vice versa. Besides server blades also storage and network blades can be installed in the blade chassis. However a combined treatment of different types of blades will be even more complex and thus may be an issue for a later stage.

Line 368 – Power supply efficiency requirements: We support the CSCI harmonisation and the proposed levels.

A significant number of servers is still operated at low loads although server and consequently workload consolidation becomes more common. Especially in cases of redundancy and in low configured blades but also for rack servers low electric load levels are encountered. This is partly also due to “relative over sizing” of power supplies. Consequently, it makes sense to address the 10 % load level as proposed by the current draft specifications.

To avoid a disadvantage for small power supplies 20 % requirements may be implemented from minimum PSU size upwards (see also further below). On the other hand we think that there is little relevance in testing efficiency at 100 % power load since this level will rarely be encountered in practise. It may be more practical to introduce a 75 % load level and to skip the 100 % level . Overall it would seem appropriate not only to address PSU efficiency but also to try to cover right sizing of PSUs in some way.

Calculations for typical server configurations (modified data from SPECpower) show that max power (at maximum workload) is often only 45 – 50 % of max rated PSU power (non-redundant situation). In case of redundancy the gap between maximum rated power and maximum required power is even bigger. However, in any case the PSU needs to be oversized compared to the on mode due to more power drawn at start up.

It is sometimes argued that power supply capacity must meet power demand also for high configurations of servers and consequently in case of upgrading. Alternatively different sizes of power supplies for different configurations could be offered. Since manufacturers offer online tools for server configuration an adequate sizing of the power supply should be possible. For the moment the major problem is the limited availability of different sizes of power supplies especially in the low-end power segment.

The following potential approaches, which could provide an incentive to avoid oversizing of power supplies could be considered for future tiers:

- Defining a maximum ratio of PSU rated power and max active power at full work-load and typical (and high) configuration.
- Introducing an additional size class for power supplies. Additional power supply classes may be introduced (e.g. <400W, 400–1000W, >1000W) and 10 % requirement would not be valid for the low-end class (smallest power supplies) but only for the other classes. Thus small power supplies would have no 10 % requirement since they anyway will be operated at higher power loads in practice.

