Leonardo

system configuration and hosting site

COMPBIOMED – all hands meeting
June 2022

Mirko Cestari - CINECA

EuroHPC Joint Undertaking

CINECA

MUR
Ministero dell'Università e della Ricerca

INFN
Istituto Nazionale di Fisica Nucleare

SISSA

MINISTRY OF EDUCATION, SCIENCE, RESEARCH AND SPORT OF THE SLOVAK REPUBLIC

Federal Ministry
Republic of Austria
Education, Science and Research

MINISTRY OF INNOVATION AND TECHNOLOGY

REPUBLIKA SLOVENIJA
REPUBLIC OF SLOVENIA
Ministrstvo za izobraževanje, znanost in šport
Ministry of education, science and sport

HELLENIC REPUBLIC
Ministry of Digital Governance

Procurement timeline

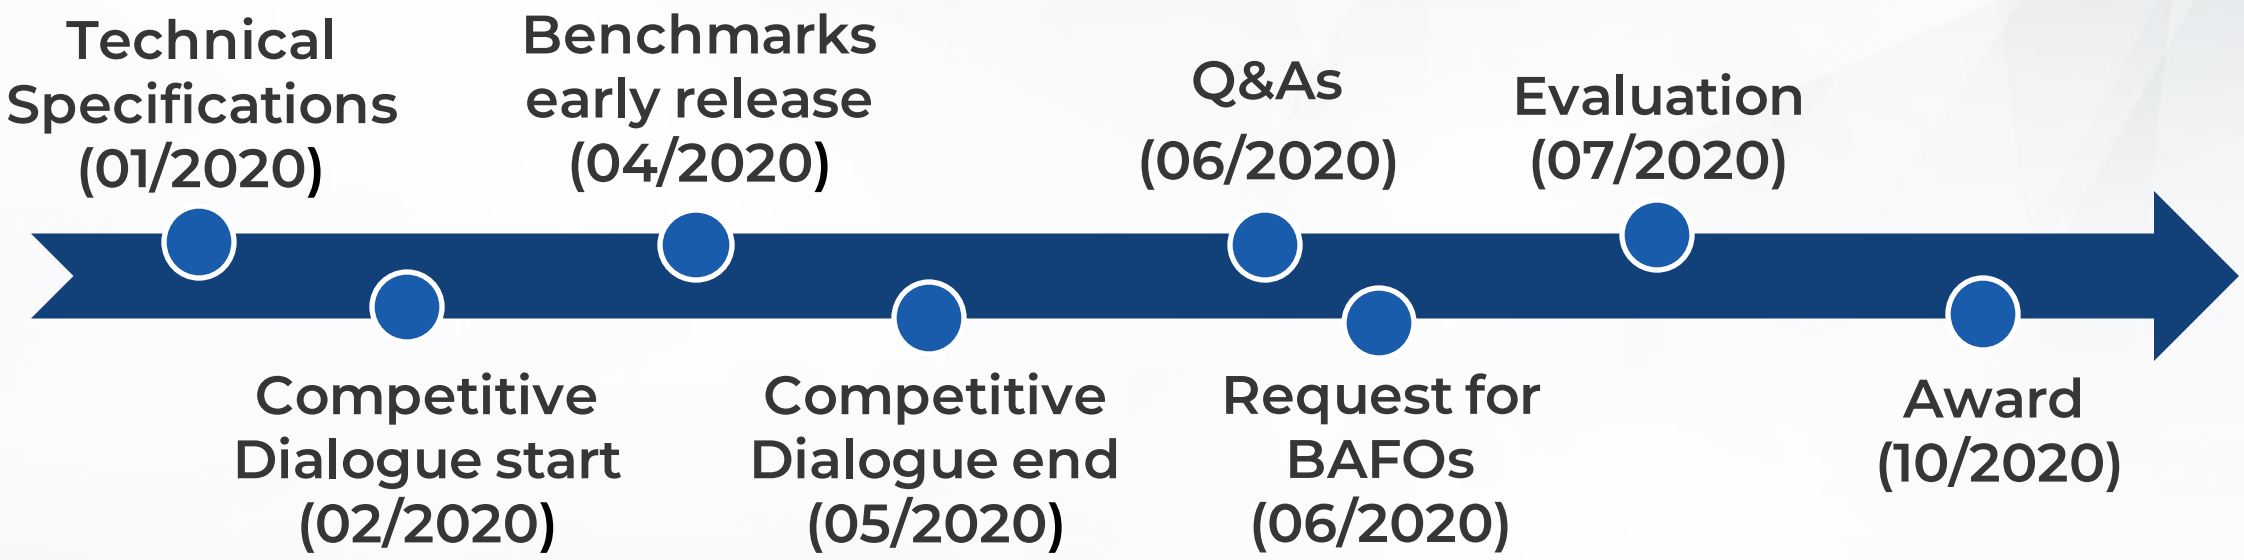

Procurement technical aspects

CINECA

- Write most of the tender **technical specifications**
- Setup of the **benchmark suite**
- **Provision** the benchmark suite to all tender participants
- Write benchmarking **rules**
- Write a **response template**
- Provide feedback on **Q&A**
- Provide template for **benchmark evaluation**

Public Gitlab

Document shipped with benchmarks

Clearly states evaluation points
Helpful in collecting data the way procurer want
Ease comparison between offers






