Position on the draft Ecodesign requirements for Servers and Data Storage Products

March 2017

We welcome the opportunity to provide comments on the draft Working Documents on Servers and Data Storage Products, as presented to the Ecodesign Consultation Forum on 17th February 2017. The Commission's proposal includes positive elements, such as:

- The three-stage power supply efficiency requirements
- The scope and idle power levels for servers that are in line with those proposed in summer 2016 by the U.S. Environmental Protection Agency (EPA) for the first draft of the ENERGY STAR version 3 specification for servers
- The resource efficiency requirements.

There are however several significant issues and improvement opportunities that should be addressed:

1. **Idle power demand requirements**: We strongly support the inclusion of idle power demand requirements in Ecodesign, which we see as critical to ensure a reduction of overall energy use.
    a. **The memory allowance is too high:** The proposed memory allowance of 0.25 W/GB for servers is much higher than required by current mainstream memory technology. This is particularly problematic for server configurations with large amounts of memory. We propose to use a different equation that better reflects the fact that power draw does not scale linearly with capacity.
    b. **Base allowance and overall stringency:** The idle power levels proposed by the Commission are already achieved by almost half (47%) the systems introduced in 2015 and 2016. By 2019, a much larger share will pass due to natural technology evolution, yielding limited savings from Tier 1. We recommend the Commission to tighten the stringency for Tier 1 to around 25% of 2015 and 2016 systems, and set lower levels for Tier 2 (2023) and Tier 3 (2026).
2. **Power Supply Requirements:** The Commission's proposal is achievable based on data from the 80-PLUS program. We recommend moving the Platinum requirement to Tier 1 and the 10% load requirement to Tier 2.
3. **Buffered DDR Channel Allowance:** Align allowance with ENERGY STAR.
4. **Inlet temperature and humidity ranges:** Require ASHRAE class A2 in Tier 1.
5. **Data Storage Products Efficiency Requirements:** Include energy efficiency feature requirements as in ENERGY STAR for data storage products v1.
6. **Scope:** Investigate network equipment in the ICT study announced in the Ecodesign Working Plan 2016-2019.
7. **Safeguard the resource efficiency requirements.**

We provide more detailed information on these issues and improvement opportunities below.

1. Idle power demand requirements

We strongly support the inclusion of idle power demand requirements in Ecodesign. A large share of servers still spend most of their time in an idle or near-idle state, particularly in small and medium server rooms and colocation data centres. During that time, they draw significant amounts of power and this state is responsible for the majority of overall energy use. Idle power limits remain critical to ensure a reduction of overall energy use.

We appreciate the Commission's and industry effort to develop a single active mode efficiency metric. We welcome future requirements using this metric whenever there is sufficient data and confidence in it, but this new metric should complement, not replace, idle power demand.

We agree with industry that Ecodesign should not exclude higher power servers, preventing virtualisation and consolidation. But there is no contradiction between low-power idle and high-performance. The proposed Ecodesign framework is performance-based, with categories and adders that scale power allowance with performance. We support adjusting the framework as needed to ensure that the most efficient products of all levels of performance are able to comply, while the least efficient are not.

a. The memory allowance is too high

The proposed memory of allowance of 0.25 W/GB for servers is much higher than required by current mainstream memory technology. This is particularly problematic for configurations with large amounts of memory. We propose to use a hyperbolic tangent equation, or an equivalent approach, to better reflect the fact that power draw does not scale linearly with capacity, and ensure that high-memory configurations are appropriately handled by the regulation.

While reduced from the current 0.75 W/GB in ENERGY STAR v2, the proposed memory adder of 0.25 W/GB is still much higher than warranted by today's memory technology. This is evident when using online server configurators like the HPE Power Advisor Utility which calculate the power increase resulting from adding various amounts of memory to a server. We chose a recent HPE server model ProLiant DL380 Gen9, and configured it with various amounts of memory in addition to the 8GB in the baseline configuration, with results shown in Table 1:

Table 1– Idle Power Impact of Adding Memory to an HP Server According to HPE Power Advisor Utility

Proliant DL 380 Gen9	Additional memory (GB)	Idle Power (W)	Incremental power (W)	W/GB
Base config (8 GB)	0	43.11		
+ 8 GB	8	43.28	0.17	0.021
+ 16 GB	16	43.28	0.17	0.011
+ 32 GB	32	43.55	0.44	0.014
+ 64 GB	64	44.04	0.93	0.015
+ 128 GB	128	44.69	1.58	0.012
Average				**0.015**

While this is only for one sample server (albeit a mainstream and high-volume product), this exercise suggests that the memory allowance should be around 0.015 watt per GB, less than one tenth that is currently proposed.

