



**Environmental Engineering (EE);
Energy Efficiency measurement methodology
and metrics for heterogeneous servers**

Reference

DES/EE-EEPS46

Keywords

energy efficiency, power, server, sustainability

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from the
[ETSI Search & Browse Standards](#) application.

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format on [ETSI deliver](#) repository.

Users should be aware that the present document may be revised or have its status changed,
this information is available in the [Milestones listing](#).

If you find errors in the present document, please send your comments to
the relevant service listed under [Committee Support Staff](#).

If you find a security vulnerability in the present document, please report it through our
[Coordinated Vulnerability Disclosure \(CVD\)](#) program.

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.

No recommendation as to products and services or vendors is made or should be implied.

No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.

In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.

The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2025.
All rights reserved.

Contents

Intellectual Property Rights	6
Foreword.....	6
Modal verbs terminology.....	6
Introduction	7
1 Scope	8
2 References	8
2.1 Normative references	8
2.2 Informative references.....	8
3 Definition of terms, symbols and abbreviations.....	9
3.1 Terms.....	9
3.2 Symbols.....	11
3.3 Abbreviations	11
4 Heterogeneous server product categories and representative product configurations.....	12
4.1 General	12
4.2 Applications and metric applicability.....	12
4.3 Product categories	13
4.3.1 General requirements.....	13
4.3.2 Form-factors	13
4.3.3 Categories	13
4.4 Product family configuration.....	13
4.4.1 General.....	13
4.4.2 "High-end" performance configuration	14
4.4.3 "Low-end" performance configuration	14
5 Metrics.....	15
5.1 Worklets	15
5.2 Formulae.....	15
5.2.1 General.....	15
5.2.2 Weightings.....	16
6 Test setup.....	16
6.1 General	16
6.2 Environmental conditions.....	16
6.2.1 Ambient Temperature	16
6.2.2 Relative Humidity.....	16
6.3 Temperature sensor	16
6.4 Input power	17
6.5 Power requirements.....	17
6.6 Energy consumption analyser.....	17
6.7 Test tool.....	18
6.8 Controller system	18
7 Equipment Under Test.....	18
7.1 Configuration	18
7.2 Test procedure	20
8 Measurement	21
8.1 Measurement for active state.....	21
8.2 Re-verification and audits	21
8.3 Measurement for power supply	22
8.3.1 Measurement for internal power supply	22
8.3.2 Measurement for test board power supply.....	22
8.3.2.1 General	22
8.3.2.2 Test loads	22
8.3.2.3 Test leads and wiring	22

8.3.2.4	Warm up time.....	22
8.3.2.5	Power measurements.....	23
9	Measurement report.....	23
History	24

List of tables

Table 1: Configuration of EUT	18
Table 2: Test configuration of EUT	20
Table 3: Measurement of active state efficiency	21
Table 4: Measurement of internal power supply	22

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: "*Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards*", which is available from the ETSI Secretariat. Latest updates are available on the [ETSI IPR online database](#).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™**, **LTE™** and **5G™** logo are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM®** and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword

This final draft ETSI Standard (ES) has been produced by ETSI Technical Committee Environmental Engineering (EE), and is now submitted for the ETSI Membership Approval Procedure (MAP).

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Introduction

Computer server and data centers have experienced a rapid increase, playing the central role in digital services. With the development of Artificial Intelligence (AI), heterogeneous servers equipped with more than one type of processors (CPU, GPU, FPGA, ASIC, etc., the abbreviations defined in clauses 3.1 and 3.3) play an increasingly important role and take a growing market share in computing industry. Heterogeneous servers gain extra performance by adding the dissimilar coprocessors, usually incorporating specialized processing capabilities to handle particular tasks [i.4]. Categories of heterogeneous servers include CPU+GPU server, CPU+FPGA server, CPU+ASIC server. The most popular heterogeneous server type is the CPU+GPU server, which is commonly used to accelerate the training process of deep learning by parallel computing.

For heterogeneous servers, the better performance the higher the energy consumption. There is a tradeoff between the performance and energy consumption for the high performance heterogeneous servers. Energy Efficiency (EE) that scales the server performance and energy consumption is one of the critical factors of the server operation and maintenance. ETSI EN 303 470 [i.5] specifies the metric and method for EE assessment of general servers. However, there is currently no standard for EE assessment of heterogeneous servers.

The present document specifies a metric for the assessment of EE of heterogeneous servers using reliable, accurate and reproducible measurement methods, which take into account the recognized state-of-the-art. The present document formalizes the tools, conditions and calculations used to generate a single figure of merit of a single heterogeneous server representing its relative EE and power consumption. The single figure EE metric is targeted for use as a tool in the selection process of heterogeneous servers to be provisioned.

For comparisons, evaluations should be conducted across similar heterogeneous server types or categories. The EE metric is targeted for identify energy saving servers by differentiating the ability of heterogeneous servers to be provisioned for mainstream market. The present document does not prescribe the levels or values for acceptance but prescribes a method of evaluation which EE programs could use to establish such criteria.

