Italian National Agency for New Technologies, Energy and Sustainable Economic Development

CRESCO HPC

Energy Technologies Department – ICT Division
http://www.eneagrid.enea.it

CRESCO: ENEA, the Italian National Agency for New Technologies, Energy and Sustainable Economic Development operates in a range of R&D sectors, including energy technologies, new materials, life and environmental sciences. In support of the institutional research activities, the ENEA IT division provides computing and storage resources. The computing core of ENEAGRID is represented by the HPC CRESCO clusters.

CRESCO HPC clusters

Cluster name (Linux x86_64)	Network	Cores/Tflops
CRESCO6 INTEL/SKL	OPA	20832/~1400
CRESCO7 INTEL/SKL	IB EDR	3072/~200
CRESCO4F INTEL/SNB	IB QDR	1024/~20
CRESCO4C INTEL/SNB	IB QDR	512/~10
CRESCO5F AMD/EPYC+NVIDIA/A100	IB EDR	1280/~100
XCRESCO IBM/PPC+NVIDIA/V100	IB EDR	21292/~1500
Total		**~48000/~3230**

CRESCO Storage

Storage Systems CRESCO	Bandwidth GB/s	network	Capacity raw/net PB
DDN EXAScaler 400NVX	50	IB/HDR	1 PB
2 x DDN SFA7790	20	IB EDR	4 /3.5
2 x DDN SFA7700	15	IB QDR	1 / 0.75
7 x RAID Servers	1.25	10 GEth	3.2 / 2.9
4 x RAID Servers	12.5	100 GEth	0.96 / 0.72
Total			~10 / ~9
Tape Library IBM 3500	LTO4/6	FC/8Gb	4.3 PB

High performance storage systems based on DDN Infiniband QDR/FDR/EDR and GPFS Spectrum Scale as parallel filesystem.

The CRESCO HPC clusters are in operations in the 3 ENEA Research Centre: Portici, Frascati and Casaccia

HPC CRESCO Usage in 2022: In the year 2022, the usage of HPC CRESCO clusters totaled 140 million core-hours, marking an interesting shift from the trend observed in the past five years. From 2017 to 2020, we witnessed a consistent annual increase in core-hours utilization, reflecting the growing demand for computational resources in various fields.

However, from 2017 to 2020, there was a notable acceleration in the usage of HPC CRESCO clusters, with core-hours increasing significantly year after year. This period saw a surge in complex simulations, data analysis, and research projects that harnessed the capabilities of HPC clusters.

Contrastingly, in 2021, we experienced a slight reduction in the usage compared to the peak observed in 2020. This decrease, while modest, indicates a possible shift in research priorities or a stabilization of computational needs in certain areas. It's essential to note that even with this reduction, our HPC CRESCO clusters remains a critical resource for a wide range of scientific and computational endeavors.

As we analyze this trend, we will continue to adapt our resources and services to meet the evolving needs of our users. Our commitment to providing world-class computing capabilities remains unwavering, and we look forward to supporting your research and innovation in the coming years.

Looking ahead, we are committed to ensuring that our supercomputing resources align with the evolving demands of the scientific and computational community. To achieve this goal, in 2023, we have already taken significant steps by deploying two additional new clusters:

- CRESCO7 based on conventional Intel Skylake CPUs.
- XCRESCO based on accelerated GPUs NVIDIA V100 in a Power9 cpu compute node

