

AN INTRODUCTION TO THE CINECA SUPERCOMPUTING INFRASTRUCTURE:

From Marconi 100 to Leonardo and beyond

Massimiliano Guarrasi: m.guarrasi@cineca.it

Agenda

- CINECA in few words
- CINECA HPC Infrastructure
- Leonardo System
- Big Data Technopole
- Casalecchio di Reno Data Center (headquarter)
- Evolution Roadmap for Big Data Technopole
- Evolution Roadmap for Casalecchio di Reno
- Evolution Roadmap for New data center in Naples
- How to have resources

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



≈ 1000

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

UNIVERSITIES



IT Systems for the
Italian Academic
System

Technological Transfer
to Healthcare
Public Administration
Industry



PA & INDUSTRY

2000

2020

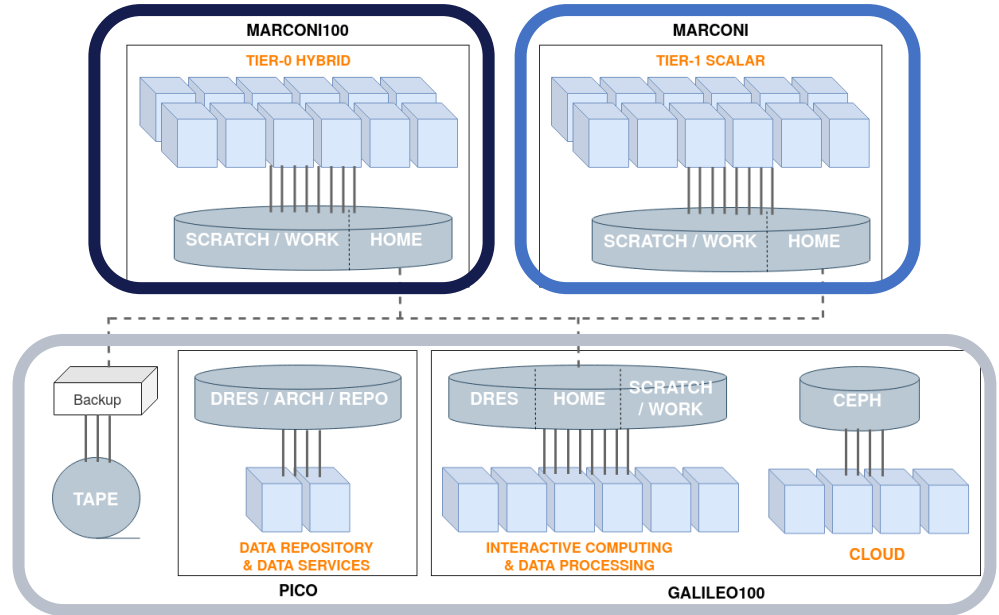
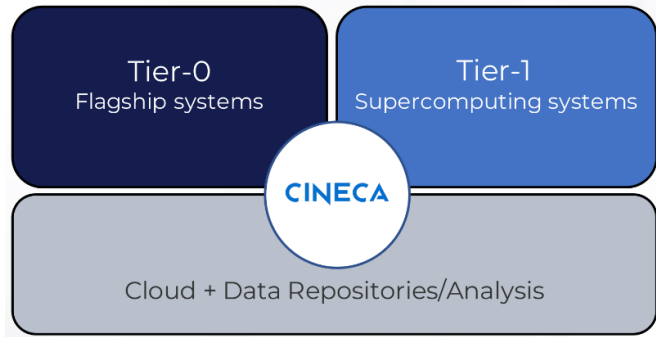
HPC PRE-EXASCALE



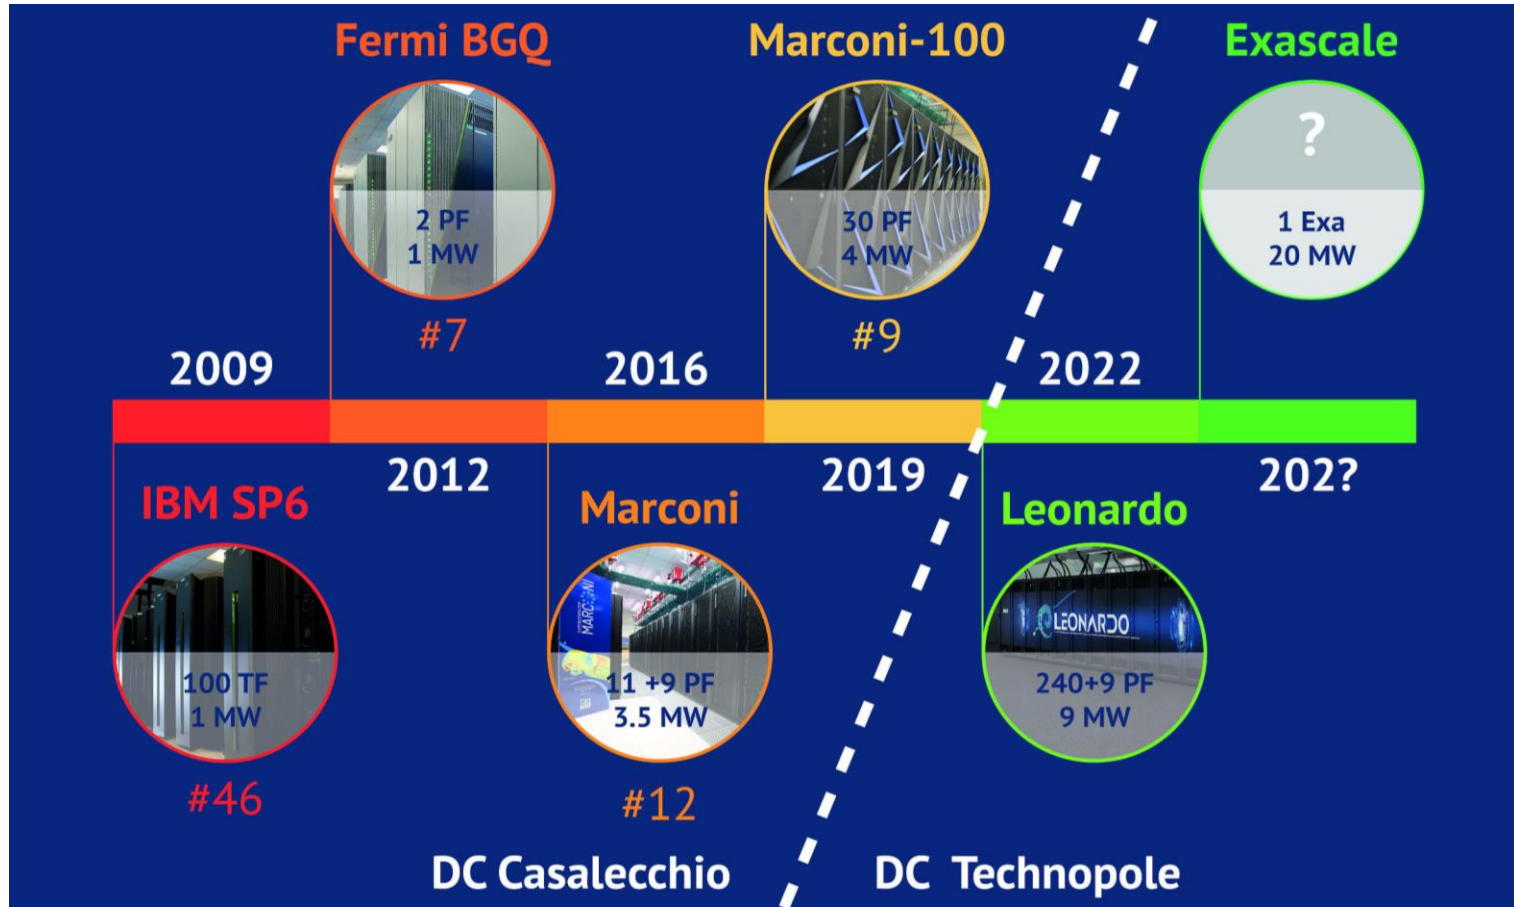
Artificial Intelligence
Big Data
Quantum Computing