As there are many operational deployments of heterogeneous servers resulting in a range of specialized equipment and configurations for a single heterogeneous server product, an EE metric that evaluates provisioning impacts to general purpose operations may not be applicable. ICT equipment, and servers in particular, are generally customized and commissioned on site for deployment. As with most ICT equipment, new technologies are regularly introduced, which may require product level customization or an industry wide tool upgrade to more appropriately represent the EE of the heterogeneous servers.

The present document categorizes heterogeneous servers to address applicability, configuration groupings to represent a diversity of heterogeneous servers to address the broad range of custom configurations, as wells as a tool revision control to ensure comparability and consistency of the resulting metric value.

The present document was developed jointly by ETSI TC EE and ITU-T Study Group 5 and published by ITU and ETSI as Recommendation ITU-T L.EE_sgpu [i.6] and ETSI ES 204 083 (V1.1.1) (the present document), which are technically equivalent.

1 Scope

The present document is based upon benchmark for server Energy Efficiency (EE) tool for general and heterogeneous servers [i.6].

The present document specifies:

- 1) test conditions and product configuration for the assessment of EE of heterogeneous servers using reliable, accurate and reproducible measurement methods;
- 2) an EE metric to support procurement requirements;
- 3) requirements for equipment to perform the measurements and analysis;
- 4) requirements for the measurement process;
- 5) requirements for the management of the EE metric calculation;
- 6) operation or run rules to configure, execute, and monitor the testing;
- 7) documentation and reporting requirements.

The metrics and methods apply to heterogeneous servers with various configurations, including type and count of APAs, CPU, memory, storage and any other add-on hardware expected to be present when deployed.

NOTE: Products whose feature set and intended operation are not addressed by active mode testing parameters are excluded from this evaluation method.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found in the [ETSI docbox](#).

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

- [1] [IEEE 802.3az™](#): "The road to energy efficient ethernet".
- [2] [EN 62623:2013](#): "Desktop and notebook computers. Measurement of energy consumption", (produced by CEN).
- [3] [IEEE 802.3™](#): "IEEE Standard for Ethernet".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

- [i.1] [Commission Regulation \(EU\) 2019/424 of 15 March 2019](#) laying down ecodesign requirements for servers and data storage products pursuant to Directive 2009/125/EC of the European Parliament and of the Council and amending Commission Regulation (EU) No 617/2013.
- [i.2] EN 60297 series: "Mechanical structures for electrical and electronic equipment. Dimensions of mechanical structures of the 482,6 mm (19 in) series", (produced by CEN).
- [i.3] ETSI EN 300 119 series: "Equipment Engineering (EE); European telecommunication standard for equipment practice".
- [i.4] Shan, Amar (2006): "Heterogeneous Processing: a Strategy for Augmenting Moore's Law". Linux[®] Journal.

NOTE: Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries.

- [i.5] ETSI EN 303 470 (V1.1.1) (2019): "Environmental Engineering (EE); Energy Efficiency Measurement methodology and metrics for servers".
- [i.6] Recommendation ITU-T L.EE_sgpu on Energy Efficiency measurement methodology and metrics for heterogeneous servers.

NOTE: Recommendation ITU-T L.EE_sgpu is at the time of publication only the draft working title. Recommendation ITU-T L.1311 should be the official number to be decided in May/June 2025.

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

active state: operational state of a computer server (as opposed to the idle state) in which the computer server is carrying out work in response to prior or concurrent external requests (e.g. instruction over the network)

NOTE 1: The work in the present document includes, but is not restricted to, active processing and data seeking/retrieval from memory, cache, or internal/external storage while awaiting further input over the network.

NOTE 2: See ETSI EN 303 470 [i.5].

Application-Specific Integrated Circuit (ASIC): integrated circuit designed and manufactured according to specific user requirements and specific electronic system requirements

NOTE 1: Processor units belong to ASIC include, but are not limited to:

- a) Tensor Processing Unit (TPU): an AI accelerator ASIC that is developed specifically for neural network machine learning by Google;
- b) Vision Processing Unit (VPU): a specific type of AI accelerator, designed to accelerate machine vision tasks;
- c) Neural Processing Unit (NPU): a specialized processing unit that is designed to accelerate artificial intelligence and machine learning applications, including artificial neural networks and machine vision;
- d) Deep Learning Processing Unit (DPU): a processor designed for deep learning that has the ability to handle large-scale neural networks.

NOTE 2: DSP is used in Embedded and Real-time Systems.

Auxiliary Processing Accelerator (APA): additional compute device installed in the computer server that handles parallelized workloads

NOTE 1: This includes, but is not limited to, Graphical Processing Units (GPUs), Field Programmable Gate Array (FPGA) or Application-Specific Integrated Circuit (ASIC) chips which can be installed in a server either on Graphics or Extension add-in cards installed in general-purpose add-in expansion slots (e.g. GPGPUs, CPU accelerators, etc. installed in a PCI/PCIe slot) or directly attached to a server component such as the motherboard.

