LEONARDO Service Description

Service name	LEONARDO Supercomputer
Service summary	LEONARDO Supercomputer service allows customers to run parallel scientific workflows through a batch queueing system. LEONARDO comprises of a large number of powerful GPU nodes and CPU nodes with varying amounts of memory.
Detailed description	The acquisition and operation of the EuroHPC LEONARDO supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as by the Participating LEONARDO Consortium States. LEONARDO is an Atos BullSequana supercomputing platform based on 13824 GPUS and 282624 CPUs. The interconnect is based on Nvidia Mellanox HDR DragonFly+ 200 Gb/s. The detailed configuration is ● LEONARDO-booster: 3456 GPU nodes with 4 NVIDIA Ampere GPUs A100 and 1 Intel Ice Lake CPU with 32 cores (2,6 GHz), with a theoretical node peak performance of 89,4 Teraflops. ● LEONARDO-Data Centric: 1536 2 x CPU nodes with 56-core Intel Sapphire Rapids processors (2.0 GHz) each. ● Computing nodes have the memory sizes: ○ 512 GB on 3456 nodes (Booster) ○ 512 GB on 1536 nodes (Data centric) ● Interconnect network Nvidia Mellanox with Dragon Fly+, nodes connected with - 2xNvidia HDR cards 2x100Gb/s (Booster) - 1xNvidia HDR cards 1x100Gb/s (Data Centric) ● LEONARDO Storage system features a capacity and fast tier: - Fast tier is based on 31 DDN Exascaler ES400NVX2 with 24x 7,68 TB SSD NVMe with encryption support and 4x InfiniBand HDR ports each. - Capacity tier is based on **31 x DDN EXAScaler SFA7990X appliances for HDD storage** each with 82 x 18 TB HDD SAS 7200 rpm and 4 x HDR100 ports and 2 x JBOD expansion per controller.

	Jobs are submitted through SLURM job scheduling system. LEONARDO can be accessed through Unix shell and X forwarding. Users can ● develop their own codes (Fortran, C/C++, python, …), ● install Linux compatible applications via compiling them or by running (singularity) containers ● utilize LEONARDO's software collection ● submit jobs and run simulations. For code parallelization MPI and OpenMP can be used. OpenMP offload, OpenACC and CUDA can be used with GPUs. Additionally, mathematical libraries are available. The LEONARDO computing environment includes tools for debugging and performance analysis. More details about the service are available at https://leonardo-supercomputer.cineca.eu/
Target audience	Academic, public, and private sector.
How to obtain the Service	To access and use the service the customer must have a CINECA user account, and a project granted on LEONARDO, which can be applied for either on an academic or commercial basis. See: How to access Leonardo resources page
Service level & availability	The LEONARDO service SLA is specified in *LEONARDO Hosting Agreement Annex III - Key performance indicators.* LEONARDO targets on 97% availability rate.
Service hours and user support	The LEONARDO service SLA is specified in *LEONARDO Hosting Agreement Annex III - Key performance indicators.* Support to the service is provided through LEONARDO UserSupport channels: ● HPC Service Desk via mail superc@cineca.it ● Documentation: Leonardo UserGuide
Pricing	LEONARDO resources accounting model: UserGuide-Accounting-Billingpolicy

Certifications	• LEONARDO User Support Services are included in the scope of ISO 9001 certification.
Data protection (GDPR) **Location of the LEONARDO service**	The service is not designed to process personal data. A user must not transfer personal data to the service. The location of the LEONARDO service production is in the European Economic Area. The LEONARDO User support is located in Italy. LEONARDO consortium members in the European Economic Area shall have the right to process the personal data of the users when providing the first level user support service to the users of the LEONARDO service. Furthermore, the supplier of LEONARDO system is situated in Europe and has access to servers based on a specific request when providing the advanced second level technical problem-solving support. See also CINECA Privacy Policy document further information.
Customer's responsibilities	Customers of LEONARDO are responsible for their data and computing. • Backups • Information security • Installations (if not utilizing LEONARDO software) • Capacity requests

Service provider's responsibilities	As a service provider CINECA is responsible for: ● Producing, developing, and maintaining the LEONARDO Supercomputer service.
Adjacent services	
Additional services	
Service producer	CINECA – Consortium of Universities and public institutions on behalf of LEONARDO Consortium