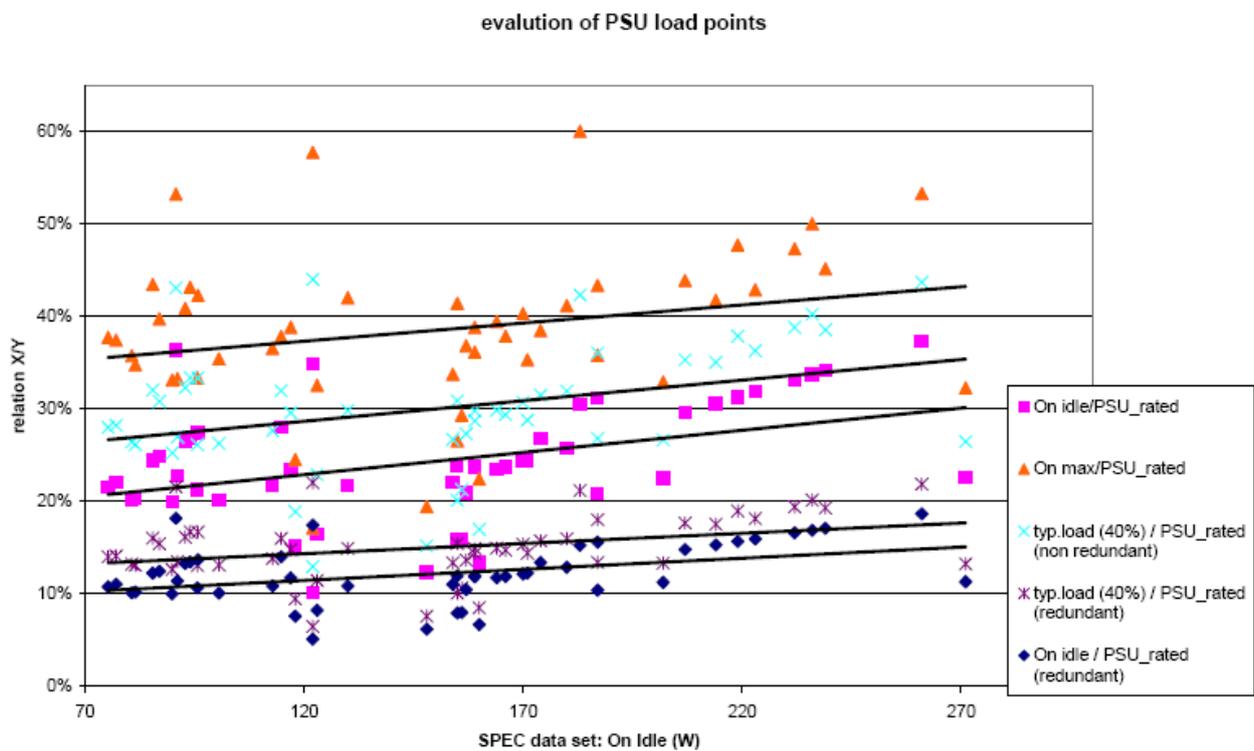


Fig. Relationship between on_max, rated PSU power etc.

Line 442 – Idle Power: The revised categorization of servers splitting subcategories by number of CPUs installed is appreciated.

Overall it may be beneficial to apply a similar sub-categorization for 2 CPU–4 CPU servers since variation in power supply redundancy and memory will also have an effect on 4P systems.

On the other hand, power demand of 2P standard availability systems will also vary with the amount of memory added. The current mix of specifications for additional memory in subcategories and in additional allowances may not be most efficient. It may be more straight-forward to address memory either only by subcategories (more subcategories) or only by additional allowances.

The levels provided seem reasonable at a first glance. Standards for installed memory will change over time and therefore have to be revised regularly. A higher number of discs (>2) within the servers today will only be encountered in very small server systems for small offices since in larger environments discs are outsourced to external storage units. Nevertheless, the total number of small servers in small offices should not be underestimated.

It was not completely clear why the additional allowance for the second and third hard disc differ by almost 100 % (8 compared to 15 W). Modern storage technology includes flash drives as well as more energy efficient hard drives. Efficient server hard drives are reported to demand between 5 and 6 W in idle mode. Flash drives show even lower power demands but limited life time and higher price are a major restrictions for broad use at the moment. There are new concepts now where servers are sold without local hard discs and the operating system is booted from a flash card.

In many larger systems as well as in most consolidated systems discs are outsourced to an external storage unit. Due to this common outsourcing of discs it seems appropriate to address storage systems as an additional product category in subsequent versions of the Energy Star requirements.

Fully buffered DIMMs now commonly used in servers (also due to the advanced memory buffer) demand more power than standard memory. The new generation of “low power dimms” designed for 1.5 V compared to previously 1.8 V is reported to allow improvement of energy efficiency by 20 to 30 %. However, we do not have sufficient data to evaluate the proposed 2 W/Gigabyte allowance.

As it seems typical configurations for SPECpower testing are partly below average. This has to be taken into account when SPECpower data is used for the specification of the required power levels.

Defining requirements based on idle power as percentage of maximum power (as discussed as one potential option in the explanatory notes) we would also not recommend at this stage since the involvement of maximum power would require a much more complex treatment of several issues (categories etc.)

Line 455 – Bins for Two or Three Installed Processors (2P & 3P): There is a potential problem with these two bins because a simple raise of the memory size, e.g. from 8 GB to 16 GB, you get additionally 52 W in idle power limit, i.e. 6.5 W/GB, which is high compared to the adder. The bins should not be characterised by only memory.

Line 510 – Standard information reporting requirements: We appreciate the standard reporting requirements especially the reporting for minimum, typical and maximum configuration.

However, a large amount of information is being required and there is therefore a bigger risk of user-input errors entering the database of qualified products. Additional review processes could help minimise significant errors.

Line 551 – Data measurement and output requirements: While the information requirements overall are welcome, some exclusions may be justified regarding low end 1P tower form factor servers which will be used in small numbers in small offices. There will often be no need for special infrastructure or cooling in these cases and prices in the low end segment may put some restrictions on technological features. On the other hand we think that all rack based servers should be covered and there should be no exclusion based on number of power supplies etc.

Line 624 – Tier 2: A statement reiterating line 324, that servers over 4 processors will be considered in Tier 2 could be inserted in this section.

We support the plan to revisit blade servers, and we would also suggest that server appliances, storage and network equipment be visited in this, or a separate Energy Star specification.