

DELIVERABLE D3.2.6

PROCEDURES FOR TRACKING ENERGY CONSUMPTION OF THE TIER-1 SYSTEM IN PLACE

TeRABIT

APRIL 2026

Mirko Cestari, Samuele Casi, Giuseppa Muscianisi, Debora Testi (CINECA)
Giorgio Bolzon, Stefano Salon (OGS)

Eccetto dove diversamente specificato, il contenuto di questo sito è
concesso in licenza con Attribuzione Creative Commons - 4.0 internazionale.



Executive Summary

This document reports on the energy efficiency design choice and on the power consumption and energy tracking of the PRACE-Italy Tier-1 HPC system (GAIA).

Table of Content

1. DESIGN AND DELIVERY OF AN ENERGY EFFICIENT TIER-1 SYSTEM.....	1
1.1. Introduction.....	1
1.2. Energy efficient design.....	1
1.3. Evaluation and execution of the offered solution.....	4
2. TIER-1 SYSTEM ENERGY TRACKING.....	5
2.1. Introduction.....	5
2.2. ISO 50001	5
2.3. CINECA energy management system.....	8
2.4. GAIA Energy monitoring	10
3. CONCLUSION.....	11

List of Figures

Figure 1. High-level architecture of GAIA.....	2
Figure 2. Power IT adsorbed profile during the acceptance stress tests.	3
Figure 3. Plan-do-check-act cycle.....	7
Figure 4. Sketch of the CINECA data centre C2 in Bologna data hosted in the DaMa Technopole premises.....	9
Figure 5. Monthly average Power Usage Effectiveness (PUE) of the data centre.	10
Figure 6. Average power consumption in kW of GAIA PRACE-Italy Tier-1 system.....	11

List of Tables

Table 1. Comparison of delivered computing performance (TFlops) and TDP of GAIA processing units, CPUs and GPUs.....	3
--	---

1. DESIGN AND DELIVERY OF AN ENERGY EFFICIENT TIER-1 SYSTEM

1.1. Introduction

The PRACE-Italy Tier-1 system, GAIA, was developed and designed according to strict best practices that aim to minimise environmental impact. This approach not only considers the proper disposal of components and infrastructure materials during decommissioning but also emphasises improving the system's energy efficiency as the primary priority.

These guidelines were maintained during each stage of the system delivery process:

- Technical specifications were designed to incorporate requirements that underscored the importance of energy efficiency.
- Offer evaluations included criteria that incentivised energy-efficient solutions.
- The awarded solution was delivered and the contract executed with a clear priority on energy efficiency.

It is essential to acknowledge that Tier-1 systems, which involve capital investments greater than 15 million euros, characteristically have substantial energy requirements. Nevertheless, as computing performance continues to increase significantly with each new hardware generation, prioritizing enhanced energy efficiency becomes increasingly critical.

A direct relationship exists between power usage and computational performance; the latter is primarily determined by the number of transistors in processing units, and their density is influenced by the employed lithography during manufacturing. Advancements in lithographic technology typically yield improved performance at equivalent power consumption levels.

Therefore, newer lithography should be prioritized to target higher energy efficiency.

1.2. Energy efficient design

Deliverable D3.2.1 contains the technical specifications. Based on these, the Tier-1 system was designed to support up to 10,000 virtual machines (cloud capacity, tender specification requirement 4.2.1-10). The goal was to create a system with more than 100,000 cores, assuming each VM would typically use 8 vCPUs and 32 GB of RAM.

This required throughput capacity needed to be delivered within a strictly defined power envelope. According to the data center tender specifications, the allowable power consumption and corresponding cooling capacity were capped at 550 kW under normal operating conditions (data hall power supply, requirement 4.1.1-7), with a maximum peak of 700 kW.

This suggests exceptionally low power consumption for each core. The intended system was expected to deliver:

$$700.000 \text{ Watt} / 100.000 \text{ cores} = 7 \text{ Watt per provided core}$$

This simple estimation does not include all the other infrastructure components that consume power, such as GPUs, RAM, network interfaces, switches, fans, transceivers, and so on.

The CPU requirements outlined are stringent. Requirement 4.4-4 specifies that the CPUs referenced in requirement 4.4-3 should prioritize achieving the maximum core count relative to power consumption (i.e., thermal design power, TDP). The proposed solution requires a CPU with at least 100 cores per processor, utilising the latest manufacturing technology. When these specifications were prepared, server CPUs had a maximum TDP of 350 W, which equates to approximately 3.5 W per core.

The result system high level architecture is reported in Figure 1.




Figure 1. High-level architecture of GAIA.

