Position on Draft Final Tasks 1-5 and Draft Task 6
Preparatory study for Enterprise servers (ENTR Lot 09)

May 2015

As a foundation for policy measures, preparatory studies need to adequately explore the policy options and technology solutions that regulation may address. The contents of Appendix 1 and 2 of this document provide a list of important considerations for the development of policy for enterprise servers and storage equipment that should be taken into account when formulating the policy options. In addition, we have the following specific comments on the preparatory study work to date.

Definitions

Whilst ENERGY STAR definitions can provide a useful foundation, they are often not sufficiently robust for use in mandatory measures which have to take into account legal considerations and outline a tightly defined scope. For example a definition aligning with ENERGY STAR may require that a server is packaged and sold with one or more ac-dc or dc-dc power supplies. This may unnecessarily exclude a number of products from scope.

➔ The study should make recommendations on how definitions should be handled within a regulation in order to ensure no loopholes are created.

Scope

Network equipment: Removing network equipment from the preparatory study risks this important product group being forgotten. In a US Environmental Protection Agency (EPA) scoping study for the ENERGY STAR programme it was found that:

- Energy use associated with network equipment is significant - A US study in 2008 estimated use at 18 TWh, or about 1% of building electricity on network equipment, with expected growth at 6% per year.
- The savings potential promising - estimated energy savings range from 20% to 50%.
- IT purchasers would benefit from access to standardised information about the energy performance of large network equipment
- Both commercial and residential network equipment need to be addressed - office building network switches and residential equipment are the two largest categories of energy use consuming 40% and 30% of the total respectively[1].

➔ As a minimum, recommendations should be made in the study as to how network equipment could be addressed in existing legislation, as well as considering the potential for information requirements (to provide the evidence base for future requirements), and recommending inclusion within any server/storage regulation of a statement of intent to bring these products under scope in future reviews.

[1] http://www.energystar.gov/sites/default/files/specs//EPA%20Letter_ES_Large%20Network%20Equipment_Launch.pdf

Enterprise storage: We strongly support the retention of storage products within scope for the following reasons:

- Basic data is already available to consider the performance/power characteristics of storage systems. For example there are already 80 storage products in the ENERGY STAR database.
- The ENERGY STAR specification has established foundation definitions and criteria for addressing these products.
- The ENERGY STAR programme states "*Data centre managers have identified storage growth as the trend having the greatest impact on their data operations. In 2012, most data centres reported data storage growth at 10 to 24% annually.*"

➔ **Enterprise storage should be retained in scope and addressed at a sufficient level of detail.**

Other boundary products: It has been implied that the following products may not be appropriate to include within the scope of a regulation:

- Resilient servers: with higher reliability, availability, and serviceability (RAS) as well as failover capacities (resilient);
- 4 socket managed servers
- High performance computing (HPC) products: designed for parallel computing and not for virtualization;
- Large mainframes and high energy density servers systems: featuring e.g. novel fluidic cooling systems as well as specialized racks and power supplies;
- Larger / higher performing embedded server and storage systems;
- Multifunctional product platforms which fit the description of a server or storage products depending on the actual configuration.

➔ **Rather than fully exempt the above groups, which should not be excluded on the basis of sales as they are subsets of a product group meeting the ecodesign sales requirements, we suggest that the application of requirements be tailored as necessary to retain these products in scope – for example, all the above products could have basic information and PSU requirements as minimum.**

Materials

Whilst Critical Raw Materials have now been addressed in the analysis of task 5, there is insufficient interpretation and conclusion on the results.

➔ **Further interpretation and conclusions on the results should be included in the report.**

BAT on SSDs

Further detail on BAT such as PCIe SSDs could be taken into account e.g.
http://www.techspot.com/review/989-samsung-sm951-pcie-ssd/

Identification of design options

The contents of Appendix 1 and Appendix 2 of this document provide a list of important considerations for the development of policy for enterprise servers and storage equipment. Many of the design options listed in the appendices to this document have not been addressed in the task 6 analysis, but have potential for improvement at low cost.

Specific comments on the options that are assessed in task 6:

SSD storage media design option: Whilst the inclusion of SSD as an illustrative scenario is interesting, realistic design options would be more constructive.

