

# CINECA

## AN INTRODUCTION TO THE CINECA (AND ICSC) SUPERCOMPUTING INFRASTRUCTURE:

From Leonardo and beyond



Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing

Massimiliano Guarrasi: [m.guarrasi@ Cineca.it](mailto:m.guarrasi@ Cineca.it)

# NOT-FOR-PROFIT CONSORTIUM

CINECA

SINCE 1969 CINECA SUPPORTS THE ITALIAN ACADEMIC SYSTEM



## 112 MEMBERS

2 Ministries, 69 Universities, 5 University Hospitals  
22 Academic and Research Institutions



## 5 OFFICES

Bologna, Milan, Rome, Naples, Chieti



≥1200

Employees



≥ 140 MLN €


Yearly Revenue



# 50+ YEARS OF IT SERVICES

1969

## SUPERCOMPUTING



4 founding Universities  
Bologna, Florence,  
Padua, Venice

IT Systems for the  
Italian Ministry of  
Universities and  
Research



## MINISTRIES

'80

'90

## UNIVERSITIES



IT Systems for the  
Italian Academic  
System

Technological Transfer  
to Healthcare  
Public Administration  
Industry



## PA & INDUSTRY

2000

2020

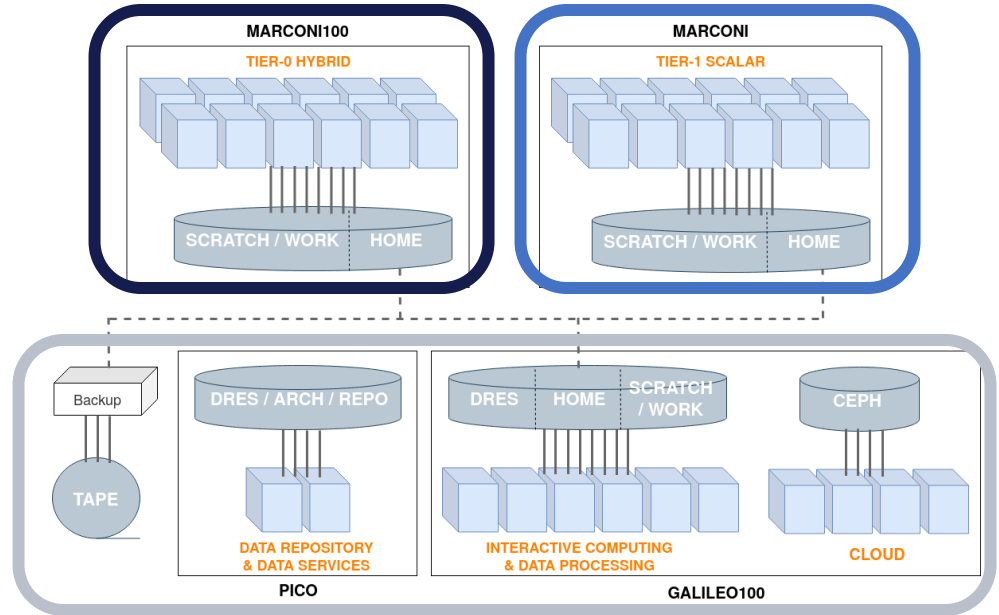

## HPC PRE-EXASCALE



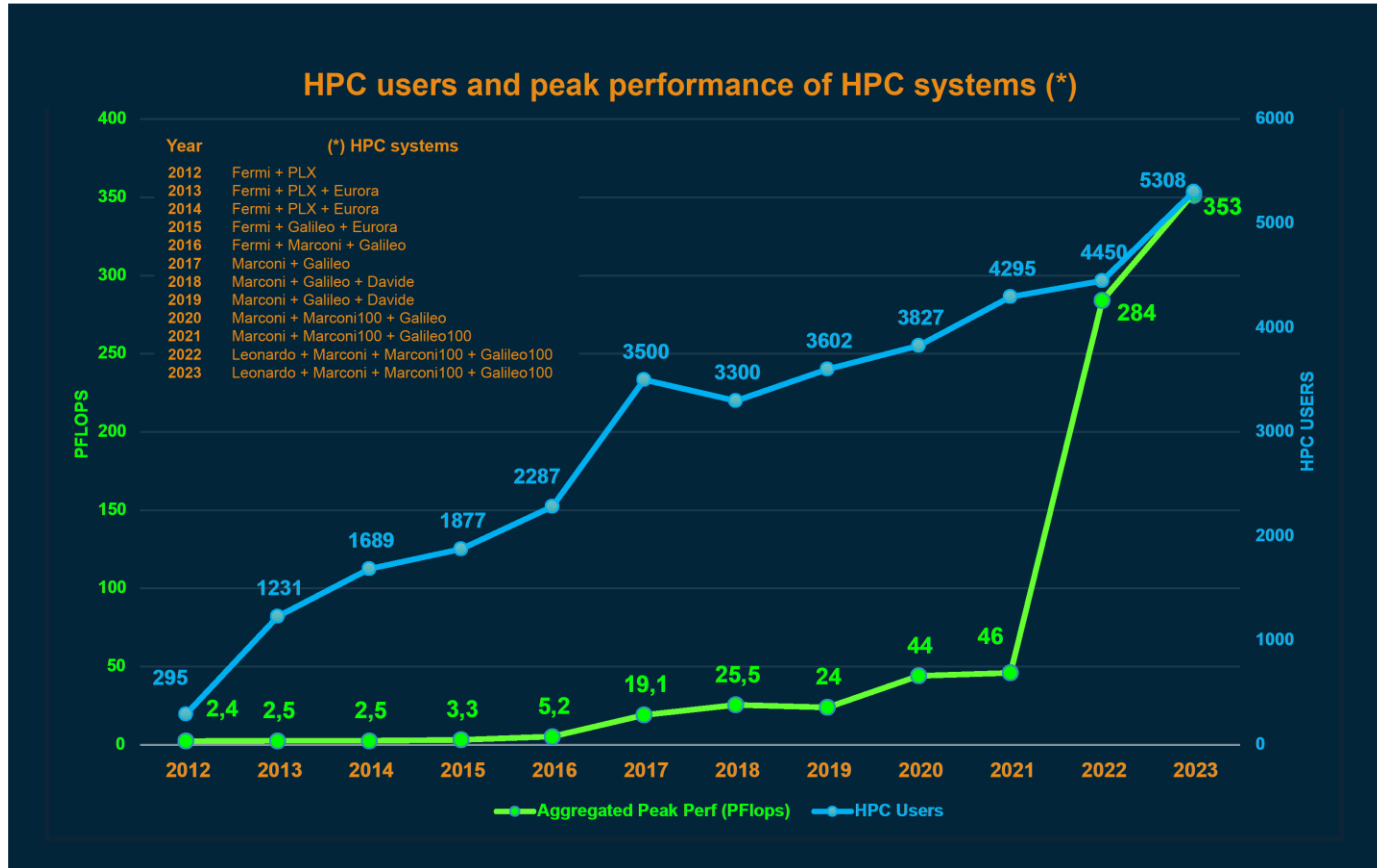
Artificial Intelligence  
Big Data  
Quantum Computing