CINECA HPC Infrastructure




Roadmap of CINECA Flagship System

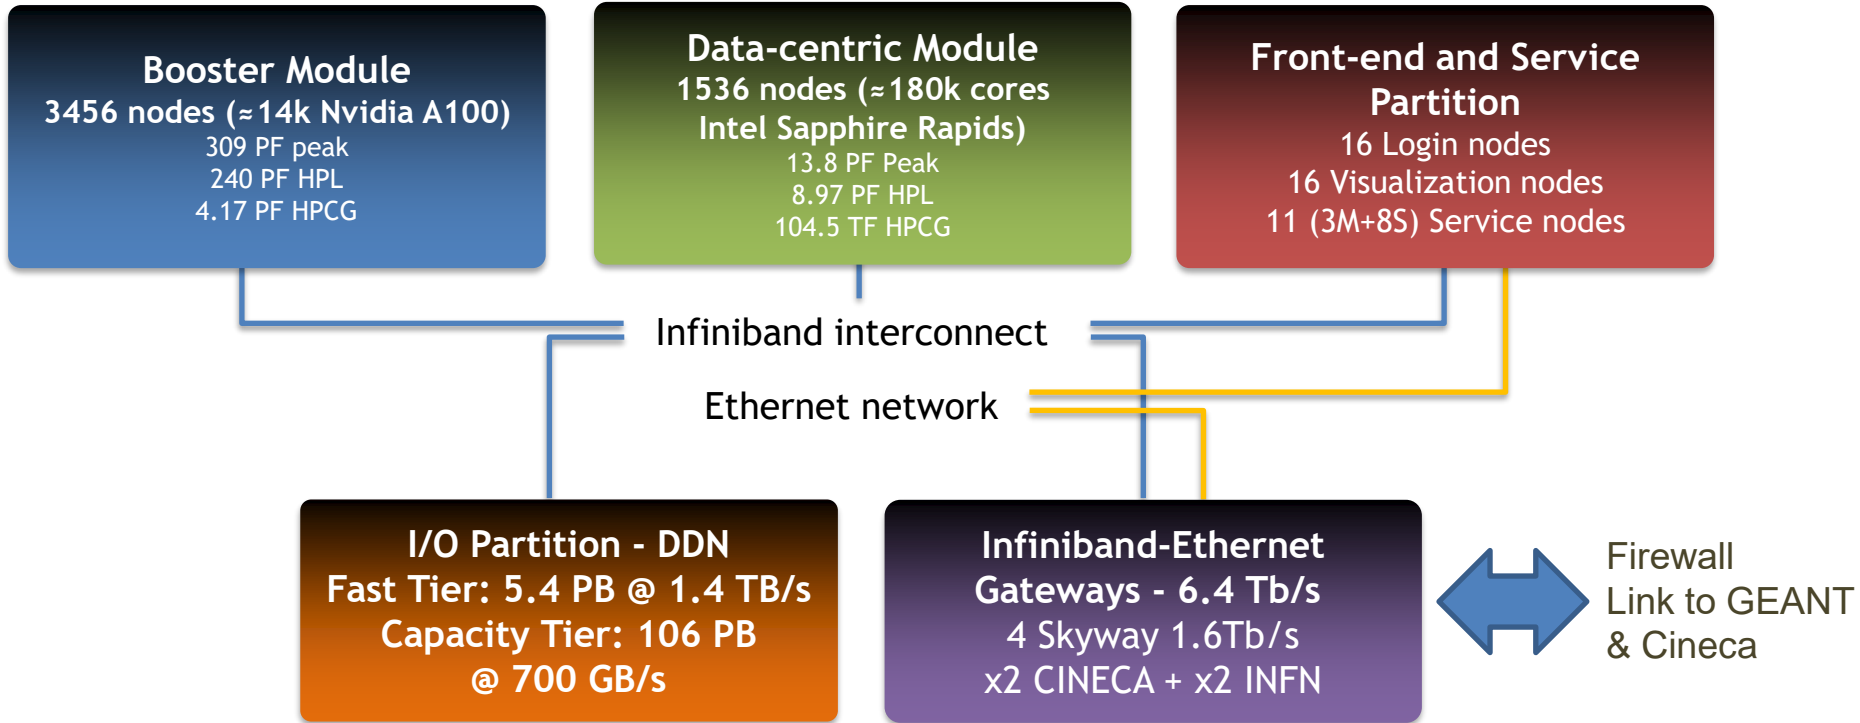


Leonardo System

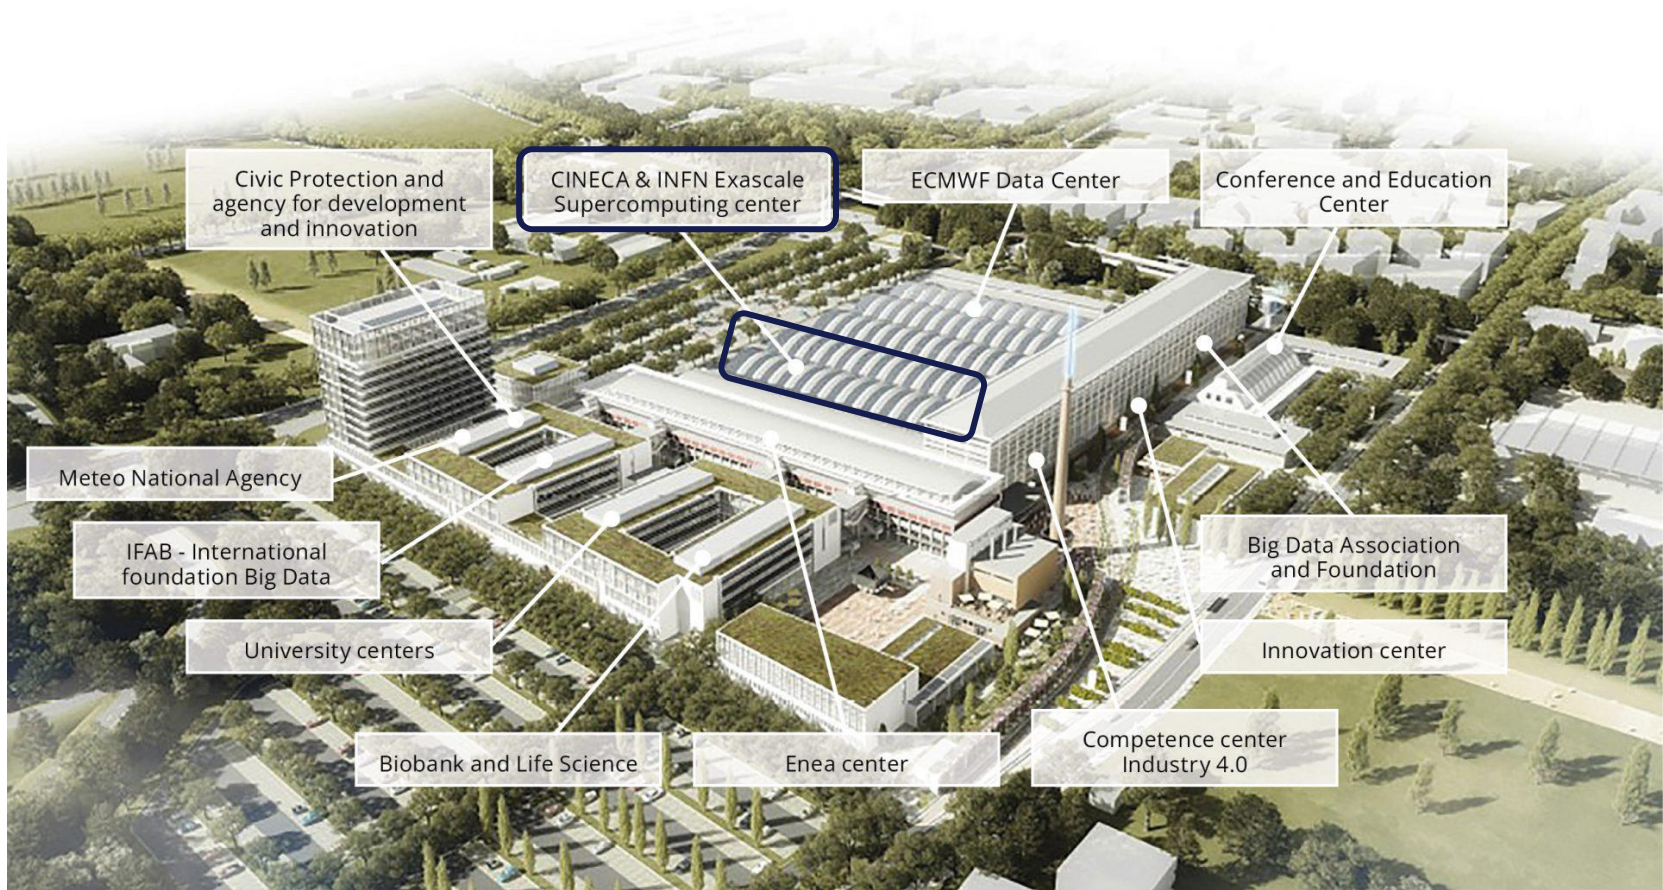
- **4th Top500**
- **HPL 240 PF + 9 PF (currently 170PF)**
- **TCO Investment: 240M€**
(120M€ Capex + 120M€ Opex)
- **5000 nodes based on BullSequana XH2000 platform technology**
(3500 GPU + 1500 CPU)
- **Computing racks: 95% Direct Liquid Cooled**
- **Data storage: > 100PB (NVMe+HDD)**
- **Warm water: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



