

Study for the review of Commission Regulation 2019/424 (Ecodesign of servers and data storage products)

Phase 1 report – Technical Analysis –
DRAFT v3.2

September 2023

Submitted to:

Davide POLVERINI

Policy Officer, DG GROW

Avenue d'Auderghem 45, 1040 Brussels, Belgium

ICF makes big things possible

ICF is a global consulting and technology services provider with more than 7,000 professionals focused on making big things possible for our clients. We are policy specialists, social scientists, business analysts, technologists, researchers, digital strategists and creatives. Since 1969 government and commercial clients have worked with ICF to overcome their toughest challenges on issues that matter profoundly to their success. Our five core service areas are described below. Engage with us at icf.com.



Study for the review of Commission Regulation 2019/424 (Ecodesign of servers and data storage products)

Phase 1 report – Technical Analysis – Draft v3.2

A report submitted by [ICF S.A.](#)

Date: September 2023

Job Number 330301272

Eirini Passia
ICF S.A.
Avenue Marnix 17
Brussels
B-1000
Belgium
T +32 (0) 2 275 01 00
www.icf.com

Document Control

Document Title	Study for the review of Commission Regulation 2019/424 (Ecodesign of servers and data storage products) – draft v3.2
Job No.	330301272
Prepared by	Laurent Petithuguenin, John Clinger, Todd Leddy,
Checked by	Eirini Passia, Tom Lock
Date	September 2023

This report is the copyright of DG GROW and has been prepared by ICF SA under contract to DG GROW. The contents of this report may not be reproduced in whole or in part, nor passed to any other organisation or person without the specific prior written permission of DG GROW.

ICF has used reasonable skill and care in checking the accuracy and completeness of information supplied by the client or third parties in the course of this project under which the report was produced. ICF is however unable to warrant either the accuracy or completeness of such information supplied by the client or third parties, nor that it is fit for any purpose. ICF does not accept responsibility for any legal, commercial or other consequences that may arise directly or indirectly as a result of the use by ICF of inaccurate or incomplete information supplied by the client or third parties in the course of this project or its inclusion in this project or its inclusion in this report.

Contents

1	Introduction.....	1
1.1	Aims and objectives of this report	1
1.2	Methodology followed	3
2	Items for review	4
2.1	Updating current Ecodesign requirements	4
2.2	Regulation definitions and scope	11
2.3	Data Storage devices performance requirements	16
2.4	Processor Power Management Function	18
2.5	Standby-Readiness for Servers	20
2.6	Parameters Information Requirements	21
2.7	Energy Label	22
2.8	Material Efficiency.....	26
2.9	Operating Conditions	38
2.10	System Performance Considerations.....	41
2.11	Liquid Cooling Systems and Solutions	44
2.12	Waste Heat Recovery Systems and Solutions	46
2.13	Direct Current Power Supply for Servers	49
2.14	Other topics	50
Part A:	ANNEXES	53
Annex 1	Installed processor metrics	54
Annex 2	Stakeholder Meeting Minutes.....	59

Table of tables

Table 2.1	Proposed Active state efficiency for servers in the EU GPP requirements ...	6
Table 2.2	Active state requirements for Block I/O Storage products on Energy Star...	17
Table 2.3	Ecodesign Requirement Standards for M573.....	27
Table 2.4	Operating condition classes.....	39

1 Introduction

The study for the review of Commission Regulation 2019/424 (Ecodesign of servers and data storage products) has been commissioned to update the Ecodesign regulation on the server and data storage product groups. The study covers a broad range of subjects, covering, amongst others, the definitions used in the regulation, the scope and exemptions granted, the energy efficiency requirements, potential labelling applications, technological improvements in the technology, usage patterns and market changes.

1.1 Aims and objectives of this report

This report covers the Phase 1 of the review study. This phase of the study seeks to answer specific questions raised in the article 8 of Regulation 2019/424, and other points of interest to DG GROW and other Commission Directorates.

Listed below are the items set out in Article 8 of Regulation 2019/424:

- a) to update the specific Ecodesign requirements on server active state efficiency;
- b) to update the specific Ecodesign requirements for servers on idle state power;
- c) to update the definitions or the scope of the Regulation;
- d) to update the material efficiency requirements for servers and data storage products, including the information requirements on additional critical raw materials (tantalum, gallium, dysprosium and palladium), taking into account the needs of the recyclers;
- e) to exempt server appliances, large servers, fully fault tolerant servers and network servers from the scope of the regulation;
- f) to exclude resilient servers, High Performance Computing (HPC) servers and servers with integrated APA from the Ecodesign requirements set out in Annex II point 2.1 and point 2.2 of Regulation 2019/424;
- g) to set specific Ecodesign requirements on the Processor Power Management Function of servers;
- h) to set specific Ecodesign requirements on the operating condition class;
- i) to set specific Ecodesign requirements on the efficiency, performance, and power demand of data storage products.

Further points of interest are:

- j) on material efficiency aspects:
 - a. the provisions on disassemblability of certain components, also considering advancements in standards (mandate M/543) since the publication of the regulation;
 - b. an analysis of the benefits of the information requirements under Regulation 2019/424 already covering cobalt in the batteries and Neodymium in the hard disks;
- k) analysis of the benefits of the information requirements under Regulation 2019/424 on the operating conditions of servers and data storage products;

- l) an analysis of the standards, and of their relevance for regulatory purposes, developed/under development under the standardisation request M/573, 'Commission implementing decision C(2021)14 of 12.1.2021 on a standardisation request to the European standardisation organisations in support of Regulation (EU) 2019/424 as regards Ecodesign requirements for servers and data storage products';
- m) Technological, market and regulatory evolutions affecting the environmental performance/aspects of data centres, and how they would reflect at product specific level, for servers and data storage products;
- n) The technical and economic feasibility and relevance of product specific requirements on liquid cooling systems/solutions
- o) The technical and economic feasibility and relevance of product specific requirements on waste heat recovery systems/solutions
- p) The technical and economic feasibility and relevance of product specific requirements on the standby-readiness of servers (for instance allowing to move to and from idle mode in a fast and seamless manner), if not covered by the analysis on the Processor Power Management Function
- q) The technical and economic feasibility and relevance of product specific requirements on DC (direct current) power supply for servers
- r) The technical and economic feasibility and relevance of product specific requirements related to:
 - a. the availability of information (temperature, (fan) speed, etc..) for open data exchange about the input/output air flow data of the server/data storage product, and/or
 - b. the capability to enable external overriding of the internal fan speed control, in view of potential synchronisation of the product cooling system with the data centre cooling system.
- s) The technical and economic feasibility and relevance of introducing an energy label for servers and data storage products, including a label and a detailed product information sheet comprising targeted indicators for the different possible uses of the product (e.g. as webserver, disk server, database server, file/disk server, etc.)
- t) Other topics, as emerged from consultations with stakeholders.

The subjects set out in this list of items, have been grouped in the report along the themes of:

- Updating current Ecodesign requirements (items a and b)
- Regulation definitions and scope (items c, e, and f)
- Data storage performance requirements (item i)
- Processor Power Management Function (item g)
- Standby-Readiness for servers (items p)
- Parameters Information requirements (item r)
- Energy label (item s)
- Material efficiency (items d, j, and l)
- Operating conditions (items h and k)

- System performance considerations (item m)
- Liquid cooling systems and solutions (item n)
- Waste heat recovery systems and solutions (item o)
- Direct Current supply for servers (item q)
- Other topics (item t)

1.2 Methodology followed

To answer the queries set out in Phase 1, the research team consulted with stakeholders through a stakeholder meeting, qualitative and quantitative questionnaires, and direct 1-to-1 calls. From this feedback, along with ICFs expertise and datasets from the Energy Star programme, the items below were answered. For each theme, the report details the background to be aware of in the theme, develops the stakeholder feedback and research results, makes recommendations to policymakers on next steps for Ecodesign and sets out in which sections of Phase 2 the subject would be further developed in the update to the Ecodesign preparatory study.

2 Items for review

2.1 Updating current Ecodesign requirements

- a) to update the specific Ecodesign requirements on server active state efficiency;*
- b) to update the specific Ecodesign requirements for servers on idle state power.*

2.1.1 Background

It was determined that to deliver the best savings, the Ecodesign Commission Regulation 2019/424 should include an active efficiency component. Although this had not been included in the original draft regulation, active efficiency was included into the final version of the regulation, with 9.0, 9.5 and 8.0 included as the minimum active state efficiency for 1-socket, 2-socket and blade servers, respectively. The regulation sets a maximum limit for idle energy consumption in servers. This limit is defined by a baseline consumption level and additional allowances for extra components, making it difficult to establish a standardised value. Idle consumption is crucial for distributed IT servers, as they often lack optimised processes to minimise idle time. Data centre servers, especially those providing cloud services, prioritise maximising efficiency by increasing server utilisation. However, tracking the market share of distributed IT servers, their utilisation rates, and ensuring compliance becomes a complex task.

2.1.2 Questions

1. What are your thoughts on the Ecodesign requirement for minimum active efficiency on servers?
2. The Ecodesign regulation sets active state efficiency requirements for 1-socket, 2-socket, blade and multi-node servers at 9.0, 9.5 and 8.0 respectively. Do you think the bulk of the market is above these requirements? As a result, do you think these regulatory requirements should be tightened?
3. What are your thoughts on the Ecodesign requirement for maximum idle state power on servers?

2.1.3 Feedback/ Research results

Regarding the requirement for minimum active efficiency on servers, the stakeholder feedback emphasised the significance of server energy efficiency in data centres, with servers being the largest energy consumers. It highlighted the importance of active state efficiency as a metric, considering it represented the common utilisation level for servers. The feedback recommended that governments prioritise active state efficiency as the primary metric for setting minimum energy performance standards (MEPs). Additionally, the feedback mentioned that SPEC SERT serves as the foundation for the ISO/IEC 21836:2020 standard. Another stakeholder highlighted an issue with the commercial implication that only "titanium" power supply units (PSUs) are compliant, while platinum/gold PSUs do not meet the requirements. The feedback suggested that the threshold for compliance is unreasonably high.

The key concern to enable appropriate energy efficiency policies for servers is to understand their utilisation rates. As servers are always "on", the utilisation is

needed to determine if policy action is required on the active efficiency side, or idle performance, where the device is “on” but not delivering any work at a specific point in time. Utilisation levels are expected to vary across the industry, notably depending on the application: if the server is operating in a hypercloud environment, it is likely to be running load optimization features which increase the utilisation. This figure is expected to be lower enterprise IT delivered in house without such features, whether hosted in a colocation or in a distributed IT environment. It is however difficult to find utilization figures for servers. Current estimates for the average utilisation rates of servers are quite low, with the Uptime Institute Intelligence survey results showing that “at least 40% of servers operate at <30% utilisation”¹. This low rate is generally justified by operators due to an abundance of caution to ensure that there is capacity to respond to peak demand times. Statements from IBM correlate this figure stating that the average rate of server utilisation is of 12-18% capacity.² Using a normal bell curve time distribution for a utilisation level of 12-18% would imply that the average server is in idle mode between 10-25% of the time.

Data requested on the utilisation rates for servers

At this point the study team had limited access to data relating to the utilisation rate of servers. This is expected to be different averages depending on the server application. Input from stakeholders is welcome.

With regards to the required levels, stakeholders mentioned that the 2019 evaluation of the ENERGY STAR database revealed that most servers met the active mode requirements of the previous ENERGY STAR version. However, there is a need to tighten the current requirements to align with the advancements in technology. The original requirements were based on a limited dataset and might not be sufficiently ambitious. The industry has consistently shown improvements in server energy efficiency. Additionally, a significant majority of servers in the market surpass the existing active state efficiency requirements of regulation 2019/424. One stakeholder suggested that imposing regulatory requirements on the active state efficiency for specific server types (1-socket, 2-socket, blade, and multi-node) at fixed levels (9.0, 9.5, and 8.0) may pose difficulties and may not yield optimal market behaviour.

The current Ecodesign Active efficiency requirement is for servers to be tested (and information provided), either as an individual product model configuration, or if part of a server product family, for only the low-end and high-end performance configurations to be tested and declared. However, the low-end and high-end performance configurations do not represent the maximum and minimum performance from an energy efficiency perspective. As there is a wide range of potential configurations within a product family, this means that testing only low-end and high-end performance products may mean some configurations in the family are allowed onto the market yet have low active efficiency values. To accommodate for this variability and provide more representative information to buyers, the Energy Star programme provides information of the “typical performance configuration”, which is defined as “*A product configuration that lies between the Low-end Performance and High-end Performance configurations and is representative of a deployed product with high volume sales.*” Each manufacturer must define which is the typical performance configuration. Providing a score for typical server configuration therefore provides a much closer representation of performance of the

¹ Transactions per megawatt-hours: Keys to increasing data centre efficiency, Uptime intelligence, 27 June 2023

² [Are Your Data Centers Keeping You From Sustainability? - IBM Blog](#)

final purchased product and hence be used for reference of the active efficiency minimum energy performance standard.

Furthermore, it was suggested to consider the Green Public Procurement minimum requirements in **Error! Reference source not found.** for active state efficiency of servers.

Table 2.1 Proposed Active state efficiency for servers in the EU GPP requirements³

Product type	Minimum Eff _{Active}
1 socket	
Rack	13.0
Tower	11.0
2 sockets	
Rack	18.0
Tower	12.0
Blade or multi-node	20.0
4 sockets	
Rack	16.0
Blade or multi-node	9.6

Regarding the maximum idle state requirements, stakeholders strongly opposed using maximum idle state power as a regulatory tool for servers due to concerns about counterproductive behaviours in data centres. Requiring idle power limits may lead to the deployment of lower idle power systems that consume more energy to complete workloads, however stakeholders have yet to provide evidence of this trend being realised, such as how other markets where idle requirements are not in place having higher active efficiency averages. Stakeholders advocated utilising the deployed power assessment to validate server energy efficiency. The active efficiency metric in SPEC SERT already incentivises energy efficiency improvements across different utilisation levels, making separate regulation of idle power unnecessary. They appreciate the holistic approach taken in the current revision and suggested reviewing the case studies in the referenced white papers. While they recommended against focusing on idle power, they are open to collaborating if it remains a focal point.

Another stakeholder view stated that the active mode alone is not sufficient for evaluating server efficiency since high-performance systems may require more electrical energy and can consume more energy when underutilised. To enhance the assessment, the SERT-2 method, which slightly favours high-performance systems, should be supplemented with an indication of average electrical energy


³ EU green public procurement criteria for data centres, server rooms and cloud services, 2020, [https://ec.europa.eu/environment/gpp/pdf/20032020_EU_GPP_criteria_for_data_centres_server_rooms_and%20cloud_services_SWD_\(2020\)_55_final.pdf](https://ec.europa.eu/environment/gpp/pdf/20032020_EU_GPP_criteria_for_data_centres_server_rooms_and%20cloud_services_SWD_(2020)_55_final.pdf)

consumption. They also suggested that a comprehensive analysis and revision of idle requirements, including the reduction of current adders for components like RAM and networking, should be conducted.

To clarify these diverging stakeholder views, the following sections describe how idle is tied to SERT and having additional idle only metrics is understood to penalise configurations with more features, which also use more energy (but typically also bring more performance or needed functionality).

Stakeholders have indicated that the maximum idle consumption requirement incentivises server designers to compromise performance enhancements in order to reduce idle power. This could have a negative impact on the EU's data centre energy consumption in aggregate as more servers would be needed to meet performance requirements of large workloads. Idle power in ICT products is gradually rising as seen in the SERT database, not shrinking, (as outlined further in the Task 3 report) but the increased performance gained by these generational improvements vastly outweigh small idle energy increases, resulting in very large efficiency gains (work/watt).

Figure 2.1 Global trends in internet traffic, data centres workloads and data centre energy use, 2010-2020⁴



Datacentre workloads have followed a similar curve of development as internet traffic from 2010 to 2020, although at a less intense rate. This amounts to a 260% increase in workloads from 2015 to 2021.