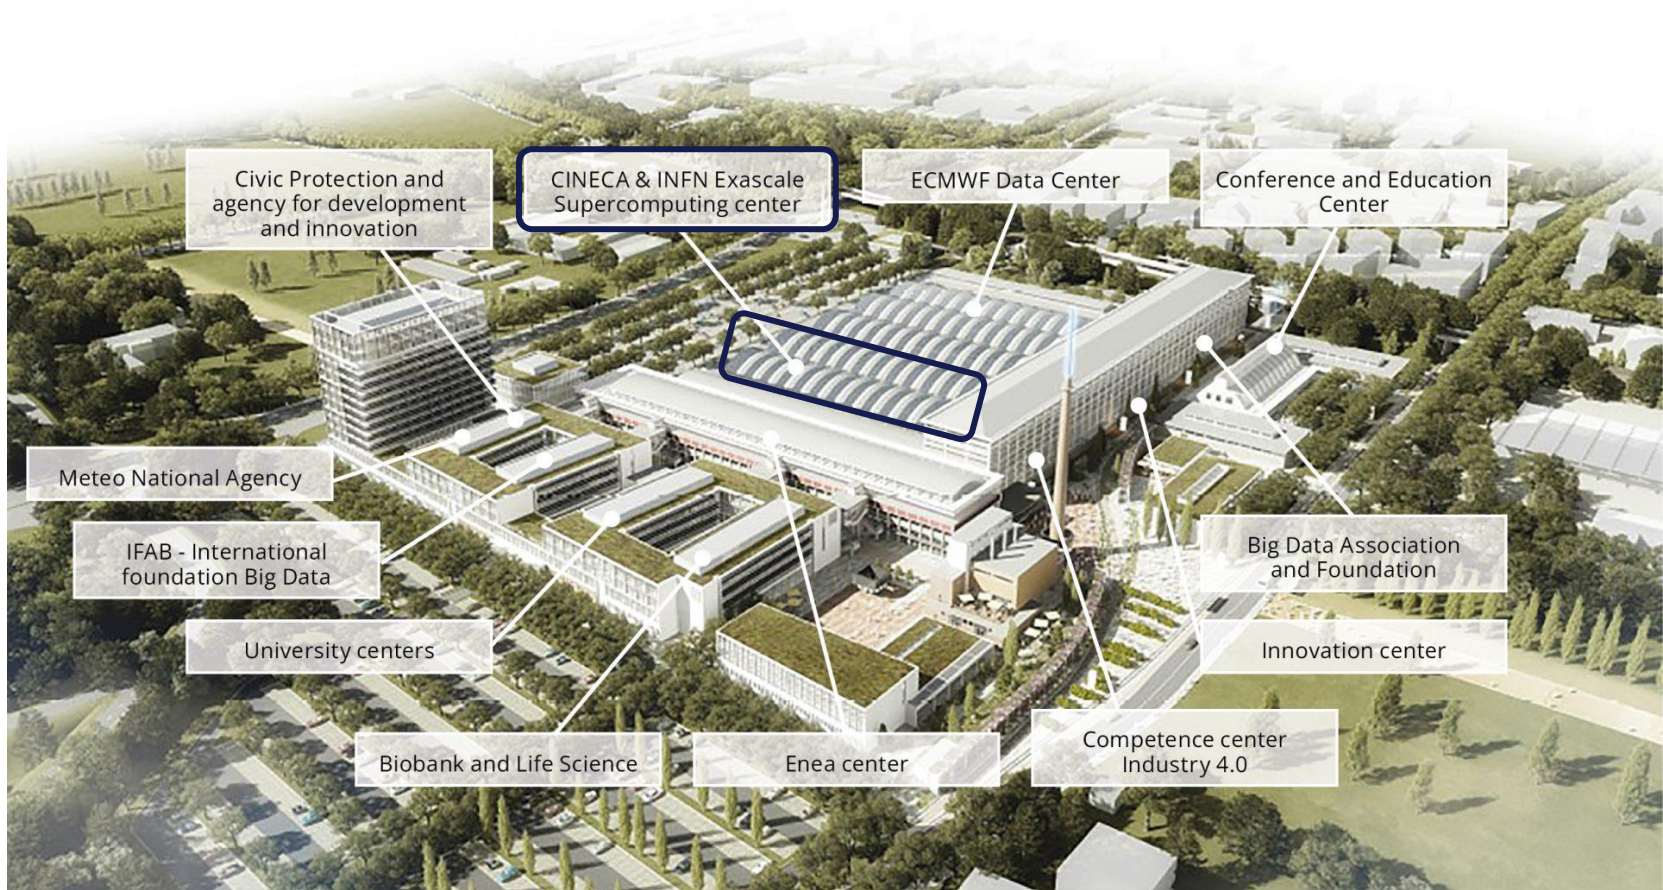
# CINECA HPC Infrastructure



# Evolution of the infrastructure

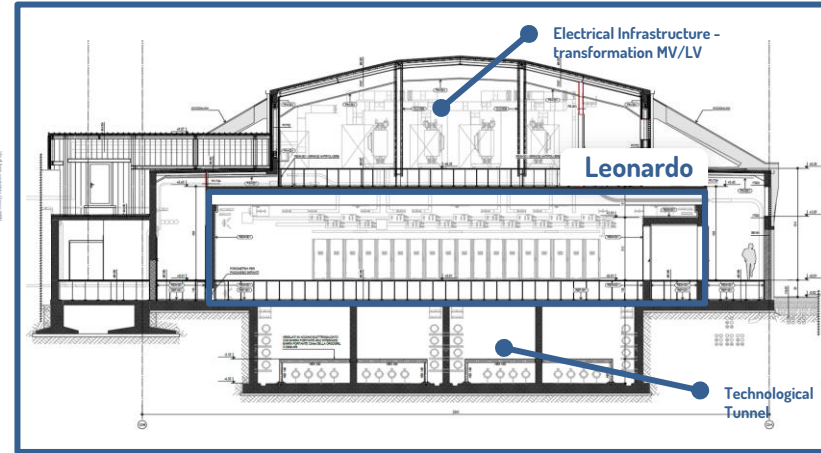


# Big Data Technopole – Bologna Science Park





# CINECA/INFN Data Center - Big Data Technopole



## Features

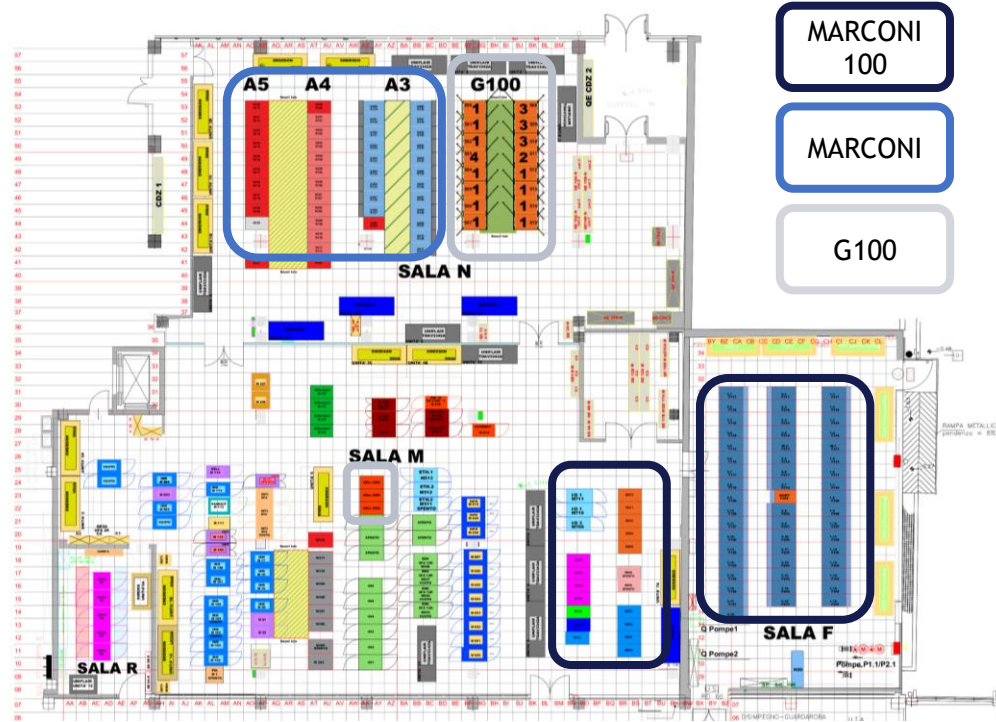
- **12 MW** and **1240 sqm** Rack Room
- **PUE < 1,10** (year based measurement strategy compliant to Level 3 Green Grid/ASHRAE)
- **Redundancy** Configuration: **3+1**, Electrical and Mechanical
- **Certify Rating 4 - TIA942 and Tier IV**

- Upgrade for **additional 12 MW** and **additional 2600 sqm** Rack Room for the Leonardo Next
- Mechanical and Electrical infrastructure able to comply with **2 different expansion strategies**:
  - **Stage 2a: Liquid Cooling** Expansion (20 MW Liquid Cooled + 4 MW Air Cooled)
  - **Stage 2b: Air Cooling** Expansion (8 MW Liquid Cooled + 16 MW Air Cooled)