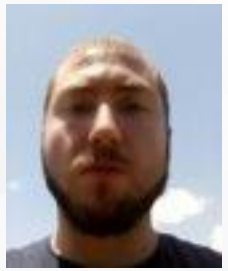


LEONARDO CINECA

Cineca HPC ecosystem

LEONARDO CINECA

Cineca HPC infrastructure

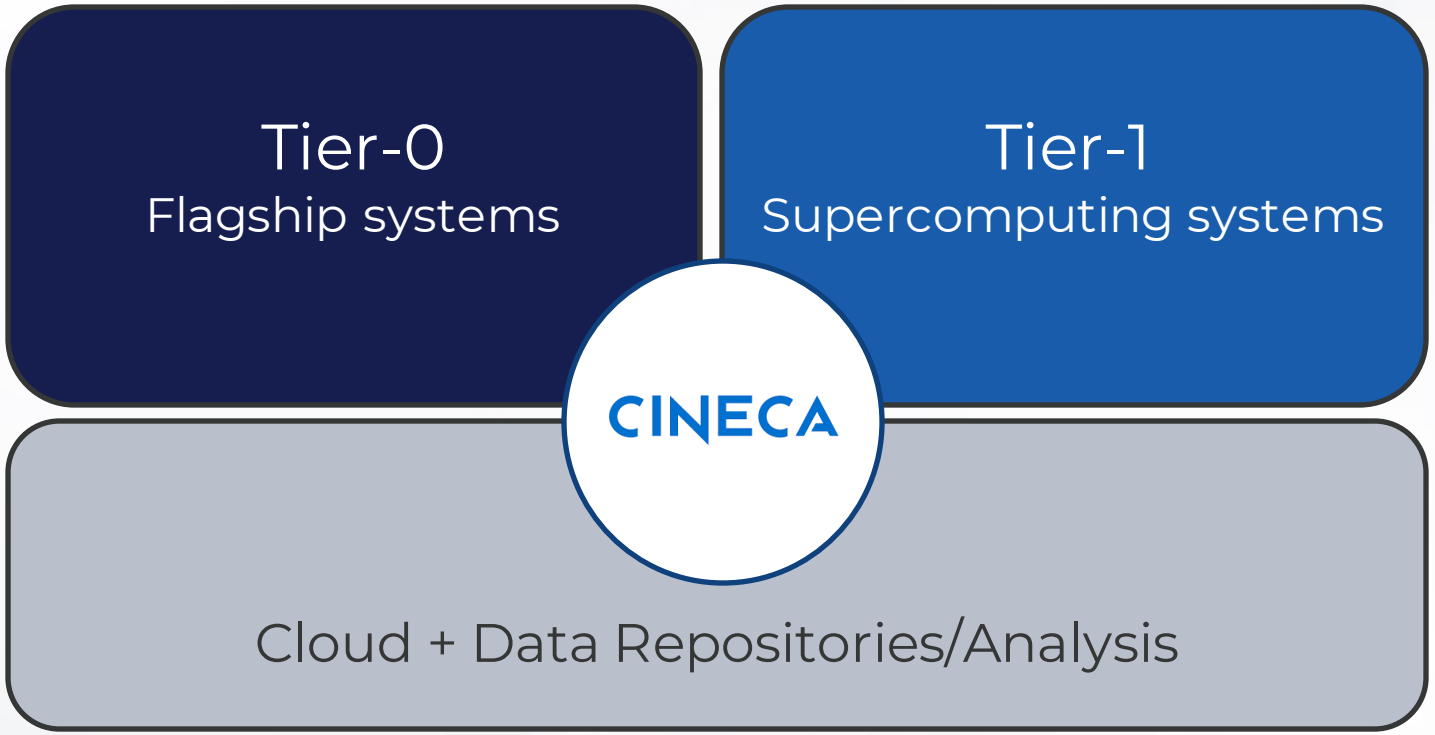

Cineca HPC infrastructure

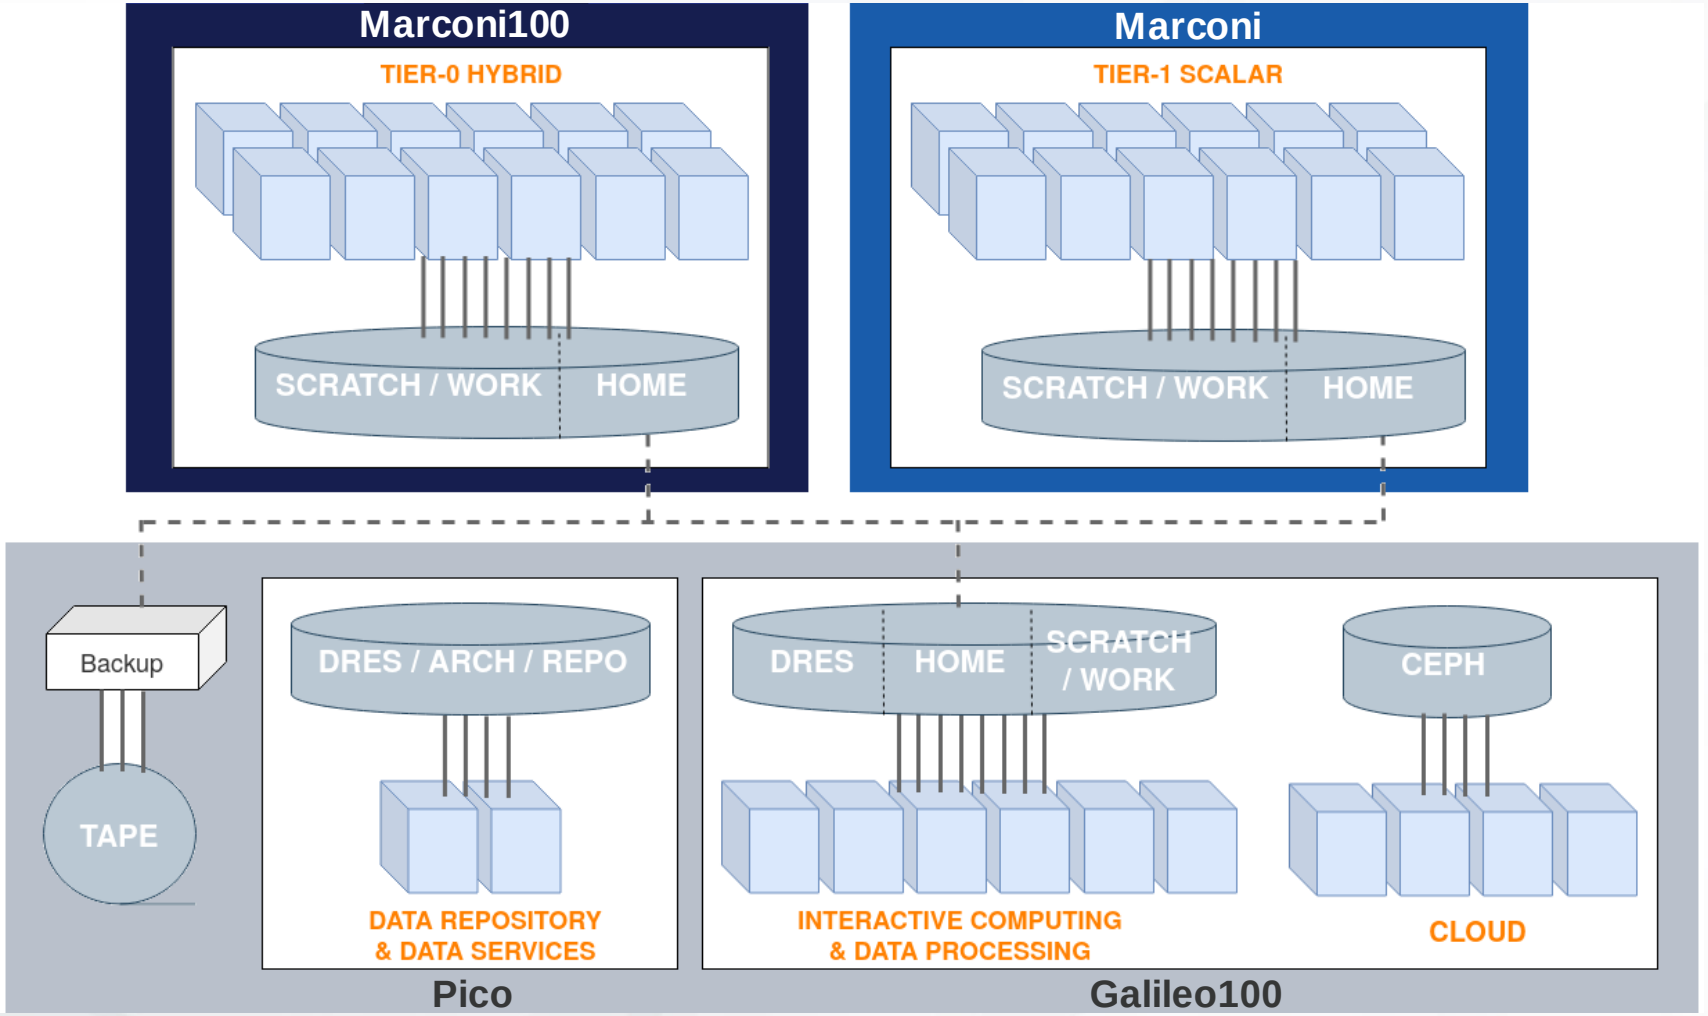