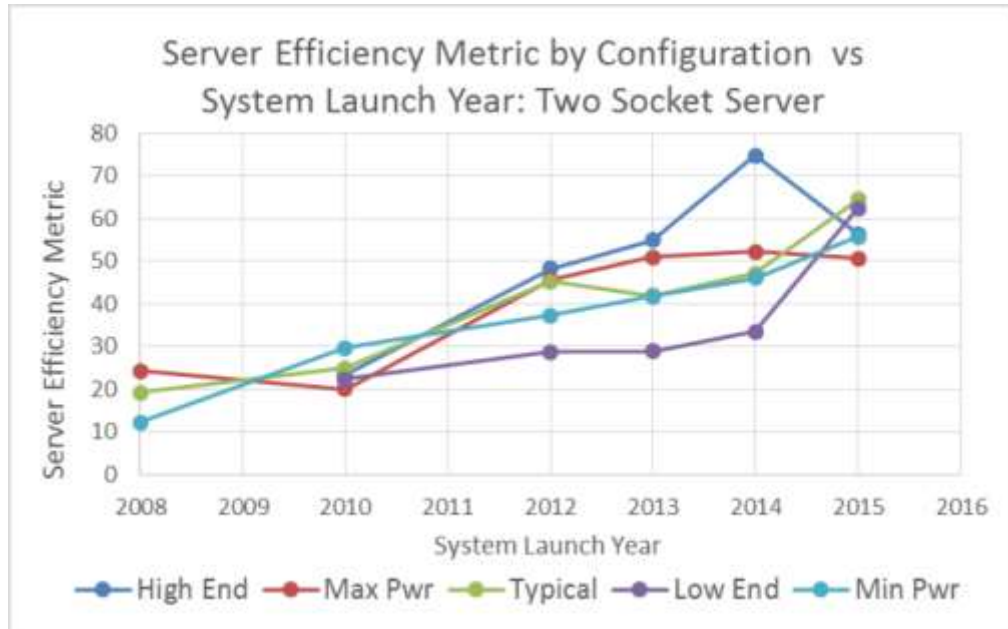
Error! Reference source not found. shows the improvement in server SERT scores from 2008 to 2015 based on the SERT SPEC database. This performance improved across entire server families, from high spec to low spec configurations.

Error! Reference source not found. evidences this point by showing the average

⁴ Nov 2021, IEA report, [Global trends in internet traffic, data centres workloads and data centre energy use, 2010-2020 – Charts – Data & Statistics - IEA](#)

SERT scores and idle consumption of 2-socket rack servers (the most common server configuration). Regulating on idle energy consumption specifically can have the unintended consequence of penalising systems with very high true efficiency (work/watt), especially when calculating energy use in a deployed power scenario in a large data room or data centre.


Figure 2.2 SERT active efficiency metric improvements with new system introductions⁵



Draft

⁵ Server energy efficiency in data centres and offices, white paper 75, The Green Grid, 2017, <https://www.thegreengrid.org/en/resources/library-and-tools/496-WP#75---Server-Energy-Efficiency-in-Data-Centres-and-Offices><https://www.thegreengrid.org/en/resources/library-and-tools/471-SERT%E2%84%A2-Active-Efficiency%3A-Demonstrating-howSERT%E2%84%A2-Active-Efficiency-Testing-Includes-Server-Idle>

Figure 2.3 Average SERT server score and the average idle consumption from 2016-2021 for 2-socket rack servers⁶



As-is, idle power is measured and reported as part of the SERT metric and should still be reported for user benefit. SERT typically provides idle power measurements in watts (W), indicating the energy consumed by the server system while in an idle state. SPEC SERT v2 testing includes a significant amount of idle time during its normal operation (roughly half the overall active test run), especially at the lower load levels, such as 25% and 12.5%. During these periods, servers go into their deepest and lowest power Core C-states often. This means that a server which has low idle power will get an improved SPEC SERT overall score as also validated by a stakeholder in the study, however quantitative evidence of the above would be needed. Evidence also lies within the SERT design document⁷. These measurements are often reported alongside other energy efficiency metrics, such as performance-per-watt or power efficiency ratings, to provide a comprehensive assessment of a server's energy efficiency characteristics.

Further information on the SERT score vs Idle Power is presented in a series of charts within Annex 1. These charts show where the majority of servers lie within each product subcategory. These data will be used to identify the relevant and updated Ecodesign requirements.

The maximum idle energy consumption requirement is currently capable of accommodating for exceptions of the highest performing server products due to the inclusion of adders within the regulation. The adders provide additional idle power allowances if the servers have extra components included. These components are: CPU performance, additional PSUs, HDD or SSD, Additional memory, additional buffered DDR channel and additional I/O devices. However, the figures for these

⁶ SPEC SERT server data set

⁷ Server energy efficiency in data centres and offices, white paper 75, The Green Grid, 2017, <https://www.thegreengrid.org/en/resources/library-and-tools/496-WP#75---Server-Energy-Efficiency-in-Data-Centres-and-Offices><https://www.thegreengrid.org/en/resources/library-and-tools/471-SERT%E2%84%A2-Active-Efficiency%3A-Demonstrating-howSERT%E2%84%A2-Active-Efficiency-Testing-Includes-Server-Idle>

adders need to be updated in line with the improvements of the technology. For a fast moving technology like servers, changing these adder requirements regularly (every few years) would be needed.

The debate around the inclusion of idle consumption requirements is often brought up due to the concerns over so called “zombie servers”. Zombie servers refer to physical or virtual servers that are running and consuming energy but are no longer actively used or serving any productive purpose. These servers often remain operational despite being outdated, obsolete, or redundant. Zombie servers can exist in data centres, cloud environments, or within an organization's own IT infrastructure. It is important to note that as the Ecodesign regulation only applies to products being placed onto the market, it will not be applicable to such legacy systems that are already in the field. The key metric to ensure in this instance, is for server utilisation to be increased, and hence avoid the “zombie servers”. The widespread adoption of virtualization and cloud computing technologies has significantly reduced the creation of new zombie servers. Virtualisation allows multiple virtual machines to run on a single physical server, enabling better utilisation of hardware resources. Cloud computing services offer scalable infrastructure, allowing organisations to dynamically provision and deprovision resources as needed. Both virtualisation and cloud computing help reduce the likelihood of having idle or underutilised servers.

2.1.4 Recommendations

We recommend using SPEC SERT as the recommended benchmark for EU Ecodesign server requirements, noting its development by the SPEC consortium and its alignment with high-quality and transparent processes.

Aligning with the stakeholder feedback, we recommend that with the comprehensive SERT database and expertise, it is possible to determine the suitable efficiency levels for each server product category currently covered within the 2019/424 regulation ICF has a large data set comprising of over 500 models since 2019 which can directly support an increase in active state efficiency requirements across all currently covered server categories based on the existing SERT metric.

The current regulation sets out aggressive internal power supply efficiency requirements for servers, to the extent that the implementation of the 80 Plus efficiency titanium levels for the internal power supplies was postponed until January 2023. This attests to the fact that there is no room for additional stringency at this time.

With regards to the maximum idle state power requirements, we recommend exploring further regulatory approaches alternative to the currently applicable idle power requirements, as it would seem – as reported so far by a few stakeholders without significant data in support -, that maximum idle power requirements might have the potential to adversely affect overall energy efficiency and data centre consumption. Instead, it could be analysed if a tighter application of the SERT Active efficiency score could provide more energy efficiency savings. However, if the idle efficiency requirement is kept, then it would require for the adder requirement criteria to be updated as the market improves. As servers are a fast-moving technology, it is recommended not to take this approach as if there were ever a delay (of a few years) in updating the regulation, the best performing products would be kept off the market.

As another potential regulatory approach to be explored, in lieu of a continued focus on idle power, the regulation could instead focus on ways to incentivise the

increased utilisation of existing and new servers, as far greater savings potential exists there but needs unlocking. The systems reporting real time utilisation enable users to evaluate the energy consumption of server systems during periods of inactivity, which is an essential aspect of overall energy efficiency. This information can help IT professionals and data centre operators make informed decisions regarding server hardware, configurations, and power management strategies to optimise energy efficiency and reduce operational costs. We recommend that all systems possess the ability to report out their utilisation in real time. Obviously, it should not be forgotten that such an approach – differently from ‘hard’ requirements on idle power - would be strongly dependent on the user/user habits, and not verifiable at the moment of the placing of the product on the market.

2.1.5 Link to Phase 2

The Task 1 report will introduce the SERT SPEC test standard.

Review the energy consumption, efficiency and idle rates of average servers in Task 4 of the preparatory study. Include these findings into the model values of Task 5. Propose and develop the above recommendations into Task 6 to be modelled in Task 7.

2.2 Regulation definitions and scope

c) to update the definitions or the scope of the Regulation;
e) to exempt server appliances, large servers, fully fault tolerant servers and network servers from the scope of the regulation;
f) to exclude resilient servers, High Performance Computing (HPC) servers and servers with integrated APA from the Ecodesign requirements set out in Annex II point 2.1 and point 2.2 of Regulation 2019/424.

2.2.1 Background

The 2019/424 Ecodesign regulation currently applies to servers and data storage products. These are defined as follows in the regulation:

(1) ‘server’ means a computing product that provides services and manages networked resources for client devices, such as desktop computers, notebook computers, desktop thin clients, internet protocol telephones, smartphones, tablets, tele-communication, automated systems or other servers, primarily accessed via network connections, and not through direct user input devices, such as a keyboard or a mouse and with the following characteristics:

(a) it is designed to support server operating systems (OS) and/or hypervisors, and targeted to run user-installed enterprise applications;

(b) it supports error-correcting code and/or buffered memory (including both buffered dual in-line memory modules and buffered on board configurations);

(c) all processors have access to shared system memory and are independently visible to a single OS or hypervisor;

(10) ‘data storage product’ means a fully-functional storage system that supplies data storage services to clients and devices attached directly or through a network. Components and subsystems that are an integral part of the data storage product architecture (e.g., to provide internal communications between controllers and disks) are considered to be part of the data storage product. In contrast, components that

are normally associated with a storage environment at the data centre level (e.g. devices required for operation of an external storage area network) are not considered to be part of the data storage product. A data storage product may be composed of integrated storage controllers, data storage devices, embedded network elements, software, and other devices;

Server appliances, large servers, fully fault tolerant servers and network servers were exempt from the scope of the regulation.

Resilient servers, High Performance Computing (HPC) servers and servers with integrated APA were exempt from the energy efficiency requirements of active efficiency and idle state energy consumption, which are set out in Annex II point 2.1 and point 2.2 of Regulation 2019/424.

The scope of the initial preparatory study in 2014 also included network equipment. However, due to the complexity and variability of the network equipment, establishing a regulation that would define and set standards for these devices was deemed too challenging at the time. Nonetheless, there is interest to review if there may be simple measures which could be applied to improve sustainability of network equipment devices, notably with regards to material efficiency requirements.

2.2.2 Questions

1. Do you have any views or concerns with regards to the definitions and scope set out in the regulation?
2. What are your views on the exclusion of server appliances from the scope of the regulation?
3. What are your views on excluding the server types listed below from the scope of the regulation?
 - a. large servers
 - b. fully fault tolerant servers
 - c. network servers
 - d. hyperconverged servers
4. What are your views on the exclusion of the following devices from the idle state and active efficiency requirements?
 - a. network servers
 - b. resilient servers
 - c. high performance computing servers
 - d. servers with integrated APA

2.2.3 Feedback/ Research results

2.2.3.1 Review of definitions

Regarding the scope, most stakeholders agreed with the need to update the regulation with definitions from ENERGY STAR, EPEAT, ISO/IEC 21836:2020 and ETSI EN 303 470. In particular, an update should be considered for the definitions of resiliency under the resilient server definition and the High-Performance Computing (HPC) servers as these do not align with ENERGY STAR. Not only should these be updated, but another stakeholder indicated that with the many definitions and exemptions, it is difficult for market surveillance authorities to regulate the product.

Storage heavy servers definition should be added to test servers with more storage capacity than normal with SERT as they cannot be tested using SNIA's Emerald tool. However, they can be tested under SERT, hence their definition as servers will

allow for appropriate testing taxonomy. It should be noted however that this product is very niche and limited, for example, HP and Dell only have a couple models each of this product type.

Data storage definitions should be updated to follow the SNIA Emerald Taxonomy to allow for harmonisation. The SNIA Emerald Storage Taxonomy defines a market taxonomy that classifies storage products or subsystems in terms of operational profile and supported features. It is structured under the 3-level hierarchy of Set, Category and classification. The main structures which need definition are the Disk Set Online and NVSS (Non-Volatile Solid State) Set Disk Access Categories as they are the dominant Data Storage technologies found in Data Centres.

2.2.3.2 Exemptions from the regulation

Regarding exemptions, multiple stakeholders stated that exclusions should only be granted for individual requirements that are not technically feasible or reasonable to deliver for a server type.

Stakeholders agree that some exemptions are required as not all products can be tested under SERT. However, Ecodesign also brings out other requirements, such as to provide information requirements, material efficiency and PSU efficiency criteria.

For information requirements, this can be applied to all servers. However, as noted above, some servers cannot be tested under SERT. Therefore, the regulation could be applied for all servers to provide information requirements, with exemptions made to information requirements linked to SERT if not applicable. We will detail these exemptions to SERT in the following section.

For material efficiency requirements, there are in principle few limiting factors to justify exempting these from servers. However, an exemption may be required for large servers (mainframes), as these have different and specialised structures to commodity servers, which may be difficult to regulate.

For PSU efficiency criteria, in principle there are no concerns with applying this criterion to all servers.

Current exemptions

The current regulation exempts large servers, fully fault tolerant servers, server appliances and network servers.

As described above, large servers should continue to be exempt from the regulation, not only as they are out of scope from SPECT SERT, EPEAT and IEC 21836:2020, but also due to their complexity compared to commodity servers, which would make them difficult to regulate, even for non-energy criteria.

Fully fault tolerant servers are out of scope for both SPECT SERT and IEC 21836:2020, however, there are no other concerns which would justify their exclusion from the remaining scope of the regulation.

Server appliances justify their continued exclusion from energy requirements as they have preinstalled, customised characteristics, where energy efficiency requirements would likely not lead to increased efficiency, however, there is no justification to exclude them from the rest of the regulation requirements (such as for material efficiency and information provision).

Network servers are currently not tested under SERT. There is some debate within industry on whether these should be included within the regulation as some would justify, they are more suited to be placed under the Large Network Equipment (LNE) test method developed by the Alliance for Telecommunications Industry Solutions (ATIS). One could therefore consider their inclusion into the server regulation under only the non-energy criteria, or for them to be included in a future regulation covering networking equipment instead. A note is made that these servers are a part of LNE which represents 10-15% of IT equipment energy use for an average data centre, which can be assumed to be low priority compared to servers.

Stakeholders also suggested exemptions should be extended to custom made servers be kept out of the scope as they are not tested under SERT efficiency metrics and they are not sold outside of the manufacturer. The regulation currently exempts these servers from the information provision requirements. A definition for custom made servers may be required due to the variability of the server family configurations.

For data storage products, “small” and “large” data storage products are exempt from the regulation. These devices are defined as follows:

Small data storage product is defined as “a data storage product containing a maximum of three data storage devices”. Due to their small size, these devices are unlikely to have online features or contain basic reliability, availability and serviceability features, making their architecture difficult to consistently and comparably test for performance.

Large data storage product is defined as a high end or mainframe data storage product that supports more than 400 data storage devices in its maximum configuration and with the following required attributes: no single point of failure, non-disruptive serviceability and integrated storage controller. These are excluded from scope as due to their size, as the testing of these systems can be in the hundreds of thousands of Euros to implement, making it excessive for manufacturers to test.

2.2.3.3 Exemptions specifically for energy efficiency requirements

Current energy efficiency exemptions

Currently the regulation exempts resilient servers, High Performing Computing (HPC) servers and servers with integrated APA from the idle state power and active state efficiency metrics.