# CINECA Data Center – Casalecchio di Reno (BO)

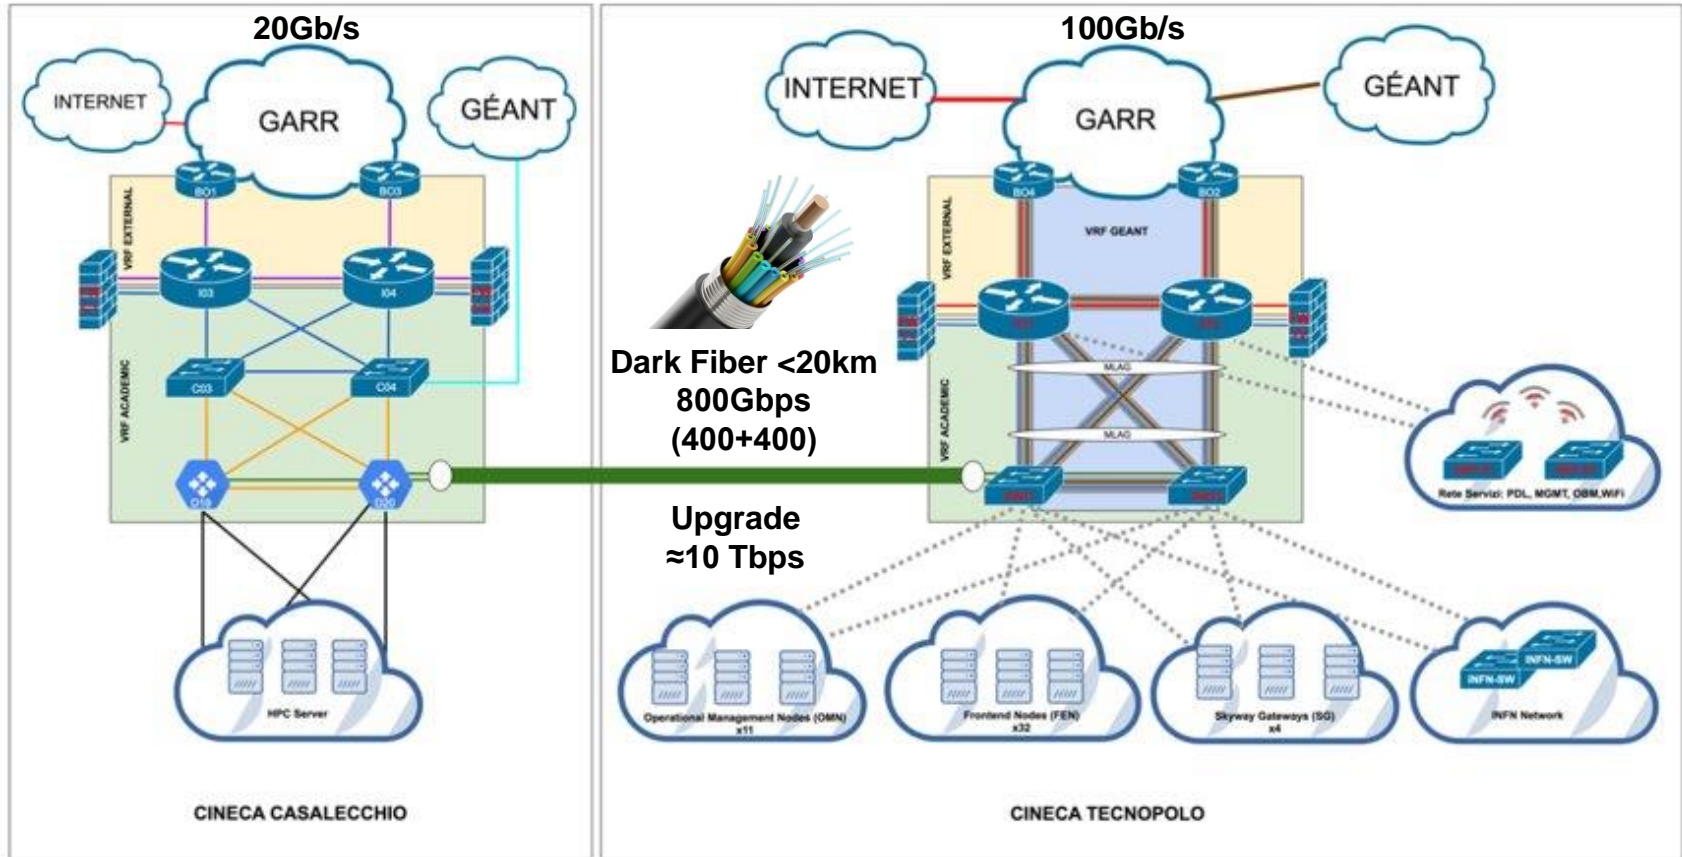
## Features

- **5 MW** and **920 sqm** Rack Room (Sala F + Sala M/N)
- Upgrade to **6 MW** in the next year
- **PUE  $\approx 1,4$**  (year based measurement)
- **Certify Rating 3 - TIA942 and Tier III**
- Mix HPC and ICT systems






# Casalecchio di Reno <-> Technopole - Interconnection



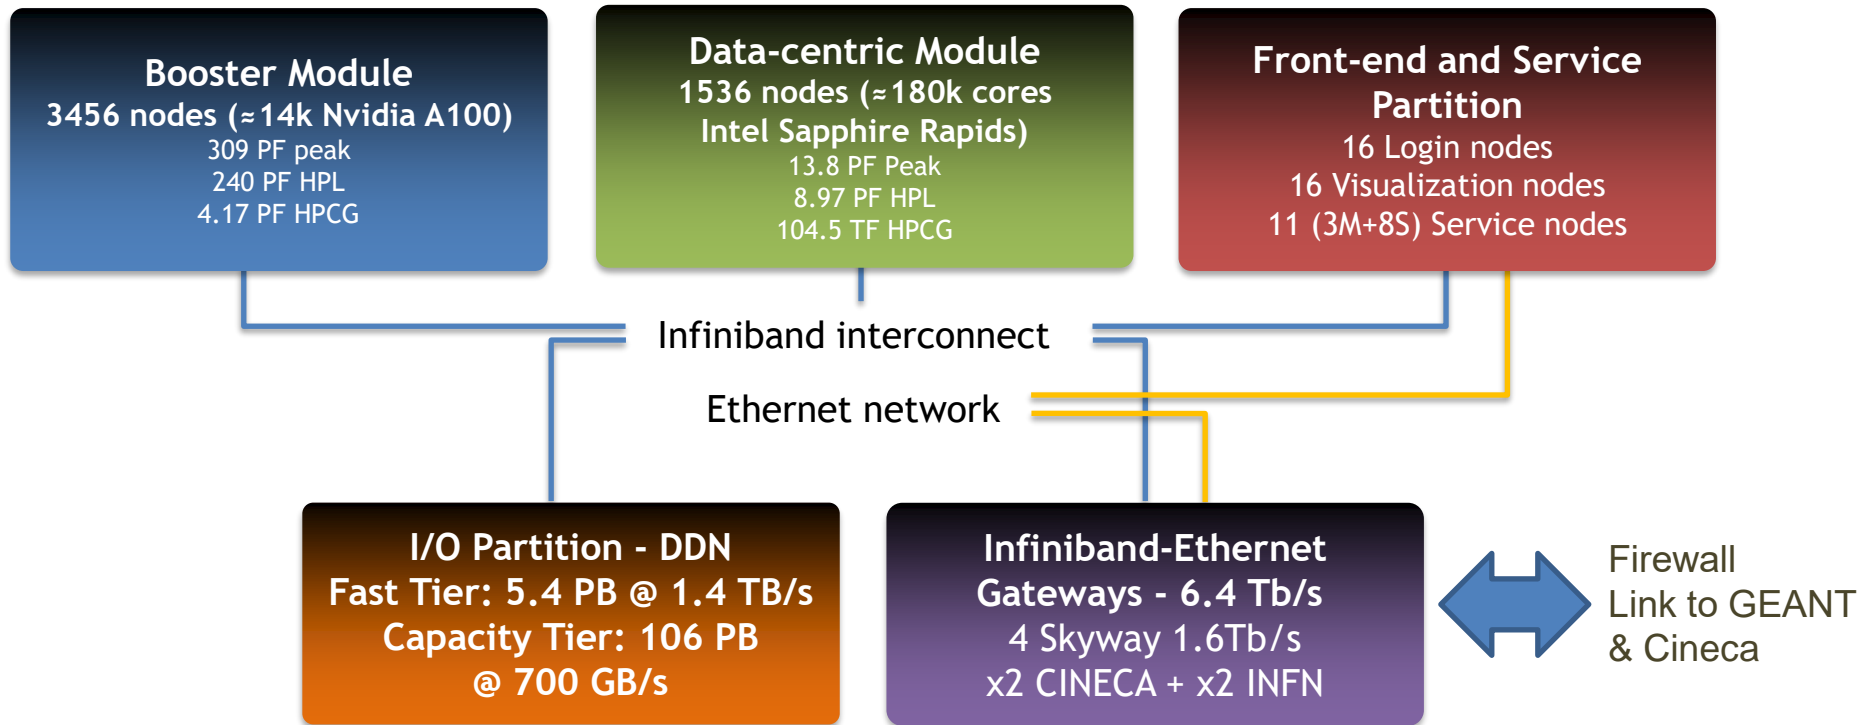
# Leonardo System

---

- **4<sup>th</sup> Top500 (7<sup>th</sup> now)**
- **HPL 240 PF + 9 PF (currently 170PF)**
- **TCO Investment: 240M€**  
(120M€ Capex + 120M€ Opex)
- **5000 nodes based on BullSequana XH2000 platform technology**  
(3500 GPU + 1500 CPU)
- **Computing racks: 95% Direct Liquid Cooled**
- **Data storage: >100PB (NVMe+HDD)**
- **Warm water: Inlet temperature of 37 degrees**
- **NVIDIA Mellanox HDR 200 interconnect**
  - **Dragonfly+ topology**



