Reporting requirements on the energy performance and sustainability of data centres for the Energy Efficiency Directive

Task A report: Options for a reporting scheme for data centres

October – 2023

Viegand Maagøe

COWI

EUROPEAN COMMISSION

Directorate-General for Energy
Directorate B — Just Transition, Consumers, Energy Security, Efficiency, and Innovation
Unit B.2 — Energy Efficiency policy and financing

Contact: Nikolaos Kontinakis

E-mail: nikolaos.kontinakis@ec.europa.eu

European Commission

B-1049 Brussels

EUROPEAN COMMISSION

Reporting requirements on the energy performance and sustainability of data centres for the Energy Efficiency Directive

Task A report: Options for a reporting scheme for data centres

Directorate-General for Energy

2023 EN

Manuscript completed in October 2023

Luxembourg: Publications Office of the European Union, 2023

PDF ISBN 978-92-68-08438-0 doi:10.2833/304891 MJ-09-23-473-EN-N

TITLE:	**Technical assistance in support of the introduction in the EED provisions (Commission proposal for the EED recast) of reporting requirements on the energy performance and sustainability of data centres.** **Specific contract No ENER/B2/2022-487/SI2.878821 implementing framework contract No OP/0021/ENER/C3/2020-724 FWC EE & RES**
DATE:	**09/10/2023**
PROJECT NO:	**2857**
VERSION:	**Final**
PREPARED BY:	**Viegand Maagøe and COWI**
PREPARED FOR:	**European Commission – Directorate-General for Energy**
QA BY:	**Jan Viegand**
APPROVED:	**Jan Viegand**

VIEGAND MAAGØE A/S

ZEALAND MAIN OFFICE
NR. FARIMAGSGADE 37
DK 1364 Copenhagen K
Denmark

T +45 +45 33 34 90 00
info@viegandmaagoe.dk
www.viegandmaagoe.dk

CBR DK29688834

JUTLAND
Samsøvej 31
DK 8382 Hinnerup
Denmark

Contents

1 Introduction

Data centres are estimated to have the fastest growing energy consumption and carbon footprint across the whole ICT sector, mainly due to technological advances such as cloud computing and the rapid growth of the use of Internet services. In the EU, from 2010 to 2018 data centre energy consumption increased by 42% and is forecast to further increase by 28.2% by 2030[1], representing about 3.2% of the EU final electricity demand. Reducing the energy demand of ICT, including data centres, is an important step in achieving the ambitious climate targets of the EU – a reduction of overall GHG emissions of 55% by 2030 compared to 1990 levels.

The European Green Deal and the Climate Law (Regulation (EU) 2021/1119) represents a major leap in that direction. As a part of the 'Fit for 55' package, the recast Energy Efficiency Directive[2] (EED recast), includes new obligations that target data centres.

To support sustainable development in the ICT sector, the Article 12 of the EED recast asks data centres to make information about their energy performance publicly available. These reporting requirements would apply to all DCs, old and new, whose IT installed power demand is above a threshold.

2 Objectives

In line with the EED recast, the Commission will adopt a delegated act to establish a common EU scheme for the reporting of the sustainability of data centres within EU.

While the negotiations were ongoing and within the frame of uncertainty about its final outcome, the Commission started preparing the necessary evidence by means of this study and whose specific objectives are to:

- Organise a consultation with relevant stakeholders and Member State representatives.
- Assess and propose the main elements that will define the scope of reporting on the energy performance and sustainability of data centres: how are data centres defined, which data centres will be required to report and possible exceptions to this obligation
- Assess and propose the main elements of the reporting scheme: which entities will be responsible to fulfil the reporting obligation for each data centre, access to data and ways to ensure the consistency and quality of the reported data
- Identify, assess and propose the possible key performance indicators that can be used to report the energy performance and sustainability of data centres, as well as their respective calculation methodologies
- Identify, assess and propose the possible data and information, which data centres will need to report along with the key performance indicators
- Propose options for an EU-wide repository that will be used to gather, keep and publish the reported data

3 Structure and objectives of this document

Task A report contains proposals for the topics covered by Task A, i.e.:

[1] Own calculations based on figures provided in European Commission (2020b). Energy-efficient cloud computing technologies and policies for an eco-friendly cloud market.
[2] OJ L 231, 20.9.2023, p. 1–111 EUR-Lex - 32023L1791 - EN - EUR-Lex (europa.eu)

- Proposal of elements that will define the scope of reporting on the energy performance and sustainability of data centres: the definition of data centres, the data centres that will be required to report and the type of information covered by the reporting scheme;
- Proposal of which entities will be responsible to fulfil the reporting obligation for each data centre and ways to ensure the consistency and quality of the reported data;
- Proposal for key performance indicators that can be used to report the energy performance and sustainability of data centres, as well as their respective calculation methodologies;

These proposals are built upon the consultation to the stakeholder in the meetings held on 6 December 2022 and 15 June 2023 and the written feedback received. Each proposal is accompanied by a rationale supported by relevant information and comments from stakeholders. Each section also includes an evaluation of the alternatives that are deemed less suitable.

4 Definition of data centre

The development and implementation of any policy measure requires a clear definition of the matter within its scope of action, in this case, data centres. The definition proposed must be unequivocal, set clear boundaries that encompass all data centres within the scope and be consistent with other EU policies. This section is mainly built on the definitions provided by standards and EU legislation, which has been supported by stakeholders that generally accept them as suitable for policy purposes.

4.1 Proposal

Eurostat definition according to Regulation (EU) 2022/132[3]

A *data centre* is defined as a structure or a group of structures used to house, connect and operate computer systems/servers and associated equipment for data storage, processing and/or distribution, as well as related activities.

Additional definitions

A structure can consist of multiple buildings and/or spaces with specific functions to support the primary function. The boundaries of the structure or space considered the data centre, which includes the information and communication technology equipment and supporting environmental controls, can be defined within a larger structure or building.

Associated equipment […] for related activities means all the facilities and infrastructures for power distribution and environmental control together with the necessary levels of resilience and security required to provide the desired service availability.

Computer room space is the area within the data centre that accommodates the data processing, data storage and telecommunication equipment that provides the primary function of the data centre.

4.2 Rationale

The amendment of the energy statistics regulation adopted by the European Commission in January 2022 contains a definition of data centre, which must be incorporated in the EED recast for reasons of consistency. The definition is *a structure or a group of structures used to house, connect and operate computer systems/servers and associated equipment for data storage, processing and/or distribution, as well as related activities.*

[3] OJ L 20, 31.1.2022, p. 208–271 EUR-Lex - 32022R0132 - EN - EUR-Lex (europa.eu)

Eurostat definition is compatible to EN 50600[4], keeping a similar concept. However, it may be too vague in some aspects, such as the terms “associated equipment” and “related activities”, which may require more details for the purpose of a reporting scheme. This is why we have provided an additional definition to complement Eurostat definition.

The definition of a data centre provided in the EN50600 Series of standards developed by the European Committee for Electrotechnical Standardization (CENELEC) is used by several standardisation organisations (ISO/IEC, ETSI, CEN-CENELEC). EN 50600 defines data centre as *a structure, or group of structures, dedicated to the centralised accommodation, interconnection and operation of information technology and network telecommunications equipment providing data storage, processing and transport services together with all the facilities and infrastructures for power distribution and environmental control together with the necessary levels of resilience and security required to provide the desired service availability.*

If necessary in the future, EN 50600 definition is proposed to complement Eurostat definition. From the definition of EN50600 and its notes, the additional definitions of “associated equipment for related activities” and “structure or group of structures” are put forward, if it is needed to clarify those terms of Eurostat definition in line with EN 50600.

For setting clear boundaries that outline the equipment within the data centre, the definition of data centre is accompanied by the definition of computer room space. The area comprised by the computer room space and the ICT equipment that this area contains define what equipment is to be considered in the scope of the reporting obligations. Other telecommunication equipment such as radio equipment that is outside the computer room space is out the scope.