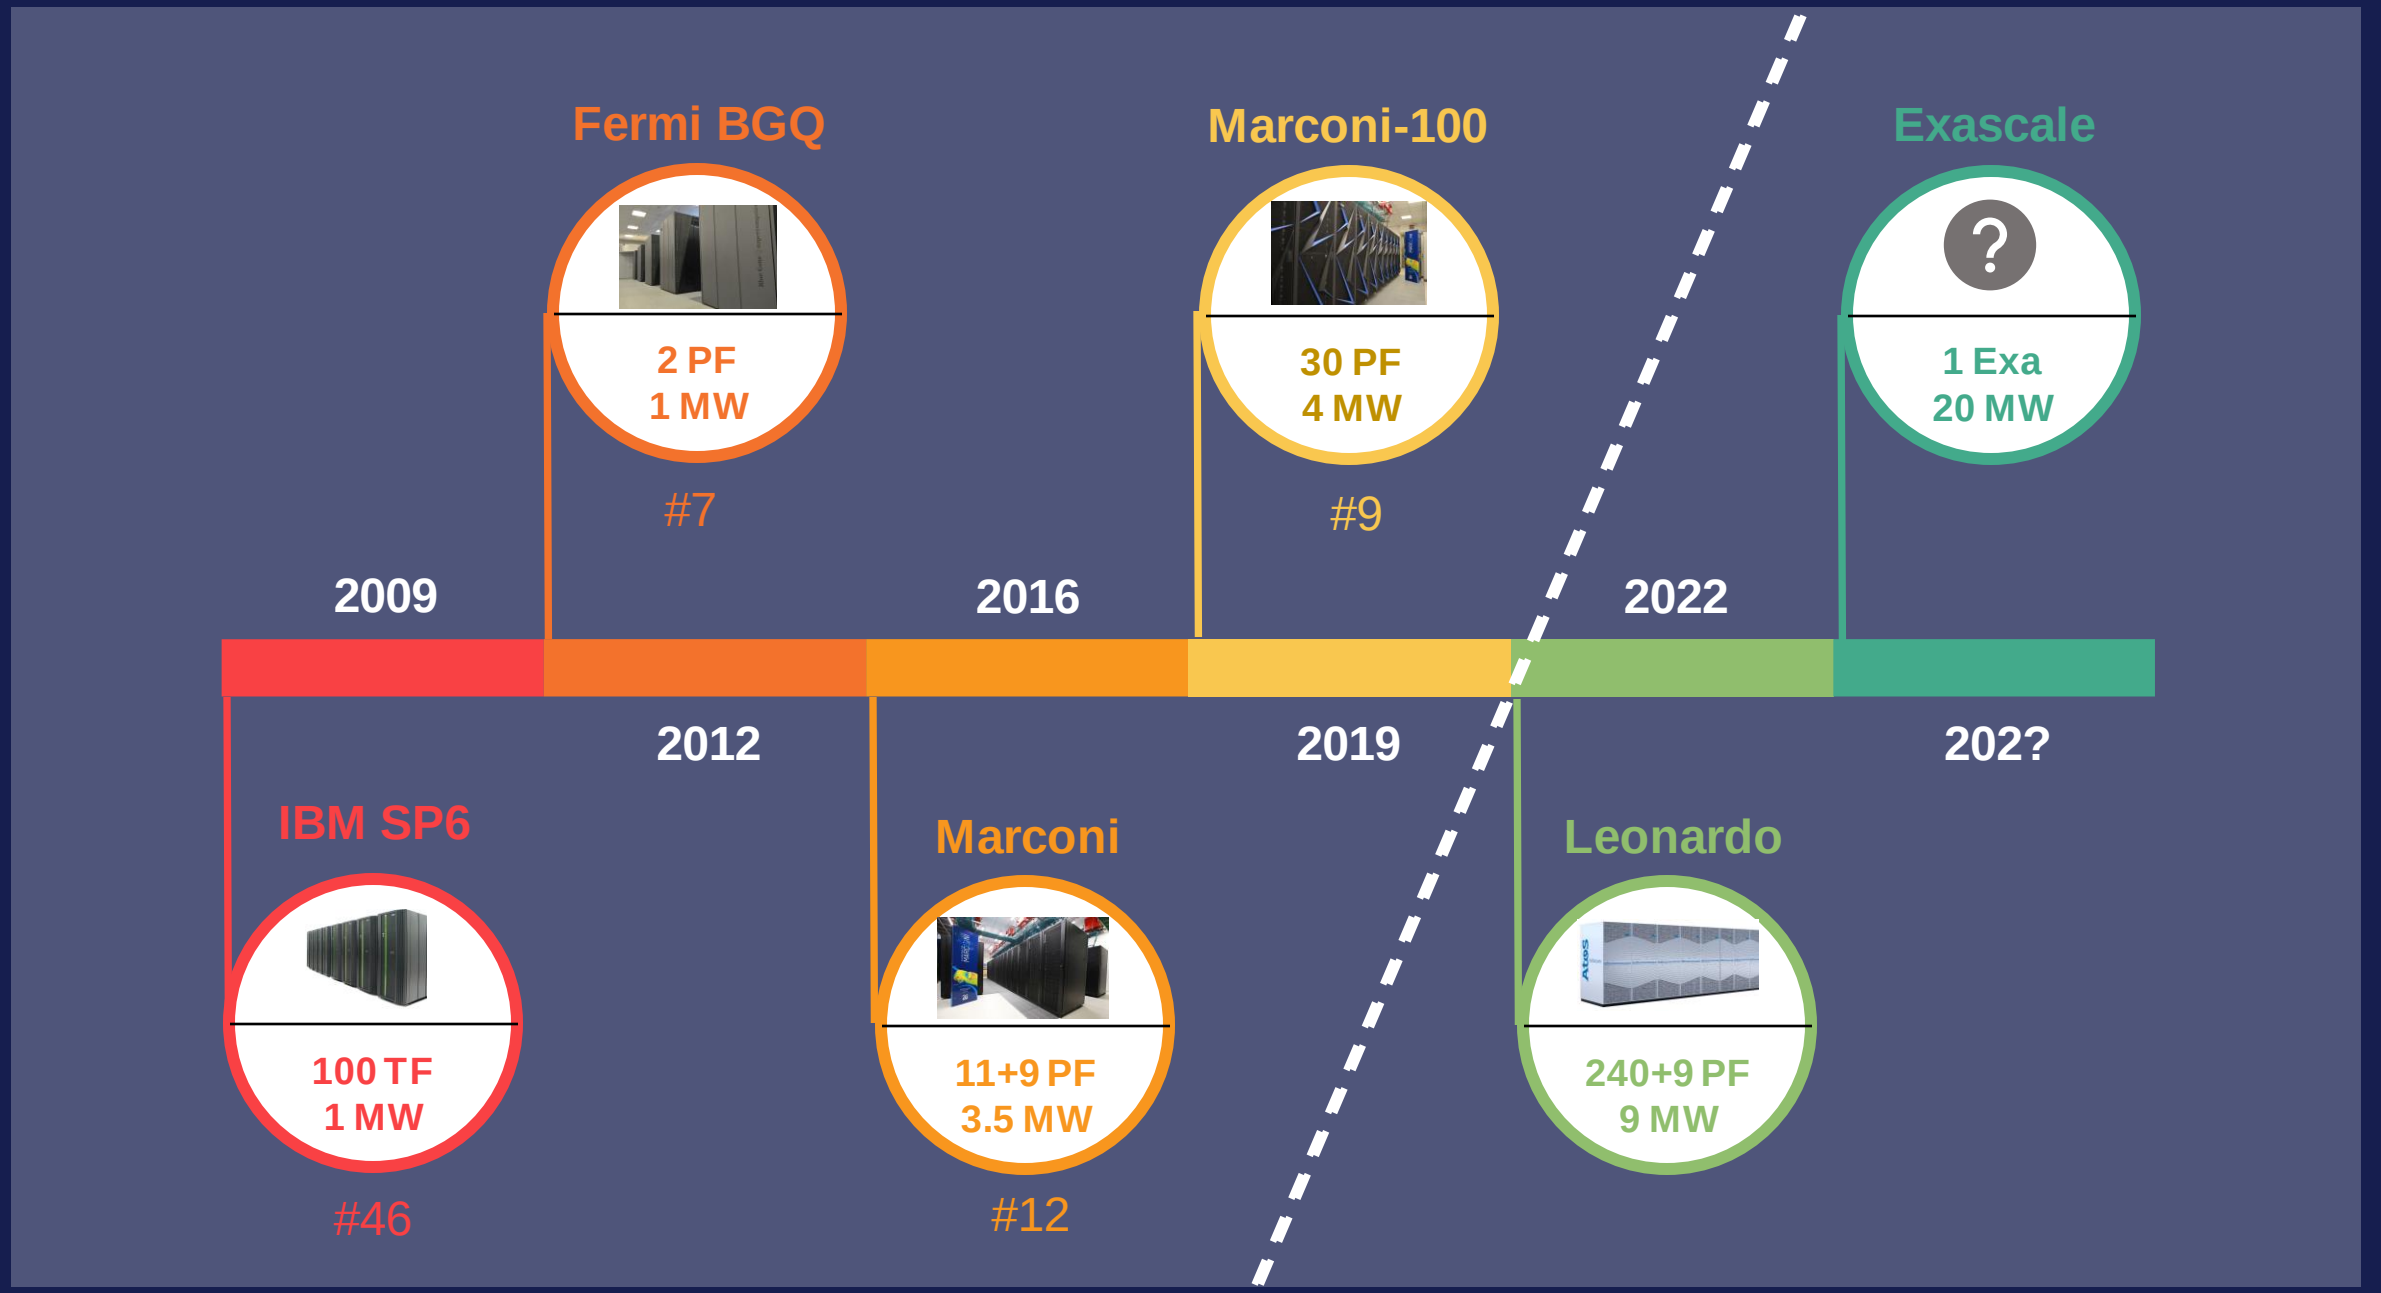
Fermi BGQ
2 PF
1 MW
#7
Marconi-100
30 PF
4 MW
#9
Exascale
?
1 Exa
20 MW
2009
2016
2022
2012
2019
202?
IBM SP6
100 TF
1 MW
#46
Marconi
11+9 PF
3.5 MW
#12
Leonardo
240+9 PF
9 MW
DC Casalecchio
DC Technopole

