CABILDO DE TENERIFE

SERVICE CATALOGUE
ITER Supercomputing Center

ITER

CABILDO DE TENERIFE

The ITER Supercomputing Center offers public and private entities a high-capacity processing resource designed to enhance and expand the scope of their research at both national and international levels.

The center makes available to its users an advanced technological infrastructure that allows the execution of complex calculations and the efficient processing of large volumes of data.

Entities accessing our services can benefit from available computing resources, optimizing their research and development projects. In addition to processing power, the ITER Supercomputing Center guarantees a secure and reliable environment for handling sensitive data.

Our commitment to excellence is reflected in the quality of our systems and the specialized technical support we provide, ensuring that each project has the necessary resources to achieve its goals.

The center's infrastructure is designed to be scalable and flexible, adapting to the specific needs of each entity and providing a comprehensive solution that covers everything from data storage to advanced analysis.

The ITER Supercomputing Center is located in the D-ALiX data processing center, a Tier III+ facility with over 2,000 square meters dedicated to equipment colocation. This data center, integrated into a renewable energy facility and connected to international submarine cables, ensures maximum reliability and energy efficiency.

The location within D-ALiX enables the creation of hybrid solutions, combining client infrastructure hosted in the data center with the advanced capabilities of the ITER Supercomputing Center. This integration provides an optimized environment adaptable to diverse business and research needs.

Additionally, companies within D-ALiX benefit from the Canary Islands' special tax regime, offering significant economic advantages. This synergy between local resources and ITER's supercomputing power ensures a robust and cost-effective solution for technological and scientific projects.

Our Staff

Our supercomputing center boasts a highly qualified team responsible for maintaining both the data center infrastructure and communications, as well as the services provided by the supercomputing center itself. This specialized team ensures that all operations are conducted efficiently and without interruptions, offering continuous support and advanced technical solutions.

The personnel responsible for our infrastructure have extensive experience in managing complex, high-tech systems. Their expertise ensures that the data center environment, including hardware and software, is always in optimal condition, maximizing uptime and the reliability of the services offered.

Our communication experts work diligently to maintain a robust and secure network, utilizing the latest technologies to ensure seamless operation for both R&D and commercial networks.

Collaboration

Access to ITER Supercomputing Center services not only enhances the competitiveness of organizations in scientific and technological fields but also facilitates collaboration and knowledge exchange at a global level.

By integrating into our network, entities can participate in joint initiatives and multidisciplinary projects, promoting an ecosystem of continuous innovation and development.

With the support of the ITER Supercomputing Center, public and private entities can transform their ideas into tangible realities, positioning themselves at the forefront of research and technology.