The definitions from Eurostat and EN 50600 were generally supported by stakeholders, considered as based on international standards and useful for policy harmonisation, and adequate to capture any building, structure or container used as data centre, even small server rooms. However, the ability of the definition to include some decentralised data centres was questioned. Some stakeholders recommended to test the definition to ensure that is expansive enough to include distributed data centres or edge data centres, data centres not hosted in buildings, and data centres related to smart cities. A peculiar case of distributed data centres was brought up by a stakeholder: servers that are marketed as “digital boilers and radiators”[5]. The digital boilers consist of computer processors equipped with a water-based heat recovery system, while the digital radiator emits the heat directly, similar to a stove. In this regard, the notes accompanying the EN definition are deemed sufficient to cover these situations, as the term “spaces” is ample enough to cover these situations. In addition, the Guidance document will provide information to help the competent authorities and data centre operators to properly interpret the definition.

4.3 Alternatives

The Scoping document presented a proposal for a modification of the EN 50600 definition that explicitly included the IT equipment contained. However, this modification is no longer proposed since Eurostat definition together with the additional definitions, including those in section 5, clearly describe the ICT equipment that is within the structure of a data centre.

5 Classification of data centres and additional definitions

Different operating and ownership models for data centres exist that are relevant for the definitions of reporting obligations and may influence the data collection practices. This section contains a proposal to clearly characterise the different business models of data centres and the roles of the organisations

[4] https://standards.cencenelec.eu/dyn/www/f?p=CENELEC:110:::::FSP_PROJECT,FSP_ORG_ID:65128,1258297&cs=1D5A74123400C0818509CE78C52EF9463
[5] https://qalway.com/fr

involved, in a way that can be linked to the responsibilities of reporting. According to that same logic, it also covers additional definitions that are needed for the correct interpretation of the following sections, e.g., the definition of “installed IT power demand”.

5.1 Proposal

Classification based on EN 50600

- Enterprise data centre is a data centre that is operated by an enterprise, which has the sole purpose of the delivery and management of services to its employees and customers
- Colocation centre is a data centre facility in which multiple customers locate their own network(s), servers and storage equipment
- Co-hosting data centre is a data centre in which multiple customers are provided with access to network(s), servers and storage equipment on which they operate their own services/applications. Both the information technology equipment and the support infrastructure of the building are provided as a service by the data centre operator.

Additional definitions

- IT installed power demand is the sum of the nominal powers in kW of the installed information technology and network telecommunications equipment in the data centre providing data storage, processing and transport services together The IT equipment covered by this definition is equal to the IT equipment covered by the calculation of the power usage effectiveness, EN50600-4-2.

  Where IT installed power demand changes frequently or cannot be determined, the rated IT load can be used as an alternative.

  Rated IT load is the maximum load, in kW, from the network(s), servers, and storage equipment that the data centre infrastructure for power distribution and environmental control is capable of handling while providing the desired service availability.
- Enterprise data centre operator is a physical or legal person, who operates the entire data centre, including the physical building through to the use of the IT services delivered.
- Colocation data centre operator is a physical or legal person, who operates the entire data centre for the primary purpose of selling space, power and cooling capacity to customers, who install and manage their own network(s), servers and storage equipment and services.
- Co-hosting data centre operator is a physical or legal person, who operates the data centre space, power, cooling, network(s), servers, and storage equipment, and part of the necessary software to deliver IT services to customers, including IT outsourcing.
- Colocation customer is a physical or legal person, who owns and manages network(s), servers and storage equipment located in a colocation data centre in which they purchase managed space, power, and cooling capacity.
- Co-hosting customer is a physical or legal person, who procures access to network(s), servers, and storage equipment in a co-hosting data centre on which they operate their own services and applications.
- Redundancy is the duplication of certain components or functions of a data centre system so that if they fail or need to be taken down for maintenance, others can take over. "N" represents the number of pieces to satisfy the normal conditions. Redundancy is often expressed compared to the baseline of "N". Some examples are "N+1", "N+2", "2N," and “2(N+1)”.

5.2 Rationale

The classification set by EN 50600 is broadly accepted and supported by stakeholders and can also be linked to the responsibilities on the information on energy efficiency and environmental performance, ensuring the consistency with the definition of responsibilities. However, some comments indicated that the definition of co-hosting data centre fails to characterise a cloud facility since the definition should consider the level of control exercised by the customer, which depends on the service purchased. In this regard, the maximum level of control by the customer is at the virtualisation layer, meaning that in all cases the cloud service provider controls the physical ICT equipment. In other words, a cloud data centres operator has the same level of control on the physical ICT equipment than an enterprise data centre operator.

This discussion links to the concerns raised by some stakeholders on the lack of a software-oriented approach, since the reporting scheme is very much focused on ICT equipment and does not propose key performance indicators beyond that limit (i.e. end-users service indicators). In this regard, the decision to keep the reporting scheme within the boundaries of the physical ICT equipment (i.e. indicators related to servers, storage and network) is based on pragmatism, which prioritises the definition of parameters that are feasible and practical for all the entities involved.

In addition to the classification of data centres, this section includes a set of definitions that are relevant for the reporting scheme. In the first place, the term “ICT installed power demand” used by EED recast is deemed confusing for stakeholders, which required a specific definition. The proposed definition clarifies that it refers to rated input power of the installed ICT equipment, not operation power demand, and that covers all the ICT equipment that is measured to calculate power usage effectiveness (PUE). In addition, the definition is presented using the same terms as in the data centre definition from EN 50600, which completes the consistency of the different elements of the reporting scheme with existing standards. In addition, an alternative is proposed in case that IT installed power demand changes frequently or cannot be determined, based on rated IT load or maximum load provisioned for the data centre to provide its services.

Stakeholders also recommended to clarify and define the roles of the various entities that can be involved in the operation of a data centre, for a better understanding of the responsibilities of reporting (section 7). The European Code of Conduct for Data Centres[6] defines the following organisations:

- Operator: operates the entire data centre from the physical building through to the consumption of the IT services delivered;
- Colocation provider: operates the data centre for the primary purpose of selling space, power and cooling capacity to customers who will install and manage their own IT hardware and services;
- Colocation customer: owns and manages IT equipment located in a data centre in which they purchase managed space, power and cooling capacity;
- Managed service provider (MSP): owns and manages the data centre space, power, cooling, IT equipment and some level of software for the purpose of delivering IT services to customers. This would include traditional IT outsourcing;
- Managed service provider in colocation space: a managed service provider which purchases space, power or cooling in a data centre in order to provide services to third parties.

Some comments pointed out that the reporting scheme does not need so many definitions and classifications. They should be as simple as possible and directly related to the reporting responsibilities,

[6] https://e3p.jrc.ec.europa.eu/communities/data-centres-code-conduct

meaning the only useful distinction is colocation data centres and other data centres, as compiling all relevant data might require collaboration with colocation customers, together with the definitions of data centre operator, co-colocation operator, and colocation customer. The definitions finally proposed are aligned to the classification of data centres based on EN 50600 and their reporting roles described in section 7.

Stakeholders suggested to also define redundancy as a relevant aspect influencing the energy consumption, hence a definition is included to clarify this term.

5.3 Alternatives

Task A scoping document proposed a classification based on installed IT power, in case that the reporting scheme established different levels according to the size of the data centre. This was generally accepted by stakeholders, while other parameters such as floor area were considered unsuitable.

The additional classification on IT installed power is not included because the reporting scheme as proposed in this report does not provide for any differentiation of that kind. However, if it were needed, the typical classification is as follows:

- Server room: <50 kW
- Very small Data Centre: 50–250 kW
- Small Data Centre: 250–1000 kW
- Medium size Data Centre: 1–2 MW
- Large Data Centre: 2–10 MW
- Very large Data Centre: >10 MW

As mentioned in the rationale, a stakeholder suggested the following modification of the definition of co-hosting data centre to better reflect the reality of cloud data centre, which may be included in the Guidance document for clarity:

A data centre offering server and data storage services where the customer pays for the service and may install and manage aspects of the middleware, operating system and/or hypervisor and applications. The managed service or cloud provider manages the required hardware, data centre equipment and/or software (as applicable to the provided service). The level of customer control depends on the service purchased.

6 Scope of the reporting obligations

The EED recast proposal established a general criterion for including data centres in the scope as those with a significant energy consumption. The rapporteur (a member of the European Parliament made responsible for handling a legislative proposal drawn up by the European Commission, both procedurally as well as with regard to the substance of the proposal) replaced that term by a threshold of 100 kW of installed IT power[7]). However, the final text of the EED recast raised the threshold to 500 kW.