# Leonardo System



# Evolution Roadmap

## Big Data Technopole

# LISA

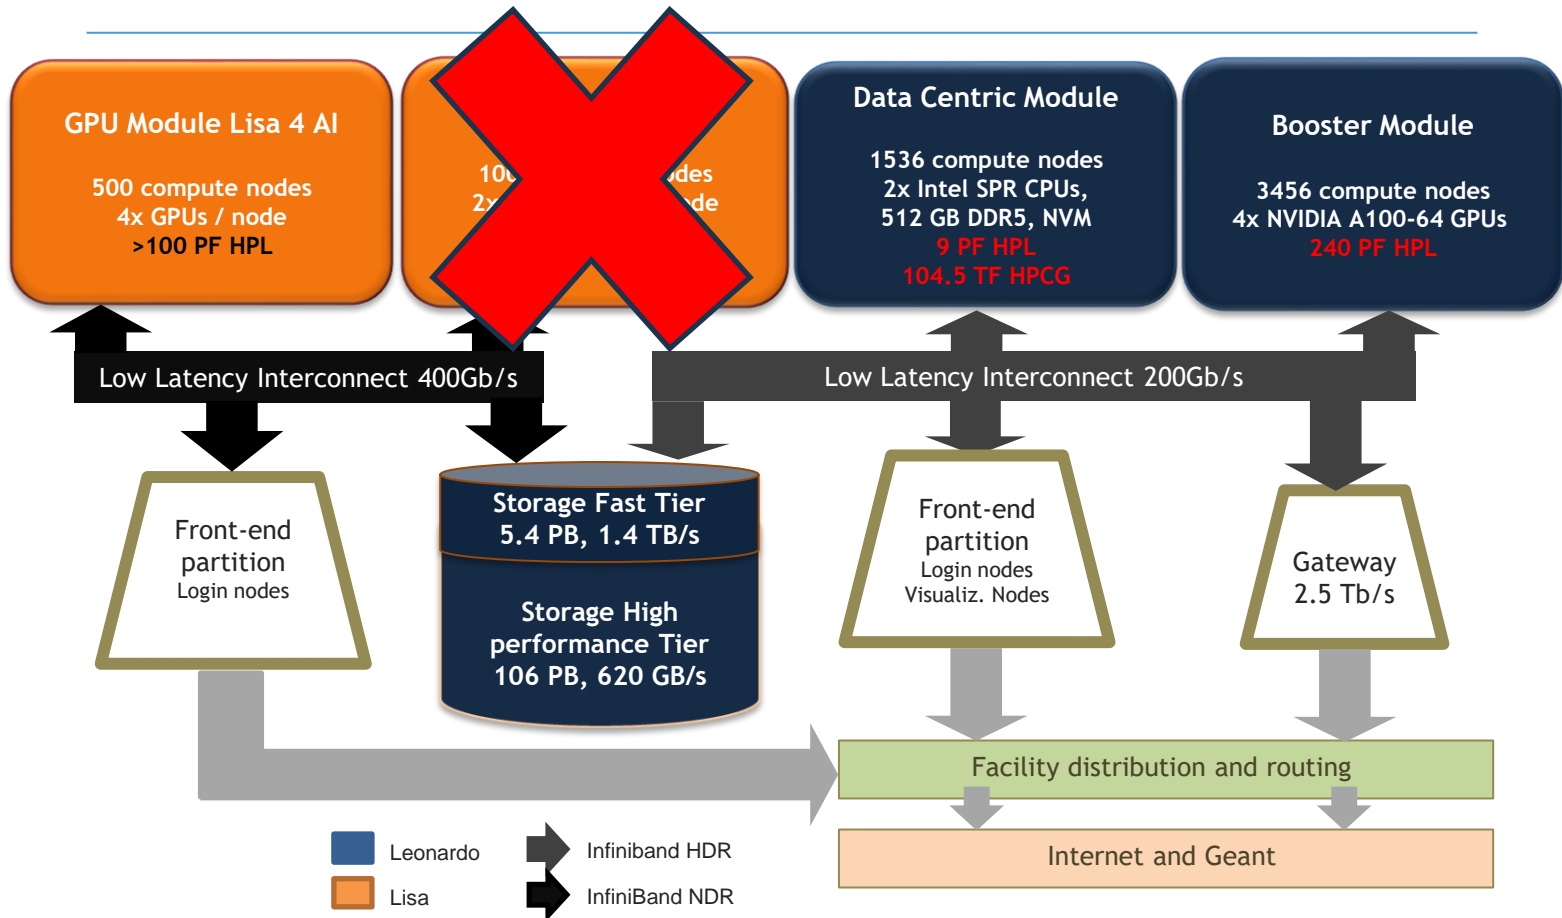
## Leonardo Improved Supercomputing Architecture

**TCO investment: 39M€**  
**65% Italy + 35% EuroHPC JU**  
**Capex: 28M€**  
**Opex: 11M€ (3 years)**





# Leonardo + LISA



# National Research Council (CNR) National Institute for Astrophysics (INAF)

CINECA in collaboration with CNR and INAF are going to install a system in the technopole focused on computational research on material solid and cosmology.

## Architecture:

- **HPL:** > 18 PF
- **GPU:** > 100-150 accelerated nodes
- **Storage:** > 1-2 PB Full-flash storage
- **Network:** InfiniBand HDR or NDR


High dense and energy efficient system

TCO Investments € >20  
M

Capex: ~ 65%

Opex: ~ 35% (5 years)

Update of Data hall 2 of Technopole: 1M€



Consiglio Nazionale  
delle Ricerche



**INAF**  
ISTITUTO NAZIONALE DI ASTROFISICA  
OSSERVATORIO ASTROFISICO DI CATANIA

# National Meteorological Agency

CINECA have a partnership with the National Meteorological Agency to acquire and operate the HPC system

## Architecture:


- **HPL:** 1.5PF CPU + 2.5PF GPU
- **CPU:** 200-300 conventional nodes
- **GPU:** 16-20 accelerated nodes
- **Storage:** 10PB parallel storage (HDD)
- **Network:** InfiniBand HDR or NDR
- **Gateway:** InfiniBand/Ethernet (>1 Tb/s)

This system will be require high availability and it will be connected with ECMWF storage with a dedicated network (100Gb/s – 1Tb/s) to support high-volume data movement.