The CPU partition uses standard dual socket server nodes, each equipped with two Intel Xeon 6780E CPUs, giving a total of 288 physical cores per node and 1 TB of DDR5 RAM running at 6400MT/s. The CPU's reported TDP is 330 W, which works out to 2.29 W per core (660 W total divided by 288 cores). For comparison, the Leonardo DCGP partition offers nodes with two processors, each with 56 cores and a TDP of 350 W, resulting in 6.25 W per core, more than twice what the GAIA Tier-1 system provides.

This outcome aligns with the system's objective to maximise core count per watt, thereby enabling optimal allocation of virtual machines. Achieving this is a primary goal for cloud infrastructure solutions.

The GPU partition features CPUs Intel Xeon 8592+, for a total of 128 cores and 1 TB DDR5 RAM 4800MT/s. Moreover, this partition offers high-end NVIDIA GPUs both for AI and 3d rendering workloads (L40s) and for scientific use cases requiring double precision FP64 arithmetic (A30 and H100). Thanks to the contract modifications, more GPU capacity was acquired – highlighted in red in Figure 1 – for a total of:

- 80 GPUs NVIDIA A30
- 112 GPUs NVIDIA L40s

- 8 GPUS NVIDIA H100 NVL

GPUs are typically more energy efficient than CPUs in providing energy efficient computing performance according to Table 1.

	GPUs			CPUs	
Version	A30	H100	L40s	Intel 6780E	Intel 8592+
TDP (W)	165	400	350	330	350
Performance FP64 (TFlops)	7	40	-	4	4.5

Table 1. Comparison of delivered computing performance (TFlops) and TDP of GAIA processing units, CPUs and GPUs.

The storage is composed of two main tiers:

- CEPH
 - CEPH high ops partition, based on full flash solution
 - CEPH high capacity, based on HDD drives.
- GARDA-B data lake
 - Full flash, multi-protocol solution

The storage infrastructure usually uses far less power than the computing infrastructure. The peak power usage of the four storage racks is under 70 kW, which accounts for about 10% of the system's total power consumption, which is around 700 kW as shown in Figure 2.



Figure 2. Power IT adsorbed profile during the acceptance stress tests.

Finally, it is worth noting that the limit of 700 kW required in the technical specifications seems to be exceeded in the plot of Figure 2. It must be reminded that 8 H100 GPUs and 32 L40s were added as a contract variation from the original technical specifications' requirements. This amounts to roughly 15 more kW of peak power consumption to the initial offered infrastructure.

1.3. Evaluation and execution of the offered solution

The proposed solutions were assessed with respect to energy efficiency according to two distinct evaluation criteria.

- **The technical value of system design** was determined by assessing offers that utilize energy-efficient components, particularly CPUs and GPUs.
- **Environmental impact**, was assessed according to DNSH criteria and, where applicable, energy certification such as Energy Star.¹

The awarded offer provided the best solution according to both evaluation criteria.

In addition, rear door heat exchangers (RDHx) were acquired, as documented in deliverable D3.2.3, to optimise the energy efficiency of the system. Rear Door Heat Exchangers (RDHx) improve energy efficiency by removing heat directly at the rack level, capturing server exhaust air before it mixes with the data-hall environment and reducing recirculation. By intercepting heat at the source, RDHx significantly lowers the thermal load on room-level air-conditioning systems, enabling reduced fan operation and lower dependence on mechanical refrigeration. The use of liquid cooling allows effective heat removal with higher chilled-water supply temperatures (19–26 °C), improving chiller efficiency and enabling free-cooling operation. In addition, RDHx limits the need to increase airflow to support high rack power densities, reducing fan energy consumption while enabling high-density server deployments without major upgrades to facility cooling infrastructure.

These cooling units help maintain a low Power Usage Effectiveness (PUE) in the data centre (see section 2.4). The monthly average PUE ranges from 1.16 to 1.12, significantly lower than the global average of 1.58 reported in the Uptime 2023 survey². According to the Uptime reference, the global average PUE improved from 2007 to 2014 but now appears to have plateaued at approximately 1.58.

¹ Verification that the electronic products purchased bear a Type I ISO eco-label in accordance with UNI EN ISO 14024, the EPA ENERGY STAR label, or other equivalent documentation

² <https://journal.uptimeinstitute.com/global-pues-are-they-going-anywhere/>

2. TIER-1 SYSTEM ENERGY TRACKING

2.1. Introduction

Considering the worldwide increase in energy demand, it is no longer sufficient for organizations with high energy consumption to just adopt energy-saving strategies. It is also necessary to implement processes and practices that can improve the efficiency of their energy consumption.

CINECA uses ISO 50001 standard practices as the foundation for its energy efficiency monitoring and improvement efforts. The ISO 50000 series includes several documents that set standards for energy management in organizations, each named ISO 5000X. ISO 50001, which is the main document of this family, outlines requirements for energy management systems (EMS) and provides guidance on how to implement those requirements.