In terms of information, the Annex VII of the EED recast sets minimum information that must be provided by data centres. In a similar structure, this section addresses the scope of the reporting obligations, in two aspects:

- The delimitation of those data centres that should be within the scope of the reporting scheme
- The identification of the information that is relevant for the reporting scheme.

[7] [7] https://www.europarl.europa.eu/doceo/document/ITRE-PR-703281_EN.pdf

6.1 Proposal

6.1.1 IT power demand threshold

Data centres with a power demand of the installed information technology (IT) of at least 100 kW

Note: this proposal is not in line with EED recast final text.

6.1.2 Elements of the data centre included in the scope of the reporting obligations

A) Information of the data centre

- **Building Information**
  - Data centre name is the name used to identify and describe the reporting data centre.
  - Owner and operator of the data centre shall include the name and contact details of the owner and of the operator of the reporting data centre.
  - Location of the data centre is the Local Administrative Unit Code (LAU code) of the location of the reporting data centre (building or site) expressed in accordance with the most recent LAU tables published by Eurostat.
  - Type of data centre is the type of the reporting data centre that matches better the operation of the reporting data centre, in accordance with the definition of data centre and the definitions of type of data centres of this Regulation.
  - The type of a reporting data centre can take one of the values ‘enterprise data centre’, ‘colocation data centre’ or ‘co-hosting data centre’, combined with one of the values ‘structure’ or ‘group of structures’.
  - Year and month of entry into operation is the calendar year and month that the reporting data centre started its operations.
- **Reporting organisation domain of control (indicate yes or no)**
  - Physical building
  - Mechanical & electrical plant
  - Data floor
  - Racks
  - IT equipment
  - Operating system / Virtualisation layer
  - Software
  - Business process
  - All the above
- **Operation**
  - Electrical infrastructure redundancy level
  - Cooling infrastructure redundancy level
  - Total number of modular capacity steps or separately provisioned halls
  - Total number of racks

    For the redundancy levels, if "N" represents the baseline number of components or functions to satisfy the normal conditions, redundancy shall be expressed compared

to this baseline of "N", e.g., as "N+1," "N+2," "2N", etc. Facility redundancy can apply to an entire site (backup site), systems, or components. IT redundancy can apply to hardware and software.

B) Key performance indicators (section 9)

6.2 Rationale

6.2.1 Threshold

The minimum threshold of 100 kW of installed IT power demand determines the scope of the obligation in terms of size and leaves out server rooms and a share of very small data centres. According to JRC Preliminary report for EU GPP of data centres[8], enterprise data centres are usually smaller than colocation and co-hosting data centres but 96% of data centres in EU are enterprise data centres. However, the proposed threshold is relevant in terms of capacity of the data centre to fulfil the reporting obligations: smaller data centre may encounter difficulties to access the necessary means for this reporting and this must be taken into account for defining the scope of the reporting scheme.

The threshold of 100 kW has been discussed with stakeholders and it is generally accepted as balanced point: it is able to capture the reality of the sector, which is characterised by a significant number of small data centres, while it is high enough for the reporting scheme to be manageable, avoiding too much administrative burden on very small data centres. In terms of energy consumption, according to feedback from stakeholders and a report by 451 Research[9], one third of the energy consumed by the data centre sector corresponds to 16000 enterprise data centres with an IT power of more than 100kW. JRC Preliminary report for EU GPP Data centres estimated that, in 2017, 60215 data centres in EU were enterprise, 2215 colocation and 152 co-hosting. Crossing the figures of both studies, the conclusion is that 2/3 of the energy consumption (67 TWh according to 451 Research) is distributed between around 55000 very small data centres, not covered by the 100 kW threshold, and less than 800 professional (colocation and co-hosting) data centres, covered by the threshold. IEA[10] estimates that in 2021 70% of the energy demand of data centres globally was consumed by cloud and hyperscale data centres. These different sources of information suggest that 100 kW is an appropriate threshold to cover a significant share of the energy demand of the data centre sector, while capturing a representative share of small data centres. In addition, private and public sector dedicated data centre rooms are the most energy inefficient, thus, the ones with a larger energy saving potential. The EURECA Project[11] provides good data on this aspect, with average PUE of 4-5 for small data centres.

Other clarifications may be also needed to interpret if specific situations are within the scope, which are to be addressed in the Guidance document. For example, some stakeholders asked to clarify what criteria would apply in those situations where in the same space, usually governments and universities, there are multiple small data centres, which may not reach the threshold separately. A stakeholder suggested to aggregate all data centres connected to the same substation.

Feedback from a Member State indicated that installed power capacity is a suitable parameter, but it can change according to modification in equipment along time. This may entail that some data centres are not consistently eligible to the application of the EED recast. In France, regulation for reduction the energy consumption of the tertiary sector is based on the IT surface area (m2), which is constant over time and can be correlated to IT installed power. It is considered essential to get a correspondence

[8] https://susproc.jrc.ec.europa.eu/product-bureau//product-groups/458/documents
[9] https://cispe.cloud/improving-datacenter-efficiency-in-europe-the-role-of-pue/
[10] https://www.iea.org/data-and-statistics/charts/global-data-centre-energy-demand-by-data-centre-type-2010-2022
[11] https://dceureca.eu/#primary

between the “installed power capacity” and the “IT surface area” to facilitate the implementation by competent authorities. Turnover was also suggested as parameter to be communicated in the first data collection, to help competent authorities identify the operators within the scope of the reporting scheme. The identification of data centres in the scope is a relevant aspect for implementation and will be addressed in the Guidance document.

6.2.2 Elements within the scope of the reporting scheme

Annex VII of the EED recast reads as follows:

The following minimum information shall be monitored and published with regard to the energy performance of data centres referred to in Article 12:

a) the name of the data centre, the name of the owner and operators of the data centre, the date on which the data centre started its operations and the municipality where the data centre is based;

b) the floor area of the data centre, the installed power, the annual incoming and outgoing data traffic, and the amount of data stored and processed within the data centre;

c) *the performance, during the last full calendar year, of the data centre in accordance with key performance indicators about, inter alia, energy consumption, power utilisation, temperature set points, waste heat utilisation, water usage and use of renewable energy, using as a basis, where applicable, the CEN/CENELEC EN 50600-4 ‘Information technology – Data centre facilities and infrastructures’, until the entry into force of the delegated act adopted pursuant to Article 33(3).*

A comprehensive list of elements of the data centre was proposed for discussion with stakeholders in Task A Scoping Document. The list included the minimum requirements for monitoring and publishing the energy performance of data centres in the EED recast, together with other elements in a comprehensive way for discussing all the possibilities. Stakeholders raised different issues:

- Confidentiality issues, particularly for critical data centres for security purposes
- Commercial sensitive information on equipment that should not be shared with competitors, and for security reasons
- Too detailed information on equipment which is not feasible to collect, in particular peripheral equipment such as lighting or supporting laptops.
- Specifically, the annual incoming and outgoing data traffic; and the amount of data stored and processed within the data centre were deemed unfeasible to measure
- No apparent reason from the energy performance perspective to monitor some elements proposed information

To understand the motives behind the request of minimum information, it is necessary to characterise that information:

- Point (a) of Annex VI requests basic information to identify the data centre and the organisations involved. Competent authorities require this basic information and the only debatable question is whether it should be publicly available in certain situations.
- Point (b) of Annex VI requests some technical details to describe the data centre such as the floor area and the IT power. Competent authorities need some descriptive details on the data centre reported, for verification purposes and for the evaluation of the policy measure. It also includes annual incoming and outgoing data traffic; and the amount of data stored and processed within the data centre, which are discussed in section 9
- Point (c) is about key performance indicators that are discussed in section 9

The information required in point (b) has different purposes, as explained above. A minimum level of detail is necessary to analyse the key performance indicators in relation to the technical characteristics of the data centre, since a figure of PUE or absolute energy consumption is meaningless for policy-making if it cannot be placed in a context for understanding the dimensions and variations of the reported data. That context is the technical description of the data centre, which is also needed to understand the issues that may encounter the different parties involved in the reporting, under the various business models. For this reason, the proposal in this report expands the minimum information required in point (b), inspired in the application form of the Code of Conduct which distinguishes:

- Building Information
- Operator Domain of Control
- Operation

The concerns on confidentiality and commercial sensitive information are already addressed by Article 12 of EED recast, which includes an exception for *information subject to Union and national law protecting trade and business secrets and confidentiality* and for those *data centres used for, or providing their services exclusively with the final aim of, defence and civil protection*.