Infrastructure

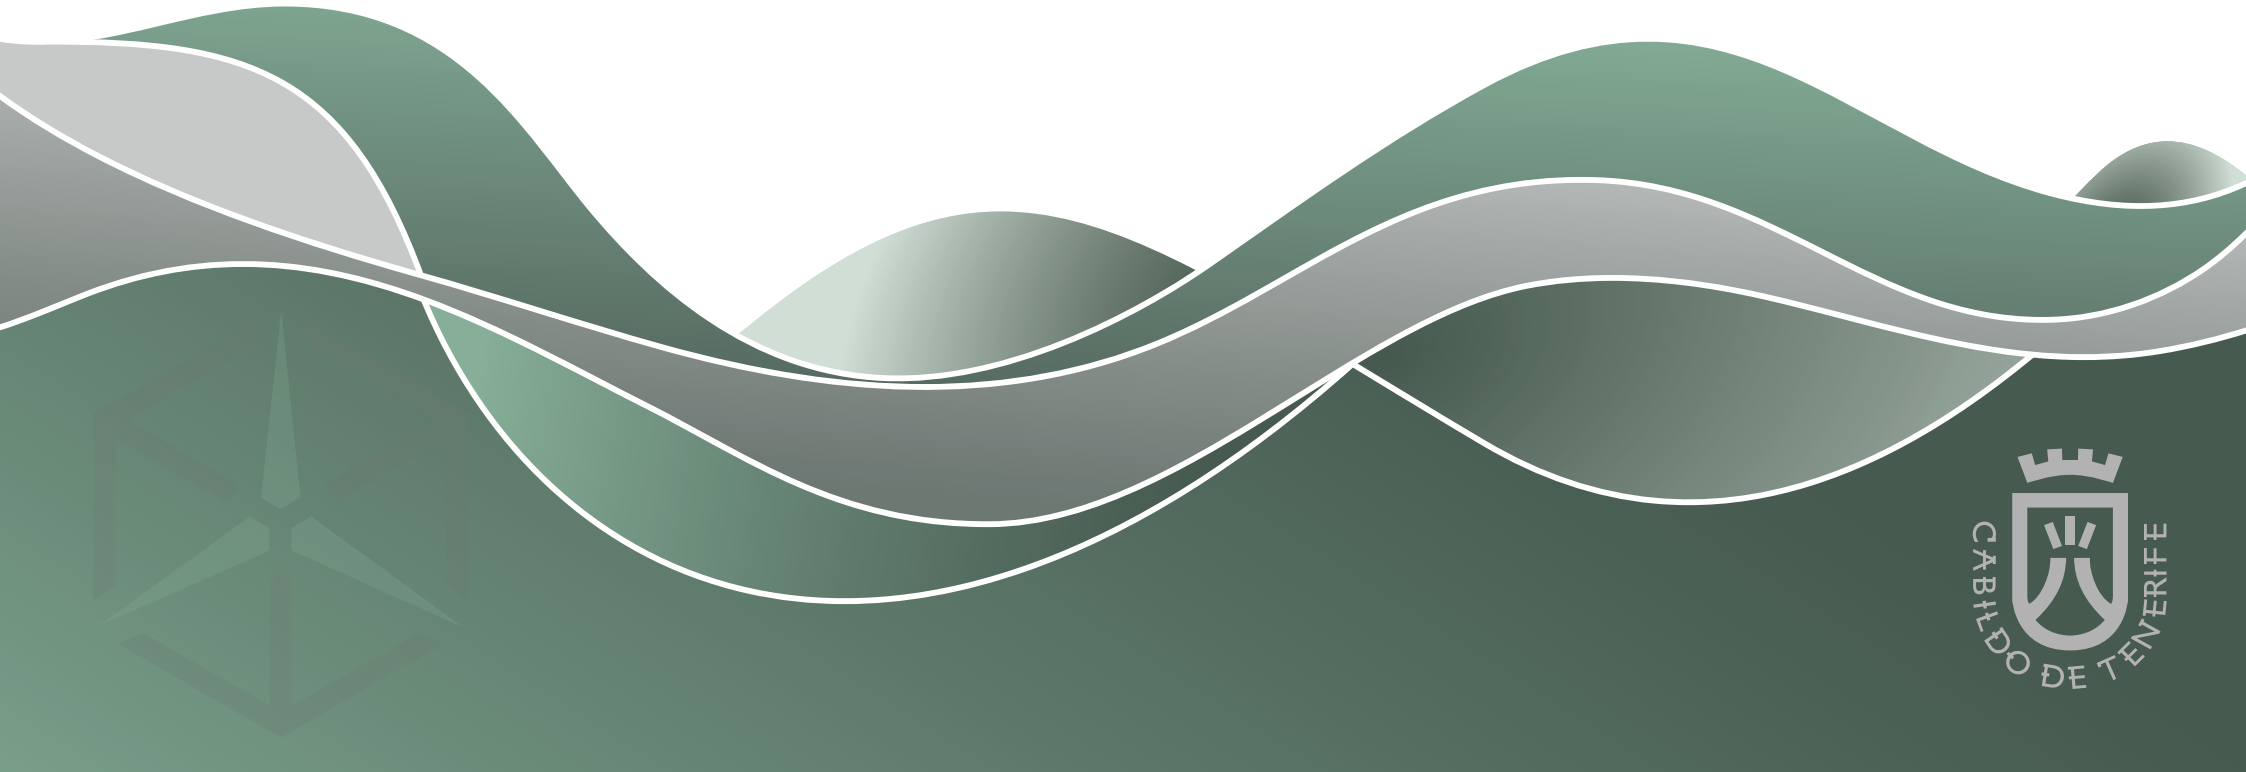

Infrastructure

At the ITER Supercomputing Center, we offer a robust and advanced infrastructure that integrates computing, storage, and state-of-the-art communication capabilities.