NOTE 2: There are two specific types of APAs used in servers:

- a) Expansion APA: an APA that is on an add-in card installed in an add-in expansion slot (e.g. GPGPUs, CPU accelerators, etc. installed in a PCI/PCIe slot). An expansion APA add-in card may include one or more APAs.
- b) Integrated APA: an APA that is integrated into the motherboard or CPU package or an expansion APA that has part of its subsystem, such as switches, included in the non-APA server configuration that would be used to run the energy efficiency test.

NOTE 3: See ETSI EN 303 470 [i.5].

blade chassis: enclosure that contains shared resources for the operation of blade servers and other blade form-factor devices

NOTE 1: Shared resources provided by a chassis include, but are not restricted to, power supplies, data storage, and hardware for d.c. power distribution, thermal management, system management, and network services.

NOTE 2: See Commission Regulation (EU) 2019/424 [i.1].

blade server: computer server, designed for use in a blade chassis, that is a high-density device and functions as an independent computer server and includes at least one processor and system memory, which is dependent upon shared blade chassis resources (e.g. power supplies, cooling) for operation

NOTE 1: A processor or memory module that is intended to scale up a standalone server is not considered a blade server.

NOTE 2: See ETSI EN 303 470 [i.5].

blade system: blade chassis and one or more removable blade servers and/or other units which provide a scalable means for combining multiple blade server or storage units in a single enclosure

NOTE 1: A blade system is designed to allow service technicians to easily add or replace (hot-swap) blades in the field.

NOTE 2: See ETSI EN 303 470 [i.5].

controller system: computer or computer server that manages a benchmark evaluation process

deployed power: average power level of the utilization applicable to the total number of servers provisioned to meet an aggregate peak load

hypervisor: supervisory system level software that establishes and manages a virtualized environment which enables multiple operating systems to run on a single physical system at the same time

pedestal server: self-contained computer server that is designed with power supply units, cooling, input/output devices, and other resources necessary for stand-alone operation within a frame similar to that of a tower client computer

NOTE: See ETSI EN 303 470 [i.5].

processor: digital circuit which performs operations on some external data source, usually memory or some other data stream

NOTE: This includes, but is not limited to, Central Processing Unit (CPU), Graphics Processing Unit (GPU), Neural Processing Unit (NPU), Tensor Processing Unit (TPU), Deep Learning Processing Unit (DPU), Vision Processing Unit (VPU), etc.

rack-mounted server: computer server that is designed for deployment in a standard 19 inch ICT equipment rack

NOTE 1: As defined by EN 60297 [i.2] or ETSI EN 300 119 [i.3].

NOTE 2: For the purposes of the present document, a blade server is considered under a separate product category and excluded from the rack-mounted product category.

server: computing product that provides services and manages networked resources for client devices

NOTE 1: 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 (OSs) 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.

NOTE 2: See Commission Regulation (EU) 2019/424 [i.1].

single-wide blade server: blade server requiring the width of a standard blade server bay

NOTE: See ETSI EN 303 470 [i.5].

worklet: synthetic software routine, using real application functions focused on a particular type of computing activity, which stresses a particular characteristic of the system

NOTE 1: A floating point and integer performance stress code is an example of a CPU worklet.

NOTE 2: An image classification training code is an example of AI worklet.

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

a.c, AC	alternating current
ACPI	Advanced Configuration and Power Interface
APA	Auxiliary Processing Accelerator
ASIC	Application-Specific Integrated Circuit
BIOS	Basic Input/Output System
CENELEC	European Committee for Electrotechnical Standardization
CPU	Central Processing Unit
CSJ	Correctly predicted Samples per Joule
d.c, DC	direct current
DDR	Double Data Rate
DIMM	Dual In-line Memory Module
DPU	Deep learning Processing Unit
DRAM	Dynamic Random Access Memory
DSP	Digital Signal Processor
EDP	Energy Delay Product
EE	Energy Efficiency
EEE	Energy Efficient Ethernet
EUT	Equipment Under Test
FPGA	Field Programmable Gate Array

GPGPU	General-Purpose computing on Graphical Processing Units
GPU	Graphical Processing Unit
HDD	Hard Disk Drive
I/O	Input/Output
ICT	Information and Communications Technology
ISO/IEC	International Organization for Standardization/International Electrotechnical Commission
NPU	Neural Processing Unit
OS	Operating System
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect express
PDU	Power Distribution Unit
PSU	Power Supply Unit
RMS	Root Mean Square
SSD	Solid State Drive
TPU	Tensor Processing Unit
UPS	Uninterruptable Power Supply
VPU	Vision Processing Unit

4 Heterogeneous server product categories and representative product configurations

4.1 General

Heterogeneous servers are sold in different types and counts of APAs, and configuration types and different groups of heterogeneous servers will have distinct performance capability and power demands.

Heterogeneous servers are categorized as defined in clause 4.3 by specific configuration parameters to enable the setting of appropriate idle power or active efficiency thresholds and assessment of like products with regards to procurement or market entry requirements.

To compare or evaluate systems, the evaluation or EE metric shall only be made among similar products. Similar products are grouped into categories. Products of other/different categories not listed shall not be compared using the EE metric of the present document.