TCO Investments **€ 8.5M**

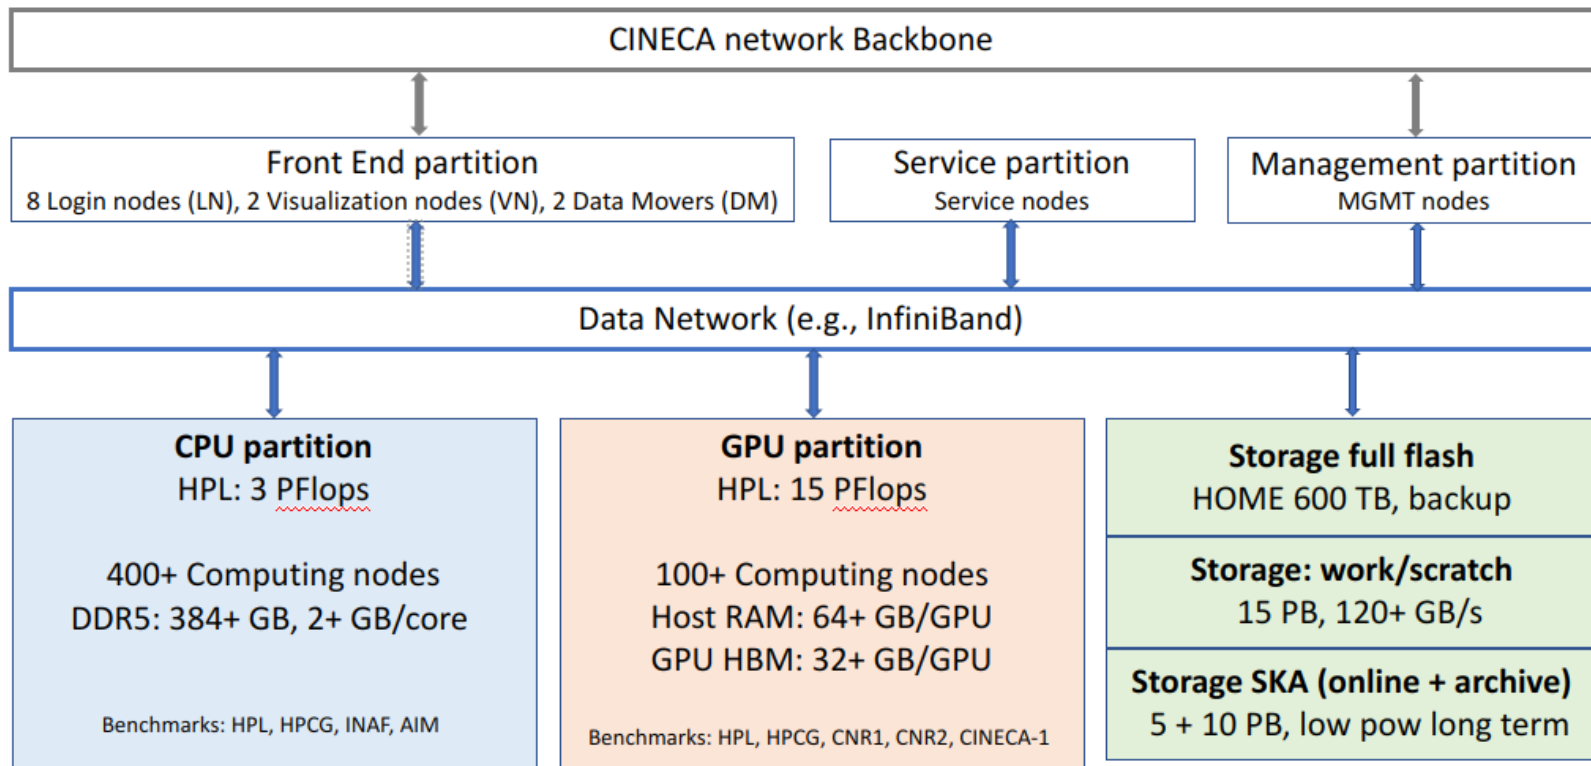
Capex: 5.25 M€

Opex: 3.25M€ (5 years)



# TIER-1 TECNOPOLO

CINECA



↔ Ethernet  
↔ Data Network

# CINECA Quantum Computing

**CINECA** plans to acquire a Quantum Computer

Initially the QC will be an experimental and dedicated system, but the idea is to use QC as an **accelerator of Leonardo**

Some **QC technologies** are under investigation, aiming for a 3-phase system upgrade:

1. “QC for “Education”
2. 500 qubit (dig)
3. 500/1000 qubit (dig)

It will be considered **QC European technologies**

Time frame: **installation H2-2024**

TCO investment **€ >25M**

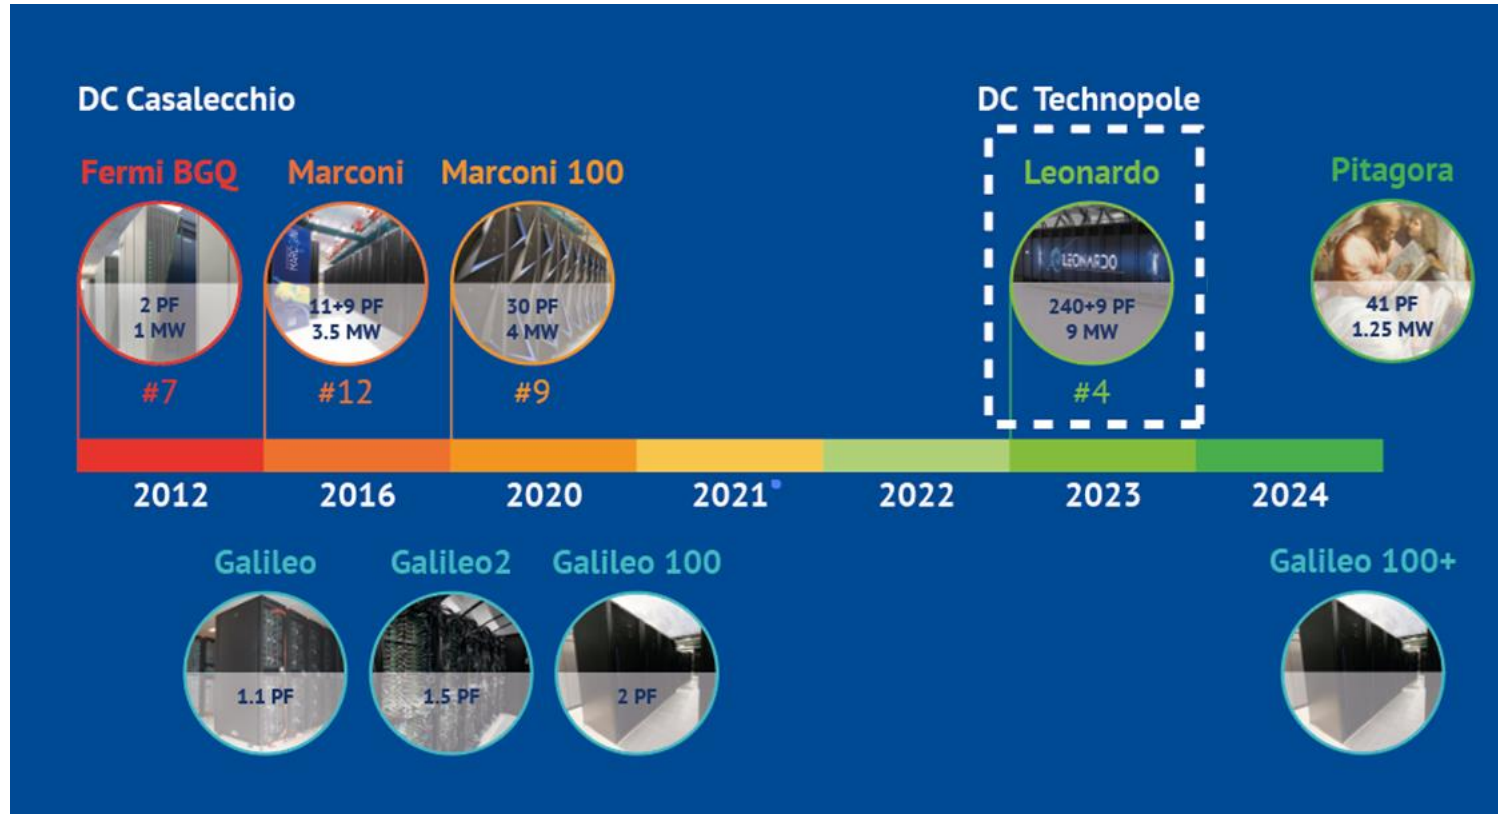




# Evolution Roadmap

## CINECA Headquarter - Casalecchio di Reno

# Roadmap of CINECA Flagship System



# EUROfusion

Italian partners (Cineca, ENEA) plans to support next phase of EUROfusion HPC provision

## Architecture:

- **HPL:** 28 PF GPU and 17 PF CPU
- **CPU:** 1008 conventional nodes
- **GPU:** 200 - 250 accelerated nodes
- **Storage:** 10PB parallel storage
- **Network:** InfiniBand NDR


Most of the investments will be used for the CPU partition (2:1 ratio)