- Partial replacement of HDD storage with SSD, rather than full replacement: A sensitivity analysis around this may identify an optimal level of SSD storage substitution that meets LLCC levels.
- Learning curves[2] should be taken into account to account for reductions in costs of SSD in future. Including a BNAT case on affordable and higher capacity SSD could be envisaged.
- Scenarios such as powering down of storage devices into Slumber or DevSlp modes.

➔ **These alternatives should be considered.**

Blade design Option 5 - Full configuration vs. reduced configuration: It is understood that the scenario is about improving upon hardware utilisation, but it is unclear why a scenario is used that is considered entirely theoretical as "hardly anyone would buy two half populated blade systems and run them in parallel".

➔ **The real-world relevance of this option could be better explained. We suggest a more realistic foundation for this scenario is used.**

Operational efficiency: It is necessary to ensure that whatever mode is in use is optimized as far as energy consumption is concerned. There are no overarching design options that consider minimum levels of performance in operation of the server. For example, designing servers to meet ratios for SERT results, building upon the Triple E programme approach using SPECpower results, memory improvements to reduce idle and sleep mode power demand (e.g. DDR4) etc.

It is important that the design options to reduce power demand in individual operational modes are adequately considered. In particular an extended discussion on the power consumption of servers and data equipment in idle mode and low load, and the technology options for reducing these power levels should be included. If stakeholders choose not to provide data, this should be based upon research and assumptions to ensure that the preparatory study addresses all the options without bias. We suggest that the EPA may have some data from early development of the server specification that may provide a useful starting point.

➔ **It is important that the design options to reduce power demand in individual operational modes are adequately considered**

Power management: The focus of the power management design option is on processor design. There are other design options that should be considered in relation to power management, such as: default enabling, presence of remote power management capability, latency and stability requirements for recovery from standby/off mode etc.

➔ **Investigate these other power management design options**

---

2 For more information on learning curves, please see the following documents:
http://www.coolproducts.eu/resources/documents/2015-ALL/Priorities-for-a-Surge-in-Savings.pdf
http://www.eceee.org/events/eceee_events/seminar-ecodesign-and-innovation/9-hans-paul-siderius
http://proceedings.eceee.org/visabstrakt.php?event=4&doc=3-042-14
http://proceedings.eceee.org/visabstrakt.php?event=3&doc=6-289-13
http://proceedings.eceee.org/visabstrakt.php?event=3&doc=2-517-13
http://www.eceee.org/events/eceee_events/seminar-ecodesign-and-innovation/7-shailendra-mudgal
http://www.eceee.org/events/eceee_events/seminar-ecodesign-and-innovation/resolveuid/a119823234514d68a4d1125993299cdf
http://www.eceee.org/events/eceee_events/seminar-ecodesign-and-innovation/resolveuid/00564e688a234f54a76733e7cbee3627

ASHRAE A1 option for servers: In their 24 April 2015 paper, Digital Europe state that existing IT equipment and data centre infrastructure are largely capable of attaining the ASHARE A2 standards for the temperature and humidity. An LLC evaluation on ASHRAE A2 would be much more insightful than evaluating ASHARE A1 which is already the industry minimum. Moreover, A2 is a preferable improvement as it provides more headroom in the event of cooling failure and allows cooling provision to be lowered, giving data centre operators more confidence to raise the inlet temperature.

➔ **Assess ASHRAE A2 rather than A1**

Reuse rate design option: We support the inclusion of a design option addressing reuse rates for storage, but suggest that more details are necessary of exactly what design changes this scenario would require. In addition, we suggest that a scenario addressing wider lifecycle impacts of servers also be included. It is recognised that there can be a delicate balance between energy efficiency improvements achieved through server replacement against the reduction of production-phase impacts through extended lifetime of existing servers. However, there are other opportunities to reuse/remanufacture server components and modules which could help considerably to reduce environmental impacts.

➔ **Develop the storage reuse rate design option.**

➔ **Include a scenario addressing wider lifecycle impacts of servers encompassing opportunities to reuse/remanufacture server components and modules**

Resource use and end of life aspects

Whilst these aspects might suffer sometimes from a lack of data, resource use and end of life considerations nevertheless have an important environmental impact that EU product policy needs to address.