Even though heterogeneous servers are classified in categories by the type of system, each server is customized by its configuration to best match the application for which it is being sold or purchased. As a result, a product is represented by a fixed set of configurations that is decided by the form-factor, category, number/type of APAs and processors, and capacities of memory. For an appropriate evaluation, the category shall be defined.

4.2 Applications and metric applicability

Heterogeneous servers are designed to be configurable to different groups of applications. The full configuration, including logical and physical elements, is optimized to deliver the most effective platform for operating those applications.

The software component of the EE metric is designed to execute typical real world applications and is designed to stress and assess the elements and associated functionality included in the server systems. By stressing the elements in a fashion that replicates real world applications, a process for measuring workload output and its associated power demand is established.

Changes in application target, technology, configuration, or elements will impact the results and applicability of the EE metric. Therefore, it is necessary to establish categories of products which describe the elements of similar products, determine representative configurations, and ensure applicability of the EE metric.

Categories or comparison groups of heterogeneous servers are formed by a combination of physical characteristics and limitations. The categories are separated based on computer architecture and physical differences that determine a different energy profile unique to that group. The server metric in the present document may not be applicable to certain categories of systems.

4.3 Product categories

4.3.1 General requirements

For the purposes of the present document, a heterogeneous server shall meet all of the following criteria:

- be marketed and sold as a computer server;
- be designed for and listed as supporting one or more computer server Operating Systems (OSs) and/or hypervisors;
- be designed such that all processors have access to shared system memory and are visible to a single OS or hypervisor;
- be targeted to run user-installed applications typically, but not exclusively, enterprise in nature;
- be packaged and sold with one or more a.c.-d.c or d.c-d.c. power supplies;
- provide support for error-correcting code and/or buffered memory (including both buffered Dual In-line Memory Modules (DIMMs) and buffered on board configurations).

4.3.2 Form-factors

The heterogeneous servers follow the same form-factors of general computer servers. Three form-factors are considered as defined in clause 3.1:

- blade;
- rack-mounted;
- pedestal.

4.3.3 Categories

The heterogeneous server metrics of the present document are applicable to the following product categories as defined and uniquely characterized in clause 3.1:

- X CPU + GPU server;
- X CPU + FPGA server;
- X CPU + ASIC server.

NOTE 1: X represents the form-factor, X may be blade, rack-mounted or pedestal.

NOTE 2: "CPU + APA" means the server is equipped with both CPU and one type of APAs. The type of APA could be GPU, FPGA, ASIC (ASIC could be anyone of the NPU, TPU, DPU, VPU, etc.).

4.4 Product family configuration

4.4.1 General

A heterogeneous server product family configuration shall:

- be from the same model line or machine type;
- either share the same form-factor or share the same mechanical and electrical designs with only superficial mechanical differences to enable a design to support multiple form-factors;
- either share processors from a single defined processor/APA series or share processors/APAs that plug into a common socket/connector/port/slot type;

- have the same number of processors and APAs;
- be dependent on common power supply unit(s).

A given product family can have in excess of tens of thousands of possible combinations of components (separate, distinct configurations). This can include many different processor models with differing APA core count/frequency/memory, as well as numerous types of components such as CPU, I/O devices, memory, and storage devices.

For the purposes of defining the configurations of clauses 4.4.2 and 4.4.3, the following apply to both heterogeneous server and general server with APA cards:

- Calculated APA processor Computing Capacity (dimensionless) = "the number of APAs" × "the number of cores per APA" × "APA frequency (GHz)".

EXAMPLE 1: 2 GPUs, each GPU has 10 496 cores, frequency of 1,7

GPU computing capacity = 35 686,4.

- Calculated APA Memory Capacity (GB) = "the number of APAs" × "the APA memory capacity (GB)".

EXAMPLE 2: 2 GPUs, GPU memory 24 GB

GPU memory capacity = 48 GB.

- Calculated CPU Capacity (dimensionless) = "the number of central processor units (CPUs)" × "the number of cores per CPU" × "the number of threads per core" × "CPU frequency (GHz)".

EXAMPLE 3: 2 CPUs, each CPU has 4 cores and 2 threads per core, frequency of 2,2

CPU capacity = 35,2.

- Calculated Memory Capacity (GB) = "the number of central processor units (CPUs)" × "the number of cores per CPU" × "the number of threads per core" which is subsequently rounded up to "the number of memory channels" × "the lowest capacity DIMM available for the product family".

EXAMPLE 4: 2 CPUs, each CPU has 4 cores and 2 threads per core

Memory capacity = 16 GB

If the product has 4 memory channels and the lowest capacity DIMM is 2 GB, then 8×2 GB_DIMMs (assuming 2 slots per channel).

4.4.2 "High-end" performance configuration

A 'high-end performance configuration' of a heterogeneous server product family means:

- the APA processor with the highest product of core count/frequency/memory;
- CPU processor with the highest product of core count/frequency;
- memory capacity equal to or greater than 3 times the product of the number of APAs;
- cores and hardware threads that represents the highest performance product model within the product family.