While overly high memory allowances have limited impact for configurations with low amounts of memory, the impact becomes significant for servers with high memory capacity. The ENERGY STAR v2 Qualified Product List (QPL) shows that half of the qualified models have more than 64 GB, a quarter more than 192 GB, with some as high as 3072 GB.

Figures 1 & 2 – Memory Configurations in QPL[1]

Installed Memory (GB)	Number of Configurations in QPL
0-64	367
64-128	98
128-192	100
192-256	13
256-320	51
384-448	24
512-576	32
768-832	13
1024-1088	13
1472-1536	1
1536-1600	7
2048-2112	3
3008-3072	2
	724

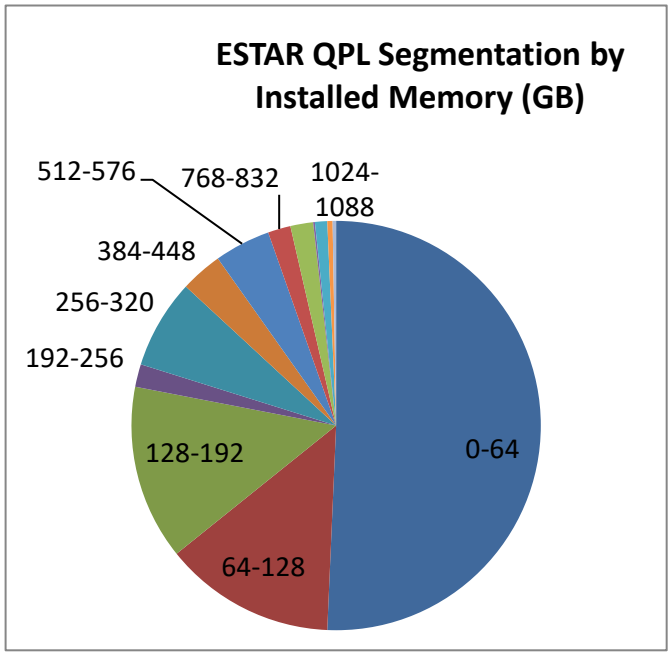

The overly high allowance of 0.25 W/GB vs. 0.015 W/GB can result in several hundred watts of extra allowance for servers with large memory capacities, which dwarfs all other allowances, effectively giving those servers a free pass even if they are very inefficient. This is shown in Table 2:

Table 2 – Impact of overly large memory allowances on total allowance

Memory (GB)	16	32	64	128	256	512	1,024	2,048	3,072
0.25 W/GB	4.0	8.0	16	32	64	128	256	512	768
0.015 W/GB	0.2	0.5	1.0	1.9	3.8	7.7	15	31	46
Unwarranted allowance (W)	3.8	7.5	15	30	60	**120**	**241**	**481**	**722**

The impact of this memory allowance issue is confirmed by Figure 3 which shows a clear correlation

[1] Including Minimum Power, Low-end, High-end and Typical, but excluding the Maximum Power configuration per ENERGY STAR v2.1

between pass margin and memory: the more memory, the easier it is for systems to pass, and by a wider margin. This is evidence that memory adders are too high. There should be no significant correlation between pass-rate and memory.