Designed to meet the most demanding requirements of research and technological development, our facilities provide the performance, reliability, and flexibility necessary to drive projects of any scale.

With a focus on efficiency and security, our solutions are prepared to support and enhance the most innovative initiatives of public and private entities, facilitating an ideal environment for scientific and technological progress.

TeideHPC

General-purpose supercomputer based on CPU technology

Sandy Bridge Nodes

TeideHPC started operations in 2013 with 1028 Sandy Bridge nodes, each equipped with 2 Intel Xeon E5-2670 processors featuring 8 cores and 16 execution threads. Each node contains 32 GB of RAM, delivering efficient and reliable performance for intensive computing tasks.

Ivy Bridge Nodes

The supercomputer also includes 72 Ivy Bridge nodes, each with 2 Intel Xeon E5-2670v2 processors featuring 10 cores and 20 execution threads, along with 32 GB of RAM per node. These nodes provide enhanced processing capacity ideal for advanced and demanding applications.

TeideHPC

General-purpose supercomputer based on CPU technology

Fat Nodes

To meet the demands for greater memory and performance, TeideHPC has 3 Fat Nodes, each featuring 4 Intel Xeon E5-4620 processors with 8 cores and 16 execution threads, and an impressive 256 GB of RAM per node. These nodes are ideal for applications that require high memory and processing performance.

Performance

TeideHPC's total computing capacity is notable, with a theoretical maximum performance (Rpeak) of 366 TFlops and a real maximum performance (Rmax) of 274 TFlops. In November 2013, this capacity placed TeideHPC at position 138 on the Top500 ranking of the most powerful supercomputers in the world.

AnagaGPU

General-purpose supercomputer equipped with GPUs | optimized for AI applications

4 GPU Nodes

AnagaGPU consists of 16 high-performance nodes, each equipped with 4 nVIDIA A100 GPUs, providing 40 GB of memory per GPU. Each node features 256 GB of RAM, offering massive and efficient processing capacity for deep learning tasks, complex simulations, and large-scale data analysis.

8 GPU Node

For applications requiring even greater processing power, AnagaGPU includes an exclusive node with 8 nVIDIA A100 GPUs, each with 40 GB of memory, and 512 GB of RAM. This unique node is designed to support extremely demanding workloads such as large-scale AI models and intensive data research projects.

AnagaGPU

General-purpose supercomputer equipped with GPUs | optimized for AI applications

Visualization Nodes

AnagaGPU also includes 4 visualization nodes, each with an nVIDIA T4 GPU and 256 GB of RAM. These nodes are optimized for advanced visualization tasks, graphical processing, and support for virtual and augmented reality environments, providing a versatile and powerful platform for various visual applications.

Performance

With an Infiniband EDR-based network infrastructure, AnagaGPU ensures high-speed, low-latency communication between nodes, optimizing the system's overall performance. The total computing capacity of AnagaGPU reaches a theoretical maximum performance (Rpeak) of 1.25 PFLOPS and a real maximum performance (Rmax) of 681.90 TFLOPS.

Common Infrastructure

Storage

The Supercomputing Center offers a cutting-edge shared storage infrastructure with a unified block and file system that provides 2.2 PB of net capacity.

This robust system guarantees fast and efficient access to large data volumes, optimizing performance for advanced research projects and critical business applications while facilitating collaboration and continuous information flow in a secure and highly available environment.

Communications

The center features a high-capacity communication infrastructure, including external connectivity via RedIRIS for R&D with a 10 Gbps link, as well as a commercial link of 5 Gbps.

This advanced configuration ensures fast and reliable data transfer, fostering collaboration in research projects and access to external resources while supporting daily commercial operations efficiently and swiftly.