Another option for ensuring confidentiality and transparency is, as recommended by a stakeholder, publishing data aggregated, after being consolidated, anonymized, and devoid of any reference to a particular data centre. This could be done aggregating the reported data per type of data centre, or per installed IT power range.

6.3 Alternatives

Notice that the final text of EED recast sets the threshold of the reporting scheme scope in 500 kW and that the considerations and suggestions mentioned hereafter were made before the final text was published.

Some stakeholders advocated to lower the threshold to 50 kW to cover very small data centres, as there is a trend in edge computing due to the proliferation of 5G, industrial IoT and the adoption of software-defined networking and network functions virtualisation (SDN/NFV) technologies. While this is a valid argument, lowering the threshold would increase exponentially the number of data centres in the scope, which will hinder its already complex implementation.

Other proposal from a stakeholder was a specific threshold of 1MW for high-performance computing, to recognise the purpose of the data centre. Along the consultation, it was proposed that the detail of reporting could be proportional to the size of the data centre, however this option is not in line with the objective of the reporting scheme, which is providing real and first-hand information on the data centre sector. This need of information is acknowledged by most stakeholders which consider that any exception or special treatment can affect the completeness and transparency of the information.

Other option consists of increasing the threshold between 250 and 500 kW, in order to focus on those data centres that are most intensive in energy consumption. This would leave small, usually non-professional data centres out of the scope of the reporting scheme, and could facilitate the identification of the data centres within the scope and implementation of the reporting scheme by competent authorities.

Another approached suggested by a stakeholder was to set the threshold on the energy consumption, instead of installed power. This will align the threshold of the reporting scheme with other thresholds in the EED recast which are based on energy consumption. However, this would be an obstacle for competent authorities to identify which data centres are affected by the threshold, who ultimately would have to estimate the energy consumption based on the installed power.

A stakeholder proposed to use installed capacity of supporting infrastructure, instead of IT: 400kW infrastructure with 99kW of installed IT power. The threshold would be more representative of the energy consumption, but it would not mean a significant change in the scope.

As for the list of elements, the rationale explains the main points of discussion with stakeholders. Additional information proposed by them is further analysed in section 9.

7 Responsibilities to fulfil the reporting obligations

The implementation of the reporting scheme requires to identify the entity responsible to comply with these obligations and who has to interact with the competent authorities. In the data centre sectors, there are business models where several organisations participate in the operation, so the identification of the responsible for reporting entity can be challenging.

7.1 Proposal

For enterprise data centres and co-hosting data centres (as defined by EN 50600, i.e. the same organisation operates ICT equipment and support equipment), the operator of the data centre is the responsible for reporting the information and data defined in the reporting scheme.

For colocation data centres (as defined by EN 50600, i.e. the organisation that operates the ICT equipment is different from the organisation that operates the support equipment), the colocation provider is responsible to declare the energy consumption of the whole data centre, both support and ICT equipment. The colocation customers are responsible to provide the productivity and functions data to the colocation provider, who will be responsible to submit it in an aggregated way to the competent authority. If there are legal or competition issues that prevent the colocation customer to provide the data to the colocation provider, the collocation provider could appoint a third independent party, bound to non-disclosure agreement, to collect and aggregate the data of the colocation providers, and if necessary, also submit it to the competent authority.

7.2 Rationale

The EED recast establishes that owners and operators of data centres are to be required to report the information about the data centre. It is important to distinguish between operation and ownership, since the responsibilities may be different depending on the type of policy measure. In the case of reporting obligations on energy consumption, the responsibility could be on the operators, or on the owners. Since operators can control the main parameters affecting the consumption, from a technical point of view it is reasonable that they are responsible for the reporting.

The different roles involved in the various DC business and operation models have been defined in section 5 and serve as a basis to further elaborate the responsibilities in reporting data.

For enterprise data centres and co-hosting data centres, there is just one entity operating both support and ICT equipment, therefore, the responsibility is clear. In colocation data centres, colocation customers operate their ICT equipment, and the colocation provider operates the support equipment. The three possible options were presented and discussed with stakeholders:

- Option 1: The colocation customer is responsible to declare the energy consumption associated to their IT equipment and their share of energy consumed by support equipment, and the annual incoming and outgoing data traffic; and the amount of data stored and processed within the data centre.
- Option 2: The colocation provider is responsible to declare the energy consumption of the whole data centre, both support and IT equipment, and the annual incoming and outgoing data traffic; and the amount of data stored and processed within the data centre.
- Option 3: Joint responsibility of colocation provider and colocation customers.

Option 2 is acknowledged as the most pragmatic solution, in particular, for the competent bodies and authorities in charge of the implementation for whom it would be unfeasible to manage the number of reporting entities (one per each colocation customer) that Option 1 and 3 would entail. Option 2 would enable only one entity bearing the responsibility of reporting the information of the whole data centre

as a unity. Option 3 on the contrary, would imply several organisations and contractual agreements, hence there is a risk of incomplete reporting if some of the customers fail to provide the information.

However, along the consultation processes some stakeholders raised objections against this option, since it would entail legal and contractual issues between colocation providers and their customers, and strongly recommended either option 3. It is acknowledged that the colocation provider will be an intermediary in the reporting of data related to the activity of the colocation customers and they cannot be accountable for the quality and accuracy of that data.

A similar logic leads to the conclusion, raised by a stakeholder, that Options 2 and 3 are essentially the same option. This is because Option 2 will also require that colocation provider and ICT operators establish contractual clauses which specify the responsibilities for measuring, aggregating, sharing, and reporting mandated information. As the stakeholder clarifies, these contractual clauses can be done within the services agreements (Option 2) or an externally established code of responsibility (Option 3), and Option 2 is considered the simplest to implement.

In that sense, some comments stressed the need of a provision in the EED recast for colocation customers to provide the required data to the colocation provider, for a common legal framework to support those contractual clauses. A Member State suggested that a draft template of contractual agreement between the two parties could be made available at EU level.

Some stakeholders have advocated for Option 1 to avoid that colocation customers were obliged to share sensitive information with potential competitors, since many organisations in the data centre sector develop different activities. However, Option 1 is not a feasible alternative for the reasons explained above, hence this issue must be addressed with a different approach. As reasoned previously, Option 2 needs to be redefined as a compromise solution between the original Option 2 and Option 3, that is, recognising the responsibility of ICT operators in the data reporting. Therefore, colocation providers are responsible for the collection, aggregation and submission of the data, but not for its quality and accuracy. In basic terms, they are intermediaries between colocation customers and the competent authority. This structure enables the introduction of another intermediary between the colocation provider and colocation customer if colocation customers are not willing to share sensitive information with the colocation providers. This additional entity must be external and independent and can act as a screen between both sides of the colocation data centre. To ensure the secrecy of the data, this third party could sign a non-disclosure agreement.

7.3 Alternatives

Task A Scoping document presented the three different option that were theoretically possible. Option 1 proposed that colocation customers were responsible of reporting, however, this option is discarded because it could create double-counting, incomplete information and would significantly multiply the number of reporting entities.

8 Ways to ensure the consistency and quality of the reported data

The consistency and quality of the information provided in the reporting scheme needs to be supported by a compliance assurance system. The EED recast sets some verification tools in Article 11, which are proportional to the energy consumption of the organisation, and therefore they can easily accommodate the wide variations in size of the data centre sector.

8.1 Proposal

The data reported by the operator responsible for reporting will be self-declared. The operator must carry out a quality assessment of the data reported according to internal management procedures which will also cover third parties involved (see colocation data centres in the section 7). The reported

data and the results of the quality assessment may be selected randomly for quality check by the competent authority.

Additionally, those data centres that are within the scope of Article 11(1) or Article 11(2) will include the data required by the reporting scheme for data centres and the internal management procedures in either their energy management systems or their energy audits.