4.4.3 "Low-end" performance configuration

A 'low-end performance configuration' of a heterogeneous server product family means:

- the APA processor with the lowest product of core count/frequency/memory;
- the CPU processor with the lowest product of core count/frequency;
- the memory capacity that is at least equal to the product of the number of memory channels;

- the lowest capacity Dual In-line Memory Module (DIMM) (in GB) offered on the server that represents the lowest performance product model within the server product family.

5 Metrics

5.1 Worklets

EE testing tools shall report energy consumption, runtime and performance data for:

- 3 AI training worklets, i.e. Image Classification, Image Segmentation, Natural Language Processing;

and correspondingly:

- 3 AI inference worklets, i.e. Image Classification, Image Segmentation, Natural Language Processing.

The selected training and inference worklets are the most general AI applications.

For each worklet, data is reported for a set of proportional performance intervals and associated, measured energy consumption and runtime values along with other test measurements. The set of individually measured performance, energy consumption and runtime values with their associated efficiency value is termed "interval data".

5.2 Formulae

5.2.1 General

The geometric mean function is used to combine the interval data to produce a worklet Efficiency score, and the worklet efficiency scores are used to create a workload efficiency metric. Using the geometric mean prevents any single performance, energy consumption, worklet runtime or workload efficiency score from unduly influencing the single metric.

In order to create a single EE metric for a server it is necessary to combine the interval Efficiency values for all the different worklets in a specific testing Scenario using the following general procedure:

- combine the interval Efficiency values for the individual worklets using the geometric mean to obtain individual worklet Efficiency values for the worklet;
- combine worklet Efficiency scores using the geometric mean function by workload type (training or inference) to obtain a workload type value;
- combine the workload types using a geometric mean function to obtain a single, total server Efficiency value.

In order to facilitate the deployed power assessment for each heterogeneous server configuration, the Performance and Energy consumption can be combined using the same process as that above.

The Energy Efficiency Metric is defined as:

$$Eff_{server} = Exp(W_{train} \times \ln Eff_{train} + W_{inference} \times \ln Eff_{inference})$$

where:

W_{train} and $W_{inference}$ are respectively the weight factors of AI training worklets and AI inference worklets;

and:

$$Eff_{train} = (\prod_{i=1}^3 Eff_i)^{\frac{1}{3}}$$

where:

$i = 1$ for normalized interval efficiency of training worklet of Image Classification;

$i = 2$ for normalized interval efficiency of training worklet of Image Segmentation;

$i = 3$ for normalized interval efficiency of training worklet of Natural Language Processing;

and:

$$Eff_{inference} = (\prod_{i=1}^3 Eff_i)^{\frac{1}{3}}$$

where:

$i = 1$ for normalized interval efficiency of inference worklet of Image Classification;

$i = 2$ for normalized interval efficiency of inference worklet of Image Segmentation;

$i = 3$ for normalized interval efficiency of inference worklet of Natural Language Processing;

and:

$$Eff_i = \frac{Perf_i}{Pwr_i}$$

where:

$Perf_i$: the normalized interval performance measurements of worklet i ;

Pwr_i : the measured interval power values of worklet i .

5.2.2 Weightings

In the present document, the weighting factor of training and inference worklets are defined as follows:

- $W_{training}$ is the weighting assigned to the CPU worklets = 0,5.
- $W_{inference}$ is the weighting assigned to the Memory worklets = 0,5.

6 Test setup

6.1 General

A single test setup shall be used to undertake the measurements to be used to determine the energy efficiency metrics.

6.2 Environmental conditions

6.2.1 Ambient Temperature

Ambient temperature shall be 25 ± 5 °C.

6.2.2 Relative Humidity

Relative humidity shall be between 15 % and 80 %.

6.3 Temperature sensor

The temperature sensor shall:

- have a temperature measurement accuracy of no greater than $\pm 0,5$ °C when measured no more than 50 mm in front of (upwind of) the main airflow inlet of the EUT;

- b) have a logging performance of minimum reading rate: four samples per minute.

6.4 Input power

During the test, the input voltage tolerance to the EUT shall be as specified below:

- a) $\leq 1,0\%$ if power consumption is $\leq 1\,500\text{ W}$;
 b) $\leq 4,0\%$ if power consumption is $> 1\,500\text{ W}$.

For a.c. input voltages, the frequency tolerance shall be $\leq 1,0\%$ and the total harmonic distortion shall be as specified below:

- a) $\leq 2,0\%$ if power consumption is $\leq 1\,500\text{ W}$;
 b) $\leq 5,0\%$ if power consumption is $> 1\,500\text{ W}$.

6.5 Power requirements

Testing shall be conducted within the following frequency and voltage ranges:

- AC Frequency: $\pm 1\%$ of 50 Hz.
- Voltage: $\pm 5\%$ of 220 V, 230 V, 240 V.

6.6 Energy consumption analyser

The following requirements apply to single-phase power analyzer.

The power analyzer shall report true root mean square power and at least two of the following parameters: voltage, current, and power factor.