Regarding resilient servers, these are covered under SERT and are included within Energy Star version 4 criteria. There is therefore no clear reason to award them this exemption. There may be a need for a separate requirement to be made to accommodate for their higher uptime availability.

HPC servers are currently not included in the testing scope for SERT. However, SPEC is actively working for the next version of SERT to include these and is likely to be available in a few years time. Therefore, it is not possible to set active efficiency requirements until the testing standard is out. Focus should however be kept to regulate these products in the future as it is expected that the demand for these servers will increase as they are used in AI processes.

Servers with integrated APAs are still waiting for SERT to develop a new test standard, and therefore should be maintained in this exemption. However, it is noted that the technology is developing quickly, such newer CPU designs with accelerator

blocks may be confused under the current server with integrated APA definition. This definition therefore needs to be updated. The Energy Star Version 4 definition can be taken as a first definition update.

Suggestions for inclusion/exemption

There was also stakeholder feedback that servers utilising liquid cooling solutions should be explicitly in the scope of the regulation to incentivise the usage of more efficient cooling systems and heat recovery. These products are currently not defined under the regulation and should therefore be better defined as in scope of the regulation. It should be noted that they are currently not testable under SERT, nor is there any datasets on their performance to set a metric for. However, it is noted that these are usually more efficient. This topic is further explored in section 2.11.

An additional definition and exemption requirement may be required in the regulation for hyperconverged servers. Hyperconverged servers are defined as: *a highly integrated server which contains the additional features of large network equipment and storage products*. These products use virtualisation techniques to combine features of a server, a storage product and network switch, which make them difficult to test for performance features and for energy efficiency in the same manner as other servers. They are also justified as having a small market presence and are not general-purpose devices. These products are currently not tested under the SERT tool and therefore are difficult to incorporate into active and idle state efficiency metrics. However, they can be considered for the remainder of the regulation standards such as for product information requirements, PSU and material efficiency.

2.2.4 Recommendations

To ensure clarity and harmonisation, it is recommended to align product definitions with Energy Star, EPEAT, ISO/IEC 21836:2020 and ETSO EN 303 470. In particular, an update should be considered for the definitions of resiliency under the resilient server definition and the High-Performance Computing (HPC) servers as these do not align with ENERGY STAR. This aligns with the stakeholder feedback. Furthermore, definitions specifically for resilient servers' recovery section and HPC servers should be updated to align with the latest industry standard. We also recommend the inclusion of the following definitions:

Storage Heavy Server (SHS): A computer server with greater storage capacity than a standard computer server. As shipped, these computer servers support 30 or more internal storage devices. These servers differ from Storage Products in that they run computer server operating systems and software stacks.

Hyperconverged Server: Hyperconverged Server: A highly integrated server which contains the additional features of large network equipment and storage products.

Data storage definitions should also be updated in line with SNIA Emerald taxonomy.

With regards to server appliances, these are out of scope of both SPEC SERT tool and ISO/IEC 21836:2020 standard which makes their inclusion into energy efficiency criteria difficult. Our recommendation is to include server appliances into the regulation, with specific exclusions for energy efficiency requirements set out in Annex II point 2.1 and point 2.2 of Regulation 2019/424.

For large servers, we recommend for their exclusion from the regulation to be maintained.

We recommend the inclusion of fully fault tolerant servers, and hyperconverged servers into the regulation, with an exemption granted for energy efficiency requirements set out in Annex II point 2.1 and point 2.2, where these products are not yet in the test standard scope.

For custom servers, we recommend creating a definition within the regulation in order to ensure that server configurations are not inappropriately deemed as “custom made”. The regulation has an exemption for the provision of energy efficiency information for custom made servers, the same exemption should be extended for energy efficiency criteria as these devices are not in scope for the SERT testing standard. The definition for custom made servers needs to be careful not to allow for market loopholes.

As networks servers may be better defined as large network equipment, they are not in scope for the SERT test standard. They can therefore not be included in the active efficiency and idle score metrics, and that exemption should be maintained. We recommend investigating if network servers can be included into a regulation for network equipment in order for them to be measured under the ATIS test method.

As resilient servers are included in the scope of SPEC SERT, we recommend removing their exemption from the energy efficiency criteria. Research can be done to consider accommodations with regards to their higher uptime availability.

We recommend maintaining the exclusions on energy efficiency metrics (from annex II point 2.1 and point 2.2) granted to High Performance Computing (HPC) servers and servers with integrated APA as these devices are out of scope of the testing standards SPEC SERT tool and the ISO/IEC 21836.2020. The definition of servers with integrated APAs should be updated for clarity.

SPEC has shared that SERT V3 will be developed over the next two years, which should address HPC servers and those with many types of integrated APAs. The commission could look to include HPC servers and servers with integrated APAs into the regulation after SERT has updated the standard.

Recommendation for liquid cooling servers are expanded upon in section 2.11.

2.2.5 Link to Phase 2

The feedback and recommendations set out above should be included in the Phase 2 Task 1 report which determines the scope of the regulation. This provides the update of how the regulation, and the study should be scoped.

2.3 Data Storage devices performance requirements

- i) *to set specific Ecodesign requirements on the efficiency, performance and power demand of data storage products.*

2.3.1 Background

Under the Ecodesign regulation, ‘data storage product’ means a fully functional storage system that supplies data storage services to clients and devices attached directly or through a network. Components and subsystems that are an integral part of the data storage product architecture (e.g., to provide internal communications between controllers and disks) are considered to be part of the data storage

product. In contrast, components that are normally associated with a storage environment at the data centre level (e.g., devices required for operation of an external storage area network) are not considered to be part of the data storage product. A data storage product may be composed of integrated storage controllers, data storage devices, embedded network elements, software, and other devices.

2.3.2 Questions

1. What are your thoughts on the regulation setting out energy efficiency requirements for data storage products? Would the SNIA Emerald Power Efficiency Measurement method laid down in ISO/IEC 24091:2019 be an appropriate metric for this energy requirement?
2. How could the use of *capacity optimising methods* be increased in data storage devices?

2.3.3 Feedback/ Research results

Regarding setting a regulation for energy efficiency of data storage products, stakeholders stated that SNIA Emerald provides a reproducible and standardised evaluation of energy efficiency for commercial storage products, specifically SAN and NAS systems in active and idle states. However, testing a significant proportion of storage configurations may not be practical. The type and selection of data storage devices can heavily influence energy consumption and performance in storage products. Providing information on these products, which are simpler to test and have a significant impact, could be more useful to users. Setting active energy efficiency requirements for data storage products in the regulation is supported, and the proposed measurement method aligns with the Energy Star specification for servers and data storage products.

Reports on active state requirements on Energy Star for storage products are as follows:

Table 2.2 Active state requirements for Block I/O Storage products on Energy Star⁸

Workload type	Specific workload test	Minimum performance/Watt Ratio	Applicable Units of Ratio
Transaction	Hot Band	28.0	IOPS/watt
Streaming	Sequential Read	2.3	MiBS/Watt
Streaming	Sequential Write	1.5	MiBS/watt

These requirements can provide a metric to measure the performance of storage products. However, as the metric is not used apart from Energy Star, there is little available data on the active state performance of data storage products at this time.

Stakeholders stated that the use of Capacity Optimisation Methods (COMs) can greatly influence the amount of data storage needed in a data centre. These include:

- 1) **Thin Provisioning:** A technology that allocates the physical capacity of a volume or file system as applications write data, rather than allocating all the physical capacity at the time of provisioning.

⁸ [ENERGY STAR Data Center Storage Version 2.1 Final Specification](#)

2) **Data Deduplication**: The replacement of multiple copies of data – at variable levels of granularity – with references to a shared copy in order to save storage space and/or bandwidth.

3) **Compression**: The process of encoding data to reduce its size. For the purpose of this specification, only lossless compression (i.e., compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed exactly) is recognized.

4) **Delta Snapshots**: A type of point-in-time copy that preserves the state of data at an instant in time by storing only those blocks that are different from an already existing full copy of the data.

Energy Star requires some of these features be included for certain storage systems.⁹

Feedback requested on data storage products performance

The study team encourages stakeholders to provide data on data storage products performance following the active state metrics set out by Energy Star set out in Table 2.2 and the availability of Capacity Optimisation Methods on the market.

2.3.4 Recommendations

We recommend assessing the potential introduction of SNIA and/or standardised energy and performance reporting requirements specifically for data storage products. ISO/IEC 24091:2019 test methods offer information on relative efficiency under various workloads. The SNIA Emerald benchmark provides valuable data on energy efficiency in storage systems. This can inform the creation of new active levels and can be designed to match the workload requirements of storage devices more accurately, allowing for more efficient power management. Given the smaller size of the storage market and limited number of unique models, care should be taken not to create requirements as aggressive as may be developed for servers which are far greater in number and variety. One can consider an energy efficiency requirement but data is required to support setting of these.

The Ecodesign PSU requirement applies to both servers and data storage products. It is stated above that for servers, the PSU maximum efficiency has been reached. We recommend keeping the power supply efficiency requirements for storage products as they are to align with servers.

Energy consumption of storage products in data centres is primarily driven by the number of physical storage devices present (HDDs, SSDs, etc.), which can be reduced by using higher capacity devices and implementing capacity optimisation methods. It is recommended to require the availability of capacity optimisation methods data, such as Thin provisioning, data deduplication, compression and delta snapshots.

Additionally, the Commission should provide users/operators with information on the benefits of utilising these methods. Educating consumers about the benefits of energy-efficient storage products and providing guidance on selecting and using such products can drive demand for energy-efficient options.

⁹ [ENERGY STAR Data Center Storage Version 2.1 Final Specification](#)

2.3.5 Link to Phase 2

Phase 2 Task 1 report will set out the test standards for data storage products.

Task 4 will set out the average and Best Available Technologies for data storage products.

2.4 Processor Power Management Function

g) to set specific Ecodesign requirements on the Processor Power Management Function of servers.

2.4.1 Background

Dynamic Voltage and Frequency Scaling (DVFS) is a technique used to reduce energy consumption in electronic devices by adjusting the operating voltage and frequency. DVFS allows for the dynamic adjustment of the supply voltage and clock frequency of a device based on its workload or performance requirements. By lowering the voltage and frequency during periods of low activity or idle states, energy consumption can be significantly reduced. This approach helps to optimise the power-performance trade-off and minimise unnecessary energy usage. There are currently no processor power management function requirements in Ecodesign. However, there is interest to review this technique as by reducing voltage and frequency, the energy consumption and associated carbon footprint of electronic devices can be lowered.

2.4.2 Questions

1. What are your views on Ecodesign requirements on the Processor Power Management Function of servers?
2. What are your views on reducing voltage/and/or frequency through dynamic voltage and frequency scaling as a power management function?
3. What are your views on enabling processor or core reduced power states (C-states) when core or socket is not in use?

2.4.3 Feedback/ Research results

Regarding Ecodesign requirements on the processor power management function of servers, the feedback suggested that servers should have processor power management enabled by default and at the time of shipment, with all processors capable of reducing energy consumption during low utilisation. SERT should account for power management and be tested in the default shipped state. Specific technology requirements addressed by SERT and active model efficiency requirements may not be necessary. New product parameters should be reliable, accurate, and reproducible, aligned with harmonised standards. It is important to align different initiatives, such as the EU Energy Star Label and Code of Conduct for Energy Efficiency of Data Centres, to avoid fragmented or overlapping policies.

With regards to reducing voltage and/ or frequency through dynamic voltage and frequency scaling a stakeholder suggested that dynamic voltage or frequency scaling can effectively reduce energy consumption in servers without impacting system performance. They recommended choosing one option (reducing voltage or frequency) instead of both to avoid diminishing server resilience and increasing implementation complexity and cost. While P-states provide energy savings, it

should be acknowledged that they introduce latency that may not be acceptable for certain end customers.

Regarding enabling processor or core reduced power states when core or socket is not in use, the stakeholder feedback highlighted that dynamic voltage and frequency scaling (DVFS) and activating processors or cores with reduced power when not in use have been established practices in the ENERGY STAR program. It is acknowledged that while C-states can save energy, they also introduce latency that may be deemed unacceptable by specific end customers.

2.4.4 Recommendations

The Commission could look to add new requirements using reliable, accurate, and reproducible methods and aligned with harmonised standards. We suggest considering incorporating the methods outlined in point 11 and 12 of the US ENERGY STAR specification for computer servers version 4 into an Ecodesign regulation.

We suggest implementing dynamic voltage or frequency scaling to reduce intrusive power management and achieve energy savings in IT devices. The Commission must consider the specific needs and requirements of colocation data centre operators who do not own the IT equipment when implementing power management solutions.

The power management features like dynamic voltage and frequency scaling provide value at all utilization levels.

We suggest continuing the implementation of DVFS and activating reduced power states for idle cores or sockets, as these practices have proven effective in the ENERGY STAR program. Despite concerns about latency, we recommend that the Commission consider the requirement of having power management features not only available but enabled as shipped as it is advisable to require the availability of P-states and C-states in order to realise energy savings, while recognising that some customers may consider the resulting latency to be unacceptable. There are various levels of power management tunability available. A manufacturer could tune basic power management with no real impact on workload or more aggressive which may impact some workloads due to increase latency. This will vary from manufacturer to manufacturer but to go into deeper levels of power management, it is understood that it is usually up to the user to configure.

The Commission should consider the requirement of having power management features not only available but enabled as shipped.

2.4.5 Link to Phase 2

Phase 2 Task 4 will include a section to detail the above information on processor power management. Task 6 can then look to include processor power management requirements which can then be incorporated into Task 7 modelling.

2.5 Standby-Readiness for Servers

p) The technical and economic feasibility and relevance of product specific requirements on the standby readiness of servers (for instance allowing to move to and from idle mode in a fast and seamless manner), if not covered by the analysis on the Processor Power Management Function.

2.5.1 Background

There has been interest in servers to gaining the ability to enter deep sleep state and recover from them in a rapid fashion to give data centre operators the option to power down groups of servers in a behaviour similar to “core parking” at the CPU level. Up till now, this functionality has not been available on the market.

2.5.2 Questions

1. Do you have any insight into the feasibility of product specific requirements on standby readiness of servers (for instance allowing to move to and from idle mode in a fast and seamless manner)?
2. Would these be economically feasible?
3. What benefits could standby readiness unlock?

2.5.3 Feedback/ Research results

Stakeholders expressed that a distinction should be made between standby readiness that requires a server to maintain the state of the system (stored in RAM) vs. a stateless example of a very fast boot after fully powering down. Stakeholders cautioned against any MEPS that dictate new technologies that customers are not requesting or willing to pay for, which this topic falls under. In addition, they cautioned that this type of behaviour can negatively impact reliability and increase jitter and latency which can negatively impact certain workload types within the data centre more broadly.

Stakeholders confirmed that lower power modes below the current idle implementations with power management enabled would require extensive development and testing to assure that servers can both enter and exit the state without causing system crashes. There are also some server components that do not currently implement functionality to allow this behaviour at this time. These stakeholders suggest that putting server systems in deeper sleep states for periods of time is already possible through data centre management software.

2.5.4 Recommendations

Given this functionality is not currently available on the market and that consumer demand for this functionality appears low, we recommend that this functionality is not required at this time. If consumer demand rises to the level that manufacturers begin to develop this functionality in the future, we recommend that this topic be revisited at that time.

2.5.5 Link to Phase 2

Phase 2 Task 4 report will include the description above of standby-ready servers.

2.6 Parameters Information Requirements

- r) The technical and economic feasibility and relevance of product specific requirements related to:*
- a. the availability of information (temperature, (fan) speed, etc..) for open data exchange about the input/output air flow data of the server/data storage product, and/or.*

b. the capability to enable external overriding of the internal fan speed control, in view of potential synchronisation of the product cooling system with the data centre cooling system.

2.6.1 Background

As electronic equipment, it is important for servers to operate within a dedicated temperature range. Servers generate waste heat during operation. Therefore, each server is typically equipped with temperature sensors, passive cooling systems (metal heat exchangers distributing heat away from the CPU), and active cooling systems (fans to increase air flow rate on the device). When a server is operating at too high a temperature, the fan will engage in order to help cool the system.