ISO 50001 aims to help organizations set up internal systems and processes that continually boost energy performance—including energy efficiency, use, and consumption. The standard outlines requirements for an Energy Management System (EMS) and encourages a culture focused on improving energy performance, which relies on strong commitment from top management throughout every level of the organization. Implementing this standard often leads to cultural changes within the company.

The process of developing and implementing an EMS includes defining energy policies, objectives, targets, and action plans related to energy efficiency, use, and consumption, while also meeting applicable requirements (such as physical and legal requirements). Implementing an EMS enables an organization to set and achieve objectives and energy targets, to take actions to improve its energy performance, and to demonstrate the conformity of its energy management system to the requirements of the ISO 50001 standard.

2.2. ISO 50001

The ISO 50001 standard is issued by the International Organization for Standardization (ISO). ISO is a worldwide federation of national bodies (known as ISO member bodies). The ISO standards are typically written by ISO technical committees that are responsible for a specific subject area, with assistance from members of other international organizations. For example, ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters related to electrotechnical standardization.

The ISO 50001 standard:

- is applicable to any organization regardless of its type, size, complexity, geographical location, organizational culture or the products and services it provides;
- is applicable to activities affecting energy performance that are managed and controlled by the organization;
- is applicable irrespective of the quantity, use, or types of energy consumed;
- requires demonstration of continuous energy performance improvement, but does not define levels of energy performance improvement to be achieved; and

- can be used independently, or aligned, or integrated with other management systems.

The EMS of an organization needs to include:

- all documented information required by the standard, namely the organization's energy policy, action plan, energy use review, method for determining and updating the Energy Performance Indicators (EnPIs), energy baseline, collected energy data, competence documentation, communication documentation, operation documentation, performance evaluation, audit program, and management review; plus
- documented information that the organization deems to be necessary for the effectiveness of the EMS and to demonstrate improvements in energy performance

To comply with the ISO 50001 standard, the EMS of an organization needs to be based on the Plan-Do-Check-Act (PDCA) continuous improvement framework and must incorporate energy management into existing organizational practices, as illustrated in Figure 3.

In the context of energy management, the PDCA approach can be outlined as follows.

- **Plan:** Understand the context of the organization, establish an energy policy and an energy management team, consider actions to address risks and opportunities, conduct an energy use review, identify significant energy uses (SEUs) and establish energy performance indicators (EnPIs), energy baseline(s) (EnBs), objectives and energy targets, and action plans necessary to deliver results that will improve energy performance in accordance with the organization's energy policy.
- **Do:** Implement the action plans, operational and maintenance controls, and communication plans, ensure the competence of the staff and consider energy performance in the design and procurement processes.
- **Check:** Monitor, measure, analyze, evaluate, audit and conduct management review(s) of energy performance and the EMS.
- **Act:** Take actions to address non-conformities and continuously improve energy performance and the EMS.



Figure 3. Plan-do-check-act cycle.

Top management must demonstrate leadership and commitment to continuously improve energy performance and effectiveness - through the EMS - by:

- ensuring that the EMS scope and boundaries are established;
- ensuring that the energy policy, objectives and energy targets are established and are compatible with the strategic direction of the organization;
- ensuring the integration of the EMS requirements into the organisation's business processes;
- ensuring that action plans are approved and implemented;
- ensuring that the resources needed for the EMS are available;
- communicating the importance of effective energy management and of conforming to the EMS requirements;
- ensuring that the EMS achieves its intended outcome(s);
- promoting continuous improvement of energy performance and the EMS;
- ensuring that an energy management team is formed;
- directing and supporting staff to contribute to the effectiveness of the EMS and to energy performance improvement;
- supporting staff in other relevant management roles to demonstrate their leadership as it applies to their areas of responsibility;
- ensuring that the EnPI(s) appropriately represent(s) energy performance; and

- ensuring that processes are established and implemented to identify and address changes affecting the EMS and energy performance within the scope and boundary of the EMS.

The energy management team has the responsibility and authority to implement an energy management system and to deliver energy performance improvement.

2.3. CINECA energy management system

CINECA serves as Italy's top supercomputing centre for scientific research, managed by a consortium that includes Italian universities, national research institutes, and the Ministry of Universities and Research. Its main data centres are located in Casalecchio di Reno and DaMa Technopole in Bologna.

The CINECA consortium adopts energy policies to guide its energy management system, aiming for better energy performance and lower energy use. These policies, outlined in an official document, demonstrate a commitment to energy objectives and reducing CO₂ emissions. CINECA intends to fully integrate these policies into its organizational culture and other procedures for lasting impact.