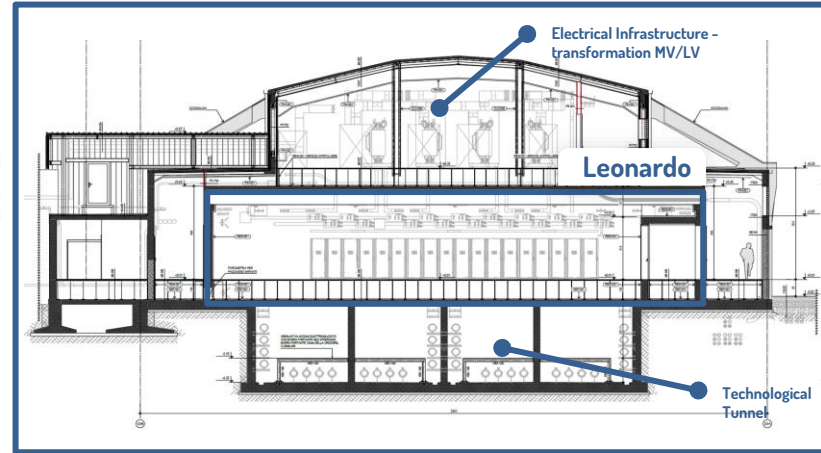
Leonardo System



Big Data Technopole – Bologna Science Park



CINECA/INFN Data Center - Big Data Technopole



Features

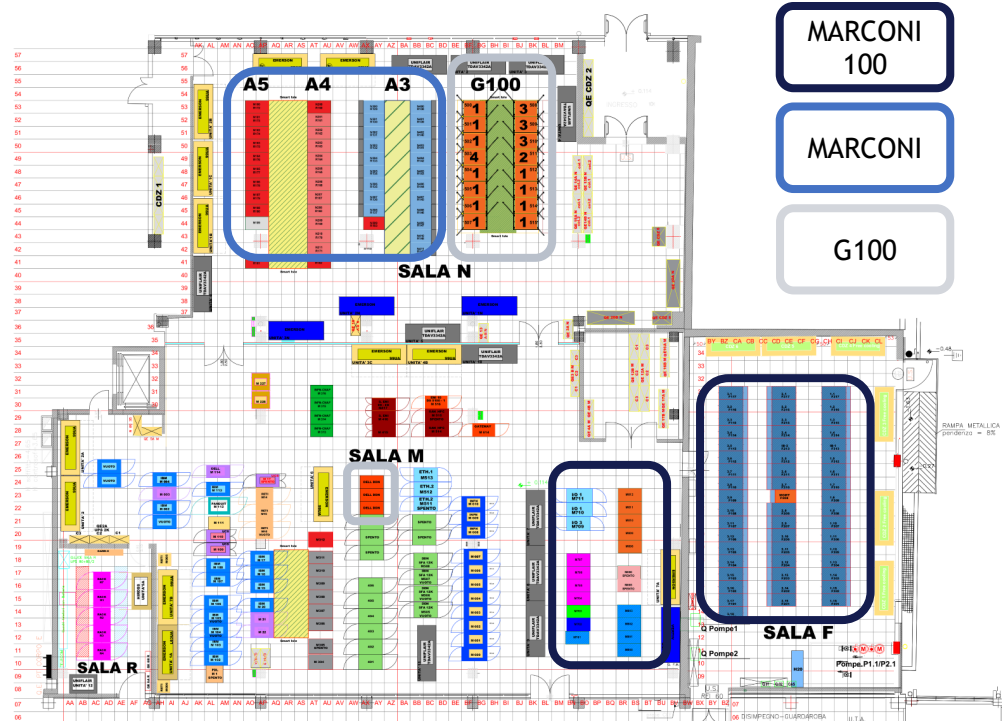
- **12 MW** and **1240 sqm** Rack Room
- **PUE < 1,10** (year based measurement strategy compliant to Level 3 Green Grid/ASHRAE)
- **Redundancy** Configuration: **3+1**, Electrical and Mechanical
- **Certify Rating 4 - TIA942 and Tier IV**

- 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)

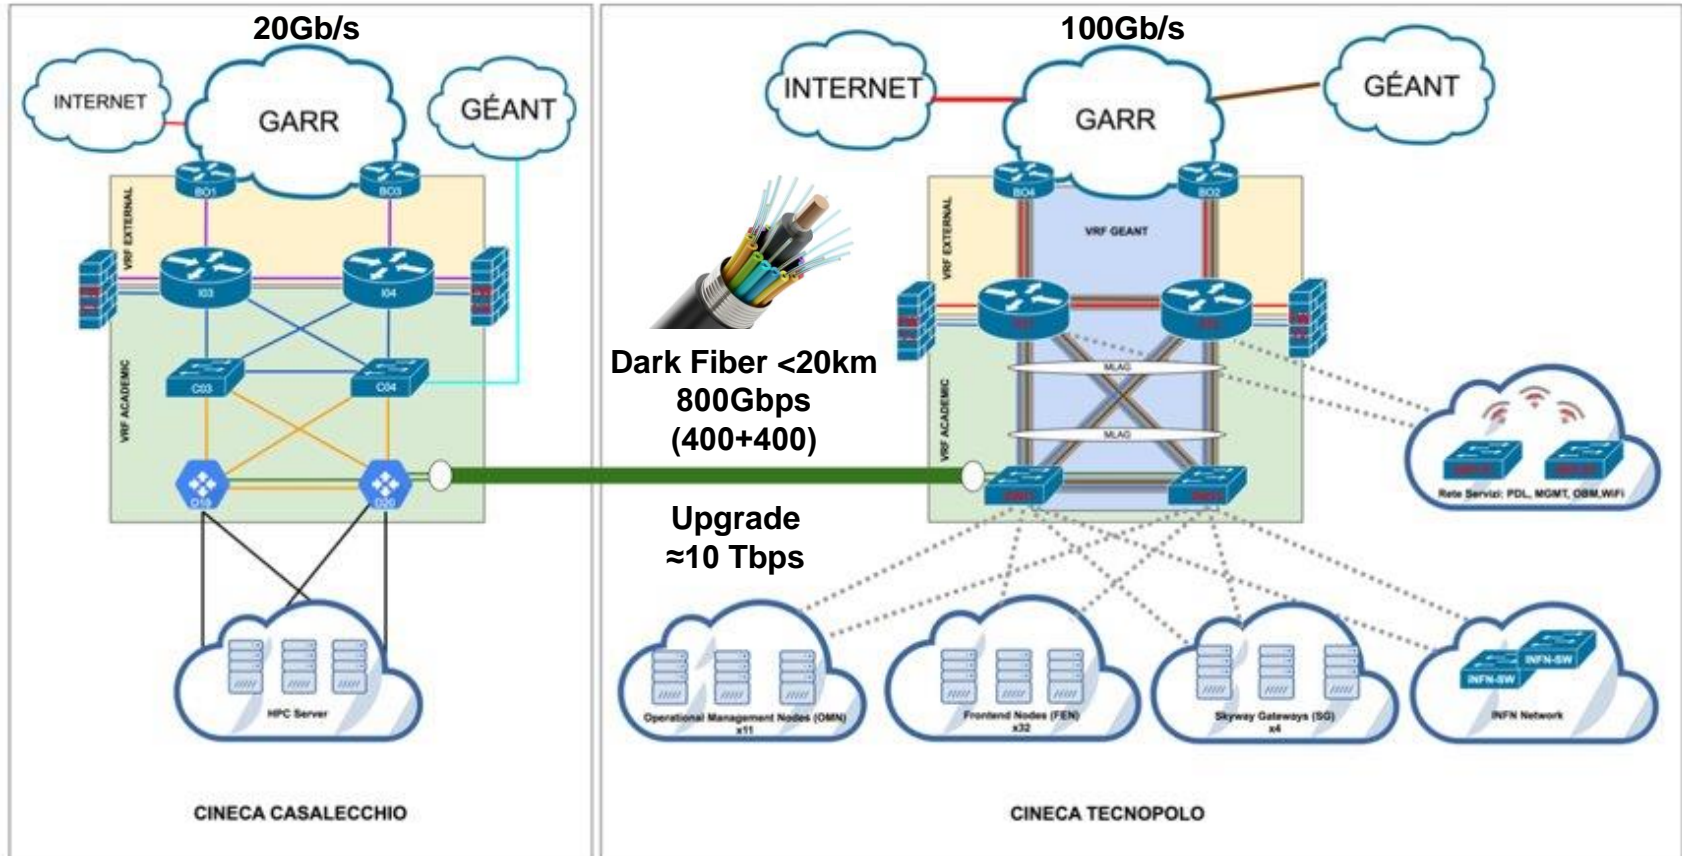
CINECA Data Center – Casalecchio di Reno (BO)

