AUSTRIAN ENERGY AGENCY

Comments on Energy Star Program Requirements for Computer Servers Final Draft

General remarks

Overall we appreciate the final draft of the first version of the program requirements for computer servers. We support the plan that the first version of the requirements is officially implemented and published not later than 15th May. This should be possible as most more complex issues can/have to be addressed for the second tier. A further delay of the implementation bears the risk that some criteria and requirements become outdated before they go into force.

In that sense several of the comments given below are rather to be seen as recommendations for tier 2 and should not further delay the implementation of the tier 1 specifications.

Specific comments

191 Including hard drive into product definitions

It has been specified that only computer servers with a hard drive installed are covered by the current requirements. This is ok in principal. However it may be beneficial to indicate in brackets that drives include solid state and flash drives and not only conventional hard disks (just to avoid any misunderstandings).

240 Multi node server

According to the explanation text has been added to indicate that dual node and multi node servers are not hot swappable. This has been done to delineate dual and multi node servers from blade servers.

We don't think that this additional text is really necessary since dual and multi node servers are already discerned from blades due to their internal power supply.

341 Revision of definition of "product families" regarding processor speed

We appreciate the approach to widen the "product family definition" by including products using the same processor but different processor speeds into one family.

367 Maximum configuration

We appreciate the new definition of “maximum configuration” targeting at the configuration with the highest power consumption.

407/408 Coverage of tier 2 and separate future specifications

It is appreciated that storage and network equipment shall be covered in separate future specifications. We believe that this product group is more relevant in terms of energy saving potentials compared to for example fully fault tolerant servers or multi-node systems. The plan to treat storage and network equipment in separate specifications makes sense as it allows to avoid overloading the server specifications as well as a different timing in the development of requirements.

410 Exclusion of blade servers in the tier 1 specifications

We support the approach to exclude blade servers from the tier 1 specifications primarily due to the fact that currently the database for the development of specifications is not broad enough. Data collection would lead to an undesired further delay of the implementation of the tier 1 requirements. Nevertheless we believe that a testing approach similar to the one already indicated is feasible and consequently idle power of blades - in combination with an approach addressing power consumption at significant workloads - should be addressed as soon as possible after publication of tier 1.

437 Power factor

From our perspective the new rule regarding loading conditions with greater than or equal to 75W on the power factor requirements is ambiguous. In Table 2 requirements are defined for 10% load at rated output power of below or equal 500 Watt. This load point refers to an output power of 50 W in the case of even 500W rated power. This implies a conflict with the above mentioned requirement. Unclarity should be avoided in this context.

471/475 Base system idle power requirements and additional idle power allowances

Power supplies

The idle power requirements and additional allowances related to the number of power supplies implemented are not completely clear. According to the definitions (190/229) both standard servers and managed servers may be equipped with one or more power supplies. The rather weak delineation regarding power supplies only implies that ***standard servers may have*** multiple power supplies or multiple power supply capability and ***managed servers must have*** this capability. According to the current definitions capability or installation of number of power supplies does not play a role for the requirements in table 3. The only distinctive feature in this categorization seems to be the installed dedicated management controller.

Despite the explanation in (478) the additional power allowance of 20 watts (table 4) most likely will be claimed for any server equipped with more than one power supply.

Furthermore the allowance for the additional power supply seems to be rather generous. 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 10–15W depending on the specific idle level category to encourage the use of better sized power supplies.

I/O devices

We appreciate that allowances for I/O devices have been have been altered to be technology neutral and based on link speed and number of ports.

Hard drives

We think that 8W is not a bad level to start with and a reduction (e.g.6 W) could be considered for a next tier.

Memory

Allowances for additional memory should deserve some attention. Power consumption is more strongly dependant on number of memory modules than on specific number of GB. 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.

Except for the issue regarding power supplies these aspects are not really critical for the tier 1 requirements but should be taken into account for a second tier.

529 Computer servers with greater than two processor sockets

The exclusion of 3S and 4S systems seems to make some sense for 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 at significant loads is available.

On the other hand there is a current trend to a development of multi-core CPUs with more than 4 cores per chip. SUN Microsystems already offers systems with 8 core chips which however belong to the mid-range and high end server segment not addressed in the current specifications. However there is also a general trend to a development of CPUs with a higher number of cores (announcements by Intel etc.) which shall become available on the market quite soon.

If this multi-core technology with more than 4 cores per chip enters the market of volume servers a new categorization for the *Energy Star* power requirements may become necessary. This will also be the case if the scope is generally extended to more powerful servers in tier 2.

EPA currently has chosen the approach to categorize servers by the number of sockets only which is a valid approach for the moment. However the current development in the area of multi-core chips may lead to a situation where multi-core systems with >4 cores/chip have the same or even higher relevance compared to multi-socket systems. Multi-core/chip has several advantages compared to multi-socket for example regarding management and co-ordination of CPUs and thus may become more abundant. Dual core systems will differ significantly from 6 or 8-core systems in performance and energy consumption.

Consequently a next tier server categorization has to consider both the number of cores and the number of sockets.

These types of systems should be covered in the requirements as soon as a benchmark applicable for the assessment of energy efficiency at significant workloads is available.

577 Reporting requirements

Overall we appreciate the standard reporting requirements and the "*Power and Performance data sheet*". The approach is comprehensive.

However we still believe that the request of power and performance data for "at least one benchmark chosen by the partner" is not very helpful.

Manufacturers have their own internal benchmarks which can be used for marketing reasons but are likely to be tuned to work particularly well with the specific hardware. Consequently the results based on different internal benchmarks provided by manufacturers will not be comparable and thus not very useful for practical purposes.

It may be more effective to define one benchmark or criterion which is mandatory.

630 Data measurement and output requirements

The accuracy required for the power consumption measurement provided by internal features of a server (service processor, embedded thermal and power meter) has been specified to +/-10W for measurements below 100W. This seems a rather week requirement. Overall also the 10% level above 100W seems rather week for example indicating an inaccuracy of +/-25W at a power level of 250W for example (225-275).

To support a reasonable monitoring by internal features, it is proposed for a next tier rather to specify accuracy to a level +/-5% with a cut-off at 5W. Thus values at levels below 100W do not have to be more accurate than +/-5W.

994 Measuring full load power

The approach to cover also conditions for full load measurements is appreciated. However it has to be considered that the nature of the benchmark applied has considerable impact on the results and may prohibit comparison of results.

Consequently it would make more sense to propose SPECpower as a standard for this type of measurements, despite of its shortcomings in the fields of memory-, network- and IO utilization.

B. Schäppi, T. Bogner, H. Teschner

Austrian Energy Agency