Cineca DC connection

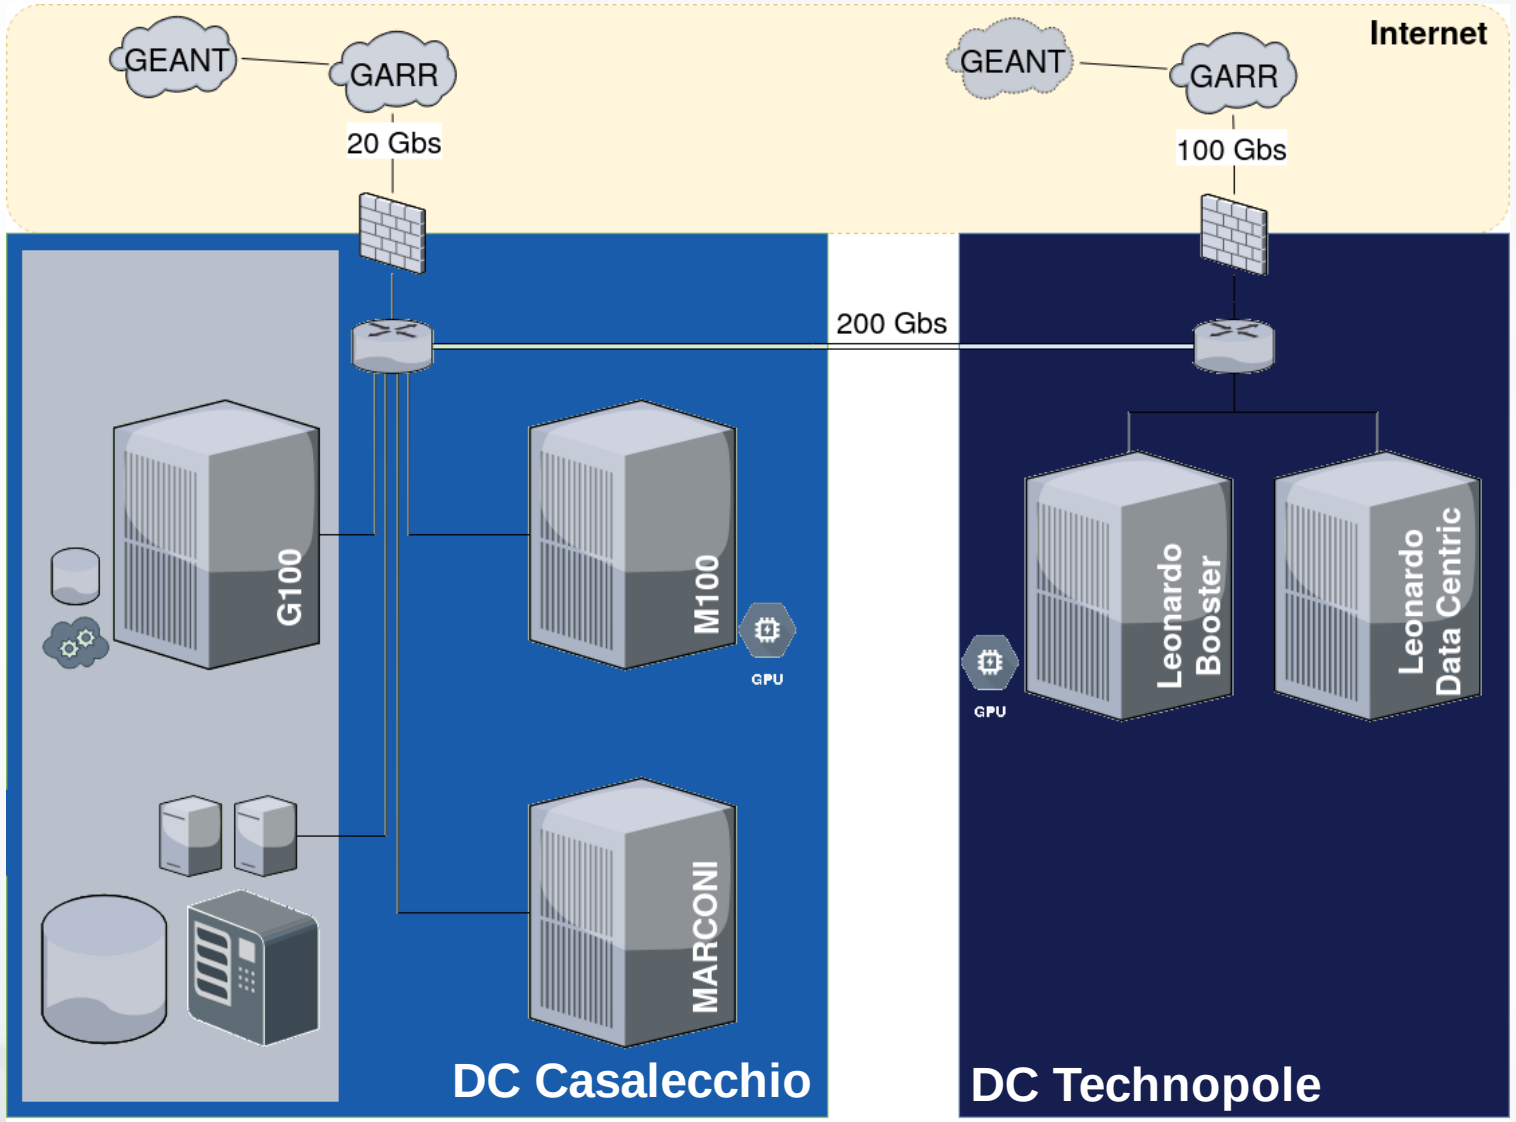

Leonardo configuration

LEONARDO CINECA

Leonardo main features

- Based on **Atos XH2000** platform technology
- **Warm water**: Inlet temperature of 36 degrees, outlet 46 degrees
- 95% of power dissipated via DLC

- **Bull Smart Energy management** suite
  - Bull Energy Optimizer
  - Bull Dynamic Power Optimizer