On a wider perspective, servers are generally operated in data centres, with many servers in one place. Having these operating together can create a lot of heat, which is why data centres have dedicated systems to cool the rooms down.

Parameter information requirements are of interest to regulation as providing this information or capability may result in a more efficient overall system operation between the servers and data centre.

2.6.2 Questions

1. What are your views on the technical and economic feasibility of:
 - a. Providing open data exchange about the input/output air flow data of server/data storage products?
 - b. The capability to enable external overriding of internal fan speed control for synchronisation with data centre system cooling?

2.6.3 Feedback/ Research results

2.6.3.1 Providing open data exchange

Servers currently measure the inlet air temperature and have metrics to activate the fan. The actual airflow is estimated from the fan speed, rather than measured. Exit air temperature is not typically measured during operation and would require additional development.

There is general agreement between stakeholders and research that the sharing of temperature and fan information is useful for data centre operators to manage their cooling facilities. Energy Star included this requirement a decade ago for this information to be shared, hence the Commission could consider aligning with Energy Star.

2.6.3.2 Capability to enable external overriding of internal fan speed control

Stakeholders agree that the capability of external overriding of internal fan speed control would result in uncertainty for who would take liability for product failure: the product manufacturer or the data centre operator. Therefore, if internal fans were being operated by an external function than the server, the server warranty would be void.

Our assessment agrees with these findings. Furthermore, it is unclear what would be the benefits of this measure from an energy efficiency perspective, as servers

are already optimised to use fan cooling to the strict minimum in order to achieve a higher active efficiency score.

2.6.4 Recommendations

We recommend considering aligning with the Energy Star requirement to ensure that servers have an open data exchange of the temperature and fan activity data. The outputs should be open source and usable by all types of energy and data centre infrastructure management software packages.

We recommend not pursuing the investigation of external overriding of server fan due to the legal risk and a lack of evidence to potential environmental benefit.

2.6.5 Link to Phase 2

Discuss in Task 6 the inclusion of open exchange of server temperature and fan speed. The measure can then be considered for modelling impact in Task 7.

Discuss in Task 6 the capability for external fan override, present the risks and conclude why the measure was not taken further.

2.7 Energy Label

s) The technical and economic feasibility and relevance of introducing an energy label for servers and data storage products, including a label and a detailed product information sheet comprising targeted indicators for the different possible uses of the product (e.g., as webserver, disk server, database server, file/disk server, etc.)

2.7.1 Background

Energy-related product policy has traditionally consisted of two parts: a minimum energy performance requirement to remove the worst performing products (MEPS); and an Energy Labelling regulation to inform and incentivise buyers towards best-in-class products. The Ecodesign 2019/424 regulation for servers is currently a minimum energy efficiency performance requirement and has no associated energy labelling component. Labelling has traditionally been more effective for domestic consumers, whereas non-domestic consumers are expected to be better informed about their purchasing and hence labels are considered less effective (with a few exceptions, such as professional refrigerators, refrigerators with a direct sales function)¹⁰. This section investigates the arguments for and against the use of an energy label for computer servers and data storage products.

The section also investigates if there should be mandatory product information requirements to inform buyers on products available on the market and compare like for like criteria.

2.7.2 Questions

1. With regards to the introduction of an energy label on servers and data storage products:
 - a. What are the feasibility considerations for their implementation?

¹⁰ [Energy efficient products \(europa.eu\)](https://ec.europa.eu/euro-observatory/energy-efficient-products/)

- b. What Information would you recommend be included in such a label?
 - c. Would It be helpful for such a label to include values such as idle energy consumption or SERT workload scores on CPU, memory and storage?
 - d. Should icons be included in the label to signify servers designed to deliver specific tasks?
2. With regards to the introduction of a detailed Product Information Sheet on servers and data storage products:
- a. What are the feasibility considerations for their implementation?
 - b. What Information would you recommend be included in such an information sheet?
 - c. Should detail such as the SERT worklet scores be included?

2.7.3 Feedback/ Research results

2.7.3.1 With regards to a label requirement

A few stakeholders set out arguments that the creation of a label for servers would not be helpful as the information provided by SERT is only tested at low and high-performance configurations and not the specific device configuration. These scores do not reflect the performance of the individual configurations sold. One stakeholder developed on this idea, by stating that one could not create a label as not only are servers highly customisable but are also used for many different applications which cannot be modelled for. Furthermore, the stakeholder stated that a label would not be useful as purchases are made using Request for Proposal (RFP) processes, which accounts for energy consumption details. There is also a comment on how creating a label for custom servers would be overly burdensome.

However, most stakeholder feedback took the opposite view that the creation of a label would facilitate customer procurement processes. They argue that procurement teams purchasing servers do not have IT expertise, and hence a simplified method for them to distinguish energy performance is needed.

Suggestions of inclusion are the idle energy consumption, SERT scores, but more basic information should also be provided (such as a rating score, or even potentially icons). The ideal mentioned by stakeholders would be to have a simple grading system giving servers a label class from A to F. Following that framework, a stakeholder proposed a label grading on the SERT SSJ worklet score, which covers both process and RAM activities. They indicated that some FTSE500 companies already use this metric for decision making for their server purchases. The concern noted in this methodology is that it does not provide a good metric for the variety of server workload applications.

Some stakeholders mentioned that labels targeted towards data centre operators could be useful, such as to include the ASHRAE operating conditions and useful lifetime of the asset. One suggested that a label score could be created with capabilities to operate at higher temperatures for waste heat recovery. This note could be problematic as it may lead to improved PUEs and heat recuperation but may perversely incentivise energy wastage at the server level.

2.7.3.2 With regards to information requirements on a detailed product information sheet

Stakeholders agreed that providing energy performance for a server product in an information sheet would be helpful. However, due to the variance within a server family, providing only the minimum and maximum performance scores is not helpful. Indeed, the minimum and maximum performance configurations do not represent the minimum and maximum energy efficiency set ups. A stakeholder indicated that variance within the family could mean that other servers within a median configuration could be up to 70% less energy efficient than the min/max performance configurations may indicate. To accommodate for this variability and provide more representative information to buyers, the Energy Star programme provides information of the “typical performance configuration”, which is defined as “A product configuration that lies between the Low-end Performance and High-end Performance configurations and is representative of a deployed product with high volume sales.” Each manufacturer must define which is the typical performance configuration. Providing a score for typical server configuration therefore provides a much closer representation of performance of the final purchased product. If not applied for MEPS, the typical configuration rating for servers could be used to inform the buyers and gather performance data to inform future regulation.

The group Techbuyer are able to provide a rating for servers using the SSJ worklet alone, and model from there performance scores for different configurations within a server family, without the need for full scope SERT testing.¹¹ Worklet scores would therefore only be useful to be provided if they relate to the as-configured server. This could therefore be a metric to be referred back to the testing standards authority. However, we should stress that this methodology relying on SSJ worklet alone is not a comprehensive measure of server efficiency. The SSJ worklet combines CPU and memory activities but is not representative of all server activities. This is why the SERT score has been developed, as it has a comprehensive weighting of all server activities.

Stakeholders also suggested that the inclusion of hardware component compatibility in an information sheet would facilitate refurbishment activities.

Other metrics suggested by stakeholders which could be considered for inclusion are: idle state consumption, SERT workload and worklet scores, power management functions, PSU efficiency, active state efficiency, repairability scores, spare part provisions, operation condition classes, secure data deletion functionality and software update provision.

Further information on the SERT score vs Idle Power is presented in a series of charts within Annex 1. These charts show where the majority of servers lie within each product subcategory. These data will be used to identify the options with regards to an Energy Label scheme.

2.7.3.3 With regards to data storage products

Stakeholders did not provide feedback with respect to labelling or information requirements for data storage products. These products have a few different families on the market, with levels that can be difficult to compare due to nearly infinite

¹¹ The concept of relying on the SSJ metric only is simpler as it only requires testing on one worklet rather than all 12 in SERT. The SSJ worklet is the only worklet that combines CPU and memory activities. SSJ is one of the original attempts of SERT to design a test for all servers, before the current more comprehensive method was developed.

storage device variability and various systems sizes impacted by scale-out vs. scale-up architecture. This makes distinguishing the energy capabilities of storage system through a label an exercise which may not be useful for purchasers.

2.7.4 Recommendations

Recommend the inclusion of energy efficiency information provided through labelling or information sheets for procurers in the EU. Engage with procurement professionals in the EU to verify labelling benefits.

Feedback requested on procurement practices

The study team encourages stakeholders to provide insight into the usefulness of energy efficiency information for server procurement teams.

Recommend the inclusion of the “typical server configuration” as a regulation tool for server families in the EU, for MEPS setting and information requirement. For labelling, this can be included, if determined to be effective with procurement professionals. We will investigate the feasibility of a label according to current efficiencies and what metrics/ granularity can be applied in the future.

Keep the provision of ASHRAE operating conditions on the information sheet and consider if a label would be useful for data centre operators.

Investigate a framework to provide users with hardware component compatibility in the information sheet to facilitate refurbishment.

For data storage products, we recommend applying an energy efficiency requirement on the products, rather than implementing a labelling scheme.

2.7.5 Link to Phase 2

Task 2 will include a section around the procurement practices of server and data storage products.

Task 3 will include a section to detail how data centres are operated and if more ASHRAE information is required.

Review in Task 6 the definition of a “typical server configuration”, its use and feasibility onto a label for servers to then model as a scenario in Task 7.

2.8 Material Efficiency

d) to update the material efficiency requirements for servers and data storage products, including the information requirements on additional critical raw materials (tantalum, gallium, dysprosium and palladium), taking into account the needs of the recyclers;

j) on material efficiency aspects:

a. the provisions on disassemblability of certain components, also considering advancements in standards (mandate M/543) since the publication of the regulation;

b. an analysis of the benefits of the information requirements under Regulation 2019/424 already covering cobalt in the batteries and Neodymium in the hard disks.

l) an analysis of the standards, and of their relevance for regulatory purposes, developed/under development under the standardisation request M/573, 'Commission

implementing decision C (2021)14 of 12.1.2021 on a standardisation request to the European standardisation organisations in support of Regulation (EU) 2019/424 as regards ecodesign requirements for servers and data storage products¹

2.8.1 Background

2.8.1.1 Information requirements for critical raw materials

With regards to material declarations, the Ecodesign server regulation currently only has two critical materials that need to be declared: Cobalt in batteries and Neodymium in the hard drive.

2.8.1.2 Material efficiency requirements for disassemblability

The Ecodesign server regulation currently has the following requirements for disassemblability:

1.2.1. From 1 March 2020, manufacturers shall ensure that joining, fastening or sealing techniques do not prevent the disassembly for repair or reuse purposes of the following components, when present:

- (a) data storage devices;*
- (b) memory;*
- (c) processor (CPU);*
- (d) motherboard;*
- (e) expansion card/graphic card;*
- (f) PSU;*
- (g) chassis;*
- (h) batteries.*

2.8.1.3 Review of standardisation request M/573

In 2021, the Commission issued mandate M573 to the European Standardisation Organisations (CEN, CENELEC and ETSI) requesting to support Regulation (EU) 2019/424 as regards ecodesign requirements for servers and data storage products¹². In August 2021, ETSI accepted the mandate.

The requirements of the specific standards to be adopted involve:

- Measurement and calculations of the power supply unit efficiency, the power factor, and its rated power output.
- The calculation of energy efficiency measurements and metrics for servers.
- Measurement and calculation of the opening condition class.
- A way to verify the compliance of products with the requirement on the secure data deletion functionality for servers and data storage products.
- Ensuring verification of compliance with the requirements for the availability of firmware and of security updates to firmware.

¹² <https://ec.europa.eu/growth/tools-databases/mandates/index.cfm?fuseaction=search.detail&id=597>

- Verifying the compliance of a server with the requirements of the supply of information on the weight range of critical raw materials
- Ensure that servers verify the compliance with requirements on its ability to be disassembled
- A deliverable on the assessment of the efficiency, performance, and power demand of data storage products.

Since the Commission issued mandate M573 to the European Standardisation Organisations in 2021, all seven standards are yet to be published. Table 2.3 below features the complete lists of standards under the mandate progress against which is published on the ETSI website¹³.

Table 2.3 Ecodesign Requirement Standards for M573

Reference	Title
DEN/ EE-EEPS44	Energy efficiency metrics and measurements for data storage equipment.
DEN/ EE-EEPS47-5	Server and data storage product disassembly and disassembly instruction
DEN/ EE-EEPS47-4	Server and data storage product critical raw materials
DEN/ EE-EEPS47-3	Server and data storage product availability of firmware and of security updates to firmware
DEN/ EE-EEPS47-2	Server and data storage product secure data deletion functionality
DEN/ EE-EEPS47-1	General for server and data storage products
REN/ EE-EEPS42	Energy efficiency measurement methodology and metrics for servers

2.8.2 Questions

1. What are your views on material efficiency requirements for servers and data storage products?
2. Do you think these requirements should be updated, to include information requirements on additional critical raw materials (tantalum, gallium, dysprosium), taking into account the needs of the recyclers?
3. What type of recyclable material can be found in servers and data storage products?
4. What is the rate of materials recovered and recycled from servers and data storage devices?
5. What are your views on disassemblability of certain components, considering advancements in standards (mandate M/543)? In particular, would the principles outlined in Annex II, Section B, Part (5) of the [draft Ecodesign regulation](#) for smartphones be applicable for servers? These principles are:
 - a. fasteners shall be removable or reusable.

¹³ [Work Programme - EWP on the Web - Query Result \(etsi.org\)](#)

- b. the process for replacement shall be feasible in at least one of the following ways:
 - with no tool, a tool or set of tools that is supplied with the product or spare part, or basic tools, or
 - with commercially available tools.
 - c. the process for replacement shall, as a minimum, be able to be carried out in a workshop environment,
 - d. the process for replacement shall, as a minimum, be able to be carried out by a generalist.
6. What are your views of the benefits of the information requirements under Regulation 2019/424 already covering cobalt in the batteries and Neodymium in the hard disks?
 7. What are the packaging materials and quantities to be considered for server and data storage products?
 8. What are your views on the relevance of developing standards under the standardisation request M/573 to regulate the material efficiency of servers and data storage products?
 9. With regards to networking equipment (such as routers and ethernet switches), are there measures that should be considered to increase their material efficiency? If so, what are they?
 10. Standards have been requested with regards to energy efficiency, disassembly, critical raw material content, firmware availability, updates, and secure data deletion functionality. Are there additional standards you believe should be considered? What other standards are you aware of that apply to server and data storage products?

2.8.3 Feedback/ Research results

2.8.3.1 Reduced material content

Stakeholders brought forward that the latest research from Interreg NW CEDaCI suggests that there is significant material waste in enterprise servers, notably in the weight of casing, number of screws and individual designs for server fans. Stakeholders inquired if a standard could be requested to be developed to ensure products are not wasting materials in design. Similarly, a standard could be considered with regards to a CO₂ emissions declaration from manufacturing processes.

2.8.3.2 Firmware support considerations

Multiple stakeholders brought concerns forward around the availability of firmware provision and updates.

There was feedback for support period from a firmware perspective to be set at six years, and for BIOS to be released as open source after that. Furthermore, a stakeholder indicated that to increase hardware life expectancy, previous versions of firmware should be made available, and not only the latest update. This is because often the new versions have increased performance requirements, which make the hardware unable to keep up and run. Furthermore, newer firmware options may be created to function for a system with particular hardware components, therefore previous versions of firmware can be useful for refurbishers wishing to replace

specific components of a server. However, consideration should be made to ensure that cybersecurity levels are maintained whilst using the previous versions.

On a complimentary note, a stakeholder indicated that contact information from users looking to upgrade firmware on their device, should be considered as restricted data, and not automatically added to contact lists marketing and sales purposes.