Figure 3 – Correlation of Pass Margin vs. Memory Capacity in QPL

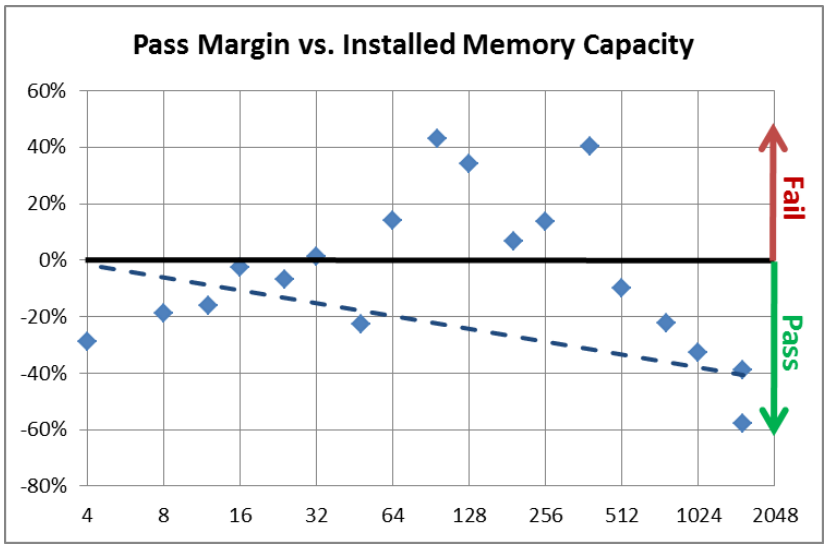

We suggest the following options to address this issue:

- Preferably use a hyperbolic tangent equation to reflect that power draw does not increase solely by memory capacity, but also by DIMM, and that large capacity DIMMs, while drawing high power than small capacity ones, are more efficient on a per gigabyte basis
- Implement the same concept in a simpler manner, use a 2-step linear allowance, with a proportional allowance of around 0.025 W/GB up to a certain limit (e.g. up to 64 GB, which would cover roughly half of the qualified models), then a linear allowance in the form of 0.015 a + b where a is the memory capacity in GB for a > 64 GB.
- Only use a linear equation 0.015 W/GB. While this may make it harder for low-memory configurations to pass compared to option 2, Figure 3 shows that low-memory configurations already pass by a wide margin.

We understand stakeholder concerns about potentially discouraging high memory configurations which are typically used in virtualisation hosts and can save energy by consolidating under-utilised servers. But this does not mean all high-memory servers should get a free pass. Both objectives can be met by setting appropriate memory allowances.

b. Base allowance and overall stringency

The idle power levels proposed by the Commission are already achieved by almost half (47%) of the systems introduced in 2015 and 2016. By 2019, a much larger share will pass due to natural technology evolution. **We recommend that the Commission sets lower levels for Tier 2 (2023) and Tier 3 (2026)** to

ensure continued savings and give manufacturers an incentive to invest in technology that reduces idle power.

The analysis of the ENERGY STAR v2 QPL shows that average pass-rates for systems launched in 2015 and 2016 are between 33% and 56%, with an average of 47%.[2] We focus on 2015 and 2016 systems because they are the most representative of current technology. The QPL contains systems that were put on the market as far back as 2010, which makes the entire QPL unrepresentative of current technology. While this analysis is focused on the ENERGY STAR v2 QPL, this dataset is quite representative of the current market. An analysis of manufacturer product configurators (like the HP Power Advisor Utility)[3] shows that almost all servers currently on the market can meet ENERGY STAR v2 energy requirements by a wide margin (but not necessarily power supply efficiency requirements).

Table 3 shows that pass rates for 2016 systems are between 33% and 56%. When the regulation goes into effect in 2019, three years after this market snapshot, it is likely that efficiency will have improved naturally, which means that the idle power requirements would have limited impact on the market as currently proposed.

Table 3 – Pass-rate for Proposed Lot 9 Levels of 2016 Servers in ENERGY STAR v2 QPL

Lot 9 Category	Fail	Pass	Fail	Pass	Total Systems
1-Socket	4	4	50%	**50%**	8
2-Socket	10	9	53%	**47%**	19
2-Socket Resilient		1	0%	100%	1
Blade and Multi-Node	4	2	67%	**33%**	6
Grand Total	**18**	**16**	**53%**	**47%**	**34**

The current proposal keeps the same idle levels for 2023 and 2026 as for 2019. It only increases power supply efficiency requirements. We encourage the Commission to tighten idle power levels for 2023 and 2026, for example by targeting 30% and 10% pass-rates for 2015-2016 models for tier 2 and tier 3, for Ecodesign to remain effective beyond 2019. Given the rapid advance of power efficiency in server technology, it is reasonable to expect that power efficiency will naturally evolve after 2019.