Services

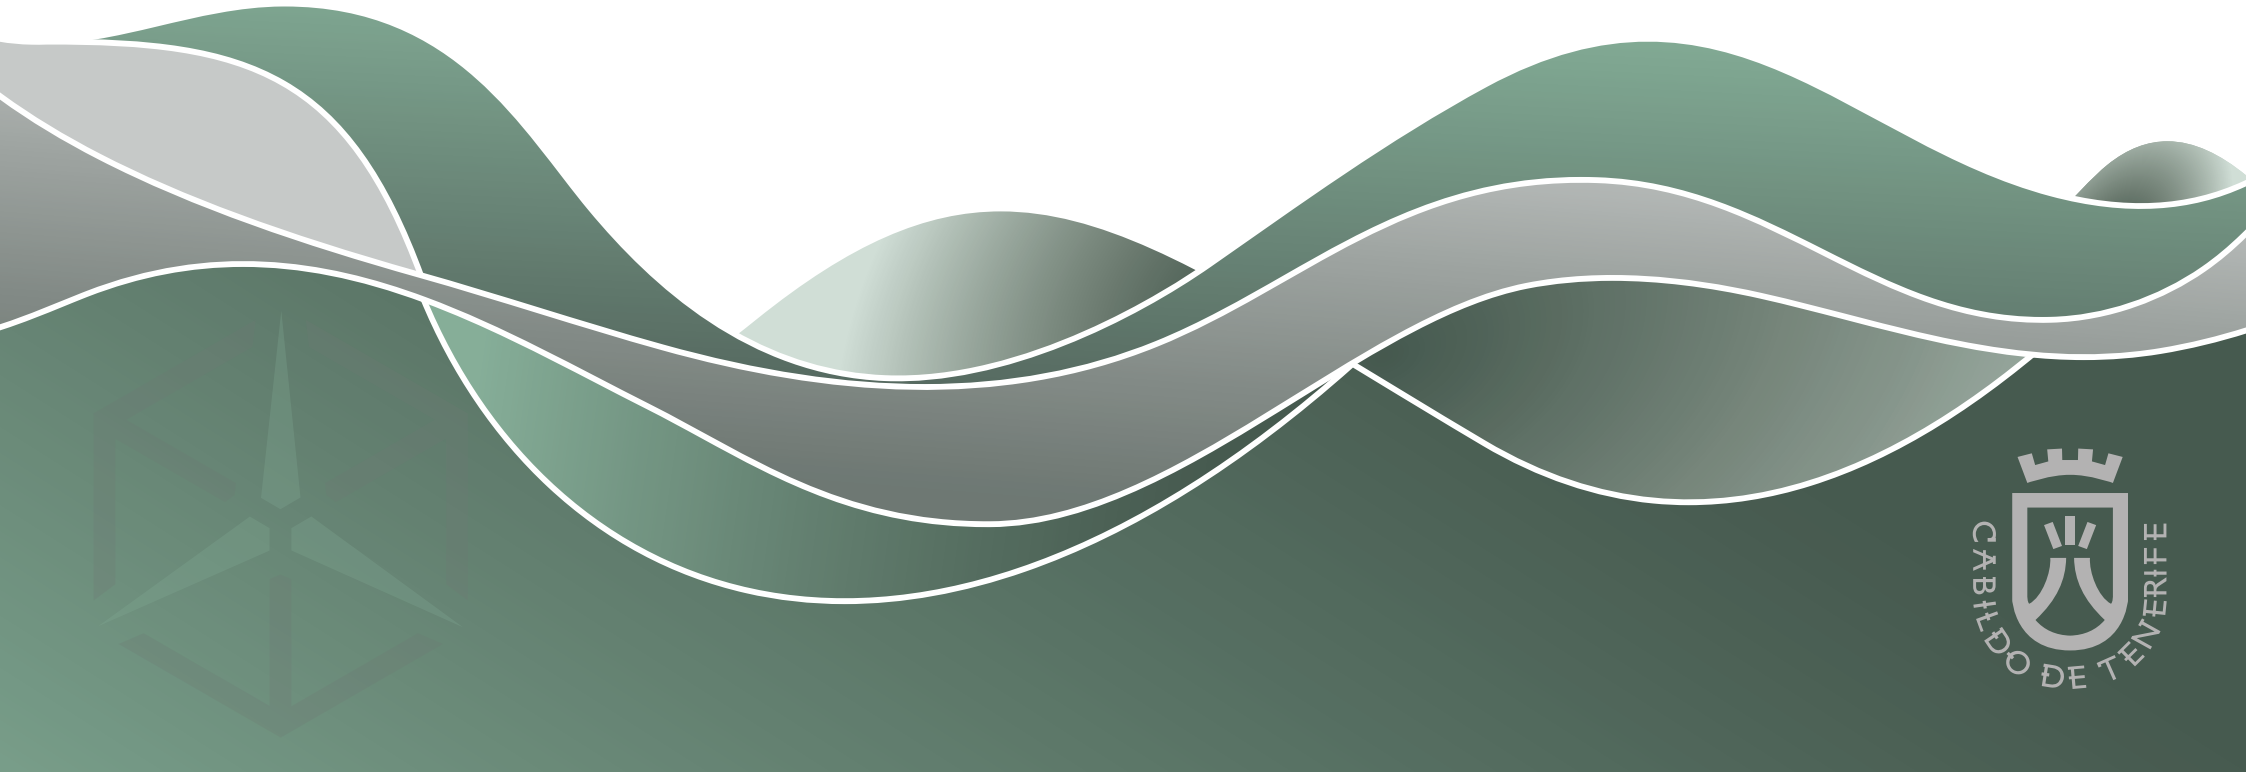

Services

The ITER Supercomputing Center offers two main services designed to meet the most demanding research and development needs: High-Performance Computing (HPC) and Infrastructure as a Service (IaaS).

Our high-performance computing service provides access to an advanced supercomputing infrastructure capable of handling complex tasks and intensive data processing with superior computing capacity. Equipped with state-of-the-art nodes and high-speed interconnection technology, this service is optimized for scientific applications, detailed simulations, and large-scale data analysis, ensuring fast and accurate results.

Additionally, our Infrastructure as a Service (IaaS) offering allows organizations to access scalable computing and storage resources on demand, without the need to directly manage hardware. This flexible solution provides a virtualized environment that adapts to the specific needs of each project, offering computing, storage, and networking resources efficiently and cost-effectively. With the Supercomputing Center's Infrastructure as a Service, entities can focus on their research and innovations while benefiting from a robust and professionally managed platform.

High-Performance Computing (HPC)

Our high-performance computing service is designed to offer both performance and flexibility. It allows job execution through an advanced task scheduler, optimizing resource usage and ensuring maximum efficiency in data processing.

Pre-installed software covering a wide range of needs enables clients to start working immediately without worrying about additional installations. Furthermore, we offer the possibility of using custom developments, allowing users to personalize and adapt their working environments to meet specific requirements. This combination of tools and flexibility ensures that our clients can tackle the most complex challenges with ease and agility.

