



Position on Draft Task 7 Preparatory study for Enterprise servers (ENTR Lot 09)

July 2015

Overall, environmental NGOs would like to welcome the publication of Draft Task 7 and the consideration which has been given to idle power and resource efficiency, as recommended in our previous comments. We find the content of sections 1.5.1, 1.5.3, 1.5.4 and 1.5.5 satisfactory, and have made suggestions on the rest of Task 7 below.

Power Supply Efficiency (section 1.5.2.1)

“The Lot 9 preparatory study recommends first, to seek for harmonisation with the existing EU Regulation No 617/2013 on eco-design requirements for computers and computer servers and as a second step to require higher standards for PSUs, such as 80 PLUS Silver in Tier 1 and 80 PLUS Gold in Tier 2, for both enterprise servers and storage. The time line could be set to 2018 for Tier 1 and 2022 for Tier 2.”

We recommend an alternative approach to requiring higher 80 PLUS certification levels in tiers 1 and 2 that could save more energy at a lower cost to manufacturers: the priority should be on idle and low load efficiency. As explained in the next comment on 1.5.2.2, reducing idle and low-load power levels will, by far, have the biggest impact on overall server energy use, because the vast majority of overall server energy is spent in idle mode or at very low-load levels. Instead of requiring higher 80 PLUS certification levels, we suggest requiring higher 10 percent load efficiency of 79 percent up to 1000 W, and 84 percent above 1000 W (which represents the average of 80 PLUS-certified 230V internal redundant power supplies), and keeping the current 80 PLUS requirements for 20, 50, and 100 percent load (Silver up to 1000 W, and Gold above 1000 W), to avoid backsliding in the higher active load zone. Focusing on the most critical load point would reduce the cost to manufacturers relative to requiring higher efficiency across the entire load curve.

Regarding the proposed timelines, 2018 seems appropriate for Tier 1, but an additional 4 years seem unwarranted for Tier 2. We therefore recommend a shorter timeframe such as 2020; the technology exists today, all that is needed is to send a clear signal to industry with enough time to adjust their product designs.

Reduction of idle power (section 1.5.2.2)

“...recommends to investigate possible idle power thresholds and specific additional idle power allowances for extra components in a next step.”

We support the draft recommendation to investigate possible idle power thresholds and specific additional idle power allowances, however it is not clear what is meant by “in a next step”. We believe that the data is available in the ENERGY STAR Qualified Products List (QPL)

and should be utilised in this study, so that it can be used to set Ecodesign regulations as soon as possible.

Reduction of idle power is critical because the vast majority of overall server energy is spent in idle mode or at very low-load levels. Multiple studies, including the most recent by Koomey and Taylor (June 2015)¹ which was based on production data from 4000 servers, have estimated the number of “comatose” servers (which have not delivered any information or computing services for six months or more) at 30 percent, and the average utilisation level between 6 and 12 percent. This situation has not changed significantly in over a decade, and this is likely to continue. It is therefore critical to optimise server efficiency for idle and low-load operating conditions.

Moreover, we believe that concerns over potential “unintended consequences” of setting idle load limits are not evidence-based, and should be researched before being included in this study. We disagree with the three concerns stated for the following reasons:

- 1) **Latency when waking up from low-power mode** – “Reactivation or synchronisation of server components shut-off to reduce idle power takes time”: it is not necessary to shut-off components completely, most components now offer low-power states with very rapid resume times (C7 states on CPUs, hybrid drives etc.). In addition, not all servers need to be instantaneously ready-to-respond all the time. For example, groups of load-sharing web servers are sized to be able to handle peak-traffic, but the majority of them can be put in low-power modes most of the time, only being ready-to-respond when traffic requires it. The power generation electricity has been using the concept of spinning and non-spinning reserves for decades, the data centre industry also uses it to some extent. Its use needs to be encouraged by ensuring that servers are designed to save energy when this technique is utilised.
- 2) **Failures to wake up from low power modes** – Failures to wake up were a known issue a decade or more ago, but recent technology no longer has this problem. An NRDC communication with a data centre software company whose solution powers down servers when unused and powers them back up when needed, indicated that the company had not seen a single failure to restart servers that are 5 years old or less, out of tens of thousands of production servers. Remnants of a legacy concern should not be a reason to dismiss a promising energy efficiency solution.
- 3) **Performance restrictions from idle power limits** – The ENERGY STAR framework includes categories and functionality adders (such as for memory), so that the limits depend on the performance capability of each machine. With appropriate categories and adders, idle power limits should ensure that the most efficient servers of all performance capabilities to comply, and that the least efficient ones do not.

| |
|---|
| Dynamic range (sections 1.5.1.3 and 1.5.7) |
|---|

“The Lot 9 study recommends to consider requirements for a dynamic range that would be an indicator for load-adaptiveness and a more efficient product.”

While the concept of dynamic range (ratio between idle and max power) is good in that it encourages energy proportionality, the metric could also provide a perverse incentive for manufacturers to increase max power rather than reduce minimum power. This would result

¹ <http://anthesisgroup.com/30-of-servers-are-sitting-comatose/>

in higher, not lower, energy consumption. For this reason, dynamic range by itself may not be a sufficient indicator of product efficiency for both product information and product labelling.

| |
|------------------------------------|
| Resource efficiency (1.5.6) |
|------------------------------------|

We welcome the proposals on dismantling, reusing and recycling (section 1.5.6.1), the requirements for technical documentation (section 1.5.6.2), and the proposal of a bonus for servers containing reused components (section 1.5.6.4)

However, we believe that the proposal for a voluntary declaration on the location of critical raw materials clearly lacks ambition. We call for an alignment with the mandatory information requirements on the presence of rare earth materials in permanent magnets which has been proposed for the review of the electric motors and industrial fans regulations. The same approach has to be taken for plastic recycling, where at least the same rules as for electronic displays should apply. Finally, spare parts availability, and the possibility to perform upgrades with reused components should be sought and we urge the study team to put forth proposals.

END

Contact:

ECOS – European Environmental Citizens’ Organisation for Standardisation

Chloé Fayole, chloe.fayole@ecostandard.org