2. Power Supply Requirements

The Commission's proposal is appropriately ambitious and achievable. The Commission's proposal corresponds to 80 PLUS Gold for 2019, Platinum for 2023 and Titanium for 2026. A large share of the market already achieves these requirements today. Figure 4 shows that **75 percent** of new redundant

[2] The pass-rates only account for idle power requirements, not power supply requirements. Power supplies are a commodity and can be changed to achieve the Commission's proposed requirements with no impact on the rest of the server.

[3] https://www.hpe.com/us/en/integrated-systems/rack-power-cooling.html

power supplies certified with 80 PLUS in 2016 achieved the Platinum level, which will only be required in 2023. Titanium models represented just short of 10 percent.

Figure 4 – 80-PLUS 230V

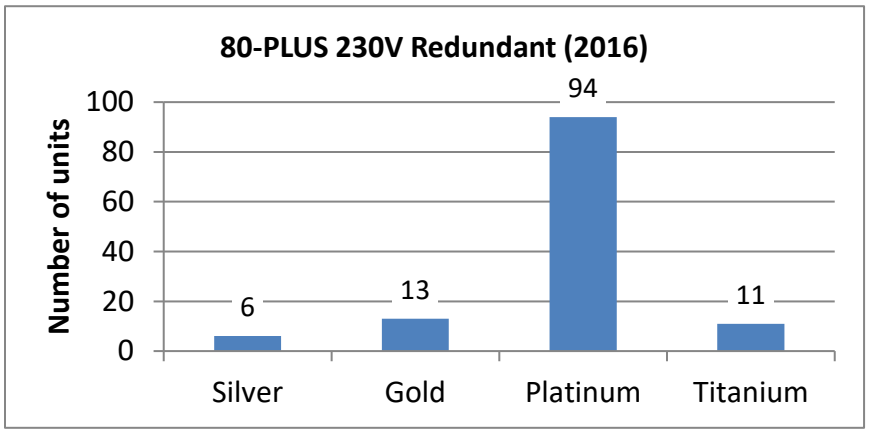

This is evidence that the requirements are feasible, and could be implemented earlier than proposed. Should the current timeline be kept, we strongly encourage the Commission to maintain the Tier 3 requirement for Titanium-level efficiency, in order to give industry a strong policy signal that will help them make the engineering investments necessary to achieve these requirements in the most cost-effective manner.

In addition, we recommend the Commission **to implement the 10% load requirement in 2023** instead of 2026, and strengthen it in 2026. Low-load efficiency is the most important load point for servers, as they spend the majority of their time and energy at low utilisation levels (per preparatory study Task 3). In addition, redundant power supplies have even lower loads (for example, they might share the load and therefore have half the load of a non-redundant PSU). At the same time, in order to reduce the burden on industry where it does not yield significant savings, the Commission could potentially slightly reduce the stringency of the requirements at full load since servers spend very little time working at more than 50% load. This would increase energy savings while adding no net burden on industry.

Strong power supply efficiency requirements are essential to guarantee energy savings from the regulation because they will continue to yield significant savings over the life of the regulation, even after technology and the market have evolved to achieve the idle state requirements with some margin and these idle power requirements no longer yield savings.

3. Buffered DDR Channel Allowance

The Commission proposes to give a 4-Watt allowance per DDR channel for installed buffered DDR channel greater than 8. This is consistent with ENERGY STAR, except that ENERGY STAR only gives this allowance to resilient servers. We encourage the Commission to align with ENERGY STAR on this.

