

UG3.0: System Specific Guides

(updated: July 2023)

In this Section, the specific features of the HPC systems available in CINECA for scientific users are described.

On these systems, you can obtain a username (access) and an account (project budget) for running programs from the Software Catalog or your own programs. Several tools for programmers are available like compilers, profiling tools and debuggers.

These systems can be reached mainly through **ssh**. All the general topics related to those systems are discussed in the Sect. "[General Information](#)".

At present the available systems in CINECA are:

- **MARCONI**: This is a Lenovo NeXtScale platform co-designed by CINECA, based on the next-generation of the Intel® Xeon Phi™ product family alongside with Intel® Xeon® processor product family (Broadwell and Skylake). It offers the scientific community a technologically advanced and energy-efficient high performance computing system. This system has been gradually completed in about 12 months in a three-step update. At present, A2 and A3 phases are available for production.
 - A1 consists of a system into production since July 2016, featuring the Broadwell processors and with the computational power of 2Pflop/s. **Starting from September 26th, 2018 the activity on Marconi-A1 has been stopped.**
 - A2 is in production since January 2017 and features Intel Knights Landing (KNL) processors and a peak performance of 11Pflop/s. **Starting from January 2020 the activity on Marconi-A2 has been stopped.**
 - A3 is in production since Aug 2017 (the first 1512 nodes), then increased to 3188 by January 2018, and features Intel Xeon Skylake (SKL) processors. With the addition of A3 partition, MARCONI reached a total peak performance of about 20 PFlop/s
- **GALILEO100**: New infrastructure is available from August 2021. This is a Cluster equipped with Intel CascadeLake 8260 nodes with 24 cores each, 2.4 GHz, 384GB RAM and Infiniband interconnection. It is also part of this cluster 77 servers for cloud computing. GALILEO100 is mainly used for special High-End projects for the technical and industrial HPC computing, as well as for dedicated projects in the field of meteo and environment. It has been expanded in November 2022 with 82 additional nodes (2 racks).
- **LEONARDO**: The new pre-exascale Tier-0 EuroHPC supercomputer hosted by CINECA and currently built in the Bologna Technopole, Italy. It is supplied by ATOS, based on a two new specifically designed compute blades, which are available through two distinct SLURM partitions on the cluster: **Leonardo Booster partition** with BullSequana XH2135 supercomputer nodes, each with four NVIDIA Tensor Core GPUs and a single Intel CPU; and **LEONARDO Data Centric General Purpose (DCGP) partition** with X2140 CPU-only blade based on Intel Sapphire Rapids processors. Leonardo also uses NVIDIA Mellanox HDR 200Gb/s InfiniBand connectivity, with smart in-network computing acceleration engines that enable extremely low latency and high data throughput to provide the highest AI and HPC application performance and scalability.
- **DGX A100** is the new system based on AMD Rome architecture and NVIDIA A100 GPUs. The DGX A100 is the universal system for all AI workloads, from analytic to training to inference. It is available for the Italian public research from January 2021.

Old HPC systems:

- **MARCONI 100**: This is the accelerated not conventional partition available from May 2020. Marconi100 is an IBM system equipped with NVIDIA Volta V100GPUs, opening the way to the accelerated pre-exascale Leonardo supercomputer. **Starting from July 2023 the activity on Marconi100 has been stopped.**
- **GALILEO** (renovated): This is a Cluster equipped with about 1000 Intel Broadwell nodes (2x18-core Intel Xeon ES-2697 at 2.30 GHz), and OmniPath interconnection. It is also part of this cluster 60 nodes equipped with K80 nVidia accelerators and 2 with V100 nVidia accelerators. This version of Galileo is available in Cineca since August 2019. **Starting from 15th March 2021 Galileo closed its production.**
- **D.A.V.I.D.E**: In production since early 2018, it is an energy aware, Peta-flops Class High Performance Cluster based on OpenPOWER servers and featuring liquid cooling and innovative technology for monitoring and capping the power consumption, developed by E4 Computer Engineering and installed at CINECA. **Starting from January, 2020 the activity on D.A.V.I.D.E. has been stopped.**
- **GALILEO**: This is a large IBM X86 Cluster equipped with 1024 six-core Intel Haswell processors, and Infiniband interconnection. Part of the nodes is equipped with Intel Phi accelerators. GALILEO is mainly used for massively parallel applications and special High-End projects for the technical and industrial HPC computing. Galileo is available in Cineca since January 2015. **Production closed Nov 20, 2017**
- **PICO**: This is an IBM Intel Cluster made of 74 nodes of different types, devoted to different purposes, with the common task of data analytics and visualization on a large amount of data. It is also the master of a large data repository (DRES) shared among all HPC systems in CINECA. **Production closed Jan 24, 2018**
- **FERMI**: IBM BlueGene/Q system, composed of 163.840 compute cores, with a peak performance of 2.1 PFlops. It has been ranked #7 in the TOP500 list (June 2012). **Dismissed in 2016. July 25th.**

- [UG3.1: MARCONI UserGuide](#)
- [UG3.2: LEONARDO UserGuide](#)
- [UG3.3: GALILEO100 UserGuide](#)
- [UG3.4: DGX A100 UserGuide](#)
- [UG3.5: ADA Cloud UserGuide](#)
- [UG3.7: Old HPC systems](#)

Outgoing links:

 Unknown macro: 'outgoing-links'