The power analyzer shall:

- a) have a valid calibration certificate or equivalent, to support its use at the time the tests are carried out;
- b) feature an available current crest factor of 3 or more at its rated range value;
- c) for power analyzers that do not specify the current crest factor, the power analyzer shall be capable of measuring an amperage spike of at least 3 times the maximum amperage measured during any 1 second sample;
- d) have a minimum frequency response of 3,0 kHz;
- e) have a minimum resolution of:
- 0,01 W for measurement values less than 10 W;
 - 0,1 W for measurement values from 10 W to 100 W; and
 - 1,0 W for measurement values greater than 100 W.
- f) have a power measurement accuracy of greater than 99,0 %;
- g) have a logging performance of:
- 1) minimum reading rate: one set of measurements (power measurement in W) per second;
 - 2) data averaging interval equal to the reading rate.

The worklet runtime is recorded, including the start and end time. The energy consumption of a worklet is:

$$E_{worklet} = \int_{t_{start}}^{t_{end}} p t dt$$

where t_{start} is the start time of the worklet, t_{end} is the end time of the worklet, p is the power.

6.7 Test tool

A server EE testing tool shall be used to determine the EE of the EUT.

6.8 Controller system

The controller system shall be capable of the following functions:

- a) start and stop each segment (phase) of the performance benchmark;
- b) control the workload demands of the performance benchmark;
- c) start and stop data collection from the power analyzer so that power and performance data from each phase can be correlated;
- d) store log files containing benchmark power and performance information;
- e) convert raw data into pdf and html format for benchmark reporting, submission and validation;
- f) collect and store environmental data (temperature), if automated for the benchmark.

The controller system may be a server, a desktop computer, or a laptop and shall be used to record input power from the equipment specified in clause 6.6 and temperature data from the equipment specified in clause 6.3.

The controller system and the EUT shall be connected to each other via one port of an Ethernet network switch.

7 Equipment Under Test

7.1 Configuration

The configuration of the EUT shall be as specified in Table 1.

Table 1: Configuration of EUT

Index	Items	Detailed requirements
A)	As-shipped condition	Products shall be tested in their "as-shipped" condition, which includes both hardware revision and system settings, unless otherwise specified in this test method. Where relevant, all software options shall be set to their default condition.
B)	Measurement location	All power measurements shall be taken at a point between the a.c. power source and the EUT. Uninterruptible Power Supply (UPS) units shall not be connected between the power meter and the EUT. The power meter shall remain in place until all the state efficiency power data are fully recorded. When testing a blade system, power shall be measured at the input of the blade chassis (i.e. at the power supplies that provide chassis distribution power).
C)	Air flow	Purposefully directing air in the vicinity of the measured equipment in a way that would be inconsistent with normal practices at the intended installation location is prohibited.

Index	Items	Detailed requirements
D)	Power supplies	<p>All Power Supply Units (PSUs) shall be connected and operational.</p> <p>For EUT with multiple PSUs:</p> <ul style="list-style-type: none"> all power supplies shall be connected to the a.c. power source and operational during the test; if necessary, a Power Distribution Unit (PDU) may be used to connect multiple power supplies to a single source (if a PDU is used, any overhead electrical use from the PDU shall be included in the power measurement of the EUT). <p>For blade servers with half-populated chassis configurations, the power supplies for the unpopulated power domains can be disconnected (see Table 2, D) for more information).</p>
E)	Power Management and Operating System	<p>The as-shipped operating system or a representative operating system shall be installed. Products that are shipped without operating systems shall be tested with any compatible operating system installed.</p> <p>For all tests, the power management techniques and/or power saving features shall be left as-shipped.</p> <p>Any power management features which require the presence of an operating system (i.e. those that are not explicitly controlled by the Basic Input Output System (BIOS) or management controller) shall be tested using only those power management features enabled by the operating system by default.</p>
F)	Storage	<p>Products shall be tested for qualification with at least two HDD or two SSD installed. Products that do not include pre-installed drives (HDD or SSD) shall be tested using a storage configuration used in an identical model for sale that does include preinstalled drives.</p> <p>Products that do not support installation of drives (HDD or SSD) and, instead, rely exclusively on external storage solutions (e.g. storage area network) shall be tested using external storage solutions.</p>
F1)	APAs	The EUT shall be tested with extra expansion APAs or add-in cards removed except for the tested APAs, when measuring the efficiency and server performance in active state.
F2)	Memory Channels	Any EUT with some memory channels unpopulated should be reconfigured so all memory channels have the same DIMMs.
G)	Blade System and Dual/Multi-Node Servers	A blade system or dual/multi-node server shall have identical configurations for each node or blade server including all hardware components and software/power management settings. These systems shall also be measured in a way that ensures all power from all tested nodes/blade servers is captured by the power meter during the entire test.
H)	Blade Chassis	<p>The blade chassis, at a minimum, shall have power, cooling, and networking capabilities for all the blade servers. The blade chassis shall be populated as specified in Table 2, D).</p> <p>All power measurements for blade systems shall be made at the input of the blade chassis.</p>
I)	BIOS and EUT System Settings	All BIOS settings shall remain as-shipped unless otherwise specified in the test method.
J)	Input/Output (I/O) and Network Connection	<p>The EUT shall have at least one port connected to an Ethernet network switch. The switch shall be capable of supporting the EUT's highest and lowest rated network speeds. The network connection shall be live during all tests, and, although the link shall be ready and able to transmit packets, no specific traffic is required over the connection during testing.</p> <p>For the purpose of testing ensure the manufacturer shall offer at least one Ethernet port (using a single add-in card only if no on-board Ethernet support is offered).</p>
K)	Energy Efficient Ethernet (EEE)	Products shipped with support for Energy Efficient Ethernet (compliant with IEEE 802.3az [1]) shall be connected only to Energy Efficient Ethernet compliant networking equipment during testing. Appropriate measures shall be taken to enable EEE features on both ends of the network link during all tests.