4. Inlet temperature and humidity ranges

The Commission's proposal requires information reporting of idle power demand at ASHRAE operating conditions. While this is a positive step forward and a minimum, we note that the realisation of the energy benefit of this requirement requires 100% of products to support ASHRAE requirements. This is because data centre operators will not take the risk of increasing operating temperatures until they are confident that doing so will not impact any of the servers and data storage products in their data centre. Per the Commission's data, almost all IT products currently on the market already achieve ASHRAE A2 requirements. Mandating A2 by 2019 would only impact very few products, while giving data centre operators the certainty that all equipment purchased after 2019 supports A2 operating condition. **We therefore urge the Commission to require ASHRAE class A2 as part of Tier 1 in 2019.** As reported in Task 7, industry experts estimate potential data centre energy savings of up to 4% in operational energy costs for every degree of upward change in the cooling temperature set point.

5. Data Storage Products Efficiency Requirements

The only efficiency requirements that apply to data storage products in the Commission's current proposal are power supply efficiency requirements. While those requirements are important, other efficiency requirements should be set for data storage products. We recommend that the Commission **includes the same energy efficiency features** currently required by ENERGY STAR for data storage products version 1 described below:

- Adaptive active cooling: Primary components of a storage product must utilise adaptive cooling technologies that reduce the energy consumed by the cooling technology in proportion to the current cooling needs to the storage product. (e.g., reduction of variable speed fan or blower speeds at lower ambient air temperature).
- Capacity Optimizing Methods (COMs): A storage product shall make available to the end user configurable/selectable features listed in Table 3 in quantities greater than or equal to those listed in Table 4.

Table 3: Recognized COM Features

Feature	Verification Requirement
COM: Thin Provisioning	SNIA verification test
COM: Data Deduplication	SNIA verification test
COM: Compression	SNIA verification test
COM: Delta Snapshots	SNIA verification test

Table 4: COM Requirements for Online 2, 3, and 4 Systems

Storage Product Category	Minimum number of COMs required to be made available
Online 2	0
Online 3	1
Online 4	1

6. Scope

Task 7 concluded that networking equipment is an important product group, but that the complexity of these products made it unfeasible for the Lot 9 study to cover them in an adequate manner and recommended that a separate preparatory study is conducted on networking equipment, based on preliminary information provided in Tasks 1-4 of the Lot 9 study.

Given the projected continuation of the exponential growth in IP traffic and internal data centre networking capacity, network equipment may be responsible for an increasingly large share of data centre energy use without policy intervention. We therefore urge the Commission to follow the Task 7 recommendation and **investigate this product subgroup as part of the general ICT study announced in the Ecodesign Working Plan 2016-2019.**

7. Safeguard the resource efficiency requirements

In the Working Plan 2016-2019, Europe has taken a strong commitment to establish product-specific requirements in order to make products more durable, reparable, upgradeable, and designed for disassembly, reuse or recycling. In this context, we strongly welcome the resource efficiency requirements put forward by the Commission on servers and data storage products, which should be maintained, considering the following suggestions to improve clarity and/or robustness.

a. Design for dismantling, reuse, recycling and recovery (Annex II section 1.2.1)

We welcome the intention to ensure the easy access of key components for dismantling, re-use, recycling and recovery. We believe that the list of targeted components should include batteries in line with the JRC report on material efficiency requirements for enterprise servers[4].

b. Built in data deletion software (Annex II section 1.2.2)

We support the inclusion of this requirement, and underline the importance of the currently included information requirement on this tool, as this is necessary for market surveillance/verification purposes. We suggest that the word “equipment” in point 1.2.2 is replaced with “component” to avoid confusion between data storage products/equipment and components included within servers.

c. Availability of latest firmware (Annex II section 1.2.3)

The requirement on availability of firmware should be refined to avoid ambiguity. The current text lacks clarity regarding:

- What party’s firmware should be made available to (i.e. “Including but not limited to brokers, spare parts repairer, spare parts providers and other third parties dealing with maintenance, reuse and upgrading of servers”).
- How long after manufacture of the product/components firmware should be available for (this could be specified as a minimum of seven years from the day on which the last model was placed on the market, for example).

d. Information to be provided by manufacturers (Annex II Section 3)

Access to information

There are two different types of information provision specified in the working document, and the way this is framed is currently unclear.