It should be noted that even with firmware updates, there are concerns that the Operating System is not included in these firmware provisions. There may therefore be limitations in product life expectancy due to operating system being no longer supported. This is out of the remit of the manufacturer.

2.8.3.3 Barriers to part harvesting

Supporting information provisions around components such as PDUs, HDDs, CPUs and Memory, for their capabilities and material content would improve part reuse processes and material recovery.

The refurbishers industry stakeholders indicate that reuse of components is at times limited by technical barriers known as “parts-pairing”. This is a software serial identification system which ensures that all the components in a device are matched to the device. Therefore, if a component is replaced (with an identical one) by a third party repairer, the system will identify that component as “other” and will not function appropriately.

2.8.3.4 Availability and usability of spare parts

Stakeholders have reiterated the need for spare parts to be made available for at least five years after the cessation of product production. This would be consistent with the material efficiency requirements of other Ecodesign regulations, such as the draft Ecodesign requirements for Smartphones and Tablets.

2.8.3.5 Material efficiency requirements for disassemblability and repair

Material efficiency requirements for disassemblability and repair are designed to ensure that components can be replaced (or recuperated) to increase product life expectancy.

Material efficiency requirements for disassembly were included in the Ecodesign regulation in Annex II, 1.2.1 . This requirement sets out that: *joining, fastening or sealing techniques do not prevent the disassembly for repair or reuse purposes of the following components, when present: data storage devices, memory, CPU, motherboard, expansion card/graphic card, PSU, Chassis, Batteries.*

Furthermore, data from Intel shows that their HPC server fleet has an overall annualized failure rate of <1.37%, with only 1.22% in the first 4 years, and up to 1.56% in the year 4 to 8 of life. The main component failures seem to be from the PSU at 0.32% across the 0 to 4th year and 4th to 8th year brackets; and the motherboard which has a high rate of failure in the first 4 years of 0.55%, then subsequently dropping to 0.16%. Drives were third placed with failures of 0.18% across their lifetime. All other components (including memory, cache battery, CPU,

fans, integrated switch, RAID controllers and Network interface Card) had an annualized failure rates below 0.1%.¹⁴

Therefore, in this section we will discuss the requirements in the scope of the following components (if included): data storage devices, memory, CPU, motherboard, graphic card, PSUs, chassis, batteries, fans, integrated switch, RAID controllers and Network Interface Cards.

According to EN 45554, there are four classifications to be aware of when discussing disassembly requirements: the fastener types, the necessary tools, the working environment, and the skill level of the disassembler. These are described below.

Skill level:

- **Layman (Class A):** person without any specific repair, reuse or upgrade experience or related qualifications.
- **Generalist (Class B):** repair, reuse or upgrade process cannot be carried out by layman (class A) but can be carried out by a person with a general knowledge of basic repair, reuse or upgrade techniques and safety precautions.
- **Expert (Class C):** person with specific training and/or experience related to the product category concerned.

Stakeholder feedback indicates that most disassembly can be done with generalist skills of **class B**. However, there is feedback that specific tasks currently require **class C** skills. Stakeholders agree that it would be preferable for all disassembly and repair operations to only require **class B** skill levels. Furthermore, when completing a refurbishment, there are server reconfigurations that need to be done. These require technical expertise to deliver but these are software based.

Feedback requested on disassembly and repair skills

The study team encourages stakeholders to provide detail on what disassembly and repair operations, or components, require expert class skills.

Stakeholders are also invited to provide feedback on Ecodesign setting a requirement for all repair operations to only require Generalist class B skill levels.

Working environment:

- **Use environment (Class A):** If a repair, reuse or upgrade process can be carried out in the environment where the product is in use without any working environment requirements.
- **Workshop environment (Class B):** If a repair, reuse or upgrade process cannot be carried out in the environment where the product is in use (class A) but does not require a production-equivalent environment.
- **Production-equivalent environment (Class C):** If a repair, reuse or upgrade process can only be carried out in an environment that is comparable with the environment in which the product was manufactured.

Class B workshop environment is likely to be the most apt description of working environment requirements. However, stakeholder feedback indicates that as many

¹⁴ IT@Intel: Green Computing at Scale, August 2021

repairs are performed on data centre premises, **class A** should be referred to under the current nomenclature. These server repairs are typically performed on tables in open areas of the data centre away from the racks. These could therefore be described as a workshop area, and hence there is a need for the **class A** definition to specify what is meant by “where the product is in use”.

Necessary tools:

- **Class A:** *feasible with no tool; a tool supplied with the product or spare part; or with basic tools (screwdriver, hex key, pliers, spanner)*
- **Class B:** *Product group specific tools*
- **Class C:** *other commercially available tools*
- **Class D:** *Proprietary tools*
- **Class E:** *not feasible with any existing tool*

Stakeholder feedback is in agreement that disassembly and repair should not require class D or E tools. As the Intel study¹⁵ revealed that the PSU, motherboard and drives are the components which are most likely to fail, it would be preferable for their access only require class A tools to facilitate repair. Therefore, tools meeting class A, B or C requirements are preferred to enable repair.

However, stakeholders indicate that Class A requirement of “tools supplied with the product or spare part” should be avoided, as many maintenance and repair activities use harvested parts that will not have the tools supplied.

Feedback requested on disassembly and repair tools

The study team encourages stakeholders to provide detail on which components (following the list set at the start of section 2.8.4.5) are most likely to fail.

Furthermore, following the same list, could detail be provided on the disassembly actions which cannot be done with basic tools (class A)? and if so, what specific tools are required?

Fastener types:

- **Reusable (class A):** *An original fastening system that can be completely reused, or any elements of the fastening system that cannot be reused are supplied with the new part for the repair, reuse or upgrade process.*
- **Removable (class B):** *An original fastening system that is not reusable, but can be removed without causing damage or leaving residue which precludes reassembly (in case of repair or upgrade) or reuse of the removed part (in case of reuse) for the repair, reuse or upgrade process.*

Stakeholder feedback agreed that fasteners should be reusable or removable. Feedback is provided that the number of different types of screws should be reduced and that no proprietary fastening systems be used in order to facilitate disassembly. This is consistent with Ecodesign requirement set out in Annex II, 1.2.1.

¹⁵ IT@Intel: Green Computing at Scale, August 2021

Stakeholders raised that requirements to allow for easy separation of the product would facilitate increased reuse and recycling. Not only does the disassemblability allow for access to parts for reuse, but also to ensure material is not lost during recycling, as making it easier to access circuit boards will ensure precious metal connections are not lost through shredding.

Feedback requested on fasteners

The study team encourages stakeholders to provide detail on what are the fasteners typically used on server and data storage product components (following the list set at the start of section 2.8.4.5). Please specify if these components are Reusable or removable.

Access to repair and maintenance information:

Information should be made available to ensure that professional repairers know how to disassemble and repair.

2.8.3.6 Current information provision of cobalt in batteries and neodymium in hard disk

Stakeholders from the IT Asset Disposal (ITAD) indicated that the provision of this information was very useful as it allows for recyclers to identify which products will be economically viable to disassemble, instead of simply shred. This logic also applies for dangerous materials (such as mercury), to know what recycling process can be applied safely. In the case of batteries, stakeholders also indicated that information on chemical content allow for recyclers to avoid mixing batteries of different chemistries, which would result in increased costs to separate work streams.

2.8.3.7 Information requirements for critical raw materials

Following the information requirements already in place and the recycling benefits they provide, stakeholders indicated the benefit of extending this measure to further materials. Materials highlighted include germanium and silicon which are present in semiconductor-based processors, memories and power amplifiers. Also mentioned are tantalum, gold, gallium and dysprosium.

Furthermore, stakeholders indicated that new recycling technologies are being developed which will be able to recover a wider range of materials. One of the drivers to this comes in the form of the EU targets for critical materials. One of these is for recycling of critical materials to include 15% of the annual EU consumption. Therefore, requiring this information be provided will enable recyclers to target the most appropriate products for recycling. Product information identifying CRM content at a component level make it simpler for recyclers to specialise their recovery efforts to the appropriate sections.

2.8.3.8 Material composition and recycling capabilities

The EU rules for the treating the waste for electronic equipment, and hence for servers, is WEEE. WEEE sets out (amongst other things) targets around the collection of electrical equipment, and the final recycling and recovery rates. It is

important to note that these targets are set out by mass and not in proportion to the most valuable or most impactful materials. The Ecodesign regulation however has the jurisdiction over the reparability and reusability of a device. This is an aspect to bear in mind, as the line between Ecodesign and WEEE is relating to if an item is deemed as “repairable” (under Ecodesign) or as “waste” (under WEEE). This is critical as recuperating component parts from a used server could be considered either as a “reuse” activity falling under Ecodesign, or a “valorisation from waste” under WEEE.

Servers are composed of plastics (Outside frame, small ABS and ABS/PC parts, cable insulation), ferrous and non-ferrous metals (chassis, metal brackets, screws, metal components from fans and HDDs and cables), precious metals (Copper, Gold, Palladium and Platinum), and batteries. Data from a European recycler stakeholder indicates that the main fractions after dismantling for a server are: the ferrous metals (66%), the power supplies (20%) then the printed circuits boards (6%).

Critical raw materials (CRMs) are located in batteries, HDDs, PCBs and connectors. The CRMs contained in the greatest amount is Neodymium in the magnets of HDDs, followed by silicon in the die of integrated circuits, and cobalt in batteries.

Traditional recycling technologies can recover precious metals, plastic, steel, copper, aluminium. For recovery of materials in smaller quantities (such as the CRMs), bioleaching and pyrolysis technologies can be used. Under the WEEE metric, approximately 80% of server materials (by weight) are recovered for recycling when a server is recycled. When including material and energy recovery, this rate increases to 99% of server mass recovery. It is important to note that WEEE operates under the metric of the total mass of the item and not the most critical materials. Therefore, this does not track if CRMs are being recovered.

However, recycling still only makes up around 15-20% of the end-of-life market which can largely be attributed to a lack of collection. Following the evidence from the 2019 impact assessment on servers and data storage products, recovery rates of servers for servers sent back to Original Equipment Manufacturers (via asset recovery services or leasing programmes) is very high. Therefore, the key to improved recovery and recycling rates is improved collection of products that are no longer repairable or reusable. This statement was echoed by stakeholders. Measures within the Ecodesign scope to facilitate this recovery relate to the provision of product data relevant for recycling such as components and material content, and disassemblability.

End of life servers are described as WEEE, which makes official statistics for these items amalgamated with other IT equipment. Therefore, tracking server and data storage end of life streams can be difficult. Specific data for servers in the Netherlands is available in GreenIT Report: Circular Data Servers. Here it is shown that out of 494,629 servers discarded in 2016 in the Netherlands, only 11% were refurbished, 24% recycled and 63% were simply exported out of Netherlands. Although it is noted that 50% of the exported data servers are intended for reuse abroad, it is unclear what happens to this equipment at their end of life, nor how long the product reuse extension is for.¹⁶

Feedback requested on recovery rates

The study team encourages stakeholders to provide figures on the recovery rates in the EU for computer servers and data storage products. To this end, the study

¹⁶ <https://www.amsterdameconomicboard.com/app/uploads/2018/06/Circulaire-Dataservers-Rapport-2018.pdf>

team seek to understand how likely it is for a product sold today to be recovered, instead of sent to landfill. Further detail breakdown regarding product reuse, refurbishment, recycling and landfill would be useful.

Another consideration to note is that servers can find a second or a third life being reused in countries outside of the EU, where recycling policies are not as strong. The inclusion of an Ecodesign expected lifetime provision in an information sheet could be considered to ensure that products being exported still have genuine life expectancy, and avoid dumping. The metric could state the expectancy for the product to perform until failure, and the further expected lifetime increase through repair. For example, as in 2.8.3.5 we describe that the component most likely to fail is the PSU. The label could indicate the life expectancy of the product to perform until first failure of the PSU, and its expected increased lifetime by replacement of the PSU (until the estimated next component failure). Providing this information can incentivise users to repair the device by providing an estimate of the potential gains from repair activities. It can also enable market regulators to check product life expectancy when exported. Market regulators can therefore be informed to avoid unsold product destruction or dumping.

Stakeholders that raised that requirements to allow for easy separation of the product into different materials should be considered to facilitate increased reuse and recycling. Restrictions on material/substance mix (e.g. choice and combination of polymers, additives) was also mentioned as useful.

2.8.3.9 Packaging

Product packaging has been categorised into three main elements:

- The cardboard packaging.
- Foam protection of the product (including bubble wrap), these are usually plastics based.
- Other plastics (films, wrappings)

Stakeholders did not provide any feedback to indicate action is required from Ecodesign with regards to packaging.

2.8.3.10 Networking equipment

Networking equipment may be capable of long technical lifespans, however the technology field is moving quickly which requires upgrades in hardware. For example, in the last eight years the technology has changed to have equipment capable of computing data from 1GB to 10GB, then 40GB and finally 100GB. The increase in bandwidth demand has forced the technology to review its hardware architecture and computing power.

Feedback requested on networking equipment

To better understand the material efficiency challenges on networking equipment, the study team is looking for feedback on what components which are most likely to fail. Data on materials content of networking equipment would also be useful.

Stakeholders have indicated that a core limitation to product reuse and repair is the licensing system. Licensing of network equipment can lock the product for usage (or reuse) without the manufacturer permission (even after the manufacturer has stopped providing product support). If the original manufacturer does not agree to

renew a license, the product is therefore forced into obsolescence. The licensing also makes maintenance and repair by independent companies difficult. For resale, or transfers of hardware (even within the EU), the original manufacturer must provide a signed transfer agreement. These become difficult to manage after multiple resales.

Improving this licensing system may improve product life expectancy, however there is a need to review if that would be impactful due to the pace of technological improvement. There are likely to be use cases where older, lower-specification systems can be successfully used in modern low-demand applications.

As with servers and data storage products, networking equipment could benefit from information provisions, such as through a data sheet, or online data portal, to inform recyclers on the components included in the device, and the material content (notably for CRMs) for increased recycling.

2.8.4 Recommendations

For all of the following measures, it is important to keep in mind how lifetime extension measures may encourage the use of older and less efficient equipment. For measures modelled in Task 7, these considerations of lifetime extension and energy efficiency will be modelled.

2.8.4.1 Reduced material content

We recommend placing a request to develop a standard to ensure materials are not wasted in the design and manufacture stages. This standard could investigate the development of:

- A reusable universal server enclosure for the most used server form factors
- Exclude any designs which prevent disassembly for repair or recycling purposes
- Reduce the material use and range, as for example, on average a server chassis contains 3-4 different steel alloys and 3 different polymers.
- Developing guidance for a modular Printer Circuit Board for improved interchangeability
- Ensure a standardised and simplified design of all parts and enclosures
- Reduce complexity by removing unnecessary fastenings, making it easy to disassemble and repair¹⁷

Furthermore, a standard to determine and declare CO₂ emissions in the manufacturing process can be investigated.

2.8.4.2 Firmware support considerations

The regulation currently states that:

1.2.3. From 1 March 2021, the latest available version of the firmware shall be made available from two years after the placing on the market of the first product of a certain product model for a minimum period of eight years after the placing on the market of the last product of a certain product model, free of charge or at a fair, transparent and non-discriminatory cost. The latest available security update to the firmwares shall be made available from the time a product model is placed on the

¹⁷ <https://vb.nweurope.eu/media/18830/cedaci-eco-design-fact-sheet-compressed.pdf>

market until at least eight years after the placing on the market of the last product of a certain product model, free of charge.

The regulation already ensures that products have access to latest firmware updates for eight years after the last product was placed on the market, therefore a request for six years is already covered. However, the availability of previous versions could be helpful for refurbishers to extend life expectancy. We therefore recommend investigating the impacts of making these previous versions available, without lessening cybersecurity norms.

Regarding releasing firmware as open-source data, this is likely to be a breach of manufacturer Intellectual Property and is therefore difficult to implement.

With regards to contact information, this is beyond the scope of Ecodesign policy, but can be referred back to GDPR regulations.