In the preparatory study work to date they are not sufficiently taken into account although relevant standardisation processes are currently being undertaken: notably Standards Development Organisations NSF and IEEE continue their discussions on a joint environmental leadership standard for servers. NSF International is beginning the initial balloting of the draft of NSF 426 *Environmental Leadership Standard for Servers*. The following options are further elaborated in the draft standard that is available via their public website here, and should be further investigated in the ongoing preparatory study.

Availability of firmware updates: For example, updates related to Security, Safety, Health and Design Flaws Corrections. These are developed by server OEMs, and are necessary to ensure reliability and compatibility of servers in the event that spare parts are replaced or maintenance is carried out. Access to these updates may be restricted by OEMs, having a significant impact on the ability of third parties to maintain machines through spare part installation.

Design for repair and reuse: External enclosures shall be removable by hand or with commonly available tools, without destruction of the enclosure. At a minimum, if present in the product, data drives or cards, processor, memory DIMMs, power supply, fans and I/O cards, shall be accessible and replaceable by hand or with commonly available tools.

Availability of manuals and tools enabling serviceability: In order to facilitate reuse and remanufacturing of components and modules, third parties should have full access to the tools that allow for thorough testing of the equipment, including access to repair manuals, technical documentation, information databases, diagnostic tools and firmware updates.

Upgradability: Some OEMs will not sell upgrades to customers using second-hand components, which forces end-users who need more resources to acquire new equipment despite the potential for upgradability within their existing equipment. OEMs should provide third parties with the activation codes required to proceed with upgrades on reasonable terms.

Removal of personal data: Provide reliable means to archive or remove personal data from electronic devices with storage capacity. The existence of personal data on storage devices and the difficulties to safely transfer and/or delete such data is considered to be an important obstacle for more widespread reuse and refurbishment.

Replacement components availability: Product replacement components and, or product service shall be made available through the manufacturer or an authorized third party for at least 5 years after the product is first placed on the market. Replacement components shall include, at a minimum, power supplies, fans, hard drives, memory, processors and printed circuit boards. Information regarding the availability of product replacement components and, or product service shall be publicly available on the manufacturer's website.

➔ **Design and policy options should include more detailed assessment of the available options to facilitate re-use and refurbishment of components and modules as well as improved maintenance, repair and upgradability[3].**

Information and reporting on disk drive magnet type and location: The manufacturer shall indicate the type of actuator/voice coil and spindle magnets in the product's hard disk drive on the external enclosure of the hard disk drive by means of a QR code. The QR code shall link directly to the magnet type and location information on a publicly available database or the manufacturer's website in at least English.

Design for plastics recycling: Plastic parts with the exception of printed circuit boards, wire and cables, shall not have molded, glued or otherwise attached metal inserts or metal fasteners, unless the metal component can be completely snapped off manually or entirely removed with commonly available tools. It also shall not apply adhesives, coatings, paints, or finishes that have a significant impact on the physical or mechanical properties of the plastic when it is recycled.

➔ **The preparatory study should explore a more detailed end-of-life characterization of enterprise servers to increase recovery of key materials and facilitate recycling of plastic parts. It should also conclude on the need for related design or information requirements such as a step-by-step disassembly instruction with required tools, product or material specifications and troubleshooting information.**

➔ **Regarding the marking of plastic parts heavier than 25g the study team should align their investigations with the current proposals in the draft Ecodesign regulation for electronic displays, including information on the presence of fillers and flame retardants.**

---

[3] See also *FREE ICT EUROPE Submission To the Institute for environment and sustainability (IES) of the European Commission - Joint Research Centre sustainability assessment unit Regarding A study about reuse and waste,* Stichting Free ICT Europe Foundation www.free-ict-europe.eu, February, 2015

Development of policy options

Regulation vs VA approach: We believe that regulatory approaches are more effective, and voluntary approaches should be adopted only where regulation is proven not to be possible. There are already a number of initiatives addressing at least server efficiency and these provide a good foundation for a regulatory approach.

➔ **Regulatory approaches should be prioritised over voluntary approaches.**

Harmonisation: There is a greater range of potential material, energy efficiency, and end of life criteria possible than what is included in the current ENERGY STAR label.

➔ **Whilst harmonisation with other initiatives such as ENERGY STAR is useful, the scope of policy proposals should aim to go further.**

Information based requirements: As stated in the study, information is important in order to influence the user towards environmentally sound use of servers and enterprise storage (users determine the application, utilisation and therefore the resulting effectiveness of the product).

