

Study for the review of Ecodesign of servers and data storage products

Feedback on Phase 1 – Technical analysis report and MEErP Task 1-4 draft reports

Brussels, 20 November 2023

ECOS welcomes the opportunity to comment on the Phase 1 – Technical analysis report and MEErP Task 1-4 draft reports shared by ICF to feed in the revision of Regulation (EU) 2019/424 on ecodesign requirements for servers and data storage products. We support the majority of these two documents and simply would like to suggest a few amendments to improve on their content.

Definitions

We support harmonisation of definitions in principle, but only through careful consideration of the risks of creating regulatory loopholes (which may be less of a concern in voluntary initiatives such as ENERGY STAR). The definitions should also consider the ease with which market surveillance authorities (MSAs) are able to clearly distinguish between products in and out of scope to assess which criteria apply and hence verify compliance.

We support the updates of the following definitions:

- **Storage heavy server definition:** These should be included as they can be tested using SERT.
- **Data storage equipment:** Definitions can be updated in line with SNIA.
- **Servers with integrated APA:** They should be included as they may be included in the next version of SERT and can currently be included in the regulation for non-energy requirements. If the definition is included, their inclusion in energy requirements should be explicitly mentioned in the review clause of the revised regulation.

Custom servers: It is stated that custom servers represent at least a third of the server market sales. Therefore, we have concerns about their exemption and consider that a new definition should only be included if it is possible to arrive at a very robust definition that does not accidentally include a large number of servers, which may be challenging.

Hyperconverged servers: The definition for hyperconverged server is insufficiently distinct from standard servers for which efficiency criteria apply. Standard servers can also be used for

hyperconverged infrastructure. If the distinction is found in the hypervisor/software, then a new definition for hyperconverged appliances may be needed. If the hyperconverged server is sold without any hypervisor/OS, it is not clear why it cannot be tested using SERT.

Scope / exemptions

We support the following suggested scope and exemptions:

	Energy requirements	PSU efficiency	Material requirements	Notes
Server appliance	exclude	include	include	/
Large servers	exclude	exclude	exclude	Fully out of scope
Fully fault tolerant servers	exclude	include	include	/
Data storage products	/	include	include	/
Resilient servers	include	include	include	/

In general, all servers should be in scope for PSU and material efficiency requirements. It should also be assumed that SERT will be updated in time to include the other types of servers for the energy requirements or energy information, rather than excluding them at such an early stage. If SERT testing is not available on time, the regulation should anticipate on that and contain a review clause triggered once testing is possible with SERT.

We propose the following changes to the proposed server requirements:

	Energy requirements	PSU efficiency	Material requirements	Notes
Hyperconverged	Include.	include	include	On energy requirements see previous comment.
Custom made	information only.	include	include	These should only be defined if a robust definition can be arrived at.
HPC	assume included	include	include	/
Servers with integrated APA	assume included	include	include	/

Storage server	heavy	include	include	include	The SERT weighting may need to be adjusted to account for different usage.
Liquid servers	cooled	assume included	include	include	Servers that only offer direct liquid cooling systems (chip or immersion), without an air cooling alternative should still be included in SERT energy efficiency testing requirements. The testing protocol would be similar to air cooling: any cooling parts (e.g. pumps) internal to the server would be included but external parts, including immersion tanks would be excluded.
Hyperconverged		assume included	include	include	/

Updates to ecodesign energy-related requirements

Product family approach / typical configuration

We consider it important to address the flaws highlighted in the product family approach that allow inefficient servers to be considered compliant. We support the suggestion to make declarations based on a "typical performance configuration representative of high-volume sales" since this information will likely be more useful to the consumer, provided that the typical configuration reflects the market. We would further suggest using a method to interpolate / estimate efficiency based on configuration. This could be either through a pre-defined methodology to be detailed within the regulations (e.g. interpolating between tested configurations based primarily on the CPU core count and frequency) or applying more sophisticated multi-variate or DL techniques across a larger dataset.

Idle metric

The analysis provided is an oversimplification of the situation due to averaging SERT scores over too many types of servers. The current SERT metric does not adequately address the idle power. A correlation of idle and SERT efficiency for current servers does not mean it is a necessary condition. While idle power criteria can be difficult to develop, we believe they are still essential (although not necessarily in the current form) for low performance servers and would not limit their performance. According to the [DC Eureka project](#), there is still a significant proportion of the market operating at low utilisation, and in idle approximately 80% of time. For cloud and similar deployments, disaster recovery sites can mean that large numbers of servers are in idle (and may be absent from declared utilisation values). In addition to idle power criteria, we consider incentivising higher server utilisation through real time utilisation reporting. This would raise awareness on the issue.

We support removing the idle state testing temperature requirement.

SPEC SERT for active efficiency

We strongly support the continued use of SPEC SERT to set active efficiency requirements but consider it important that the ambition of the criteria is increased. In addition we recommend the development of a threshold based on the performance of the server as well as other significant variables. This analysis should be easily achievable by correlating the SERT performance and efficiency test data already collated by the study contractor.

We urge the project team and European Commission to coordinate with the European standardisation organisation ETSI, who is currently working on the redraft of the EN 303 470 standard in line with the M/573 standardisation request (to address key aspects relevant to the server regulation). The changes requested by the EC have been made in the draft document and it was (until recently) close to being finalised. However, recent discussions may delay delivery of the standard and are, we consider, at risk of derailing the revision of the server regulation (2019/424).