CINECA has completed all necessary steps to achieve ISO 50001 certification. The consultancy firm RINA assisted CINECA in this process. RINA is accredited by Accord Certification in India according to the ISO 50003 standard for every technical area—including light and medium industry, heavy industry, buildings, complex structures, transport, mining, agriculture, and energy supply. Since there are no geographical restrictions, RINA can offer services to any type of organisation worldwide. This result was possible thanks to the work of the Energy Management Team (which is composed of specifically trained internal staff). CINECA developed its own EMS. The certification was issued by the accredited body RINA SERVICE. The current roadmap includes the maintenance of the certificate through periodic audits carried out by the certifying body.

CINECA ISO 50001³: energy management system certification is reported on the main website of the organization. All the systems hosted in the DaMa Technopole data center (see Figure 4) are covered in the certification, including the GAIA PRACE-Italy Tier-1 system.

³ https://www.cineca.it/sites/default/files/2025-07/_Certificatelo50001.pdf




Figure 4. Sketch of the CINECA data centre C2 in Bologna data hosted in the DaMa Technopole premises.

In general, the energy review represents the core of the ISO 50001 standard. For each energy use, the analysis focuses on elements of interest, breaking down the various uses according to specific needs. For example, Leonardo is divided into DCGP and Booster, while thermal plants are split according to tempered water services, chilled water services, and pumping systems.

Once direct energy consumption is tracked, it is necessary to compare it with production data in order to establish meaningful correlations between power/energy data and the output associated with a given use or process. For HPC systems, production is measured in computing performance (Flops), while for Cloud ICT servers SPEC performance is typically used⁴.

With regard to measurements in CINECA data centers, and in particular at the DaMa Technopole:

- On a monthly basis, electrical consumption measured by ENEL meters is collected.
- Each HPC system has multiple power supplies fed by different electrical branches. A multimeter is installed to measure the consumption of the four busbars, which in turn supply distribution boxes feeding individual racks.
- It is not possible to measure consumption at the individual rack level, as there are no multimeters installed on the PDUs or on the small distribution boxes feeding individual racks. Therefore, only total measurements at the busbar level are available.
- The data from the various busbars are subsequently summed. The multimeters are installed on the 1000 A main feeder heads; no multimeters are present on the 250 A distribution boxes.

⁴ <https://www.spec.org/serf/>

- The same approach applies to other computing systems and uses (chillers, adiabatic coolers, pumps, etc.).

CINECA also tracks PUE values, which at the Technopole consist of several components. The numerator includes IT load + cooling + losses, while the denominator includes IT load only. PUE is calculated at different levels: for the entire site: data center B5 and C2, for B5 only, for C2 only, and, starting from November 2025, also split by Hall 1 and Hall 2, and in the future by Hall 3. It is not possible to calculate a PUE for individual computing systems, as it is physically impossible to allocate specific cooling contributions from perimeter computer room air handlers (CRAHs) to single systems, since the airflow is distributed throughout the entire hall.

All measurements are recorded within a software system that manages both measurement instruments and plant automation systems (mechanical, electrical, special systems, fire protection, etc.). On a monthly basis, a report containing the main measurements is extracted and the data are recorded according to the data collection plan (ISO 50001 documentation).

2.4. GAIA Energy monitoring

GAIA is included within the ISO 50001 certification perimeter. Last year, as it was not yet in production, only absolute energy consumption was tracked.

Plots of Figure 5 and Figure 6 report examples of the energy and power consumption measures tracked in CINECA EMS and part of the ISO 50001 standard certification. The plots show an outstanding average PUE value, which is a result of a combination computing infrastructure design and data center cooling and power infrastructure adaptation.




Figure 5. Monthly average Power Usage Effectiveness (PUE) of the data centre.

GAIA system reports a baseline of roughly 100 kW of power consumption. The system will be opened to users and start the production phase in May 2026. Power consumption will be monitored and cooling infrastructure will be tuned to maximize energy efficiency.




Figure 6. Average power consumption in kW of GAIA PRACE-Italy Tier-1 system.

3. CONCLUSION

The document highlights the comprehensive energy management and monitoring practices employed at CINECA, with a particular focus on Power Usage Effectiveness (PUE) tracking across multiple facility levels. The approach ensures that data is collected and analysed through integrated software systems, aligning with ISO 50001 standards to support continual improvements in operational efficiency. Notably, challenges in attributing cooling loads to individual computing systems are acknowledged, reflecting the complexity of airflow distribution in the data centre environment.

The introduction of the GAIA system within the ISO 50001 certification perimeter marks a significant development, with its baseline power consumption monitored and future optimisation planned as it enters production. Graphical analyses further underscore the centre's commitment to achieving outstanding energy performance, facilitated by robust infrastructure design and ongoing adaptation of cooling and power systems. These efforts collectively demonstrate a proactive strategy for energy efficiency and sustainability in high-performance computing operations.