8.2 Rationale

The EED recast sets energy management systems and energy audits for enterprises in Article 11, which distinguishes two levels of verification dependent on the energy consumption:

1. Member States shall ensure that enterprises with an average annual consumption higher than 85 TJ of energy over the previous three years, taking all energy carriers together, implement an energy management system. The energy management system shall be certified by an independent body, in accordance with the relevant European or international standards.

Member States shall ensure that the enterprises referred to in the first subparagraph have an energy management system in place at the latest by 11 October 2027.

2. Member States shall ensure that enterprises with an average annual consumption higher than 10 TJ of energy over the previous three years, taking all energy carriers together, which do not implement an energy management system are subject to an energy audit.

The feedback from stakeholders suggested to integrate the quality assurance of data centres reporting scheme in Article 11, avoiding the duplication of verification systems. Therefore, the proposal established that the elements of the reporting scheme must be in the scope of the energy management system and energy audits. Their scope should also include any third party involved in data collection, i.e. the specific procedures that colocation data centres have in place for collecting, aggregating and submitting the data from colocation customers.

According to stakeholders, the threshold of 10TJ will not suffice to cover all data centres in the scope of the reporting scheme. In addition, the energy audits are to be carried out every four years, and the third-party audits for the certification of management systems may not be annual. In contrast, the reporting scheme must have a defined reporting period for all data centres that should be ideally one year. For these reasons, it is proposed that all data centre operators report the data as self-declaration (supported by internal management procedures), including those affected by Article 11(1) or (2). This allows to have a common and basic way to ensure the consistency and quality of the reported data applying to all data centres, combined with the other two stricter levels of compliance assurance already provided by the EED recast.

8.3 Alternatives

Some stakeholders were in favour of third-party verification as universal compliance assurance tool and indicated that the costs could be minimised since the data audited can be collected by electrical meter readings. In addition, that a single audit could suffice for ensuring compliance with several EU policies, if they were implemented by means of the same international standards.

This option has been discarded because the variation in size of data centres is too wide, hence the impact of the cost of a third-party verification would differ too much in the range of sizes within the scope of the reporting scheme. According to the input from stakeholders, third party verification audits cost between 5000 and 10000 euros per day depending on the size of the organization and the level of assurance provided by the audit.

9 Key performance indicators and methodologies for energy performance and sustainability of data centres

The EED recast includes an obligation for the reporting of the energy performance of data centres with the aim of eventually establishing further measures to improve energy efficiency. It also empowers the Commission to adopt delegated acts to define the data centre sustainability indicators and to set out the key performance indicators and the methodology to measure them. This section looks into various options for key performance indicators and their methodologies.

There are many different ways to present the different components or layers of the data centre. A simplified description is shown in Table 1, which corresponds to the most common data centre categories and ownership/operating models. Each layer influences the data centre sustainability in some way and because there are many different operating and business models, ideally, key performance indicators to measure the activities and energy consumption at every layer would be included in an EU-wide reporting scheme.

Table 1: Layers and their functions in a data centre

Layer	**Function**
Software	Applications
Platform	Binaries, libraries
Virtualisation	Containers, hypervisor, operating system
ICT equipment	Network, servers, storage
Infrastructure	Power, cooling, space, security equipment
Location	Site selection, climate, land, power source, water source, heat re-use feasibility

Data centres are often described as similar to a factory, but this underplays their complexity. A single data centre can be analogous to an entire industrial zone supporting many factories producing a wide range of unrelated goods, including the necessary logistics in the form of warehouses and transport to store and move the inputs and final products as and when they are needed.

While it is simple to identify what we would like to measure to determine sustainability, i.e. energy consumption and productivity, in practice this may be not be technically feasible, or even possible to define in a way that is applicable to every data centre. Key performance indicators (KPI) therefore provide a proxy of the characteristics we are interested in.

The KPI and methodologies proposed are assessed against the following criteria:

1) Representativeness of desired characteristic
2) Ease of implementation, including access, cost and resource required:
    a) Access to measurement. Measurements can be taken at different layers and on different pieces of equipment. Ideally measurements should be taken as close to the components as possible. However, when ownership is split between many parties this can be complex

to achieve both contractually and practically. The addition effort and risk of errors may outweigh the marginal increase in accuracy.

b) Data sampling. Many of the KPIs are dynamic, their value changes continuously and generally need to be captured over a year to account for seasonal and business cycles. This could be achieved by (near-) continuous measurement or periodic sampling. Continuous sampling is most representative but can create too much data and overhead, with a measurable impact on energy consumption. Periodic sampling can be more practical but the sampling methodology should ensure that it is sufficiently accurate.
It is also possible to sample a subset of the equipment rather than monitoring all devices, especially if access to the devices is restricted, but it will also be necessary to determine how much this affects accuracy for each indicator.

c) Cost/resources. measurements can only be made if the equipment has the correct sensors/monitoring (correctly calibrated) and reporting capabilities. Without these, equipment may need to be replaced or additional monitoring equipment purchased. Any measurement made over many devices or very frequently will most likely require some sort of automated system or process to record the data. The most common approach is to use DC infrastructure management software, assuming access to the data is available. The costs incurred and resources available to implement the monitoring should be considered when selecting a KPI.

The most suitable KPI will require a compromise of the above points, while considering both the current and future state of the DC as well as the resources and capabilities of different sized DCs to measure them.

International standards already specify a variety of KPI, methodologies and sustainability indicators[12] (i.e. metrics[13] that enable comparison by combination of different KPI). The standards development process will have already discussed these issues and been agreed by technical experts. Using KPI from existing standards is the preferred approach unless there are significant issues with the representativeness or ease of implementation.

KPI will be required to be reported and in the first instance are not intended to be used for comparing data centre efficiency. An KPI can be defined as a measurement or value which gives an idea of what something is like. It is not necessarily a direct measurement of what we what to know, and KPIs are distinct from any sustainability indicator that may be developed for the purpose of comparison and are likely to be calculated from a combination of different KPI. It is acknowledged that there are also many factors related to the location and the indicators should be interpreted within that context when making comparisons. These factors will be considered in Task B.

Based on a review of existing research and standards, there are three types of KPI that are relevant to this project, which can be measured at different layers:

- Energy consumption
  - Data centre (total DC energy)
  - Infrastructure layer (power, cooling, security etc)
  - ICT equipment layer (ICT equipment energy)
  - Virtualisation layer (idle state energy)
- Productivity and data functions (processing, storage and transport)
  - Data centre (DC productivity)
  - ICT equipment layer (ICT equipment capacity)
  - Virtualisation layer (ICT equipment utilisation)

[12] In this report, sustainability indicator is the term use to refer to those metrics that are calculated from the proposed KPI
[13] In this report, metric is the term used to refer to those indicators that are calculated from two or more basic indicators.

- Ancillary KPIs
    - Water use
    - Energy/heat reuse
    - Renewable energy use
    - Energy related carbon emissions

None of the energy consumption KPIs in DC provide an energy efficiency figure in terms of energy consumed per work delivered by the DC, for which productivity and data functions parameters should be integrated in performance metrics.

The following sections provide proposals, rationale and questions to stakeholders for the three types of KPI. Data centre is usually abbreviated as DC in the terminology of the KPI, in line with EN 50600 standards.

9.1 Energy consumption

9.1.1 Total DC energy consumption

9.1.1.1 Proposal

E_{DC} total data centre energy consumption (annual) in kWh as defined by, and using the methodology, in EN 50600-4-2.

In addition, all E_{DC} contributions from backup generators shall be reported separately.

Data centres that provide additional electrical grid functions such as peak demand shifting and firm frequency response (FFR) should report the functions being utilised. DCs providing FFR should report the average capacity available (in MW) and the amount of time it is made available.

9.1.1.2 Rationale

The total energy consumed by the DC is the most basic KPI and is needed to track energy consumption of the data centre and the industry as a whole. It also needed for many sustainability indicators, such as PUE. Total DC energy consumption is already very widely measured and reported.

The specific version (year) of the EN 50600-4-2 and category of measurement have not been specified because future updates may be published before or during the EED recast comes into effect. The latest version of EN 50600-4-2 should be used unless there are specific reasons preventing this.