➔ **Opportunities to improve upon information provision should be explored whilst keeping in mind that it is essential that Ecodesign policy recommendations go beyond this and also include tangible performance requirements.**

Appendix 1: Policy option considerations for servers

The following options should be assessed for development of policy measures for servers, and it should be considered to which server type / usage scenario each option could be applied (those in grey are already assessed as design options in task 6, most justifications are based on evidence from the preparatory study):

Aspect	Potential policy option	Justification
Component: PSU	PSU-level power management / load adaptiveness	Power supply losses are considered to account for 10-20% of server energy consumption. A standard for the system to identify and communicate with the PSU and therefore throttle power consumption could improve upon this.
Component: PSU	Efficiency and power factor requirements: Internal PSU at 80 PLUS gold standard or better.	See above. According to DIGITALEUROPE, PSU losses are expected to decrease to 5-10% due to the migration to gold and higher ranked power supplies (task 4 report).
Component: PSU	PSU rightsizing and modular implementation for racks feeding multiple servers.	Rightsizing means not over-dimensioning the PSU. There is considerable potential for power savings with right sizing. Larger, modular PSUs can be built into server racks which feed multiple servers so that only the necessary PSUs need be activated at the same time as providing redundancy. Because these PSU can be reused, higher investment to achieve greater efficiency and quality can be justified. Creating a standard for this would facilitate interoperability
Component: PSU	Redundant PSU configuration – no load sharing	Load sharing between redundant PSUs should be avoided as it can be inefficient (both PSUs operating at below 50% load). To minimise losses of redundant PSUs they should be turned off until there is a failure of the primary PSU. This is particularly relevant to lightly-configured systems at low utilization levels and where redundant PSUs are not powered by redundant power feeds.
Component: Storage	Replacement of HDD with SSD	HDD replaced by SSD SSD can have lower average power demand compared to the same storage being provided by HDD, although costs can be higher.
Component: Storage	Optimum SSD storage capacity	There may be an optimum proportion of storage that can be replaced with SSD balancing costs against savings.
Component: Memory	Minimum memory / power ratio (GB/W or similar)	The power consumption of the memory in the use phase is influenced by many factors. DIGITALEUROPE stated that there can be up to a 50% difference in energy use and thermal dissipation for memory chips with the same GB capacity but made by different manufacturers / fabrication methods. The new DDR4 memory generation has improved power management capabilities and shown power reductions in the range of 10 to 15%. New technologies such as 3D-memory expected on the market in 2016, could reduce energy consumption to a tenth of current values and extend the lifetime by a factor of 10 (task 4 report).
Virtualisation	Presence of a software management system that enables virtualization capacity.	Virtualisation ensures more efficient operation by increasing server utilization. Memory resource issues may also need to be resolved to ensure that utilization over 50% can be achieved.
Active mode	Minimum performance power ratio	Ratios could be suggested for SERT results, building upon the Triple E programme approach using SPECpower results.
Idle mode	Modal efficiency requirement	Servers that need to respond quickly spend much time in idle, so this is an important mode in which to ensure efficient operation. Low idle power demand can be achieved by e.g. lower power memory, partial powering of memory in idle, power management of storage media, advanced processor power management such as adaptive voltage operation (also called adaptive clocking), improved power gating, and smart processor power management featuring an "intelligent boost" microcontroller which reduces processor power at times of low