2.8.4.3 Barriers to part harvesting

Provision of more detailed component level performance and material information would allow for improved part reuse and materials recuperation. This can be reviewed in scope of inclusion as Ecodesign measure.

In order to facilitate repair and extend life expectancy of devices, we should investigate the inclusion of a repair clause which allows for server components to be replaced by legitimate repair professionals. This may require for a more innovative approach to part pairing – potentially from parts pairing to be banned, or for parts pairing firmware to be made available to certain repairers to allow for replacement or for manufacturers to provide a free-of-charge parts pairing system to authenticate legitimate repairs. Legitimate repairers could be identified through subscription to a database, as is done for the provisions of indicative weight and disassembly operations under Annex II clause 3.3. Such system should be available for longer than the availability of spare parts.

2.8.4.4 Availability and usability of spare parts

As is the case with other Ecodesign regulations such as the draft Ecodesign requirements for Smartphones and Tablets, there is a requirement for spare parts to be made available by suppliers for a certain amount of time after a product is placed on the market and for those parts to be fully useable for legitimate repairers. This measure can be investigated to be brought forward into the server regulation.

The spare parts considered are: data storage devices, memory, CPU, motherboard, graphic card, PSUs, chassis, batteries, fans, integrated switch, RAID controllers and Network Interface Cards.

2.8.4.5 Material efficiency requirements for disassemblability and repair

We recommend including a clause within the Ecodesign regulation to require servers and data storage products to be disassemblable by a generalist (class b), in a workshop environment (Class B), using tools meeting Class A, B or C nomenclature.

Feedback requested on disassembly tools

The study team invites the stakeholders to provide insight into which components cannot be disassembled with class A tools.

The definition for Class A work environment should be updated to ensure that workshop environments consistent with class B can be hosted at in use sites (which implies class A work environment). Class A tools should be reviewed to not include “tools supplied with the product”, as these are rarely available during product repair on site. Furthermore, fasteners should all be reusable (class A) or removable (class B).

Information should be made available to ensure that professional repairers know how to disassemble and repair.

2.8.4.6 Current information provision of cobalt in batteries and neodymium in hard disk

We recommend keeping these requirements within the regulation as stakeholders have mentioned they are useful to encourage recycling rates.

2.8.4.7 Information requirements for critical raw materials

As with Cobalt and Neodymium, we recommend extending this list of CRM reporting requirements to germanium, silicon, tantalum, gold and dysprosium. Providing the indicative weight of these materials will facilitate recovery and recycling activities of these critical raw materials.

2.8.4.8 Material composition and recycling capabilities

We recommend investigating further the rates of recyclability and recuperation of servers in the EU.

Within the scope of the Ecodesign regulation to improve the recycling rates of servers, we recommend investigating the inclusion of a more extensive product datasheet requirement which would track material content. This product information datasheet should include a list of the components, their number codes and their material content (both bulk and targeted CRMs: Cobalt, neodymium, silicon, germanium, silicon, tantalum, gold, dysprosium). This list will allow refurbishers to recuperate required components, source appropriate replacements for repair, and incentivise recycling. We also recommend the inclusion of an expected lifetime provision on the information sheet. This information could also be provided via a digital product passport.

The regulation can also consider measures to ban the use of particular polymer combinations, which are difficult to recycle, and a requirement for easy separation of materials.

2.8.4.9 Networking equipment

Currently, networking equipment is beyond the scope of the existing server and data storage product Ecodesign regulation 2019/424. We recommend the Commission investigate the potential to facilitate the licensing system for product reuse and refurbishment.

There could also be a requirement for the information sheet of networking equipment to provide material content information.

2.8.5 Link to Phase 2

Task 5 modelling will cover the modelling of server base case life cycle analysis. This will include the effects of materials in the manufacturing process. Therefore, for all the measures which are carried over from Task 6 to 7 for modelling, the benefits from lifetime extension compared to the benefits of newer more efficient servers will be considered.

Task 4 will include a discussion on what is firmware and its updates. Task 6 will include a review of a measure to make previous versions of firmware available to wider public, and its potential effects.

Task 6 will include a review of measures to improve part harvesting and its effect on product repair.

Task 6 will include a review of the requirement for spare parts to be made available by suppliers for a certain amount of time, and the effective deployment of spare parts in repair situations.

Task 6 will include a measure for inclusion of improved disassemblability for repair, reuse and recycling requirements.

Task 4 will review material content of servers, including of components with high CRM content. Task 6 can then include a measure for the declaration of CRM content in components.

Task 6 can include a measure for increased product information provided in datasheets to track material content and product expected lifetime.

Task 6 can include a measure for the ban of particular polymer combinations and requirements for easy separation of materials.

2.9 Operating Conditions

- h) to set specific Ecodesign requirements on the operating condition class;*
- k) analysis of the benefits of the information requirements under Regulation 2019/424 on the operating conditions of servers and data storage products.*

2.9.1 Background

The current Ecodesign regulation for servers has an information requirement to state the operating condition class of the product to be from ASHRAE A1 to A4 (which detail the temperature and humidity operating capabilities of the server). Under the ASHRAE A1-4 conditions provide a window of operation capability around the optimum temperature at approximately 27 Degrees C. This provides an implicit requirement for products to be operated within the A4 window. Temperature ranges for Operating classes are defined in Table 2.4.

Table 2.4 Operating condition classes

Operating condition class	Dry bulb temp °C		Humidity range, non-condensing		Max dew point (°C)	Maximum rate of change (°C/hr)
	Allowable range	Recommended range	Allowable range	Recommended range		

	Dry bulb temp °C		Humidity range, non-condensing			
A1	15- 32	18-27	- 12 °C Dew Point (DP) and 8 % relative humidity (RH) to 17 °C DP and 80 % RH	- 9 °C DP to 15 °C DP and 60 % RH	17	5/20
A2	10-35	18-27	- 12 °C DP and 8 % RH to 21 °C DP and 80 % RH	Same as A1	21	5/20
A3	5-40	18-27	- 12 °C DP and 8 % RH to 24 °C DP and 85 % RH	Same as A1	24	5/20
A4	5-45	18-27	- 12 °C DP and 8 % RH to 24 °C DP and 90 % RH	Same as A1	24	5/20

2.9.2 Questions

1. What are your views on setting specific Ecodesign requirements on the operating condition class? Should there be a requirement for servers to operate at higher temperatures to limit HVAC energy consumption in data centres? Consider in your answer if the operating temperature range should align with what is recommended by ASHRAE.
2. The Ecodesign server regulation requires products to provide operating information on ASHRAE ranges A1-4. In your opinion, is providing this information useful for reducing energy consumption of data centres?

2.9.3 Feedback/ Research results

2.9.3.1 Setting Ecodesign operating condition requirements

A few stakeholders indicated that setting a required operation conditions at A2 ASHRAE class could be beneficial. This is because most colocation data centres

operate under A1 conditions, as they will operate under the lowest capability of servers in their facility. Therefore, setting a requirement for operation under A2 class for allowable range may be beneficial in order to allow for data centres to align their temperature conditions for all servers at A2, instead of A1. This should still however, have data centres aim for temperature to be kept within the recommended range.

Most datacentres are assumed to operate at 18°C in order to ensure systems do not overheat. However, operating at higher temperatures, would result in lower energy costs for datacentres, as cooling consumption would decrease. This is partially true, and has been one of the drives to improve Power Usage Effectiveness in the datacentres operation. It is estimated that for every 0.56°C (or 1°F) increase in server inlet temperature, a 4% saving can be delivered on the cooling.¹⁸ However, although increasing temperature settings results in a reduction in the HVAC consumption, it also has the opposite effect on servers themselves, where the consumption increases as internal fans blow harder. This results in a sweet spot of temperature setting of 27°C, where the sum of HVAC consumption and server internal fan consumption is at its lowest. Operating at higher temperatures than 27°C would therefore result in an increase of datacentre energy consumption as internal server fans increase their output.¹⁹

This can also be counterintuitive and reduce the PUE score for a data centre as the ICT equipment would have the energy consumption increase.

There is also a concern that products designed to operate to A3 or A4 are capable of running at higher temperatures, but these are usually meant to be for short periods of time during an unplanned event or outage. Products should not be run at these temperatures consistently or it would affect their lifetime. Most stakeholders agree that operation under high temperatures should be considered as a temporary activity rather than a long-term operation.

Feedback requested on operating temperatures

The study team invites the stakeholders to provide feedback on the operating temperatures of their datacentres.

2.9.3.2 Provision of ASHRAE operating conditions

There is general consensus that operating condition information should be kept under the ASHRAE system in order to provide consistency with other markets, operator clarity and keep warranty conditions consistent. This information provision is useful for data centres to manage their operations, to select the appropriate server capabilities for their facility temperature, and in determining if additional free cooling hours are available within the environment (free cooling refers to when cooling systems can rely on low external air temperatures to cool systems instead of mechanical refrigeration).

¹⁸ [Raise the Temperature | ENERGY STAR](#)

¹⁹ [data-center-operating-temperature-white-paper.pdf \(dell.com\)](#)

2.9.4 Recommendations

2.9.4.1 Setting Ecodesign operating condition requirements

We recommend investigating a requirement for servers to operate at a minimum ASHRAE range of A2. However, the requirement for A2 allowable range is meant to be for temporary temperature fluctuations. Operation should be kept long term within the recommended range. These should also be reviewed in order to ensure that servers can operate reliably at the higher end of the recommended operation range.

2.9.4.2 Provision of ASHRAE operating conditions

Recommend keeping the information provision of operating conditions under the ASHRAE nomenclature.

2.9.5 Link to Phase 2

Review the range of servers advertised as operating under A1 to A4 conditions. This can inform the effect of a potential requirement of removing the A1 range of products.

In task 6, review the effects of the removal of A1 range servers and the incentive to operate on the higher end of the recommended range. Task 7 covers modelling of product lifecycle assessment, and hence will not have an in-depth model of data centre consumption (which the temperature operating condition change would affect). However, an estimation can be provided in Task 7 of the impact of the measures on the wider data centre system.

2.10 System Performance Considerations

m) Technological, market and regulatory evolutions affecting the environmental performance/aspects of data centres, and how they would reflect at product specific level, for servers and data storage products.

2.10.1 Background

Computer servers and data storage products require supporting equipment to operate, notably to provide power, temperature control and internet connectivity. Many servers are therefore often hosted in data centres: buildings dedicated to computer server usage, which provide optimal conditions for operation. As the servers and data storage products operate, they consume energy which is wasted as a heat output. For this reason, data centres need to have temperature control systems which cool the building. They also need to have power converters and uninterruptible power supplies to stabilise power to the facility. Therefore, server energy usage is not only related to their direct usage, but also to the extended environment. As servers consume more energy, the system will need to proportionally increase their energy consumption to deliver additional cooling. Power Usage Effectiveness or PUE is a metric used to analyse data centre efficiency, which is calculated by dividing the total data centre consumption by the consumption of the servers, data storage equipment and network equipment. PUE therefore has a theoretical minimum of 1.

2.10.2 Questions

1. How are technological, market and regulatory evolutions influencing the environmental performance of data centres? How are these changes affecting the product-specific level performance of servers and data storage products?
2. How has the average data centre Power Usage Effectiveness (PUE) evolved over the last few years, and how is PUE expected to develop in the future?
3. What are your views on the advantages and concerns around the PUE metric? Should another data centre performance metric be considered?
4. Are there any market developments of data centres facilities to be noted?

2.10.3 Feedback/ Research results

One stakeholder stated that the Ecodesign regulation on servers should align with the Energy Efficiency Directive which applies to data centres.

Stakeholders provided insight that servers activity accounts for 60-70 percent of the energy consumption of data centres, and therefore needs to be reduced. As IT equipment consumption is reduced, so would proportionally be the energy required to maintain and cool the data centre.

PUE has been used to advise on the energy efficiency of data centres. It is used as a reference KPI at a data centre level. It has significantly improved in the last 20 years, but improvements are now stalling and reaching physical limits of improvement. These efficiency gains have mostly been achieved using improved cooling technologies. The average data centre performance is now a PUE of 1.55²⁰. The best PUE scores tend to be achieved through the use of local water sources and evaporative technologies, which can therefore have an impact on local water reserves.

Some countries have made requirements for data centres to operate at low PUE values (such as Netherlands where the minimum average PUE should be 1.2). Therefore, stakeholder feedback is that the focus on PUE is now no longer required, as gains in efficiency on this metric are now near to none. Moreover, the PUE metric ignores the total energy consumption of the data centre, and hence doesn't acknowledge high consumption value of the IT equipment. One stakeholder suggested also that PUE needs to be certified according to ISO/IEC 30134-2.

One metric to reduce the energy requirements of a data centre is to allow for the facilities to operate at a higher temperature. This is only effective if the servers are not compensating by accelerating their fans, as otherwise PUE score would reduce only because IT equipment is using more power than the cooling infrastructure. Industry stakeholders note that this has recently been improved by having servers capable of operating at the higher end of the recommended temperature range of 18-27C, without the increase in IT fan speed.

The PUE does also not acknowledge system efficiency gains if waste heat is captured and reused for other facilities. The Energy Reuse Factor (ERF) as defined in EN 50600-4-6 could be used as a metric to encourage data centre efficiency developments. On a product specific level, servers which are ready for direct liquid cooling would be more effective for waste heat recovery.

²⁰ <https://www.statista.com/statistics/1229367/data-center-average-annual-pue-worldwide/>

One stakeholder has stated that cooling efficiency could be improved if, for liquid cooling systems, the energy reuse of chilled water was set between 30 and 60 degrees C.

Other data centre performance metrics include: Water Usage Effectiveness WUE (particularly critical when relying on water cooling technologies), Cooling Efficiency Ratio (CER as defined in EN50600-4-7) and Carbon Usage Effectiveness CUE.

However, the metrics of PUE, WUE, CER and CUE do not consider the effects of servers themselves. Metrics are proposed to compensate for this gap:

- ITEESV which measures the energy efficiency of servers,
- ITEUSV which is used to determine the processor utilisation rate, and
- SIEC or Server Idle Energy Coefficient relates to the amount of idle time (and hence “wasted” time and power for a server).

It is noted that data centres and their trade associations have launched the Climate Neutral Data Centre PACT (CNDP), where signatories have pledged to be carbon neutral by 2030. This is following metrics on energy, carbon, water usage, circular economy and heat reuse.

2.10.4 Recommendations

There are two parameters the server Ecodesign regulation can consider to directly affect the energy consumption of wider data centre activities. The first is to act to reduce the energy consumption of the servers themselves. This is being considered under the energy efficiency criteria in section 2.1. The other aspect is to influence the operating conditions of servers themselves, which can be to mandate a higher operating temperature range or encourage liquid cooling requirements.

Our recommendation is to encourage operation of servers at higher temperatures (up to 27°C), by providing clear information through labelling or information sheet that states that the servers would operate well at these temperatures whilst reducing energy costs. However, this should be stressed not to increase servers operating temperatures above 27°C, as this can result in an increase in total data centre consumption (due to an increase in individual server fan consumption). This follows the advice from ASHRAE which shows how for every degree increase in the air inlet temperature from 17.7°C, 4% can be saved on cooling costs. However, beyond 27°C, although cooling costs continue to decrease, and PUE value decreases, the consumption of the IT equipment increases, resulting in a total data centre energy increase.²¹

2.10.5 Link to Phase 2

Include the findings above into the Phase 2 Task 3 report, which illustrates the usage phase of servers and data storage products. Task 3 to detail how increased operating conditions temperature would increase server fan consumption. The different criteria for data centre performance KPIs can be included in the report as reference, with a mention for the Energy Efficiency Directive to refer back to the Ecodesign report when setting their regulatory action on data centres.

In task 4 of the report, temperature conditions of servers can be reviewed along with the Best in Class for operating at higher temperatures. This can also be applied for a

²¹ [ENERGY STAR Ask the Experts | Products | ENERGY STAR](#)

review of the water-cooled server temperature range. These can then be considered as modelling options in Task 6.