The additional requirement to also report the contribution of backup generator use is included because they are highly polluting and there is some evidence of excessive use. This information is needed to calculate E_{DC} and therefore is not considered to create any significant additional effort to report.

There was a concern that the measurement within a mixed use building would include all other uses including shared HVAC systems, however, the data centre size IT power threshold should exclude data centres where the related activities are not independent and do not have some level of redundancy (as recorded in elements of the data centre included in the scope of the reporting obligations).

All stakeholders' comments are in agreement with the need to report this measurement using EN 50600-4-2.

Data centres are also required to report electrical grid functions. This can improve the efficiency of the wider grid but may come at the cost of the PUE and be needed in the long term for future metrics. In the near-term, this requirement will affect only a very small proportion of data centres and is expected to be very simple for those data centres to report. A list of grid functions and corresponding definitions will need to be developed.

For FFR specifically, the average capacity and availability should be reported to enable better analysis (and comparison between data centres). These values are required to enter a contract to provide FFR and therefore should be readily available.

9.1.2 Data functions and ICT equipment energy consumption compared to total DC energy consumption

This is a measure of the energy consumed to perform the data functions, i.e. storage, processing and transport. As a simple energy measurement, it does not indicate how much work is being done with the energy.

Ideally, all the energy consumed should be used solely for the data functions (i.e. storage, processing, transport). All energy used for non-data functions (cooling power etc), should be as close to zero as possible.

In this formulation, total DC facility energy consumption is expressed as:

$$E_{Total\ DC} = E_{data\ functions} + E_{non-data\ functions}$$

The obvious place to measure the energy consumed for data functions is the power input to the ICT equipment that carry out the data functions, and the infrastructure.

$$E_{Total\ DC} = E_{ICT\ equipment} + E_{infrastructre}$$

However, this is not completely accurate because some ICT equipment have internal power and cooling components whereas other, newer designs have shifted this energy consumption externally. To account for this inaccuracy metrics have been proposed, including the ICT power usage effectiveness, to separately determine the internal power and cooling energy consumption within the ICT equipment.

From a practical point of view, there are many more pieces of ICT equipment than power equipment that also have built-in power monitoring functions. Colocation data centres also do not have physical access to the ICT equipment to make direct measurements. Measurements therefore often occur at the power equipment and some of the energy 'wasted' for power and cooling is being misallocated as energy utilised for performing data functions.

By far the most widely used and accepted sustainability indicator to compare energy used by ICT equipment is PUE, EN50600-4-2 (see Annex A). This standard establishes three categories of PUE and three methodologies to measure the ICT equipment energy consumption, including the frequency of measurement and the area the ICT equipment measurement is made:

- Category 1. UPS;
- Category 2. PDU; and
- Category 3. the ICT equipment input.

Work is also underway to update the PUE methodology to stay in line with technology changes.

9.1.2.1 Proposal

(Annual) Information technology equipment energy consumption (E_{IT}) in kWh as defined in EN 50600-4-2.

Data centres should measure the annual energy consumption at the UPS output according to the Category 1 PUE methodology in EN 50600-4-2.

For data centres that do not have a UPS, e.g. direct current data centres, E_{IT} can be measured at the PDU in line with category 2, or the specify the measurement point.

9.1.2.2 Rationale

The E_{IT} based on Category 1 is a practical approach since many data centres do not have the equipment to measure energy consumption past the UPS. E_{IT} is proposed rather than PUE since this is easier to verify during any potential audits and reduces the likelihood of errors in reporting.

PUE and hence E_{IT} is already being measured and reported against the EN50600-4-2 standard by many larger data centres and therefore would not require additional effort. For data centres not yet measuring PUE, this is the most easily measured at the UPS. Stakeholder feedback also indicates this is reasonable for all sizes and types of data centre.

Although measurement in line with PUE Category 1 may not be as representative as the measuring closer to the ICT equipment (a minority of stakeholders suggested that measurement should be made at the PDU) the difference when comparing infrastructure energy consumption between several data centres can be still very high. While the most efficient data centres are implementing the latest technologies and achieving additional reductions of a few percent, the most significant energy savings are likely to be in data centres which could in theory reduce infrastructure energy by 50-90% or more. These differences are still easily identified by PUE methodology category 1. However, flexibility may be required in future if the data centre technology changes substantially.

Alternative measurement in the guidance are designed to allow flexibility as data centres evolve and for data centres, which are designed without UPS or use direct current power.

9.2 Productivity and data functions (processing, storage and transport)

9.2.1 ICT equipment capacity

ICT equipment capacity is a measure of the total amount of work the data centre can perform. The total equipment capacity is simply the sum of the capacities of each piece of equipment.

Combined with the utilisation level of the ICT equipment, it produces a KPI of the amount of work done, rather than try to measure it directly.

$$work\ done = \text{ICT capacity} * average\ utilisation * time$$

Or

$$utilised\ capacity = \text{ICT capacity} * average\ utilisation$$

While standards exist to enable data centres to track the actual work being done, this is not practical for an EU-wide reporting scheme, since those standards set minimum requirements for data centre to develop their own indicators for their own characteristics. Instead, the KPIs proposed correspond to different types of ICT equipment and/or data functions.

The main data functions are the processing (compute), storage, and transport (network). Memory capacity has also been identified by other work.

9.2.1.1 Proposal

The proposed method to measure functional is to use the SERT performance rating and efficiency rating as required by the Regulation on ecodesign requirements for servers and data storage products EU 2019/424[14], or latest revision to the regulation in effect. Server types, which are not covered by the regulation, e.g. network servers and fully fault tolerant servers would not be required to report.

SERT measures the energy consumed by the server undertaking various worklets designed to measure the performance of the processing, RAM and storage at various utilisation levels. A weighted geometric mean of the performance/Watt is calculated to produce the SERT value. This value will vary depending on the specific configuration of the server.

[14] OJ L 74, 18.3.2019, p. 46–66 EUR-Lex - 32019R0424 - EN - EUR-Lex (europa.eu)

It may not be possible to gather this information for all legacy equipment installed in the data centre because the information has not been published. These issues are discussed in Annex B and will be addressed in guidance documents.

9.2.1.2 Rationale

The capacity KPIs are focussed on the compute and storage because these represent the main source of additional computing capability and energy consumption. Memory is included because this can be a major consumer of server energy for some types of workload, although it tends to scale closely with processing. Other indicators such as the transport (networking) are configured to support the processing and storage of data and to scale with compute and storage capacity. In addition, the network equipment consumes a relatively small proportion of the ICT energy consumed in the data centre. Including these would create more work with more limited benefit. Similarly, fully fault tolerant servers and other servers not covered by the regulations are expected to comprise a small fraction of the total number of servers and energy consumption and can be excluded.

The SERT rating includes the compute, memory and storage performances of the server, thus combining and simplifying the reporting requirements. SERT is currently reported by manufacturers for the ecodesign regulations but only the low-end and high-end server configurations which are not representative of real life use. Auxiliary processing accelerators such as GPUs and specialist hardware used for AI and other HPC are becoming increasingly important market and should also be addressed.

Updates to SERT and the ecodesign regulations have been proposed to expand the reporting of configurations as well as include performance measurements for the floating point capacity which should address these shortcomings. Ecodesign regulations could also add requirements to help reporting of storage equipment capacity.

Furthermore, SERT could also be applied to estimate utilisation as discussed in 9.2.2.3.1.

There may be situations in which these capacities will not be representative of computing capability of the data centre, particularly when the hardware is mismatched to the workload. However, it is intended for this to be reflected in the utilisation, and it is not possible to develop a perfect indicator.

Most types of data centre operator will have access to the ICT equipment and therefore can get the capacity data. Colo operators, however, will need to ask the colocation customers for this information according to the responsibilities to fulfil reporting obligations.

Stakeholders were generally in support of a capacity KPI as a practical indicator, while acknowledging its inherent limitations. Some stakeholders expressed a preference for using alternative indicators described Table 2 which are more widely available for the cost common components used in servers.

9.2.1.3 Alternative proposals

1) New indicators are currently under development and could result in a more suitable indicator by various parties such as the Climate Neutral Data Centre Pact (CNDCP) which have been suggested by several stakeholders. The alternative does not discount this as a long-term option but in the absence of a proven indicator and timelines for their completion they cannot be proposed.