Features

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



Casalecchio di Reno <-> Technopole - Interconnection



Evolution Roadmap

Big Data Technopole

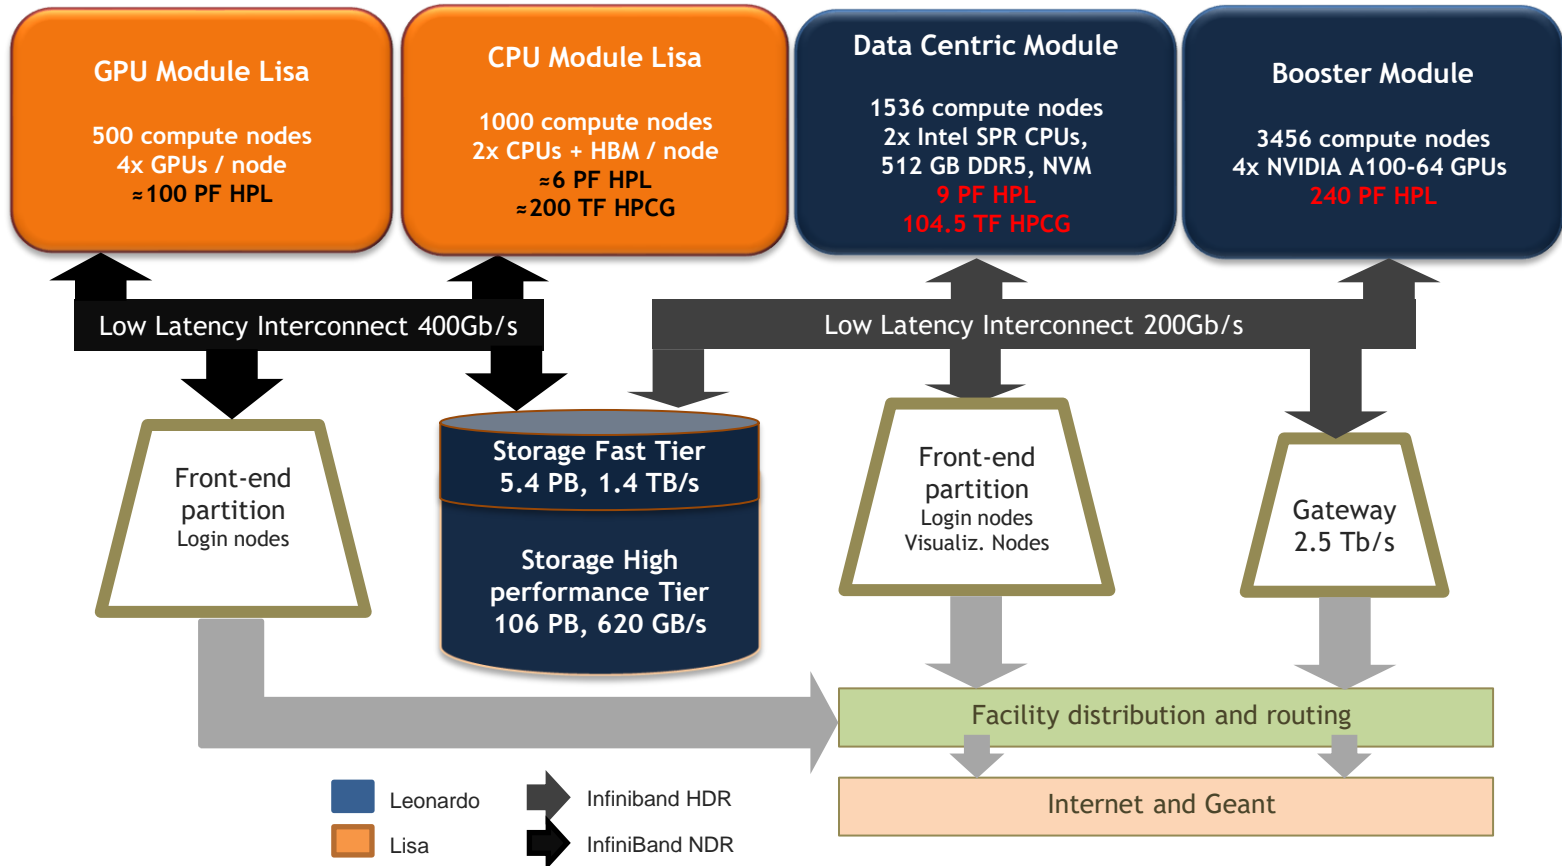
LISA

Leonardo Improved Supercomputing Architecture


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



Leonardo + LISA



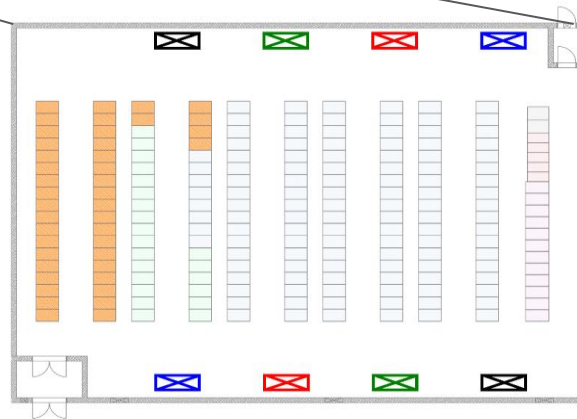
Leonardo + LISA



Building preparation + power distribution and piping (~600 k Euro)

	Leonardo	Leonardo + LISA
Power	7700 kW	10000 kW
Cooling capacity, tempered water	7100 kW	9700 kW
Cooling capacity, chilled water	600 kW	800 kW
Floor space	570 m ²	720 m ²

Leonardo upgrade
roughly 36 racks and
2.5 MW power IT.