7.2 Test procedure

The EUT test configuration shall be in accordance with Table 2.

Table 2: Test configuration of EUT

Index	Detailed requirements
A)	The EUT shall be tested with all processor connectors populated.
B)	The EUT shall be installed in a test rack or other static location and shall not be physically moved until testing is complete.
C)	For a multi-node system, the power consumption per node of the EUT shall be measured in the fully-populated blade chassis configuration. All multi-node servers installed in the blade chassis shall be identical, sharing the same configuration.
D)	<p>For a blade system, the blade server power consumption of the EUT shall be measured in the half-populated blade chassis configuration.</p> <p>For blade systems, the blade chassis shall be half populated as follows:</p> <ul style="list-style-type: none"> • The number of blade servers required to populate half the number of single-wide blade server slots available in the blade chassis shall be calculated. • For blade chassis having multiple power domains, the number of power domains shall be chosen that is closest to filling half of the blade chassis. In a case where there are two choices that are equally close to filling half of the blade chassis, test with the domain or combination of domains which utilize a higher number of blade servers. <p>EXAMPLE 1: A blade chassis supports up to 7 single-wide blade servers on two power domains. One power domain supports 3 blade servers and the other supports 4 blade servers. In this example, the power domain which supports 4 blade servers would be fully populated during testing, while the other power domain would remain unpopulated.</p> <p>EXAMPLE 2: A blade chassis supports up to 16 single-wide blade servers on four power domains. Each of the four power domains supports 4 blade servers. In this example, two of the power domains would be fully populated during testing, while the other two power domains would remain unpopulated.</p> <p>All user manual or manufacturer recommendations shall be followed for partially populating the blade chassis, which may include disconnecting some of the power supplies and cooling fans for the unpopulated power domains.</p> <p>If user manual recommendations are not available or are incomplete, then the following guidance shall be followed:</p> <ul style="list-style-type: none"> • Completely populate the power domains. • If possible, disconnect the power supplies and cooling fans for unpopulated power domains. • Fill all empty bays with blanking panels or an equivalent airflow restriction for the duration of testing.
E)	<p>The EUT shall be connected to a live Ethernet (IEEE 802.3 [3]) network switch.</p> <p>The live connection shall be maintained for the duration of testing, except for brief lapses necessary for transitioning between link speeds.</p>
F)	<p>The Controller System required to provide the workload harness control of server EE testing tool, data acquisition, or other EUT testing support shall be connected to the same network switch as the EUT and satisfy all other EUT network requirements.</p> <p>Both the EUT and Controller System shall be configured to communicate via the network.</p>
G)	The power meter shall be connected to an a.c. voltage source set to the appropriate voltage and frequency for the test, as specified in Table 1.
H)	The EUT shall be connected to the measurement power outlet on the power meter following the guidelines in Table 1.
I)	The data output interface of the power meter and the temperature sensor shall be connected to the appropriate inputs of the Controller System.
J)	It shall be verified that the EUT is configured in its as-shipped configuration.
K)	It shall be verified that the Controller System and EUT are connected on the same internal network via an Ethernet network switch.
L)	Using a normal ping command, It shall be verified that the Controller System and EUT can communicate with each other.
M)	Server EE testing tool shall be installed on the EUT and the Controller System.

8 Measurement

8.1 Measurement for active state

The measurement shall be in accordance with Table 3.

Table 3: Measurement of active state efficiency

Index	Detailed requirements
A)	The EUT shall be re-booted. System caches and any stored information that may affect the active state efficiency metric shall be flushed.
B)	All steps outlined in the User Guide of server EE testing tool shall be followed to successfully run server EE testing tool. There shall be no manual intervention or optimization of the Controller System, EUT, or its internal and external environment during the execution of the server EE testing tool.
C)	Once server EE testing tool, procedures are completed, the following output files shall be included with all testing results: <ol style="list-style-type: none"> 1. Results.xml 2. Results.html 3. Results.txt 4. All results-chart png files (e.g. results-chart0.png, results-chart1.png, etc.) 5. Results-details.html 6. Results-details.txt 7. All results-details-chart .png files (e.g. results-details-chart0.png, results-details-chart1.png, etc.)