2) Alternative ICT equipment capacity indicators are the total compute, memory and storage capacity.

The capacities are measured using the five indicators shown in Table 2.

Table 2: Proxies for the functional capacity of data centres

	Capacity Indicator 1	Capacity Indicator 2
Compute	SPECrate INT	Floating point
Memory		Total capacity
Storage	Data throughput (IOPS)	Raw capacity (TB)

There are many performance benchmarks variations and setups and the methodology will need further development.

9.2.2 ICT equipment utilisation

ICT equipment utilisation refers to the how much of the ICT capacity is used, which will change over time.

9.2.2.1 Proposal

IT equipment utilisation for servers ($ITEU_{sv}$) – ISO/IEC 30134-5 (See Annex A)

The differences in utilisation methodology and sampling frequency, not covered in ISO/IEC 30134-5 are discussed in Annex B.

9.2.2.2 Rationale

As discussed above, a work done KPI requires a measurement of the capacity and utilisation.

KPIs based only on capacity could produce a misleading indicator whereby low utilisation reduces power demand which is then rewarded because capacity is being achieved at lower overall energy consumption. Conversely, it is not believed that this will lead to data centre operators artificially increasing the utilisation because this will lead to increased energy costs and a similar result has not occurred when PUE was reported.

Utilisation will vary between workloads and there are technical reasons, which limit the achievable utilisation using current technology and an indicator of utilisation will help identify this, as well as new technologies able to overcome this. CPU utilisation is more easily measured than utilisation of all the other ICT equipment. A standard already exists for measuring CPU utilisation and all major operating systems and hypervisors provide tools to report utilisation.

Server utilisation can vary from below 5% to above 30% and it is not necessary to achieve very high levels of accuracy for the reporting of utilisation to be able to distinguish between the more and less efficient data centres. This should simplify the sampling and measurement methodology.

There is not technical barrier to collating this information but it is not widely implemented. For a colocation, the utilisation data would need to be provided by the tenant who operates the ICT equipment. For tenants that provide dedicated servers to clients and do not operate the hypervisor or operating system, the situation becomes more complex and the tenant would have to first request the utilisation data from each client. This is discussed further in Annex B.

9.2.2.3 Alternative proposal

9.2.2.3.1 Total number of servers and storage equipment. Average weighted server performance and efficiency as reported by SERT

Based on the server efficiency and ICT power consumption, testing has shown it is possible to estimate utilisation with sufficient accuracy to be useful. However, it is not clear at this stage if this holds for all data centres and how accurate the efficiency data must be. This indicator has been proposed by a stakeholder currently developing a potential metric.

SERT is currently reported under the EU ecodesign regulation for enterprise servers, which means there is a growing body of available data. This also avoids any requirement to continuously measure utilisation of the IT equipment.

Ecodesign regulations currently only requires SERT information for two server configurations, high and low-performance. This means that the majority of servers with a 'typical' performance configuration will not have a representative SERT score. SERT also provides no measurement of floating point performance/efficiency which would need to be taken into account for any possible metrics. However, there is an opportunity within the current review of the ecodesign regulations to require manufacturers provide more information to consumers that could be used for reporting under EED recast, or to work with server manufacturers to provide more data voluntarily.

9.2.2.3.2 Server idle energy consumption

The server idle energy consumption is a measure of the energy consumed while the server CPU is idle (or very low utilisation for servers with idle disabled).

The formula for server energy consumption can be constructed similarly to the data centre energy consumption:

$$E_{Total\ Server} = E_{data\ functions} + E_{non-data\ functions}$$

Where $E_{data\ functions}$ occurs when the server is active and doing work

and $E_{non-data\ functions}$ is the energy consumed during idle.

This KPI can be used for the server idle coefficient or DC idle coefficient (See Annex A) metric which has been shown to be useful in limited pilot tests and is analogous to the PUE. By measuring the energy consumed when the processor is idle, it indicates how much energy is being wasted doing no useful task. As a result, it should encourage users to minimise the amount of time a server is idle, or ensure that servers in idle states demand as little power as possible. Idle energy consumption however, does not measure how efficiently the server is utilising energy when performing the data functions themselves.

The server idle energy consumption is more complex to measure and report than utilisation but is entirely possible with support from OS/hypervisor manufacturers. A possible solution is to require reporting only for servers installed after a certain date and/or running newer versions of the OS/hypervisor and simply wait for older servers to be eventually replaced. Similarly to utilisation reporting, there may be a small difference between OSs but is unlikely to be significant.

A small number of stakeholders supported reporting of this indicator.

9.2.3 Overall data centre productivity indicator

This indicator is designed to assess the overall output produced by the data centre.

9.2.3.1 Proposal

This indicator was proposed for discussion in the Scoping documents, but it is discarded in this final proposal.

9.2.3.2 Rationale

No useful measure of productivity was identified that could be practically measured and would provide a useful indicator of productivity across all different data centre types, including sufficiently clear measurement methods. Adding contextual data could help to understand the differences in efficiency but this would require reporting from customers of the colo customers, making the information too difficult to request, collate and weight to be meaningful. Stakeholder comments suggest that new indicators are being developed but these are too immature to be included.

No stakeholder comments suggested a practical productivity indicator or metric.

9.2.4 End-user Service efficiency indicator

9.2.4.1 Proposal

No proposal.

9.2.4.2 Rationale

No practical methods were identified to measure the quantity and quality of the end-user service provided by a data centre, because of the way a service is delivered within a data centre and the wide range of services and clients within a single data centre.

Work has been done on measuring the efficiency of specific services eg streaming, but it was necessary to exceed the boundary of a data centre to be representative and therefore falls outside the scope.

9.3 Ancillary KPIs

9.3.1 Water usage

Water usage measures the volume of water consumed by the data centre. Reporting of the water use is based on the methodology in the EN 50600-4-9 standard

9.3.1.1 Proposal

Annual water input (w_{in}) and potable water input ($w_{in,pot}$) as defined and measured using EN 50600-4-9 WUE Category 1

9.3.1.2 Rationale

EN 50600-4-9 for water usage effectiveness is an existing standard that is already being utilised to report the volume of water used as a proportion of the total data centre energy consumption. This includes a methodology to measure the water input. Water input is the preferred KPI because it can be more easily verified and reduces errors arising while calculating WUE. Since potable water is generally a significantly more valuable resource, the separate reporting of this source is also required. Water consumption is not widely reported but is always metered at the building. However, it can be more problematic for data centres in mixed use buildings.

Water use or reuse is not required because of its additional complexity for measurement and is not considered to add significantly more useful information since water reuse is also limited. In addition, stakeholder comments suggest that water scarcity should also be reported. However, this can be ascertained from the reported data centre location data and therefore is not required for the data centre to report.

9.3.2 Energy/heat reuse

9.3.2.1 Proposal

- Energy reused in kWh as defined and measured by EN50600-4-6
- Temperature of waste heat for all data centres (monitoring method to be defined in guidance)

9.3.2.2 Rationale

Reusing heat is an effective energy efficiency measure and should be reported as part of a wider picture of the data centre energy use. EN 50600-4-6 is an existing international standard that defines how the heat reuse should be measured and is supported by the majority of stakeholders.

Reuse of data centre heat is currently concentrated in a few regions but is rare across EU as a whole. Reporting energy reused does not give any indication of the *potential* for the data centre heat to be reused. While opportunities do not always exist to reuse the heat for reasons out of control of the data centre, more transparency on the potential heat sources available could encourage new development of local reuse of heat via district heating or many other potential applications.

Work from Byte2Heat suggests that reporting both the IT equipment energy and temperature of the waste heat are the key KPIs needed to assess the data centres potential to supply useful heat and therefore have been proposed for all data centres. There is currently no standard for measuring the waste heat temperature and this will be addressed in the guidance for Task D.