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: ~ 13 M€

Opex: ~ 7 M€ (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.

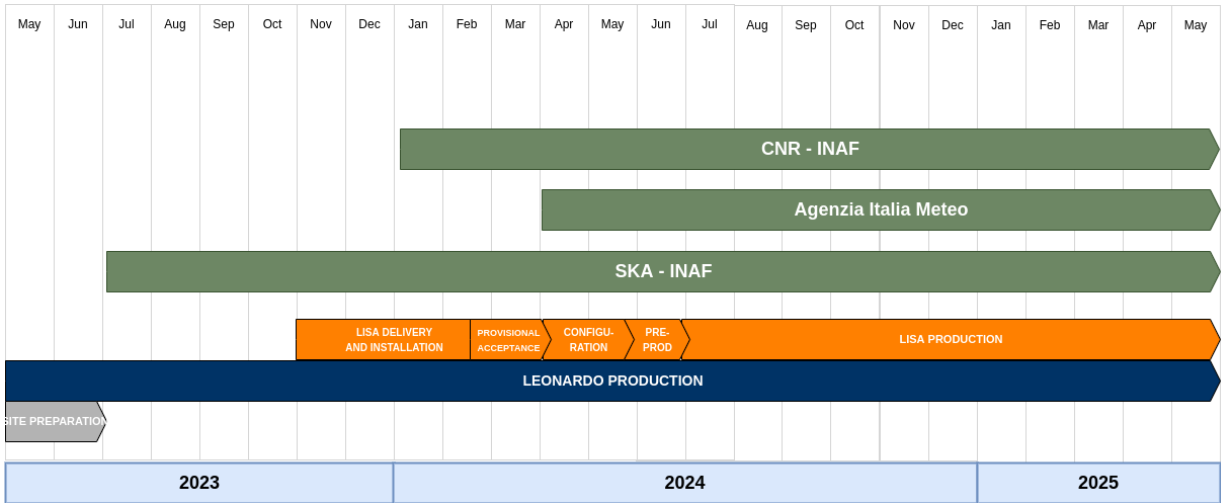

TCO Investments **€ 8.5M**

Capex: 5.25 M€

Opex: 3.25M€ (5 years)



Timeline for new HPC systems at Technopole



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. 100 qubit
2. 200 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

EUROfusion

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

Architecture:

- **HPL:** 34 PF GPU and 14 PF CPU
- **CPU:** 2000 - 2500 conventional nodes
- **GPU:** 200 - 250 accelerated nodes
- **Storage:** 10PB parallel storage
- **Network:** InfiniBand HDR or NDR


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

Operative in **Q1-Q2 2024**

TCO investment **€ 44.2M**

Capex: 28.4 M€

Opex: 15.8 M€ (6 years)



Upgrade G100

CINECA plans to **upgrade G100**

Leonardo will remain a **conventional HPC system**

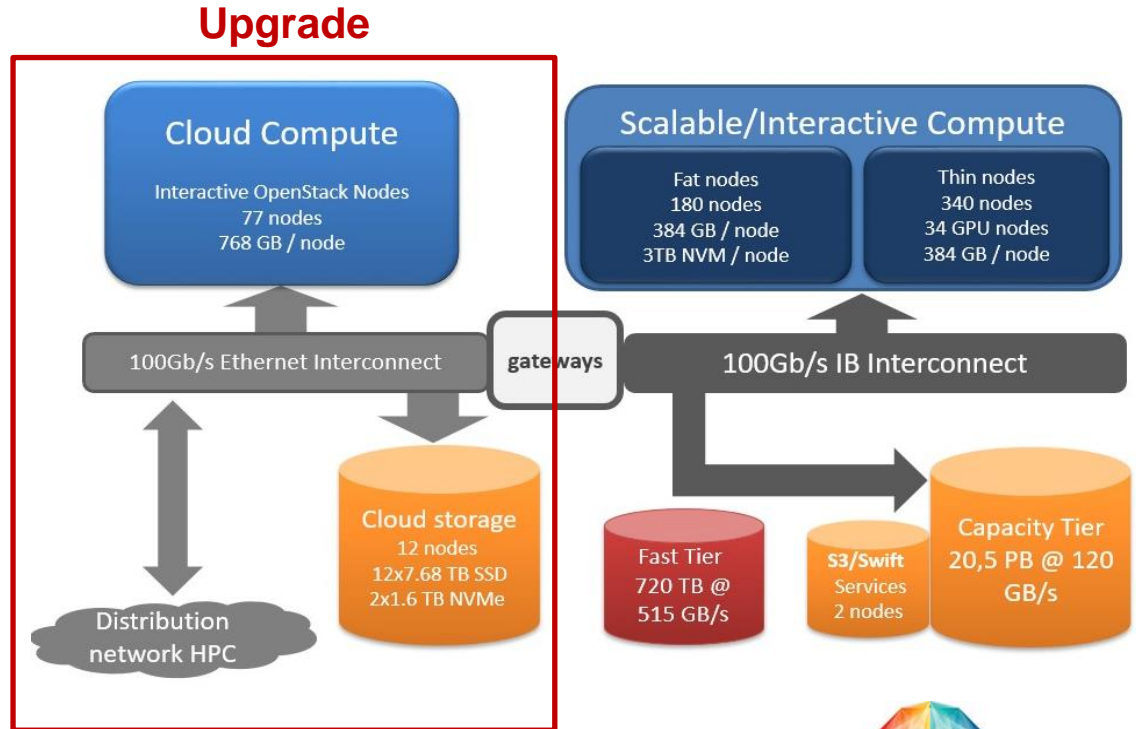
G100 will be upgraded to become an **important cloud and storage asset**

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)