Aspect	Potential policy option	Justification
		utilisation. (task 4 report)
Sleep mode[4]	Modal efficiency requirement	Large amounts of memory are still energised in this mode, and there are opportunities for improvements. This mode is especially relevant for servers not needed to respond to immediate work requests.
Standby/Off mode[5]	Latency requirement / max resume time	Would encourage greater use of this mode which is often disabled by users due to concerns about slow reaction times or stability. Especially relevant for servers that are required very infrequently e.g. used for a few annual or monthly activity peaks, not expected to be needed for several hours etc. Will be facilitated by greater use of SSD in future. IEEE 1680.4 may address resume time.
Standby/Off mode	Reliability / stability requirement	As above. According to industry stakeholders enabling of power management options is very low. Customers are said to disable power management in order to ensure high availability and continuous operation. However, traditional server utilization patterns are often low and cyclical – suggesting an opportunity for power management if the technology is designed to be sufficiently stable and quick to react.
Standby/Off mode	Remote power management capability	As above.
Power management	Default power management enabling User Information on stability, latency etc	As above. If not a requirement for default enabling, provision of information to encourage enabling by data centre operators.
Power management / load adaptiveness	Advanced processor power management (adaptive clocking)	Industry reported that this can lead to 10-20% less power consumption and improvements in work per unit of energy consumed
Higher inlet temperature	ASHRAE A1	As a consequence of the reduction in cooling capacity or cooling provision, the power usage effectiveness (PUE) of a data centre can improve at the system level.
Internal airflow and cooling efficiency	Efficiency improvements e.g. specifying minimum fan sizes	Bigger fans are more efficient. Due to data centre space constraints, servers can be designed to minimise footprint, combining more, hotter components in a smaller space which compromise airflow and require higher powered fans.
Maintenance/ Repair/ Reuse / Remanufacture of components or modules	Availability of firmware updates to third parties and for second hand parts, availability of servicing tools. And information facilitate reuse / refurbishment / remanufacture of key components, design for modularity and RRR, Ease of extraction of key components, Measurement and declaration of product and key component durability, ease of reparability of product / key components	The preparatory study observes that the re-use and recycling rates of servers is relatively high. All major companies have asset recovery services and/or leasing programmes that record resale rates of around 80%. However, these programmes do not cover the entire market and other sources discussed in Task 3 show lower figures. Lack of third party access to OEM firmware updates and software tools prevents third party maintenance/upgrading using remanufactured or reused components / modules. Improved design and information could facilitate greater component / module reuse and remanufacture The HP Moonshot is an example of modular design whereby the CPU RAM, HDD can be swapped out on a small daughterboard whilst the main chassis (a lot of metal) and network / communications fabric is retained. EPEAT IEEE P1680.4 will deal with design for maintenance, repair, lifecycle extension, end of life management. NSF 426 *Environmental Leadership Standard for Servers, issue 1, revision 1* also provides more detailed provisions and tools for verifications on those aspects. The draft standard is available via their public website here.

[4] Preparatory study mentions "Inactive power state 1: 90% of idle power consumption with up to 10 seconds recovery time", which could be considered equivalent to this mode.
[5] Preparatory study mentions (Inactive power state 2: 30% of idle power consumption with up to 20 minutes recovers time) which could be considered equivalent to this mode.

Aspect	Potential policy option	Justification
Increased recovery of materials to facilitate increased resource circularity	Design for dismantling of key components, marking of parts, declaration of the recyclability index for plastic parts, declaration of use of hazardous materials that inhibit recycling (e.g. mercury), Requirement for minimum content levels of recycled materials	Similar requirements have been considered in the review of the television and electronic display regulation. The preparatory study observes that regarding end-of-life of plastics, only a minor share goes to landfill, while the biggest part is incinerated with heat recovery. There are also significant considerations relating to critical raw materials. EPEAT IEEE P1680.4 will consider some of these aspects. NSF 426 *Environmental Leadership Standard for Servers, issue 1, revision 1* also provides more detailed provisions and tools for verifications on those aspects. The draft standard is available via their public website here.

Appendix 2 - Policy option considerations for enterprise storage equipment

Many of the important considerations for development of policy measures highlighted in the previous table for enterprise servers can also be applied to enterprise storage. In addition, some storage-specific considerations are listed in table below:

Aspect	Potential policy option	Reason
Storage efficiency	Presence of data de-duplication, compression, thin/virtual provisioning and/or array virtualisation	Such functionality optimises storage efficiency.
Storage virtualisation	COMS (Capacity Optimization Methods Software) / software-based optimization of storage capacity	The ability of storage products to manage more data per unit of energy consumed is dependent on the software capabilities enabled on the specific storage product as well as on the storage data network within the data centre. The use of COMS on a storage system can result in an increase in power consumption for the individual storage system, but reduce the overall power required to store the data. Alternatively software defined storage serves the same function as virtualization on servers.
Idle power management	Capable of intelligent power-down (e.g. to slumber/ DevSlp modes) and drive spin-down or slow spin (MAID 2/IPM) when idle.	Load adaptiveness would result in improved efficiency.
Disk tiering	Presence of a disk tiering strategy capable of supporting storage media with multiple power / capacity points with a factor of at least 2X between the slowest and fastest	Could reduce power demand per unit storage.

END

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