9.3.3 Renewable energy consumption

9.3.3.1 Proposal

Annual renewable energy in kWh owned and controlled by the DC (E_{ren}) as defined and measured by EN 50600-4-3. In addition, the renewable energy should be broken down by:

- Guarantees of Origin (GoOs) purchased and retired by data centre
- Power Purchasing Agreements (PPAs) (any GoOs created must be owned and retired by DC)
- On-site renewables (any GoOs created must be owned and retired by data centre)

9.3.3.2 Rationale

REF is an existing standard that is already being utilised to report the proportion of total DC energy consumed that comes from renewable sources. However, REF does not distinguish between where the renewable comes from, and what route the data centre used to own the renewable energy. This is important because it could reward data centres located in areas with a high proportion of grid renewables which may not be within the data centres control. Therefore, a breakdown is also required in the reporting. This is not expected to increase the complexity since these must be measured to report the total E_{ren}. Reporting the energy rather than the REF also reduces the chance of reporting errors and can be verified more easily.

Annex A: Sustainability indicators and other metrics referred in the proposals[15]

Power Usage Effectiveness (PUE) EN 50600-4-2

The PUE[16] metric is associated with the data centre infrastructure, and measures the relationship between the total facility energy consumed and the IT equipment energy consumed.

PUE is defined as the ratio of total facilities energy to IT equipment energy, as shown in Equation 1 below.

$$PUE = \frac{Total\ facility\ energy\ consumption}{IT\ energy\ consumption}$$

Server idle coefficient and data centre idle coefficient[17]

The study was requested by EDNA, aiming at exploring the KPIs quantifying 'wasted energy' in servers and data centres and thereby reducing the IT energy consumption without reducing the workload.

The Server Idle Coefficient (SIC) and the Data Centre Idle Coefficient (DCIC) metrics take a different approach than most other data centre metrics, since they are ineffectiveness metrics and not efficiency metrics. The source considers that determining ineffectiveness rather than (in)efficiency is grounded in the fact that there is no known generic metric indicating the amount of work that is done by a data centre. Therefore, it is impossible to define an efficiency metric that would by definition be of the format Unit of Work per Unit of Energy. For servers however, there is a single identifiable process that is common to all servers, regardless of make, model and even architecture, that indicates that the server has no useful workload to run. This process is known as 'Idle'. It is possible to calculate the energy used for running these idle cycles and express this as a percentage of the total energy use of the server:

SIC = Server Energy (idle) / Server Energy (total) x 100%

Expanding this for the total data centre to:

DCIC = ∑ Server Energy (idle) / ∑ Server Energy (total) x 100%

According to the source, the data needed is simple and does not need application data. The pilot data showed values for the SIC that ranged from 50% (best case) to over 90% (worst case) indicating the relative size of the savings.

IT Equipment Utilization for servers (ITEUsv) ISO/IEC 30134-5:2017

The IT Equipment utilization for servers (ITUEsv) is the average utilisation of all servers, or group of servers in a data centre.

$$ITEU = \frac{1}{a}\sum_{i=1}^{a} ITEUsv(t_0 + e * i)$$

Where

[15] https://www.iea-4e.org/wp-content/uploads/2022/10/EDNA-Studies-Metrics-for-data-centre-efficiency-Final.pdf
[16] The Green Grid (2012). PUE™: A comprehensive examination of the metric https://www.thegreengrid.org/en/resources/library-and-tools/237-WP#49---PUE:-A-Comprehensive-Examination-of-the-Metric
[17] The idle coefficients. KPI's to assess energy wasted in servers and data centres. Certios for EDNA:

ITEUsv(t) is the average utilisation of all servers or a group of servers in a data centre at time *t*

a is the number of *ITEU(sv)* measurements intervals over a year

t_0 is the starting time of measurement

e is the interval of measurement, where *e x a* = one year

$$ITUEsv(t) = \frac{\sum_{i=1}^{N} CUSi(t)}{N}$$

CUS_it is the CPU utilization ratio of server *i* at time *t* (%)

N is the number of servers in a data centre or in a group running at time t

Annex B: ICT capacity and utilisation KPIs

Compute

Integer computer capacity should be reported for all CPUs and integer processing units. It should be reported using SPECrate INT 2017 base metric.

Results are available for a wide range of server models on the SPEC website. Where the model is not available on the SPEC website, a similar model with the same CPU should be used. However, developing a standardised reporting scheme for the CPU and component manufacturers is preferred. The information could then be made available on a webpage and/or provided by the server manufacturer.

For older servers it is still possible to find results, although they may be based on SPEC 2006 and an equivalence will need to be developed. If no result can be found, a proxy value will need to be developed.

Floating point operations per second (FLOPS) is a theoretical performance measure for GPUs and other accelerated processing units. It is common to calculate the FLOPS based on different levels of floating point precision, most commonly 16-bit, 32-bit but also 16-bfloat, and other metrics designed for deep learning calculations (particularly inference engines). Theoretical FLOPS are also not achievable in real life workloads, and can be misleading, particularly for FPGAs. More development of this is needed to finalise the indicator.

Storage

Storage should be reported for both HDDs and SSDs. This increases the amount of reporting but minimises confusion and mistakes from having different requirements for different types of drive. Raw capacity is easily to obtained from component manufacturer drives or interrogating the ICT equipment.

Drive IOPS can be measured in a wide variety of ways, the most common being sequential read/write and 4K random read/write. In addition, the random read/write can be performed with a different queue depths. A higher queue depth is more representative of a data centre workload and 32/64 is proposed. Component manufacturer declaration can be used and would be easiest but many different benchmark software are available. More development will be needed to determine if there is significant difference between benchmark software and eg drive pre-conditioning.

Memory

The memory capacity is very straightforward and can be easily obtained by interrogating the server.

Equipment replacement

There is constant churn ICT equipment capacity, and the most accurate measure would be to weight the equipment capacity by the proportion of time it is installed and in operation during the reporting year. This should be possible with asset management software but may not be necessary. Taking a snapshot of the capacity at a single point, or a few points of time during the reporting year could be sufficient but will need further development. This would also simplify the data collation for colo operators.

Asset management software, which can record ICT capacity is recommended. It is expected that industry will develop solutions for this.

Utilisation

Utilisation is perhaps the most difficult indicator proposed to monitor and report. It will require frequent sampling for all servers. For data centres with many different types of servers, OSs or hypervisors, the initial set up could be time consuming. Ongoing, this could also result in a large amount of data being

collected. The ISO standard requires at a minimum hourly sampling which would result in a minimum of 8760 datapoints per server a year. Because diurnal, weekly and seasonal use patterns are very common, reducing this below every hour could result in unrepresentative indicators. There is also a concern that a more frequent sampling rate is needed because CPU utilisation can swing widely over very short time frames. More development is needed to define the sampling e.g. the average utilisation over 1 minute every hour.

Utilisation measurement should be required for all virtualised servers and more recently installed servers. For the other older, unvirtualized servers it may be necessary to use average values based on other factors such as server age.

The differences in utilisation reporting methodology between operating systems and hypervisors is unlikely to cause a significant difference, but further work is needed to confirm this. Discussion with stakeholders indicate that utilisation is easy to monitor for large cloud operators but there is a concern about the volume of data that will be collected and if this could have a measurable impact on the DC energy consumption.

Average utilisation is reported independent of the performance of the server. Weighting would be beneficial but would increase the reporting complexity, especially for colo. For servers installed or removed, the server can be treated as a separate server and the utilisation averaged for the period of time it was operating.

Software tools and services that help colocation customers to monitor and report utilisation could be very useful, and we expect that OS and hypervisors will be able to provide these with minimal effort.

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct centres. You can find the address of the centre nearest you online (european-union.europa.eu/contact-eu/meet-us_en).

On the phone or in writing

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

— by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),
— at the following standard number: +32 22999696,
— via the following form: european-union.europa.eu/contact-eu/write-us_en.

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website (european-union.europa.eu).

EU publications

You can view or order EU publications at op.europa.eu/en/publications. Multiple copies of free publications can be obtained by contacting Europe Direct or your local documentation centre (european-union.europa.eu/contact-eu/meet-us_en).

EU law and related documents

For access to legal information from the EU, including all EU law since 1951 in all the official language versions, go to EUR-Lex (eur-lex.europa.eu).

EU open data

The portal data.europa.eu provides access to open datasets from the EU institutions, bodies and agencies. These can be downloaded and reused for free, for both commercial and non-commercial purposes. The portal also provides access to a wealth of datasets from European countries.

Publications Office
of the European Union

ISBN 978-92-68-08438-0