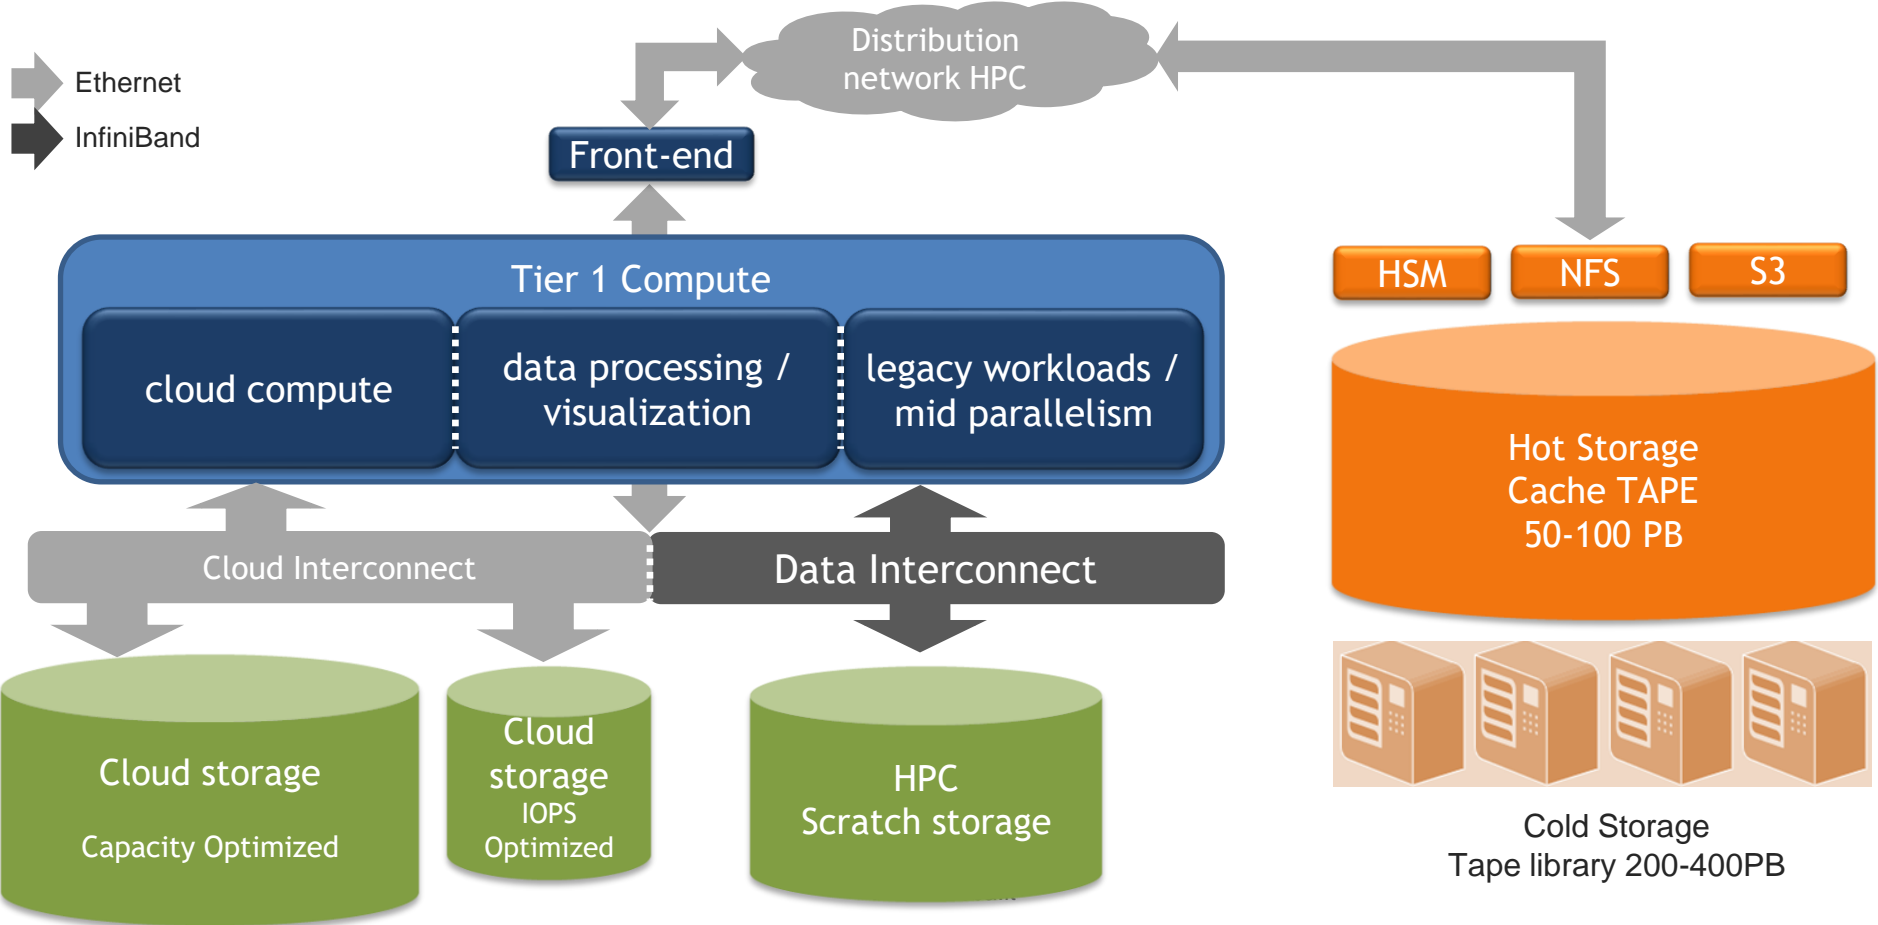


FENIX
RESEARCH INFRASTRUCTURE



Human Brain Project

Upgrade G100



Prototipi Modulari (PRiMO) – TestBench system

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

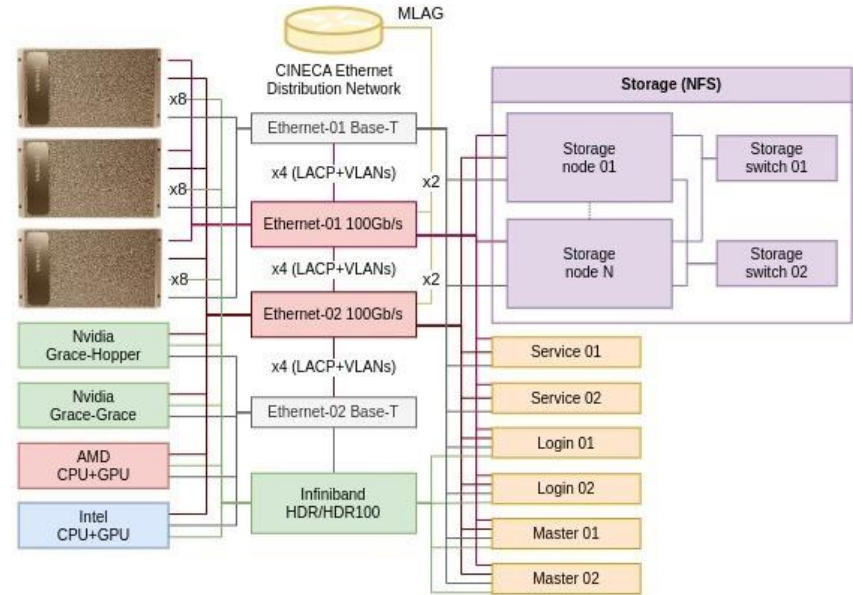
The architecture will be very heterogeneous and based on high-performance ethernet.