2.11 Liquid Cooling Systems and Solutions

n) The technical and economic feasibility and relevance of product specific requirements on liquid cooling systems/solutions.

2.11.1 Background

As water has a high heat capacity, liquid cooling is considered for servers as an effective measure for temperature control. Liquid cooling would also make for an improved medium for waste heat recovery.

In this section, when discussing liquid cooling systems, we are referring to solutions that bring liquid for cooling to the rack, server, or component level.

2.11.2 Questions

1. What are the different types of liquid cooling solutions available for servers and data storage products? How do these vary in terms of efficiency, cost, ease of installation and maintenance? What are the environmental considerations to be aware of?
2. Please expand on the previous question based on the following:
 - a. Liquid cooling to product.
 - b. Liquid cooling to the rack.
3. What do you expect is the market share of servers and data storage products that include a liquid cooling solution?

2.11.3 Feedback/ Research results

There are multiple types of liquid cooling solutions considered in this section:

- The first is indirect cooling. This is a system of liquid cooling where no liquid flows through the servers, but rather to the rack. These are usually set up as rear-door hybrid cooling, where a cooling liquid is set to flow at the back of the server rack, and allows for air to be cooled as it is removed from the server by the liquid flow. This liquid flow would then need to be cooled elsewhere, using technologies such as a chiller, or free cooling. These systems have a low to medium efficiency, require outlet water temperature to be under 50 degrees Celsius, are simple to install and have a medium total cost of ownership.
- Another to be considered is immersion or submersion cooling. These systems have the entire server submerged in liquid which grants a high efficiency due to the liquid heat capacity. However, the installation is complex, and servicing the servers is difficult.
- The last is a direct-to-chip liquid cooling system. These will have liquid cooling flow provided directly in thermal contact with the hottest components of a server. These are highly efficient and allow for higher outlet temperatures, 60 – 75 degrees C, which allows for more efficient waste heat recovery applications and the use of additional free cooling. These have low maintenance and low total cost of ownership. There are some installation concerns, but the main drawbacks are in the design of the specialised server to accommodate for the liquid heat exchange. Liquids used can be in a single phase (1P) or dual phase

(2P). Internal analysis from Asetek shows that the electrical consumption for a per-node cooling system can be reduced by 40%. Examples of such technologies and companies are Denmark based Asetek, and Canada based CoolIT.

As detailed above, liquid cooling options usually result in more efficient cooling systems compared to air-based solutions. A more efficient system would hopefully be capable of avoiding the use of evaporative towers in order to lower WUE scores, notably for water scarce regions. They can also enable waste heat recovery systems as the outlet liquid can be set at higher temperatures.

Some stakeholders commented that the liquid cooling definition requires better definition. In particular, these should be defined as meaning liquid cooling “to the product”. This would be more appropriate as indirect liquid cooling to the rack is outside of the regulatory scope of Ecodesign.

There are multiple trials underway of data centre immersion to increase data centre efficiency. However, these are out of scope of the Ecodesign regulation, as the entire system is immersed rather than the servers specifically.

Almost all high-performance computing (HPC) installations are utilising direct liquid cooling to achieve the performance and power density requirements of that market. Unfortunately, these products are out of scope of the Ecodesign regulation.

Direct cooling technology application to other servers is currently limited, however one can expect this technology to develop in the next few years due to the following trends:

- an increase in Thermal Design Power requirements for both CPUs and GPUs,
- the need for higher rack power density (for example 50% reduction in rack U-space for GPUs when using liquid cooling instead of air cooling)
- the need to increase aisle temperature in server rooms to conserve energy for HVAC
- an increase in demand for waste heat

These applications are currently limited but there is a long-term potential of their increase in demand in years to come. These technologies are currently not capable of being tested under SERT, which is designed for air cooled configurations.

2.11.4 Recommendations

As liquid cooled systems can take different forms it is important to be clear on the definitions covered.

For indirect cooling systems applying to the rack, these are simple and cost effective, but are sold as separate equipment to the server and are therefore out of scope of the Ecodesign regulation. No action is required here.

Liquid cooling to the chips is largely limited to HPC systems, which are out of scope of the regulation. Therefore, no action is required here.

Regulated server products are currently tested under SERT in an air-cooled configuration. As the implementation of liquid cooling onto those servers is likely to increase efficiency and their low market share, no action is required on this front at this time. The market share is very small and direct liquid cooling is only used in high performance computing currently due to cost and complexity issues. However, this technology should be monitored to ensure that if market share grows, the next

Ecodesign revision may need to review regulatory action. A definition for liquid-cooled servers is needed, to ensure that servers that are only manufactured in liquid cooled versions are exempt from energy efficiency criteria of the regulation.

2.11.5 Link to Phase 2

To develop a definition for liquid-cooled servers in Task 1.

Task 3 should have a section to detail the above findings of the indirect and submersion cooling technologies, as these are part of the usage environment of the servers.

Task 4 to detail the cooling to chip technology in the BAT section for HPC servers, and in the BNAT section for remaining servers.

2.12 Waste Heat Recovery Systems and Solutions

o) The technical and economic feasibility and relevance of product specific requirements on waste heat recovery systems/solutions

2.12.1 Background

As ICT equipment processes information, electrical energy is consumed which releases waste heat as a byproduct. In data centres, HVAC facilities are set up in order to keep ICT equipment operating at the appropriate temperature, which requires extensive cooling capabilities. A potential solution to increase the sustainability of the data centre system is for this waste heat to be recuperated and used for other applications such as district heating.

2.12.2 Questions

1. To your knowledge, how many data centres have had their waste heat recovered? What are the uses for recovered waste heat from data centres?
2. What are some of the most effective waste heat recovery systems/solutions used? How do they help to reduce energy consumption and costs?
3. Are there any negative environmental considerations to be aware of with regards to waste heat recovery? Do heat recovery systems disincentivise other energy efficiency efforts?

2.12.3 Feedback/ Research results

Waste heat recovery is a practice which can be applied to all industry sectors, not just data centres. This technology is nascent for most industries, except for data centres for which waste heat recovery is a well-established practice globally. The most effective and economical cases of waste heat are for district heating, but is also applied to greenhouses, swimming pools, sports complexes, schools, hospitals, etc. There are many examples of successful waste heat re-use in the EU already, as the technology is readily available, not cost-prohibitive and district heating providers are aware and asking for data centre waste heat. Data centres are increasingly built in close proximity to urban areas, where regulatory requirements are incentivising the use of sustainable energy (notably pushed by the EU Energy Efficiency Directive). Asetek bring examples of their data centres in Aalborg Forsyning and

University of Tromsø providing district heating. Other examples include Frankfurt Mainova²² and Technical University Dresden²³.

2.12.3.1 Data centre was low grade heat value

For most data centres, the waste heat to be recovered is considered “low grade” as the liquid recuperated from the cooling systems will reach a maximum of approximately 30°C. This is because the data centres aim to keep their premises at the low end of the ASHRAE A1 recommended temperature conditions 18 – 27°C. This limits the direct application of this heat to residential, commercial areas or agricultural facilities: providing direct floor heating to residential or commercial buildings using 30°C water temperature (offices, houses, common spaces, etc.) or within agricultural facilities (greenhouses, fish farming). This requires for these applications to be close as the transport of such low-temperature water would incur losses as the distance grows.

For other applications, the waste heat can be upgraded from 30°C to up to 95°C, with heat pumps. The waste heat from the data centres provides preheating, which facilitates the heat pump. However, the heat pump is still required to run to increase the value of the heat, bringing an additional energy cost to the process. This heat pump may be running an ammonia refrigerant, which will increase in efficiency if the data centre waste heat temperature is increased. Distance is again a concern to ensure that heat losses over distance are minimised. For data centres established outside of city premises (where space is premium), the transporting of the waste heat may not be feasible, meaning many data centres have decided not to invest in waste heat recovery. Far from city premises providing the most easily accessible heat sinks, the feasibility of waste heat reuse in other industrial cases, such as for example the food and beverages processing plants, still needs to be further investigated.

2.12.3.2 Seasonal concerns

The reuse of waste heat from data centres can have strong benefits and synergies with district heating. However, there are limitations to these depending on the season. In the summer, the demand for heating goes down, which means the data centres still require high efficiency chillers installed in order to expel heat in the summer months. One could also argue that the winter months are the time when data centres can use local weather conditions for free cooling, hence making the benefit district heating as a heat sink less urgent for data centres. Therefore, there needs to be a good business case for heat to be purchased from the data centres in the winter to justify the waste heat recovery system installation.

2.12.3.3 High grade heat through Direct to chip liquid cooling

As detailed in the section for liquid cooling, servers are mostly cooled with fans in an air medium. However, if servers (or specific chips) were in direct contact with a liquid cooling medium, then this liquid could be brought to a higher temperature than the air systems as it has a higher heat capacity to take heat from the servers.

²² <https://www.faz.net/aktuell/rhein-main/wirtschaft/mainova-und-digital-realty-machen-data-centre-abwaerme-zufernwaerme-18613563.html>

²³ <https://www.saechsische.de/dresden/dresdner-pilotprojekt-zum-energiesparen-5843803.html>

Direct-to-chip liquid cooling would enable waste heat re-use above 60°C. At these temperatures, the waste heat from data centres could be used directly without the need for a heat pump cycle. Examples of such technologies and companies are Denmark based Asetek, and Canada based CoolIT. Not only would this technology provide a more efficient waste heat recovery solution but could also lower the data centre cooling energy consumption from a PUE of 1.2 to 1.12.

2.12.3.4 Incentives for waste heat recovery

Under the Energy Efficiency Directive, the EU has targets to reduce energy consumption by 9% by 2030 compared to the 2020 reference scenario.²⁴ Therefore Member States are pushing for measures to reduce their energy consumption. Some Member States (such as Germany) are requiring data centre operators to conduct feasibility studies on the viability of waste heat recovery. There is, however, a risk with these measures that to incentivise waste heat recovery may increase total energy usage. To improve the waste heat recovery measure, one could improve the heat output by decreasing the efficiency of server activities, leading to an overall increase in ICT equipment.

Stakeholders provided feedback in agreement with the analysis above.

2.12.4 Recommendations

The recuperation of waste heat is beyond the scope of the Ecodesign product regulation. However, it is noted how this technology could reduce the footprint of the overall data centre system. It is noted above that the best way to improve the waste heat recovery measures are to increase the grade of the heat. This could be done either by increasing the operating temperatures of data centres, or by using liquid cooling for improved heat transfer.

Under the first option, as was discussed in the operating conditions section 2.9.4, we do not recommend pushing for server operation at higher temperatures, as that may result in higher overall energy wastage, but rather to operate in the higher end of the ASHRAE recommended range (approximately 27DegreesC) and aligning to ASHRAE A2.

The second option to have direct liquid-to-chip cooling systems would be the most effective for heat recovery. As was discussed in the section 2.11.4, we recommend continuing to monitor the development of this technology to mainstream usage. Once developed further, a test standard under SERT should be developed to determine their active efficiency.

2.12.5 Link to Phase 2

The above content should be detailed in Task 3 as it relates to the usage environment of the server products.

In Task 4, revert back to detail the cooling liquid to chip technology in the BNAT section for servers.

²⁴ https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive_en#documents

2.13 Direct Current Power Supply for Servers

q) The technical and economic feasibility and relevance of product specific requirements on DC (direct current) power supply for servers.

2.13.1 Background

Servers and data storage products have Power Supply units, which convert the alternate current (AC) from the grid into a stable Direct Current (DC) supply source which then powers the device. There is a perception that as the power supply unit has losses associated with this conversion, having DC operated servers which are provided with DC power, can save on the PSU material cost and energy losses.

2.13.2 Questions

1. What are the technical and economic feasibility considerations for implementing direct current (DC) power supply for servers? How do these impact the design, installation, and maintenance of such systems?
2. Does the use of DC power supply for servers reduce the overall consumption of servers? How does it impact consumption at product level and data centre level?
3. What are some of the key challenges and opportunities associated with integrating DC power supply for servers?
4. What do you expect is the market share of servers and data storage products that are powered by direct current?
5. Can the existing SERT V2 tool be modified to cover DC servers? How feasible is this solution for meeting the regulatory requirements for energy efficiency and sustainability in servers?

2.13.3 Feedback/ Research results

The DC server market share is currently small and not expected to increase in the short term due to the concerns described below.

The use of DC ICT equipment is considered for the energy efficiency saved from the PSU units. The servers themselves would not require a AC to DC power conversion. However, DC servers would still be equipped with step down DC converters to bring voltages to appropriate levels in the server. It is noted that although savings can be made at the PSU level, efficiency losses are typically passed on to another location in the data centre. Indeed, as utility power is primarily in AC, the AC-DC conversion will need to happen for the ICT equipment, which may mean a shift of the problem at the PSU server level to a centralised data centre conversion unit.

Another consideration is with regards to the distribution losses. EUDCA acknowledges that the technical and economic feasibility considerations for implementing direct current (DC) power supply for servers could have an impact depending on the voltage provided to the data centres (AC/DC) and the distribution distance. For example, if it is AC, it will need a conversion from AC to DC through the use of fuel cells (700 V DC) to have a conversion advantage. To minimise some of these losses on the distribution stage within the data centre, a high voltage can be applied. However, this incurs new concerns around the safety.

SPEC SERT currently does not support testing of DC servers. SERT could look into developing this, but would require comparison of DC and AC servers, to add a comparison weighting if DC losses are being accounted for elsewhere in the

conversion system (otherwise DC servers would seem to be much more efficient than AC servers as the PSU is not included).

2.13.4 Recommendations

The current evidence does not support the argument that DC servers are more efficient than AC servers. Therefore, we recommend not to incentivise the deployment of DC servers.

DC servers currently are not tested under the SERT testing metric. SPEC could be contacted to develop the testing metric for SERTs next iteration. As DC servers are not covered under the SPEC SERT methodology, they should be exempt from the SPEC SERT active efficiency requirements.

2.13.5 Link to Phase 2

In Task 4, including the discussion above on DC servers.

2.14 Other topics

t) other topics, as emerged from consultations with stakeholders.

1. Are there other topics not raised in this questionnaire that this regulation review should consider?

Stakeholders have raised the following concerns. We have provided our view and action recommendation for each.

2.14.1.1 ICT real time operating condition provision

One stakeholder requested for servers and data storage products to provide the following data in real time:

- the energy consumption [W];
- inlet temperature of the cooling medium (e.g. air/water) [°C];
- Data transmission via the network interface [Mbit/s];
- in the case of servers, the load condition for each logical CPU [%],

They also made recommendations for the data format such SNMP (simple network management protocol), IPMI (intelligent platform management interface), or XML (extensible markup language).

Recommendation:

Under section 2.6, we have included that the requiring information for the inlet temperature and fan speed of the server would be advantageous for data centre operation, and hence recommend for this to be included in the regulation. This aligns with best practice from the Energy Star programme.

With regards to energy consumption, data transmission and load condition of the CPU, this information can be provided in real time by servers, as it is a requirement under the Energy Star programme. Therefore, a similar requirement could be set for increased performance transparency.

2.14.1.2 Carbon emissions of the production process

One stakeholder suggested that an information sheet should include the CO₂ emissions associated with the production of the server.

Recommendation:

In principle, this would be good for customers to be aware of. However, this is unlikely to have a high impact on purchasing decisions as material composition between servers are not likely to change due to similar materials used. The main motivation of purchasers is also on performance instead of embedded carbon.

Furthermore, this metric is complex to calculate and implement due to the complexity of the product composition. From a policy perspective, these are being developed with the launch of the EU Carbon Border Adjustment Mechanism for select products and is not ready for other categories at this point.

The study team will provide estimations of carbon emissions for base case representative products in the Phase 2, Task 5.