Operative in **Q3-Q4 2024**

TCO investment **€ 44.2M**

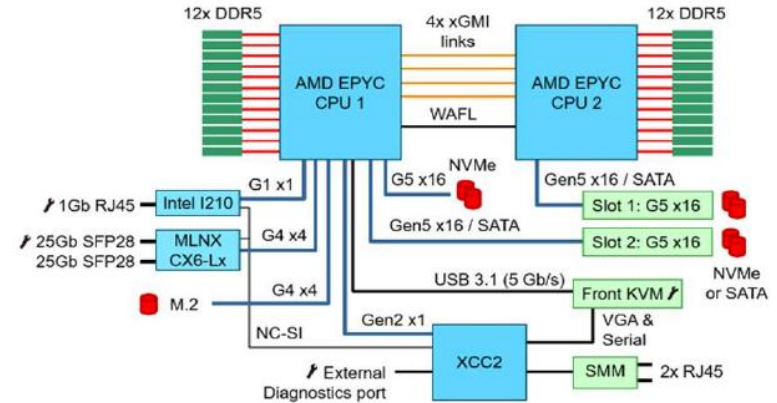

Capex: 28.4 M€

Opex: 15.8 M€ (6 years)



# CPU partition

- 14 racks
- 17 PFlops (Rmax)
- 1008 Compute nodes
  - 72 CN x rack (6x6 tray enclosures)
  - 2x AMD Turin 128c (Zen5 microarch) 2.3 GHz
  - **1.7X performance over Genoa**
  - 768 GB DDR5 6400 MT/s
  - 3GB DDR5 per core, **1.33x higher bdw**
  - Integrator w/ earliest availability from AMD Turin
  - 1x NDR200 adapter
- Full DLC (97% heat removal)









# Upgrade G100

CINECA plans to **upgrade G100**

Leonardo will remain a **conventional HPC system**

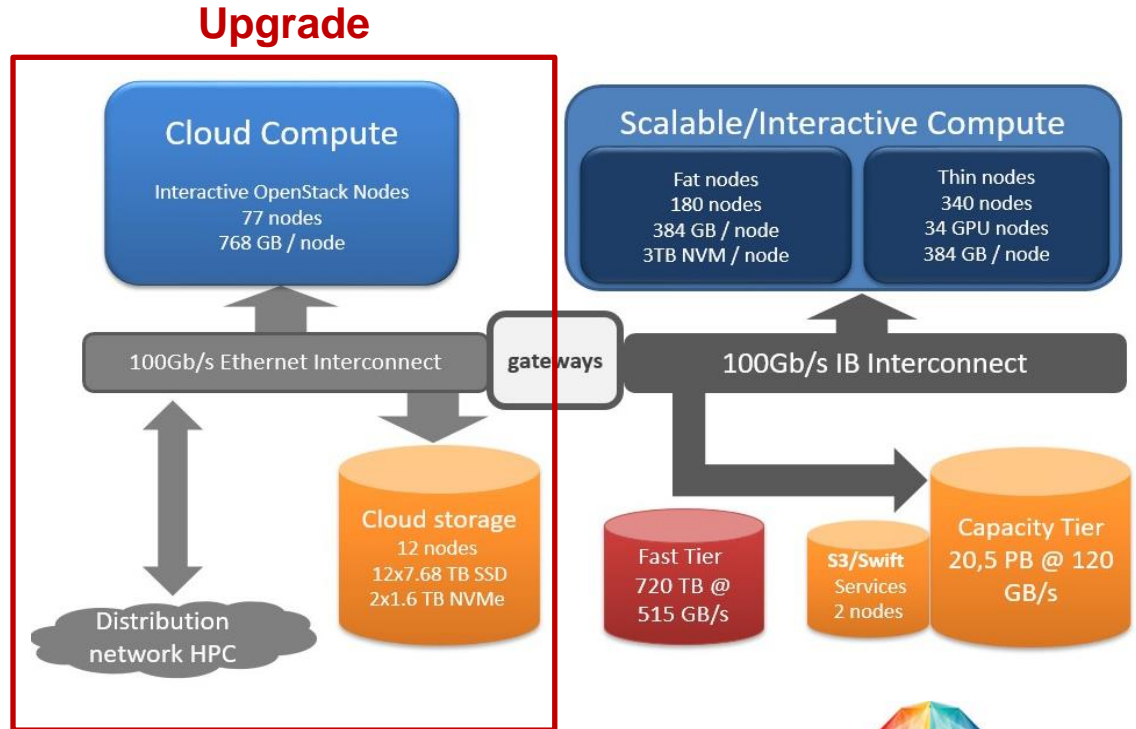
G100 will be upgraded to become an **important cloud and storage asset**

Most of the investments will be used for the **cloud partition** and a new **archive/data lake storage**

TCO investment € 23,5M€

Capex: 16,5 M€

Opex: 7M€ (5 years)