There is considerable pressure from some stakeholders within ETSI to specify or develop an alternate benchmarking tool to SERT (please see attached documents for details). ECOS considers that the development of a new benchmarking tool could derail the revision of the server regulation for the following reasons:

- Little or no established data on how the market performs in relation to an alternate benchmark.
- Lack of comparability of alternate tools tool with SERT worklets enshrined in the regulation.
- Significant delays likely in the delivery of the revision of EN 303 470
- Significant delays likely in regulatory discussions as references to SERT worklets would need to be revised and analysis reworked after the alternate tool becoming available.
- Lack of clarity on the robustness of the alternate tool and how representative of energy efficiency the results would be.
- Lack of support for an alternate tool from many manufacturers.
- Legal uncertainty that would occur for companies and MSAs where two tools are referenced that are not directly comparable.
- Considerable additional resources required in the development and maintenance of 2 testing tools.

We ask that the study contractors work with the Commission to urgently provide clear guidance to ETSI on their expectations for the delivery of the EN 303 470 revision and the benchmarking tool that should be referenced in this standard.

Note: There are similar issues for data storage in relation to the existing SNIA approach and a new Chinese standard on storage energy efficiency.

Processor power management function

We support the enabling of processor power management as default. However, given the high rate of disabling when the server is configured, we also suggest that the review team consider requiring all the processor management in the BIOS to be configured as fully OS controlled. Some hypervisors and OS now recommend this option to balance performance and efficiency without affecting stability. This should result in representative test results and optimal performance for software designed to be efficient, without compromising the reliability.

In addition, a requirement should be introduced for manufacturers to ensure that any support and guidance given to users about optimal settings for running the OS and hypervisor on their servers is based on the latest recommendations.

Parameters information requirement

We support aligning with ENERGY STAR to require open data exchange of the temperature and fan activity. We also recommend that open exchange of the power consumption and utilisation is included as these enable the operators to understand and manage the efficiency operation of the server.

Energy label & GPP

We support investigating the process of server procurement, including SMEs and local government, to ascertain the optimal method to provide information for EU GPP for data centres. Based on feedback from Futuretech we believe that any information should include an estimate of the efficiency of the actual configuration being procured. We consider that this can be determined through interpolation with sufficient accuracy.

For custom (non-typical) configurations a more flexible approach than a formal energy label may be required, as it could be challenging to define an energy label based upon interpolated efficiency data, especially in terms of compliance and registration of models in the EPREL database. An energy label for a typical configuration may either add value or create confusion for procurers looking at custom solutions - we recommend the study investigates how the label might influence procurers in such situations.

Updates to ecodesign material efficiency requirements

We support the following proposals of the review study:

- **Materials:** A standard on material intensity of servers. Restrictions on material/substance mix. Requirements to allow for easy separation of the product into different materials. A more extensive product datasheet to track material content (e.g. info to recyclers on chassis content to facilitate recovery).
- **Parts pairing:** Requirements preventing parts pairing.
- **Spare parts:** Requirements on spare part availability of at least 7 years. We also support a framework to provide users with hardware component compatibility on the information sheet to facilitate refurbishment. As this would need to be updated as new processors and firmware appear, a QR code may be the best solution.

- **Technical lifetime:** Provision of an information sheet or label with the technical lifetime of the product.
- **CRM:** Extended information requirements for critical raw materials.
- **Licensing:** Information and other requirements to tackle the issue of licensing limitations, as these can drive product obsolescence. We recommend this is also explored under product lifetime in the task 2 report.
- **Firmware:** Extended firmware availability period and then making firmware available as open-source BIOS. Requiring the availability of previous versions of firmware (any cybersecurity issues can be handled by those running the data centre). As per mandate M/573, classes for firmware (spare parts availability, class B) and deployment skill level (class C) should be specified in the revised regulation.
- **Ease of disassembly classes:** Including a clause within the Ecodesign regulation to require servers and data storage products to be disassemblable by a generalist (class B), in a use environment (Class A), using tools meeting Class A, B or C nomenclature. We recommend that the study contractors coordinate with ETSI who are working on the ease of disassembly standard to ensure that the fastener classes are consistent.
- **Repair and Maintenance costs:** We recommend more information is gathered on this in the task 2 report as "cost of ownership is one of the top criteria in purchasing decisions".

Updates to other ecodesign requirements

Operating conditions

We support the continued provisions on the ASHRAE operating conditions, including a requirement to operate at A2 at a minimum. We would also propose information provision in the form of a QR code that would take server procurers directly to the server information sheet online, based on SKU or even serial number.

System Performance Considerations

We support the policy to increase operation of air-cooled servers to higher temperatures to optimise the efficiency between the server and cooling system.

ICT real time operating condition provision

We support the requirements to increase performance transparency and believe it will become a requirement under future reporting, including the Energy Efficiency Directive.

Corrections

In Section 1.2.1.3 on standards in the task 1 report, ETSI have planned a separate deliverable (a tr) on CPU power management. Also, CRM requirements will be included in the CEN-CLC EN 45558 standard as a Z-annex, rather than as an ETSI standard.

For more information, please reach out to mathieu.rama@ecostandard.org.