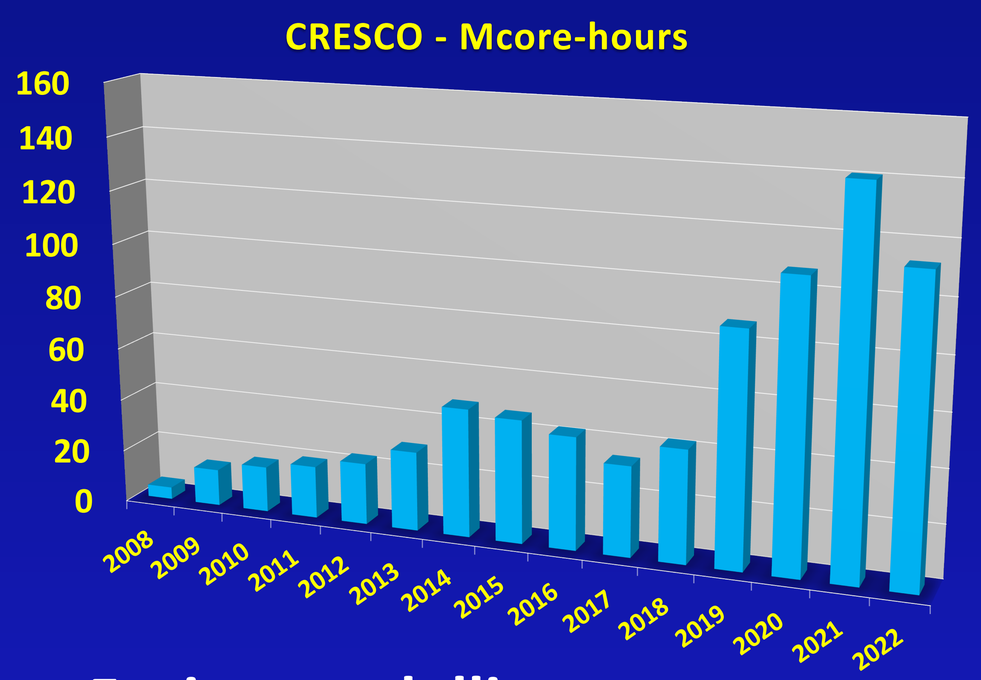

CRESCO8: ENEA HPC TIER-0 to support Nuclear Fusion modelling (operations: Q1-2024)

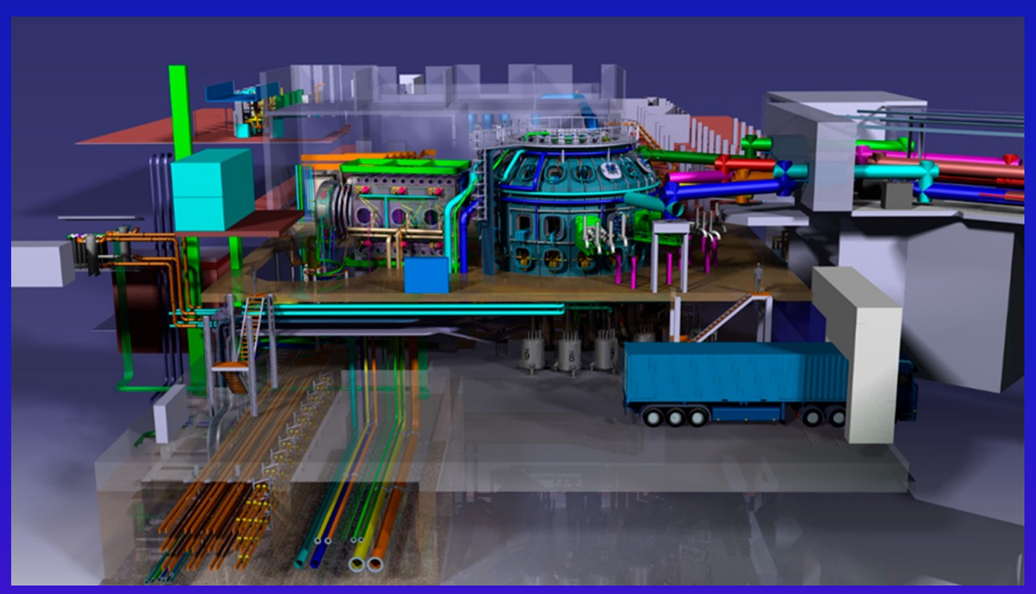

DTT, Divertor Tokamak Test facility, is a fusion experiment under construction at the Frascati ENEA Research Center. DTT plays a leading role in fusion research, the main aim being to explore alternative solutions for the extraction of the heat generated by the fusion process.

To make DTT exploitation more attractive to the international scientific community, an upgrade of the baseline design consisting in a supercomputer for nuclear fusion modelling is being to deploy.
DTTU was included among the high-priority facilities in the NextGenerationEU 'Piano Nazionale Infrastrutture di Ricerca' of the Italian Ministry of Research

HPC CRESCO8 Technical Specifications

	Specification	Technical requirement
Conventional partition	Technology	X86_64
	Peak performance	Minimum 6-7 PFlops/s
	Memory/node	Minimum 192 GB
	Memory/core	Minimum 4 GB/core
	Interconnection network	Low latency / High bandwidth (state of the art)
Accelerated partition	Technology	NVIDIA/AMD
	Peak performance	Maximum 3-2 PFlops/s
	Memory per node	At least 64 GB multiplied by the number of GPU per node
	Memory per GPU	At least 32 GB (Internal GPU memory)
	Interconnection network	Low latency / High bandwidth (state of the art)
Fast storage (scratch, work)	Technology	High Performance Parallel Filesystem targeting high availability (GPFS, LUSTRE or BeeGFS)
	Capacity (raw)	At least 10 PB
	Bandwidth	> 50 GB/s
Service/login nodes		management nodes login nodes for terminal sessions Advanced Graphic cards and NX session