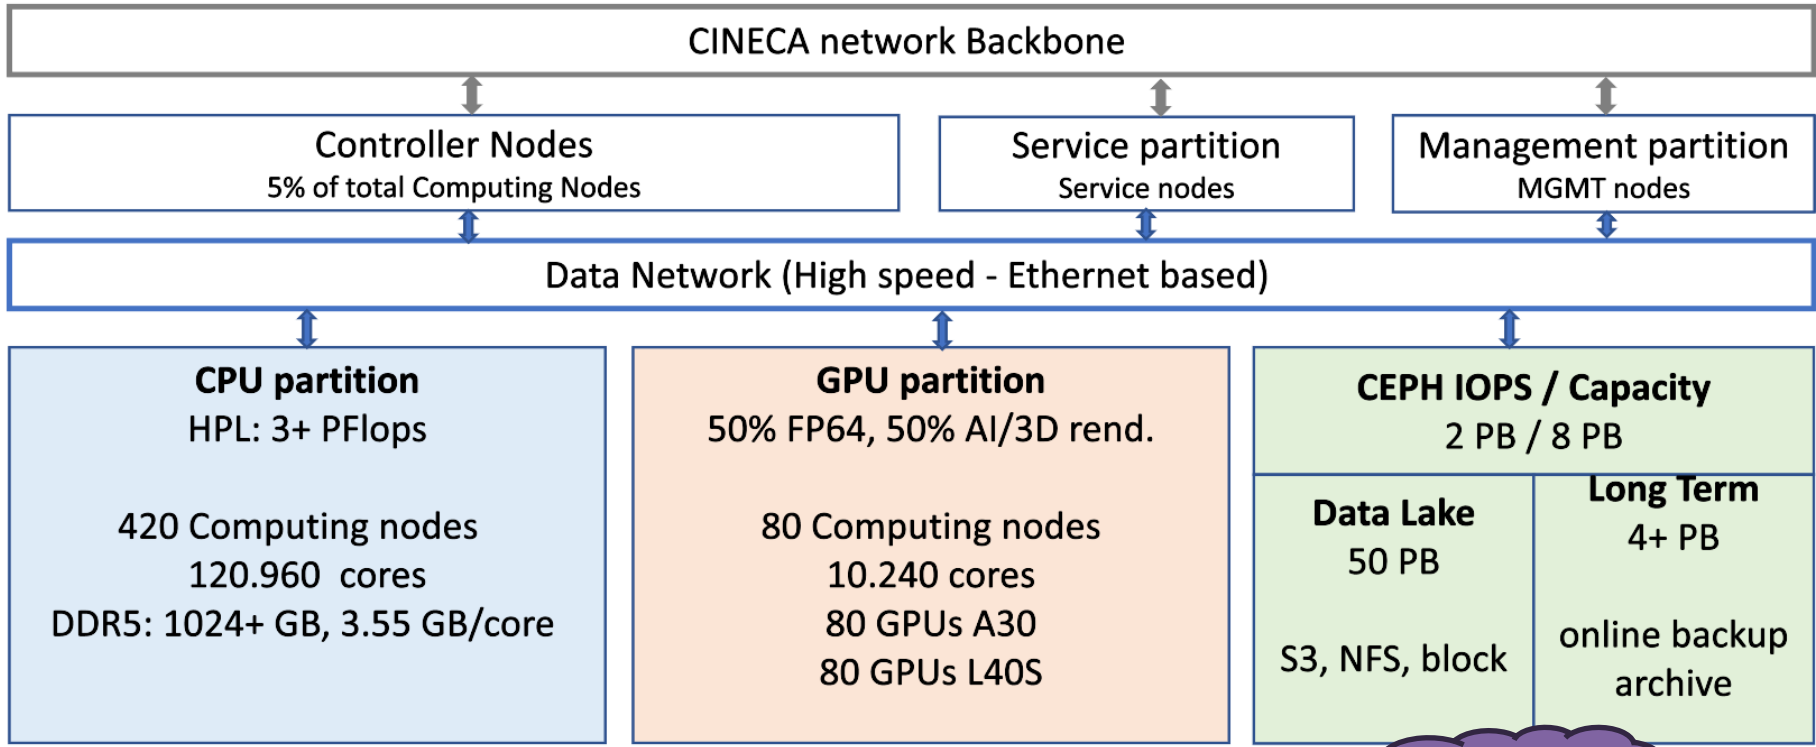
**FENIX**  
RESEARCH INFRASTRUCTURE



Human Brain Project

# CLOUD (GALILEO100-ADA) ++

CINECA



↔ Service network  
↔ Data Network

**National Cloud for research**



# CLOUD targets

- ✓ Data processing and analysis services complementing High Performance Computing services
- ✓ Bridging HPC infrastructure (Cloud to host front-end services for the user)
- ✓ Data management services Exposing simulation data to the web via community-defined services
- ✓ Workload requiring ISO27001 certification (processing sensitive data)
- ✓ Kubernetes cluster and containerized workflows
- ✓ Templated infrastructure deployment (Infrastructure as code)

# Prototipi Modulari (PRiMO) – TestBench system

CINECA is realizing an evaluation platform to host evaluation systems, prototypes and engineering samples

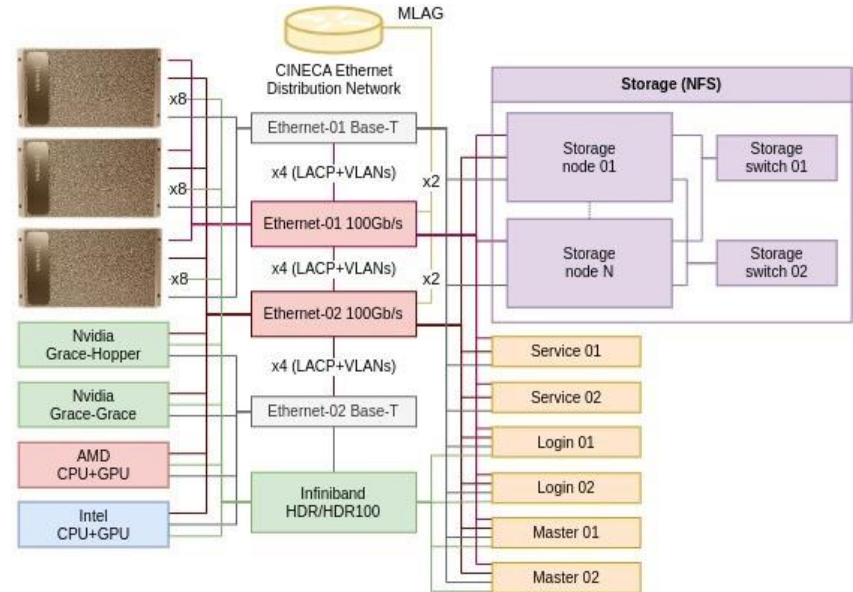
The architecture will be very heterogeneous and based on high-performance ethernet.

It will also host x3 A100 DGXs of CINECA

Full-flash storage: 100TB

Only few weeks for the publication of the open tender

Capex investment € 500k€



# Evolution Roadmap

## New CINECA Data Center in Naples



# CINECA Data Center in Naples

CINECA will open a new Tier1 data center in **Naples area**

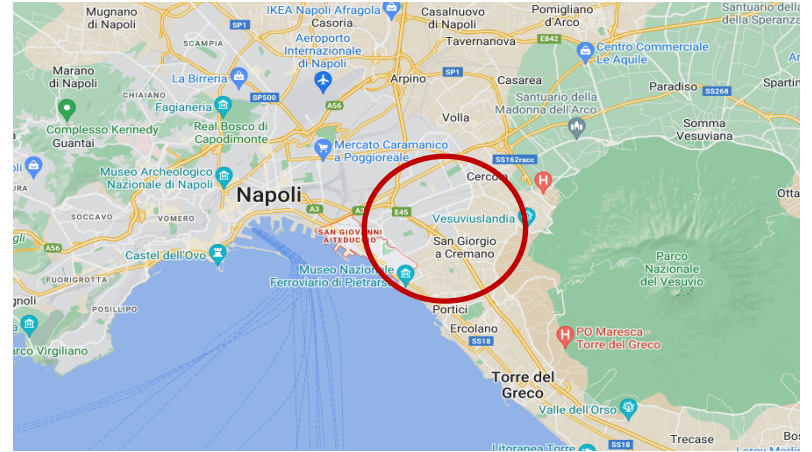
**San Giovanni a Teduccio** is the interested area

In collaboration with **University of Naples Federico II**

Time Frame: **operative 2024-2025**

Data center: **2MW – 500 sqm**

CINECA's data center investment **€ 5M**



Consiglio Nazionale  
delle Ricerche

# CINECA & CNR

CINECA and CNR plan a HPC system for basic and applied research in the southern Italy.

## Architecture:

- **HPL:** TBD
- **CPU:** TBD
- **GPU:** TBD
- **Storage:** TBD
- **Network:** InfiniBand HDR or NDR

TCO investment € 14.5M

Capex: 9.5M€

Opex: 5M€ (5 years)



Consiglio Nazionale  
delle Ricerche

# Bio-pharmaceutical Dompé

CINECA and Dompé have a long partnership to support the research on bio-pharmaceutical applications.

CINECA and Dompé are involved on several EU research projects

## Architecture:

- **HPL:** TBD
- **GPU:** The system will be fully accelerated
- **Storage:** 1-2 PB (full-flash under evaluation)
- **Network:** InfiniBand HDR or NDR
- **Long-term storage:** 8-10PB