2.14.1.3 Utilisation metric for servers

A stakeholder raised the concern that the industry is too reliant on PUE as a measure of data centre efficiency and does not have metrics for energy consumption of the ICT equipment. They propose the datacentre industry and policymakers use the Server idle coefficient as a proxy to measure server effectiveness. This coefficient was proposed by dividing server energy consumed in idle state, divided by the total server energy consumed. This would provide a ratio of the server “useful” energy consumption.²⁵

Recommendation

The active efficiency metric from SERT is a measurement of server efficiency. However, it is not widely understood and implemented. We recommend investigating options under labelling and information requirements to make this metric more ubiquitous to the market. This may also incentivise manufacturers to have more comprehensive SERT testing in their server families, instead of focusing on min/max performance.

With regards to server idle coefficient, this is a metric of how users utilise their servers and is therefore beyond the scope of the Ecodesign. It is a useful metric for users to ensure they optimise their assets productivity. However, placing a regulation requiring minimum server idle coefficient levels would be counter-productive, as it could be simply deviated by users running useless operations to avoid being idle. The result would be an increase in energy consumption. Therefore, the regulation could follow a requirement for servers to provide utilisation data, to ensure that users can monitor and maximise their utilisation rates.

2.14.1.4 Idle state testing temperature

It was reported by stakeholders that the regulation currently requires idle state power to be measured at the higher boundary temperature of the declared operating condition class. Although this requirement is only an information requirement, it places additional testing requirements on manufacturers, with limited benefit for the information provided. Indeed, SERT testing is not done at these temperatures,

²⁵ [Server-Idle-Coefficients-FINAL-1.pdf \(iea-4e.org\)](#)

therefore delivering the high temperature test needs additional time after the SERT test to bring the system to this temperature. Furthermore, it is unclear what the benefit of this information is to the end users, as servers are rarely run beyond their recommended temperature ranges.

Recommendation

In order to ensure alignment with the SERT methodology, it would be recommended to remove this requirement for testing at the higher boundary temperature.

Draft

Part A: ANNEXES

Draft

Annex 1 Installed processor metrics


The figures and charts below indicate the SERT score over idle power for the entire data set. This includes models from 2016-2021.



Table A1.1 servers as a key and shows the number of models in each of the charts presented further below.

Table A1.1 Data set (SERT)



Entire Data set				
Server Type	# installed Processors	Count in data set	# Families	# of families with Typical Configs
Rack	1	152	57	37
Rack	2	247	90	70
Rack	4	58	19	17
Tower	1	30	10	7
Tower	2	23	7	6
Blade	2	77	25	24
Blade	4	22	8	7
	Total	609	216	168

A1.2 Rack servers







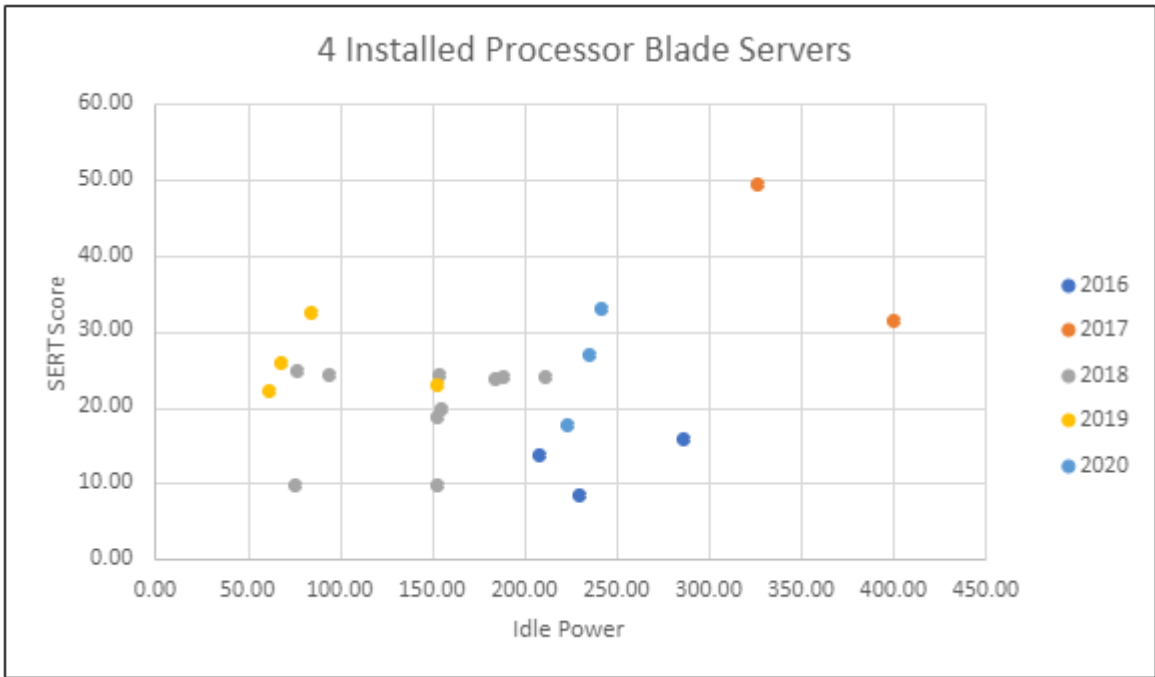
A1.3 Tower Servers



A1.4 Blade servers







Annex 2 Stakeholder Meeting Minutes

Date/ Time: Wednesday 29th March 2023

Attendees – Project team

Davide Polverini (GROW; Study Policy Officer), Eirini Passia (ICF; Project Coordinator), Tom Lock (ICF; Project Director), John Clinger (ICF; Technical Lead), Tom Ramsson (ICF; Modelling Lead), Laurent Petithuguenin (ICF; Technical Lead), Abhishek Jathar (ICF, Modeller), Todd Leddy (ICF, Researcher)

Minutes

Introductions

Introduction to the project and a round table of introductions from the study team and DG GROW.

Setting the scene (DG GROW)

DG GROW set the scene explaining the steps involved in the review of the existing Ecodesign regulation. The study has two phases: firstly, covering the Technical Review process; secondly, the preparatory study and the consultation forum. The Consultation Forum is expected to take place in the next 12 months. The aim of this stakeholder meeting is to demonstrate that the review study for this regulation has begun and collect feedback.

The scope of the review stems from the previous preparatory study and Article 8 of the Ecodesign Regulation for Servers and Data Storage products. Article 8 lists the items that require analysis within this study. Overall, the requirement is to explore the feasibility to either increase the tightness and the coverage of the regulation or to leave them as they are in light of the analysis conducted during the study.

DG GROW described how challenging it is to put exact dates on when any new regulation will be published. It was estimated that the publication date would not be before Q2 2026. ICF have been tasked to work within the framework of the current Ecodesign Directive 2009/125.

DG GROW raised the aspect of product level vs system level, which is very relevant for servers and data storage products. However, DG GROW stressed that the current regulation (2019/424) is a product specific regulation with two product categories: servers and data storage products.

Digital Europe – Q1) Asked DG GROW if they could expand on the ESPR (Ecodesign for Sustainable Products Regulation) aspects which could be considered for this review? Q2) At a system level, are there any implications from the review of the Energy Efficiency (EE) Directive e.g., for a Data Centre.

DG GROW provided the following responses: A1) ESPR expands the scope and types of requirements which can be enacted by the regulation. ESPR foresees a more specific list on material efficiency (ME), environmental footprint and energy efficiency. For servers, EE and ME parts are well covered. Carbon footprint is not covered, so far. A2) The second question will be addressed via the technical discussion during the rest of the presentation.

Delivery Plan

ICF presented the delivery plan for the study, which included the role of ICF, the study objectives, the study team, key project milestones, project deliverables, stakeholder involvement and a reminder of the study website. The first milestone will be in August 2023

when the 1st Interim Report (draft phase 1 and draft phase 2) will be published. The 2nd Stakeholder meeting will take place shortly after this (around September 2023). The finalised report is expected to be published in the summer of 2024.

Technical Analysis – Phase 1 – Review items a-t (split into themes):

Updating current Ecodesign requirements (Review of items a & b): ECOS queried whether the test methods underpinning the requirements are going to be reviewed. ICF stated that one of the objectives will be to check if the references to SERT need to be updated and to see if any new versions should be adopted.

Regulation definitions and scope: (review of items c, e & f): ICF explained that the study team are aware of the recently updated definitions for integrated APA's. ICF encouraged stakeholders to add viewpoints on this and any potential revisions to the APA definitions. DG GROW reiterated that references to new SERT versions will be included if work on the methodology part of the new version is finalised in time. Techbuyer Europe raised a question on whether other benchmarks besides SERT will be looked at. ICF replied that they are not aware of any other comprehensive benchmarks with regards to server efficiency.

Performance requirement of Data Storage devices (review of item i): ICF requested any additional context from stakeholders on the continued impact of Capacity Optimising Methods (COMs). Additionally, how COMs are operating in the market and if there any new COMs that need to be considered. ICF described how the study will also look at the SNIA Emerald Benchmark as the basis of active level performance.

Processor Power Management Function (review of item g): Interact mentioned that current research from IEEE shows that when processor power management functions are enabled this can reduce the power consumption by 18–50% for no loss of performance. Intel discussed how older models have greater savings and how scalable CPUs are less impactful when power features enable BIOS levels, whereas pre-generation CPUs are affected by these features. It was suggested that Ecodesign could regulate shipping in “balanced configuration in BIOS” to enable savings.

Standby-Readiness for Servers (review of item p): Dell acknowledged that to go into a standby mode, all components would need to save their state. This is not possible without a complete reboot. Even after many years of testing there is a risk of system crashing. Therefore, there has been little incentive to work on this because the risk is too high.

DG GROW explained this item relates to the maximum consumption in idle power state. In the previous preparatory study, there were comments raised late on that there should be more granularity on information provided on the standby state (deep stand by, etc..). Thus, the idea of this item is to consider if anything should be done at a regulatory level or not.

Parameters Information Requirements (review of item r): Dell explained that most fans do dynamic fan speed control. However, they discussed that for Dell products once an external system controls the fan speed of a server, the warranty of the server would become void.

DG GROW emphasised that part A of this item would cover the disclosure of information only on parameters related to the cooling. The idea is to see if there is any feasibility/relevance in improving the type of fan speed information provided by manufacturers to the user.

ECOS stated that they currently report the power of servers in a standardised way to allow the measurement of it. Hence, this would help create synergy with the EE Directive. Additionally, ECOS suggested that reporting utilisation would also enable the server idle coefficient to be reported.

Dell expressed that they would not be able to measure all items that are requested as a product manufacturer because there are requirements will need to be conducted by the operating system.

Energy Label (review of items s): Dell and Intel stated that it is not possible to create an energy label for the many configurations they have. They also question who the label will benefit. DG GROW replied stating that this is a new idea thus feedback is sought. During the first regulation there were efforts to include the intrinsic variability of servers. Therefore, it was decided we make these declarations at two ends of the families/configurations spectrum of servers. The low-end and high-end performance configurations. DG GROW preliminary view is that an energy label that has to be tailored for every configuration would be difficult to impose. The approach of calculating the energy label for the low-end and high-end performance configurations would for instance appear as a suitable one.

Following this, Interact described that they have created a grading system for servers. This calculates a rating based on the SSJ Ops/Watt, using banding A+ to F. It uses SERT and SPEC power to validate this. Furthermore, it would be possible to use machine learning to fill in the gaps of the other SERT worklets. Interact described in detail how data centre operators have used this information to help with reaching decarbonisation goals. They stated that the energy performance between the minimum and maximum performing products can range between 70–80%. Therefore, their label helps the end-user choose a more efficient product.

The German Environment Agency and TechBuyer Europe expressed that the label would be useful. In particular for IT procurement and reporting. IBM argued that the usefulness of a label comes before a product is bought, however, it is “how” this information is communicated that could be an issue. DG GROW stated that the label information could be displayed via the EU's EPREL database, as it is done for other product groups. Information could also be provided via the Product Information Sheet.

ECOS felt that an energy label is restrictive in what information it contains. It was suggested that it may be more useful to provide this information separately. They discussed how an energy label could be provided by a machine learning algorithm, if the regulation states how this should be carried out.

DG GROW explained that due to the highly customised nature of servers, they may not be suitable for the Energy Label. The study team requested feedback on whether a tighter scope might focus which servers are labelled.

Material Efficiency (review of items d, j, l): Digital Europe expressed doubts about material efficiency. They queried if the information about Critical Raw Materials (CRM) is useful to recyclers. Techbuyer Europe stated it would be good to capture this information, because this will support future reporting and assist current/future recycling efforts. It was also discussed how the development of technologies (such as pyrolysis and bioleaching which increase the number of materials that can be recovered simultaneously), has meant that recycling processes have become more efficient. Thus, meaning more CRMs from the server and data storage products can be reused. Cisco highlighted that there could be issues when the products get upgraded, because the original information provided will no longer be accurate. Cisco inquired if the regulation would demand that information need to be updated after every upgrade to the product. Techbuyer Europe suggested that it would be helpful to have a banding of material contents.

DG GROW clarified some details on the aspects of material efficiency. It was mentioned that in the first regulation, material efficiency requirements were introduced by the regulation. The Commission have requested ICF analyse the transition towards a language that is similar to the pending smartphones regulation.

DG GROW state there is a standardisation request from the Commission which calls for product specific standards on material efficiency. The goal is to have everything aligned: the product specific standard and the requirements in the regulation. During the last study, stakeholders requested that we explore whether the requirements of the regulation could be imposed on the circular economy aspects of networking equipment, and if this could be done in a standardised way for all networking products. DG GROW requested feedback on the possibility of whether there could be a horizontal material efficiency requirement for networking equipment.

Juniper Networks felt like this was a bolt on approach. Juniper are concerned that material efficiency measures could be missed. In response to DG GROW's request on networking equipment they mentioned that there should be energy efficiency requirements for networking equipment.

Operating Conditions (Review of items h, k): DigitalEurope raised the question of whether we are looking at ASHRAE Requirement level A5 or A6. ICF provided a reply that the A1 class is generally accepted as the standard which Data Centres are typically designed around. Following this, Kao Data provided context around the thermal guidelines within the ASHRAE 9.9 standard. They anticipate a further version of the ASHRAE standard to be published soon.

System Performance Considerations (Review of item m): The German Environment Agency described how there are other metric tools to consider for server performance, such as: year/year renewable energy factor (REF); Water Using Effectiveness (WUE); and the cooling efficiency ratio (CER). Kao Data described how WUE is coming to the forefront, however, it is geographically limited. All three are being considered by the Climate Neutral Pact. The German Environment Agency stated that PUE is a useful parameter, but it should not be considered in isolation. Only together with the other indicators (CER, WUE, REF) can we get a whole picture.

Liquid Cooling Systems and Solutions (Review of item n): ECOS stated liquid cooling is one of the beneficial factors for more heat re-use in servers and should be considered. Kao Data raised awareness that ASHRAE have published fluid and liquid temperature guidelines for liquid cooling with temperature classes. DG GROW raised that the possibility to apply liquid cooling is very feasible but wanted to understand if this was a niche market.

Waste Heat Recovery Systems and Solutions (Review of item o): ECOS stated that there could be synergies between waste heat and the EE Directive. ICF explained that from a system level, they are trying to understand this, in order to inform regulators.

Direct Current Power Supply for Servers (Review of item q): AMD explained how the finalised scope and timing of SERT V3 release is not yet finalised. The DC power servers are likely to be in scope, but this is not guaranteed.

Other Topics – No comments raised.

AOB

Intel asked for clarification on when the regulation will be published. DG GROW reiterated that taking the pending smartphone regulation as a guide, it is unlikely the updated regulation will be published before Q2 2026.

Closing statement

DG GROW shared their appreciation for everyone that attended the meeting today. They encouraged all attendees to participate and follow the process for this regulation. DG GROW wanted to know if any of the manufacturers who attended the session could help the project team to engage with recyclers. DG GROW are looking to expand the audience of the review study. ICF closed the presentation, requesting that all attendees complete both the

[qualitative questionnaire](#) by **28th April 2023**, and the quantitative questionnaire by **12th May 2023**. Presentation slides and meeting minutes are uploaded on the study [website](#).

Draft