8.2 Re-verification and audits

Due to the manufacturing variance in components, number of significant power elements in the system, and run to run variations, retest of an individual server, a server from the same configuration, or a server family will result in values that may vary from the initial testing.

Adjusting for a 90 % confidence interval, retesting values that are within 15 % of the passing level shall be acceptable under re-test. Re-test or audit resulting in greater than 15 % error to the designated passing level shall be re-evaluated as a product family to determine compliance.

For audits and re-verification of Power Supply Units, re-testing of PSU efficiency shall be not be lower than the declared value by more than 2 % and the power factor shall not be lower than the declared value by more than 10 %. Variance beyond these levels shall require re-evaluation of the Power Supply Unit.

8.3 Measurement for power supply

8.3.1 Measurement for internal power supply

The measurement shall be in accordance with Table 4.

Table 4: Measurement of internal power supply

Index	Detailed requirements
A)	For all types of internal power supplies, the efficiency and the power factor shall be measured at 10 %, 20 %, 50 % and 100 % of the rated [nameplate] output power.
B)	Test setup, test conditions, and measurement instrument specifications shall comply with clause 6.2.
C)	This test procedure assumes that the internal power supply meets the following criteria: <ul style="list-style-type: none"> Detailed input and output ratings are available on the name plate or in manufacturer's literature, specifying the maximum loads that can safely be placed on each individual d.c. output voltage bus and, where necessary, groupings of those voltage busses. The power supply has connectors that allow the d.c. output voltage busses to be connected and disconnected from the powered product non-destructively. The power supply can be easily detached from the housing of the product it powers, without causing harm to other circuits and components of the product.
D)	In the event the above criteria C) are not met, a test board (see clause 8.3.2) shall be provided to enable testing.
NOTE 1:	The power supply can be easily detached from the housing of the product it powers, without causing harm to other circuits and components of the product.
NOTE 2:	Such data could already be available from the manufacturer of the power supply; in such cases, the manufacturer could decide to use them. However, where 3 rd party test results are used, it is the responsibility of the manufacturer to assess the trustworthiness of the sources.

8.3.2 Measurement for test board power supply

8.3.2.1 General

Tests specified in this clause shall be made on either:

- the power supply of the computer under test, after it has been disconnected from the powered parts and extracted from the housing; or alternatively;
- another unit, representative of the built-in power supply.

8.3.2.2 Test loads

Active loads such as electronic loads or passive loads such as rheostats may be used as d.c. test loads. They shall be able to maintain the required current loading set point for each output voltage within an accuracy of $\pm 0,5$ %.

8.3.2.3 Test leads and wiring

Appropriate wires shall be used to avoid excessive overheating and reduce voltage drop across the wires. If measurements are not taken directly at the connector pins, voltage drop against the additional wires shall be taken into account.

NOTE: If applicable, voltage drop across the input wires will be subtracted from the measured input voltage, and the voltage drop across the output wires will be added to the measured input voltage for the measurement of power efficiency.

8.3.2.4 Warm up time

Whereas internal temperature of the components could impact its efficiency, the power supply under test shall be loaded up to the test load for a period of at least 15 minutes or until the reading over two consecutive five-minute intervals does not change by more than $\pm 0,2$ %.

8.3.2.5 Power measurements

The true RMS root mean square wattmeter used to carry out a.c. input power measurements shall meet the requirements of clauses 5.7 and 5.8 of EN 62623[2]. Input power shall be determined using an averaging technique over a minimum of 32 input cycles utilizing the measurement instrument averaging function.

For appliances connected to more than one phase, the power measurement instrument shall be equipped to measure the total power of all phases connected.

D.c. output power measurements shall be made either with a suitably calibrated voltmeter and ampere meter or with a suitably calibrated power meter.

9 Measurement report

The following metrics/measurements shall be listed in the measurement report:

- 1) Efficiency metric;
- 2) Reported maximum power.

It is accepted that the measurement accuracy of the metrics is $\pm 10\%$. Any subsequent assessment within this range shall be considered to be consistent with the quoted value.

The following additional information shall also be reported/provided under the technical documentation:

- 1) Server EE testing tool, test report and supporting data;
- 2) Author, site, and date of the testing;
- 3) Product type;
- 4) Server configuration including:
 - a) component manufacturer, product ID, number of units;
 - b) component product type for APA, CPU, memory, drive (HDD or SSD);
 - c) system test model number;
 - d) manufacturer name;
 - e) server product ID;
- 5) Power Supply Unit test information or reference to previously conducted test report;
- 6) The extremes of server inlet test temperature during the test;
- 7) Revision numbers for each of test software elements used;
- 8) Software (e.g. Java™) revision and source used;
- 9) Controller product model ID;
- 10) Test equipment (power meter, thermal sensor/meter): manufacturer, model, ID, and calibration date;
- 11) Server EE testing tools, suite revisions;
- 12) Server EE testing tools, product configuration revision.

History

Document history		
V1.0.0	February 2025	Membership Approval Procedure MV 20250408: 2025-02-07 to 2025-04-08