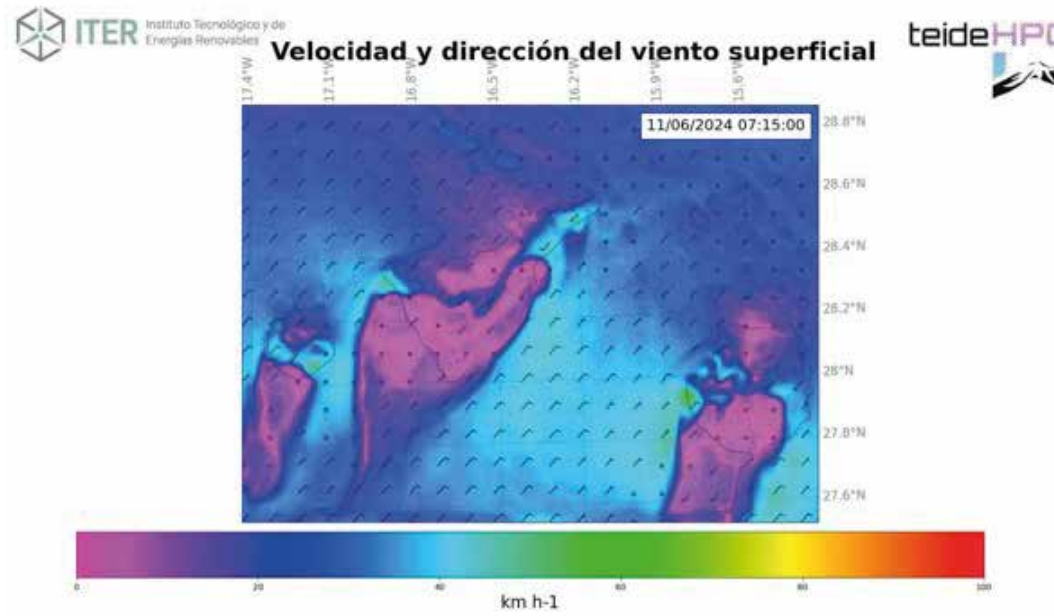

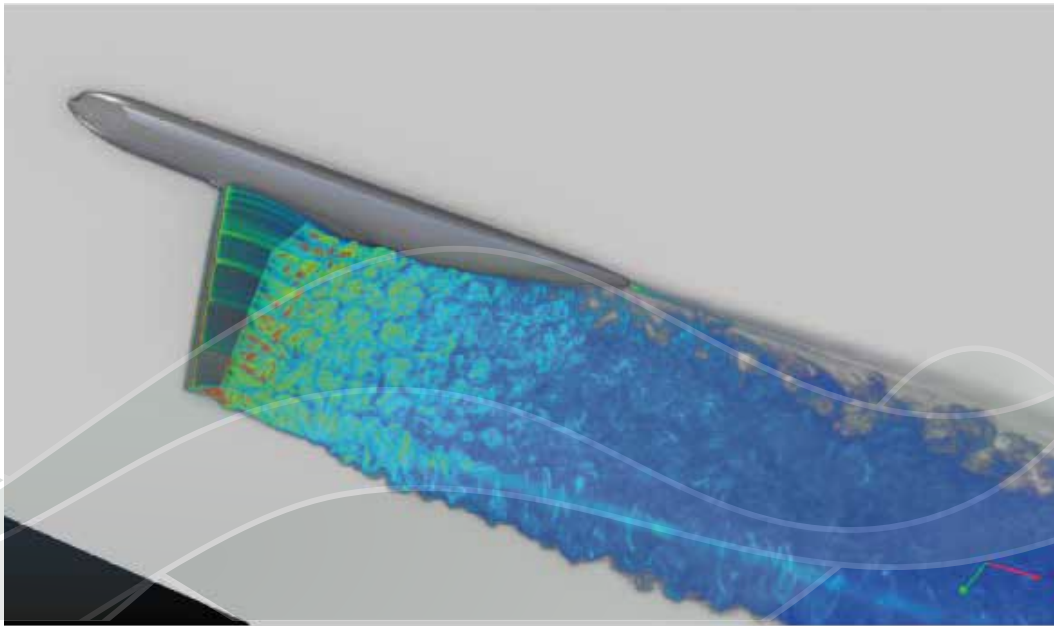

Infrastructure as a Service (IaaS)

Our infrastructure as a service (IaaS) offering provides an integral solution for scalable and flexible computing needs. Through infrastructure virtualization, we provide virtual machines with different capacities tailored to the specific requirements of each project.

Each user can benefit from private virtual networks, ensuring a secure and isolated environment for their operations. Additionally, we offer customized architecture design, enabling clients to create configurations that optimize performance and efficiency. External communication configurations are also integrated, facilitating seamless connectivity and data exchange.

This comprehensive approach ensures that our solutions perfectly adapt to the demands of each user, providing a robust and customized platform for the success of their initiatives.