[4] Page 79 of JRC Science and Policy Report on Environmental Footprint and Material Efficiency Support for product policy, Analysis of material efficiency requirements of enterprise servers, Laura Talens Peiró, Fulvio Ardente
September 2015

- [3.1 and 3.2 relating to information relevant to disassembly, recycling and disposal at end-of-life and information on the data deletions tool(s)
- [3.3 relating to CRM, instructions on sequence of operations for component removal and the firmware version]

For clarity, the wording of these requirements should be re-framed as follows:

- *From 1 January 2019, the following product information on data storage products shall be provided **to** installers and end-users, **by** manufacturers, their authorised representatives and importers **in** the instruction manuals and on the free-access websites*
- *3.3 From 1 January 2019, the following product information on servers and data storage products shall be made available **to** third parties dealing with maintenance, repair, reuse and upgrading of servers (including brokers, spare parts repairer, spare parts provider and third party maintenance) or recycling of waste electronic and electrical equipment **by** manufacturers, their authorised representatives and importers **via** free of charge website registration.*

The following further edits are also proposed:

- **Information channels:** The rationale for these different levels of information access is unclear. We suggest that all the information in 3.1 to 3.3 be made available in both instruction manuals and via free access websites.
- **Duration of availability:** It should be further specified (as described in the JRC report) that website-based information should be kept available for seven years from the day on which the last model was placed on the market.
- **Exploded diagram availability:** Further, an exploded diagram of the product should be required in section 3.3. This should illustrate the location of components/materials with special handling as identified by WEEE Directive 2012/19 EU annex VII, and the location of targeted components listed in section 1.2.1, and include the specifications of each of these components. This is necessary in order to facilitate refurbishment, maintenance and recycling.

Information on Critical Raw Materials

Information on total weight per product of cobalt, neodymium and palladium is requested in section 3.3. The inclusion of information on critical raw materials is considered appropriate in line with the potential impact CRM can have on the choice of recycling approach and with respect to recent requirements included for fans and electronic displays. It is important that the requirement for "an indication of the components in which the critical raw materials are present" is retained.

e. Other resource efficiency aspects that could be included

Spare parts availability

A requirement on information provision regarding the duration (in years) of availability of spare parts should be included. Like the rest of the information, this information could be provided for at least seven years after the last product is placed on the market. It has been suggested that spare parts availability can sometimes be a proxy for product lifetime. Therefore, provision of information on spare parts availability to installers and end users can be very useful, and declaration and increased attention on this aspect could contribute to improved reparability of products.

Incentive for enterprise servers with reused parts

A refurbished product could have higher annual total energy consumption but still generate overall environmental benefits. In principle, there is potential for regulatory approaches to account for this in the setting of Ecodesign requirements, although the specific approach needs careful consideration. Suggestions to set additional idle power allowances for servers that include reused components were detailed in the JRC Science and Policy Report on Analysis of material efficiency requirements of enterprise servers. However, it is essential that incentives for re-used parts account for the following:

- **Appropriate level of allowance:** Whilst the JRC analysis suggests a range of 7 to 20%, the allowance should be capped at a lower level to ensure that net environmental benefits are facilitated via the measure (rather than a break-even).
- **Definition of type of product/reuse to which the allowance will apply:** The JRC study suggests the allowance should apply to products that include reused components, and are sold "as-new". However, a study of reuse in the server market should inform a definition of the products to which such requirements would apply. This needs to appropriately take into account the status of the refurbished product and the types of components being reused, in order to incentivise a change in the market rather than rewarding status-quo design.

Material complexity

To facilitate plastics recycling, a requirement should be included that all plastic parts greater than 100g should be made of single polymer or directly recyclable polymer blend (to limit the variety of materials used). Examples are available in the voluntary agreement for imaging equipment.

Durability

To facilitate improved product durability, a requirement should be included for compliance with certain levels of standard MIL-STD810G (or IEC 60068/60529) relating to shocks and other damages. Examples are available in the JRC technical report for the Ecolabel of computers (2015).

Contacts:

ECOS – European Environmental Citizens' Organisation for Standardisation
Chloé Fayole, chloe.fayole@ecostandard.org

NRDC – Natural Resources Defense Council
Pierre Delforge, pdelforge@nrdc.org