

ENERGY STAR Program Requirements for Servers 1.0 – Draft 4 – Comments from the European Commission

We present in the following comments from the European Commission (EC) on ENERGY STAR Program Requirements for Servers 1.0 – Draft 4.

Comments on Tier 1

Line 239 – Definition of Computer Server: We should be certain there are no obvious gaps between the definitions of servers in this server specification and of small-scale servers in computer specification. E.g. it is uncertain which specification a web managed server without ECC and/or buffered memory and without a dedicated display interface belongs to.

Line 309 – Managed servers: The capability of redundant power supplies and an installed dedicated management controller are not very special features but found in many server models today. Thus the statement that these servers are “primarily used for mission critical applications” may not be strictly true. A different wording may be appropriate.

Line 429 – Product family: The approach of defining product families is appreciated in general. It is clear that different mainboards and different types of processors (regarding model and number of cores) also automatically would mean different product families.

However, one aspect which is less evident is the segmentation by processor speed. Typically the same server model is offered on the market with processor versions running at several different speeds. Consequently it is not clear why a number of models slightly differing in processor clock rates should not be allowed to belong to one product family in case they all comply with the Energy Star idle criterion. It should be expected in most cases that maximum clock speed of the CPU is related to idle power demand. At the moment it seems that the definition of the product family is a bit more tightened than necessary.

Line 491 – Definition of qualifying products: We agree with current definition to include up to 4 socket system servers and only specify idle power limits for 1 and 2 processor systems independently of number of processor cores for Tier 1.

Line 521 – Blade servers: We support the proposal to include blade servers as long as it will not delay the effective date of Tier 1. Alternatively, blade servers should be moved to Tier 2.

The approach for blades may need criteria respectively information for different chassis configurations (numbers of blades per chassis). Considering a 100 % and 50 % configuration may be a good starting point. The

proposed one blade per chassis configuration is irrelevant in practise but can add some information for assessing overall efficiency and defining criteria. We also support the idea of additionally assessing the power draw of the chassis.

Line 574 – Power supply efficiency requirement: We agree with the new efficiency requirements for power supplies and appreciate the approach to have a stronger differentiation of wattage levels and an approach which avoids penalizing small power supplies.

Line 632 – Base system idle power requirements: The data set used for the definition of these levels probably also reflects other differences which have not been identified and specified. The difference between the levels for C and D most likely can not be attributed to the service processor and chipset alone.

Considering only the levels specified it would also not be clear why the additional allowance for the service processor and redundant power supply capability should be 10 W for single processor servers and 50 W for double processor servers. It rather has to be assumed that the “managed server category” compared to standard servers also implies higher configuration in terms of more powerful CPUs.

Furthermore, the additional allowances from A to C (respectively single processor standard to dual processor standard, 45 W) and B to C (respectively single processor managed to dual processor managed, 85 W) are quite different. Overall, the differences do not seem to be explained by the few factors considered here (standard versus managed) but by other configuration criteria as possibly processor type etc. so the current coarse classification also involves other factors not explicitly considered.

Nevertheless, we support the current approach as best possible current option based on the dataset.

Please notice that there is a mistake in the categorization since category B should be “managed single installed processor”.

Line 633 – Additional allowances for components:

The allowance for an additional power supply seems to be rather generous and it is not adapted to the categories A to D. Thus, there is a lack of incentive for appropriate dimensioning of power supplies. A rough evaluation based on the data in the “Draft 4 Idle data set” would support the recommendation to reduce this allowance to 10–15 W depending on the specific idle level category to encourage the use of better sized power supplies.

The 8 W allowance for additional hard drives seems appropriate for Tier 1. A reduction to e.g. 6 W could be considered for a next tier.

There is currently some trend that volume servers are also offered without an internal hard drive in the basic configuration, where the system is booted with a flash memory. It seems that these types of servers are included, because according to the definition it is not required that the server has one or more hard drives in-

stalled. Therefore it could also be an option to define the base level for the categories without internal drive and adding 8 W for any drive installed.

Allowances for additional memory should deserve some attention. Power consumption is more strongly dependant on number of memory modules than on the gigabyte size. To base additional allowance on GB only is therefore not an optimized approach. FB-DIMMS have normally an idle power of 6–10 Watt relatively independent of the capacity (1 – 8 GB). Consequently, a lower number of modules is more energy efficient and it would therefore be more effective to set requirements in a way that supports the use of fewer modules. However the issue also has to be reconsidered as soon as advanced power management at memory level will be available.

The additional I/O devices should be able to power down during the test and therefore should only need marginal additional power consumption. However, it might be needed during Tier 1 to allow the proposed levels if the I/O devices on the market do not support power management.

Line 793 – Exclusion of 3S and 4S Systems: The exclusion of 3S and 4S systems seems to make sense at the moment especially since currently only idle power is addressed and availability of data is limited. However, 3S and 4S servers should be addressed more comprehensively as soon as a benchmark/criterion to address active power is available. 1S and 2S systems are still the broad mass market. It is unclear however why the power management criteria defined for 3S and 4S servers are not applied for 1S and 2S servers as well.

Line 830 – Power and Performance Data Sheet: We welcome the detailed requirements regarding provision of Energy Star related information for complying products on the website and in printed documentation of manufacturers. We also appreciate the clear specifications regarding the power and performance data sheet. This implies that the relevant technical data and power data will be available on the websites as soon as the Energy Star label is used for a product or product family.

However, it is indicated that the provision of benchmark data based on SPECpower data is not mandatory but data from at least one benchmark which can be chosen by the manufacturer has to be provided. It may be more effective to define one benchmark or criterion which is mandatory for the assessment of the full power load (100%). Otherwise results based on different benchmarks provided by manufacturers will not be comparable and consequently not very useful for the buyer.

Comments on Tier 2

Line 307 – Addressing fully fault tolerant servers in Tier 2: It should be considered if the market sale of fully fault tolerant servers is large enough to make it relevant to include them in Tier 2.

Line 632 – Power requirements: We support that the idle criterion is only considered as a first step in the process and a further development of criteria/benchmarks addressing significant workloads should take

place in order to be included in Tier 2. This will also support the development towards servers designed to operate at higher loads due to consolidated and virtualized system).

Line 633 – Additional power allowances I/O devices: Tier 2 should support powering down inactive I/O devices and no additional power allowances should be included.

Line 929 – Storage and network equipment: We propose to start considering to address storage or network equipment in Tier 2.