It will also host x3 A100 DGXs of CINECA

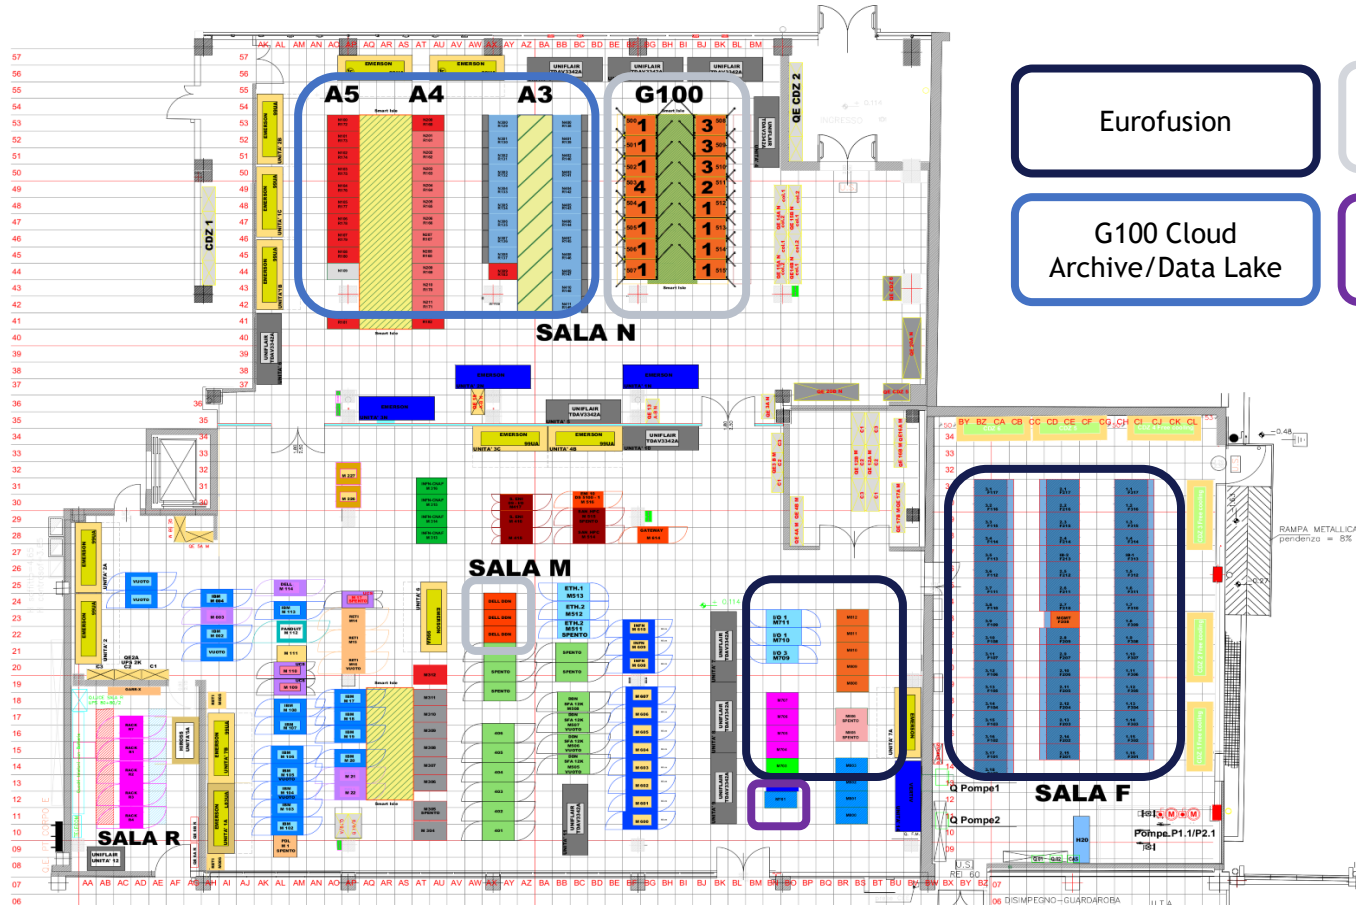
Full-flash storage: 100TB

Only few weeks for the publication of the open tender

Capex investment € 500k€



Layout Data Center Casalecchio di Reno



Eurofusion

G100

G100 Cloud
Archive/Data Lake

PRiMO

RAMPA METALLICA
pendenza = 5%

Q Pompe1
Q Pompe2
Pompe.P1.1/P2.1

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



National Cybersecurity Agency

CINECA have started a joint collaboration with the National Cybersecurity Agency.

It will be used to analysed the cybersecurity logs/data coming from the perimeter of national cybernetic security.

It will implement high-security HW/SW features, we are evaluating secure enclave technologies, full memory encryption and full-encrypt storage.

Architecture:

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

TCO investment € 17.5M

Capex: 11.5 M€

Opex: 6 M€ (5 years)



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)



Conclusion

CINECA with its partnerships will invest several millions euro in the next two years for a total TCO investment of around:

€ 200M

The total CAPEX investments in HPC equipment in the order of: **€ 110M**

The CAPEX investment will be split in the following areas:

- Big Data Technopole: **€ 42M**
- Casalecchio di Reno: **€ 45M**
- New data center in Naples: **€ 23M**

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);
- 3 calls per year (1 call every 4 months);
- -10 pages proposal;
- Next calls:
 - 3 March 2023 - 10:00 AM Luxembourg time;
 - 7 July 2023 - 10:00 AM Luxembourg time;
 - 3 November 2023 - 10:00 AM Luxembourg time;
- 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;
- Next calls on september;
- Last call closed on December 15th;
- Active on Pre-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.



Thank You

Massimiliano Guarrasi: m.guarrasi@ Cineca.it