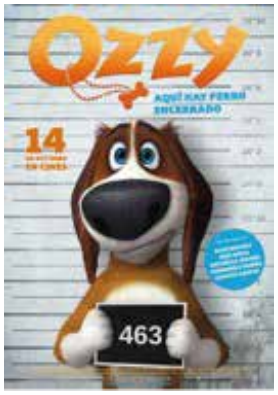

Commitment

Our supercomputing services are designed to provide maximum support with Next Business Day (NBD) service, ensuring fast and efficient attention to all your needs.

Regarding communications, we offer access to both R&D and commercial networks at no additional cost, providing smooth, high-quality connections tailored to your projects.

For those preferring customized options, we provide the possibility of contracting independent communication services with operators hosted in our data center, ensuring full flexibility and control over your connections.

Our objective is to deliver first-class service quality, ensuring that all operations run smoothly and reliably, supported by our top-tier infrastructure and a dedicated support team.

Pricing

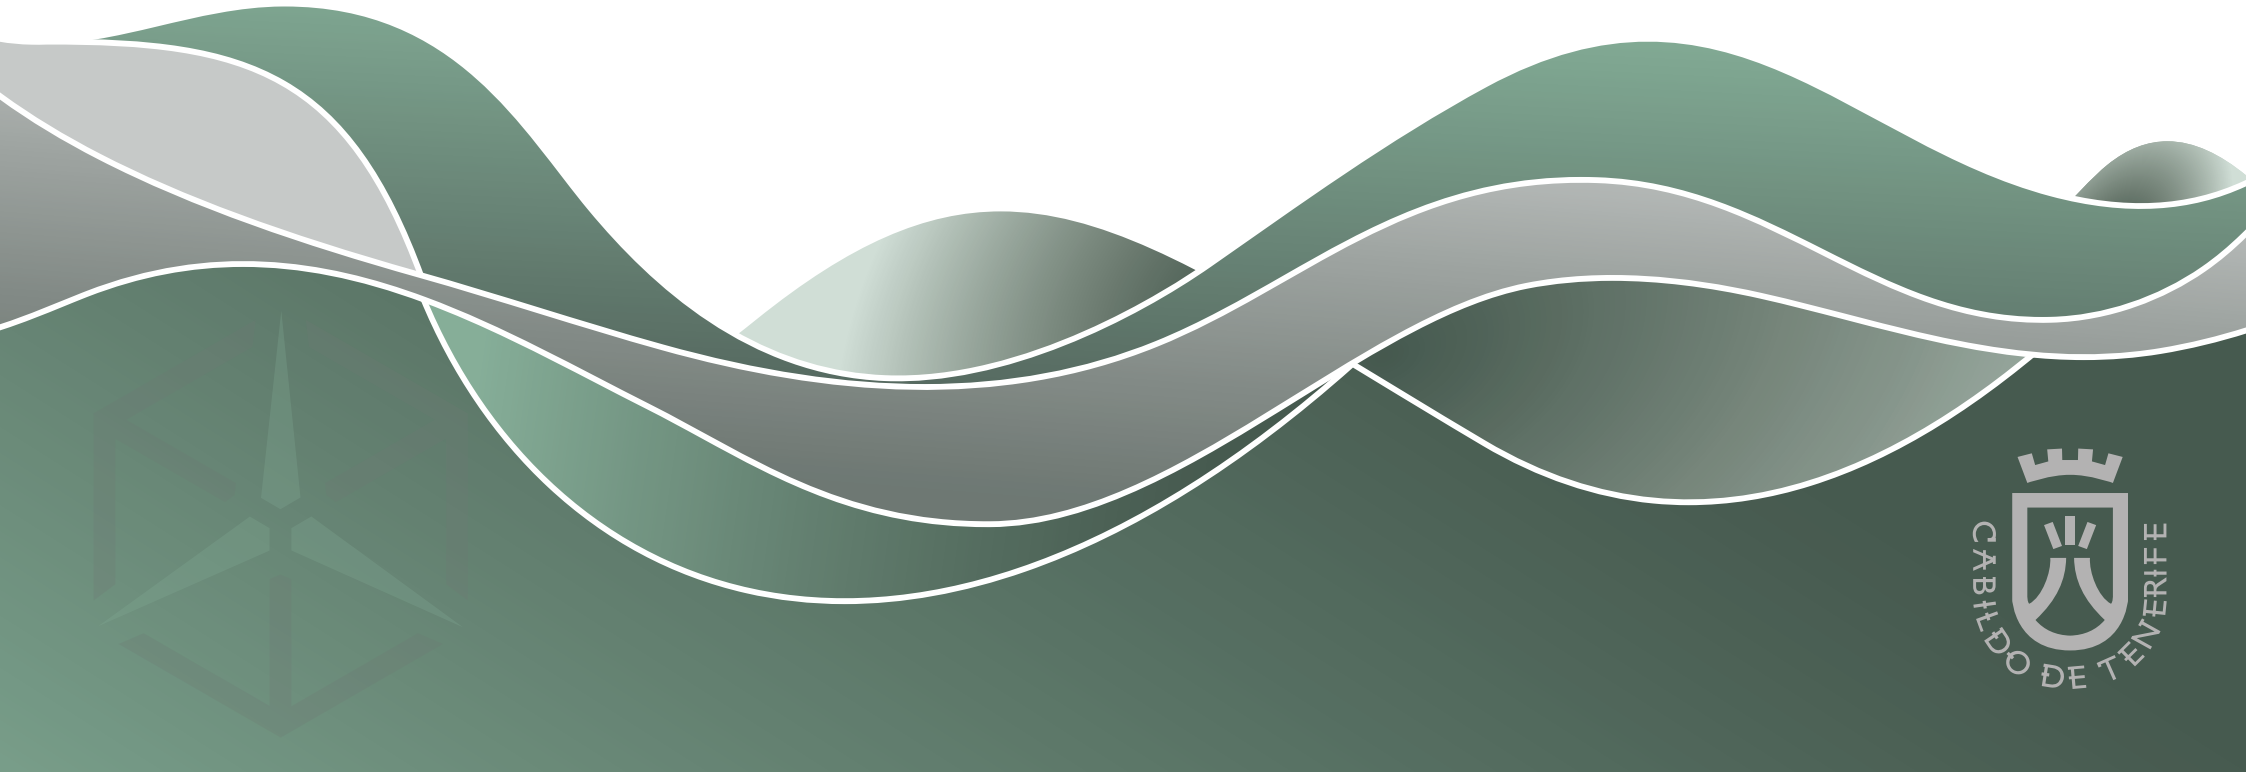

Pricing

The following are the rates for the various services offered by the Supercomputing Center.

For the high-performance computing service, pricing is detailed based on resource usage, providing a clear and transparent view of the associated costs.