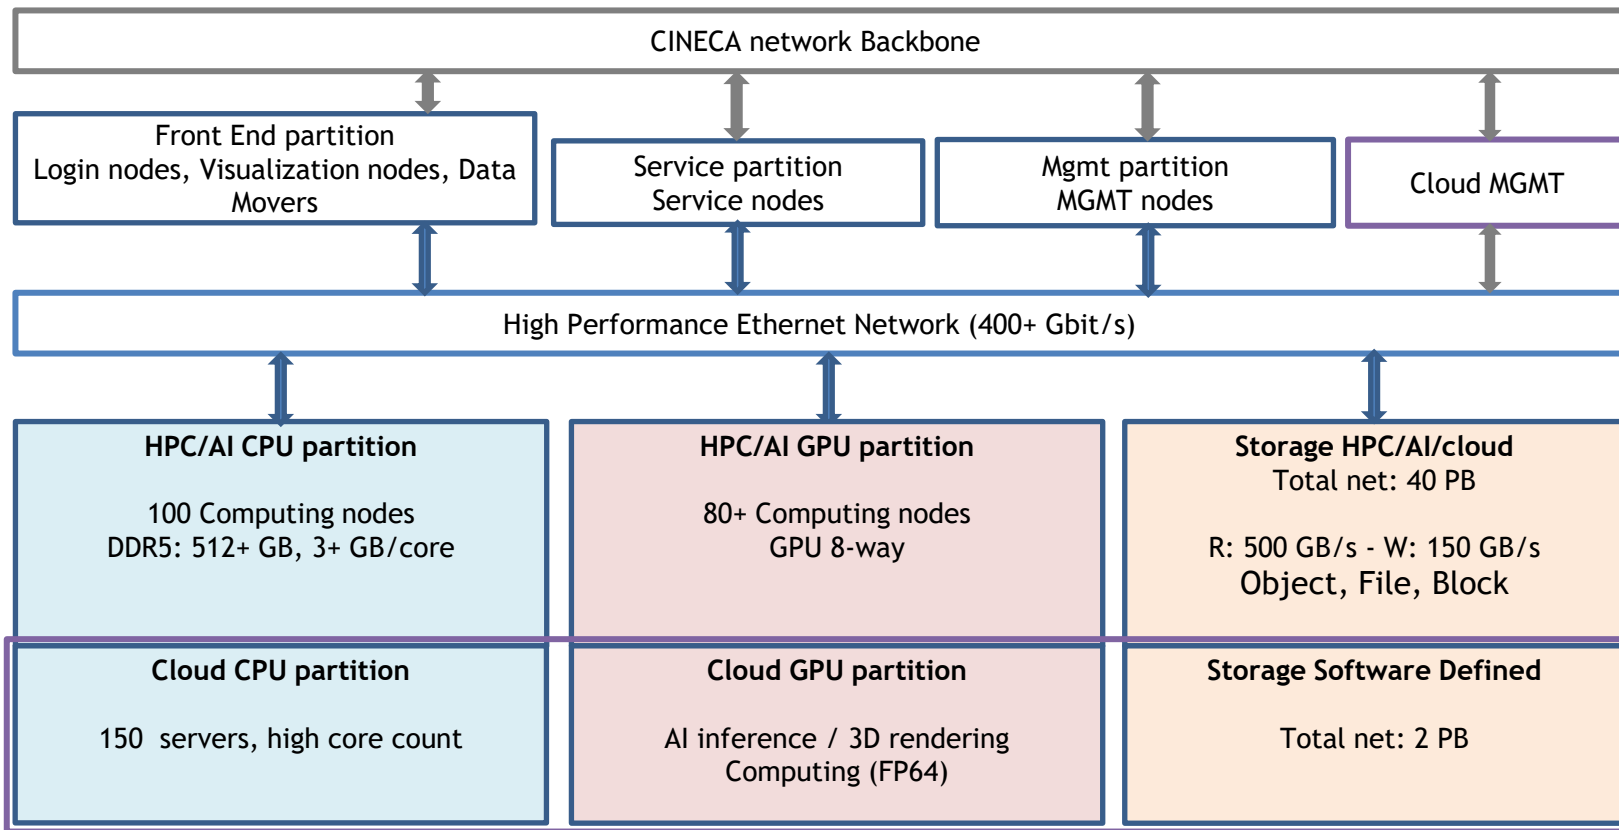
TCO investment € 2.8M

Capex: 2M€

Opex: 0.8M€ (5 years)



# TIER-1 NAPOLI



↔ Ethernet

↔ High Performance Ethernet Network

# Italian SuperComputing Resource Allocation - ISCRA

- Open to all scientific researchers affiliated to an Italian research organization needing large allocations of computer time, supporting resources and data storage to pursue transformational advances in science.
- Projects' Principal Investigators are expected to be affiliated to an Italian institution, while no restriction is applied for the Co-PI and collaborators.
- Further information ad for applying:  
<https://www.hpc.cineca.it/services/iscra>

- **ISCRA C:**
  - Small Project (8'000 GPUh on Leonardo, 3000 Nh on G100, 20K Ch on DGX)
  - Duration: 9 months
  - Easy to submit (1 page, few data needed)
  - Only technical evaluation
  - Continuous submission, 1 cut off per month
  - Several types of the project (HPC, Cloud, HPC + Cloud, Development & Benchmark, Quantum Computing) and the project focus (General purpose, Special Focus: AI & ML, Big Data/Bioinformatics, COVID-19, ...)
- **ISCRA B:**
  - Mid size projects (up to 250K Nh on Leonardo, up to 75K Nh on G100)
  - Duration: 1 year
  - More detailed proposal (some pages, scalability plot, detailed budget estimation, technical and scientific details needed)
  - Tech and Scientific evaluation
  - 2 call per year
- **ISCRA D:**
  - Only for long term storage resources
  - Up to 50 TB (To be revised) on FS and/or 200 TB on Tape Library
  - Maximum duration 36 months (+ 6 for move the data)
- **ISCRA A/Key projects (TBD)**



# EuroHPC – Regular Calls

- For Medium-size projects. Eg, Minimum request:
  - 10M Core hours on CPU machines;
  - 0,5-2M Core hours on GPU machines (depending from the machine);
- 2 calls per year (1 call every 6 months);
- ~10 pages proposal;
- Active on Petascale and Exascale machines
- 3 Access modes:
  - **The Scientific Access Track:** Open to all fields of science, will call for applications with a case to enable progress of science in the domains covered. These applications are expected to be able to justify the need for large allocations in terms of compute time, data storage and support resources because they are significantly contributing to the progress in their domain. The Scientific Track prioritises 75% of the total resources available at each cut-off period;
  - **The Industry Access Track:** Prioritises 20% of the total resources available for this cut-off period for proposals with a Principal Investigator from industry;
  - **The Public Administration Access Track:** Prioritises 5% of the total resources available for this cut-off period for proposals with a Principal Investigator from the public sector.

# EuroHPC – Extreme Calls

- For Extreme-size projects. Eg, Minimum request:
  - 80M Core hours on CPU machines;
  - 20-55M Core hours on GPU machines (depending from the machine);
- 2 calls per year (1 call every 6 months);
- Very demanding proposal;
- Active on Pre-Exascale machines (soon also on Exascale)
- 3 Access modes:
  - **The Scientific Access Track:** Open to all fields of science, will call for applications with a case to enable progress of science in the domains covered. These applications are expected to be able to justify the need for large allocations in terms of compute time, data storage and support resources because they are significantly contributing to the progress in their domain. The Scientific Track prioritises 75% of the total resources available at each cut-off period;
  - **The Industry Access Track:** Prioritises 20% of the total resources available for this cut-off period for proposals with a Principal Investigator from industry;
  - **The Public Administration Access Track:** Prioritises 5% of the total resources available for this cut-off period for proposals with a Principal Investigator from the public sector.

# EuroHPC – AI & Data Intensive Calls

- For AI & Data Intensive projects. Fixed Size:
  - 50K NODE hours (1.6M Core Hours) on GPU machines;
  - Max 11 proposals approved per call
- 2 calls per year (1 call every 6 months);
- Not So hard to submit;
- Active on Pre-Exascale machines (soon also on Exascale)

# ICSC calls

- Currently open only to flagship projects of ISCS
- When more resources will be available they will be opened to all the Italian public and private researcher for open science proposal
- Same modality and sizes (?) of the ISCRA calls

# Agreements

- Subscribing an agreement with Cineca based on some shared interests is always possible



# Training

Course	Date	Modo	Where	Participants	Interests	Average score
Python for Scientific Computing	6-8 FEB	ONLINE		63		8,9
Introduction to Python programming	6-8 FEB	IN PRESENZA	BOLOGNA	22		
Julia High Performance	1-3 MAR	ONLINE		25		7,8
Introduction to Parallel Computing with MPI and OpenMP	6-7-8 MAR	ONLINE		47		9
Debugging and Optimization of Scientific Applications	13-15 MAR	IN PRESENZA	BOLOGNA	10		8,5
Introduction to Fortran for Scientific Computing	20-23 MAR	ONLINE		30		8,7
Introduction to Scientific and Technical Computing in C	27-29 Marzo	ONLINE		32		9
Introduction to Deep Learning and Tensorflow	3-4 APR	ONLINE		63		9
Introduction to Python programming	17-21 Aprile	ONLINE		60		9
Introduction to Scientific and Technical Computing in C++	3-5 Mag	ONLINE		35		8,8
HPC Molecular Modelling	10-12 Mag	IN PRESENZA	BOLOGNA	16		9
Introduction to Quantum Computing School	5-9 Giu	IN PRESENZA	BOLOGNA	21		8,6
Summer School on Parallel Computing	3-14 LUG	IN PRESENZA	BOLOGNA	27		8,6
19th School of Computer Graphics for Cultural Heritage: AI and Cultural Heritage	9-13 OCT	IN PRESENZA	BOLOGNA	15		9,8
Introduction to Leonardo HPC cluster, for users and developers	27 OCT	IN PRESENZA/ONLINE	BOLOGNA		150	
Introduction to Python programming	8-10 NOV	IN PRESENZA	ROMA	19		9,5
Evento Mathworks	45240	IN PRESENZA	BOLOGNA		31	
Introduction to Fortran for Scientific Computing	14-17 NOV	IN PRESENZA	BOLOGNA	16		8,4
Advanced School on HPC Computing with GPU Accelerators	20-24 NOV	ONLINE		30		9,4
High Performance Bioinformatics	4-5-6 DIC	IN PRESENZA	ROMA	24		9,4
Introduction to Scientific and Technical Computing in C	11-13 DIC	ONLINE		29		8,6
High Performance Computing and Quantum Computing (HPCQC)	14-15 DEC	IN PRESENZA - ONLINE	BOLOGNA		89	



Thank You

Massimiliano Guarrasi: [m.guarrasi@ Cineca.it](mailto:m.guarrasi@ Cineca.it)