Regarding Infrastructure as a Service (IaaS), we offer a variety of generic configurations as a reference, designed to cover a wide range of common needs. However, we understand that each client has specific requirements, so our configurations are fully customizable. We are available to tailor our offerings to individual needs, ensuring that clients receive a perfectly tailored solution for their projects and objectives.

High-Performance Computing

Service	**RPR**
Core-hour CPU	0,03€
GPU-hour	0,50€

Storage	**RPR**
Storage TB/month	20,00€

Infrastructure as a Service (IaaS)

CPU

Configuration	Monthly Price
VM 4vCPU 8GB RAM	50,00€
VM 8vCPU 16GB RAM	100,00€
VM 16vCPU 32GB RAM	200,00€
VM 32vCPU 64GB RAM	400,00€

GPU

Configuration	Monthly Price
VM 16vCPU 64GB RAM 1 nVIDIA A100 40GB	1800,00€
VM 32vCPU 128GB RAM 2 nVIDIA A100 40GB	3600,00€
VM 48vCPU 192GB RAM 3 nVIDIA A100 40GB	5400,00€
VM 64vCPU 256GB RAM 4 nVIDIA A100 40GB	7200,00€

Other services

Service	RPR
Initial Configuration Services	220,00 €
Hourly Support	22,00 €/hour

Supercomputing Center

For more information: teidehpc@iter.es