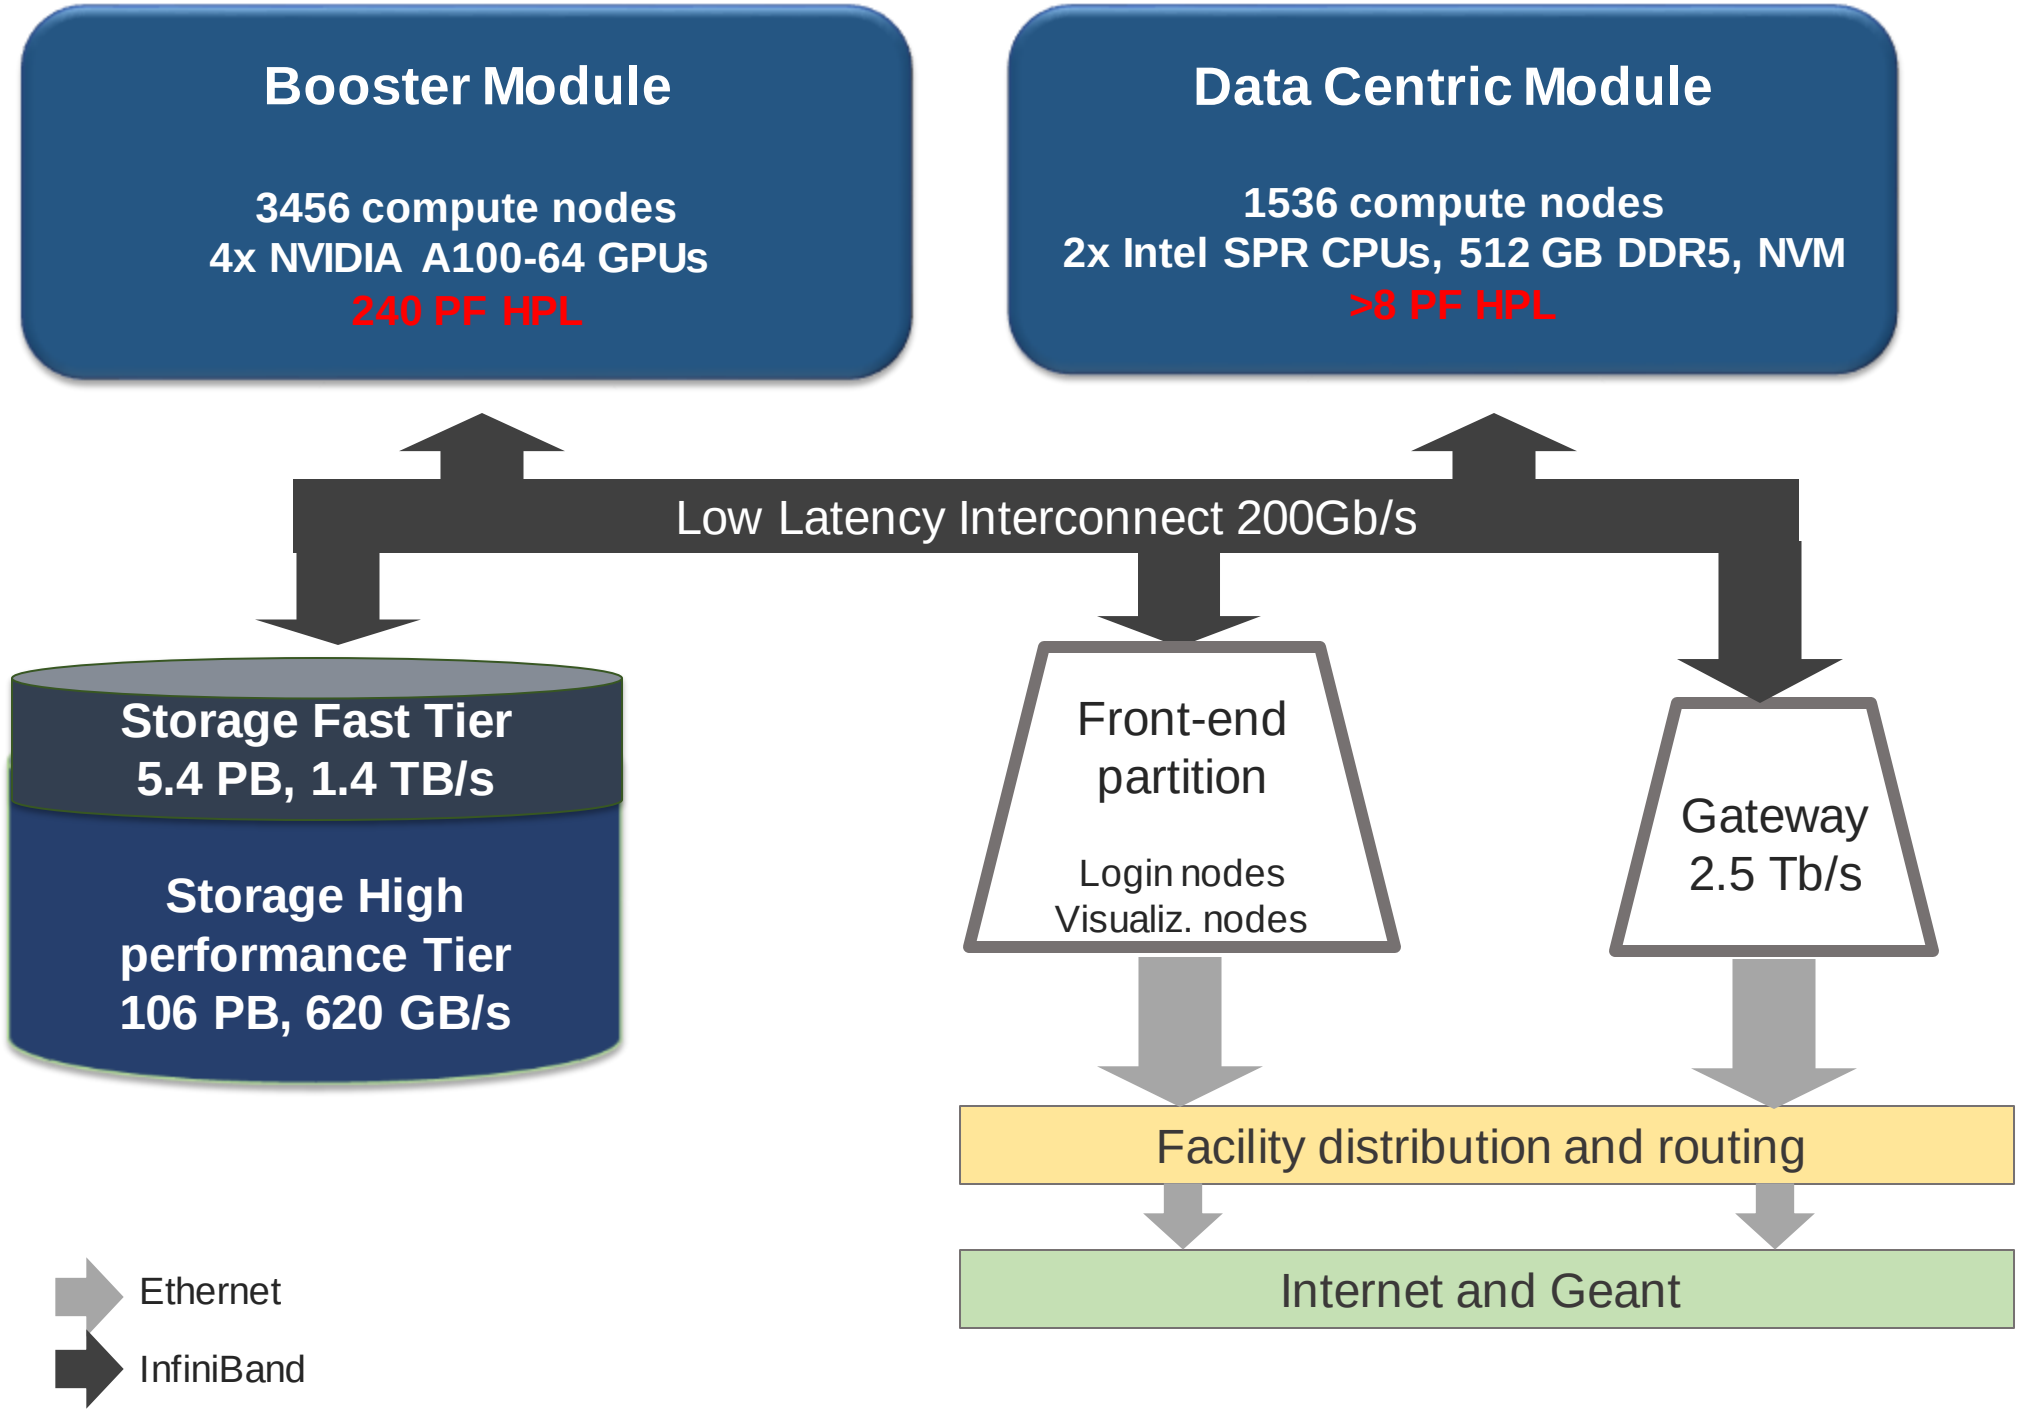
Booster Module
3456 compute nodes
4x NVIDIA A100-64 GPUs
240 PF HPL
Data Centric Module
1536 compute nodes
2x Intel SPR CPUs, 512 GB DDR5, NVM
>8 PF HPL
Low Latency Interconnect 200Gb/s
Storage Fast Tier
5.4 PB, 1.4 TB/s
Storage High performance Tier
106 PB, 620 GB/s
Front-end partition
Login nodes
Visualiz. nodes
Gateway
2.5 Tb/s
Facility distribution and routing
Internet and Geant
Ethernet
InfiniBand

Booster compute node

BullSequana X2135 “Da Vinci” single-node GPU Blade

- 1 x CPU Intel Xeon 8358 32 cores, 2,6 GHz
- 8 x 64 GB RAM DDR4 3200 MHz
- 4 x NVidia custom Ampere GPU 64GB HBM2
- 2 x NVidia HDR cards 2x100 Gb/s

Performance per node: 89,4 TFLOPs peak

Accelerated nodes comparison

	Marconi100	Leonardo-Booster
CPU	2 POWER9	1 IceLake
Cores	32 (16 per P9)	32 (2.4 GHz – 250 W)
Memory	256 GB	512 GB
CPUs : Accelerators	2:4	1:4
Accelerators	4 Volta V100	4 Ampere based GPU
GPU-GPU bandwidth	150 GB/s	400 GB/s
Accelerator DP Flops	28 TF (4x7)	NDA
Accelerator Memory	64 GB HBM2	256 GB HBM2e (4 x 64 GB)
Accelerator Memory Bandwidth	3.6 TB/s (900 GB/s x 4 GPUs)	6.5 TB/s (1.6 TB/s x 4 GPUs)
SSD Capacity	1.6 TB	-

Booster compute node

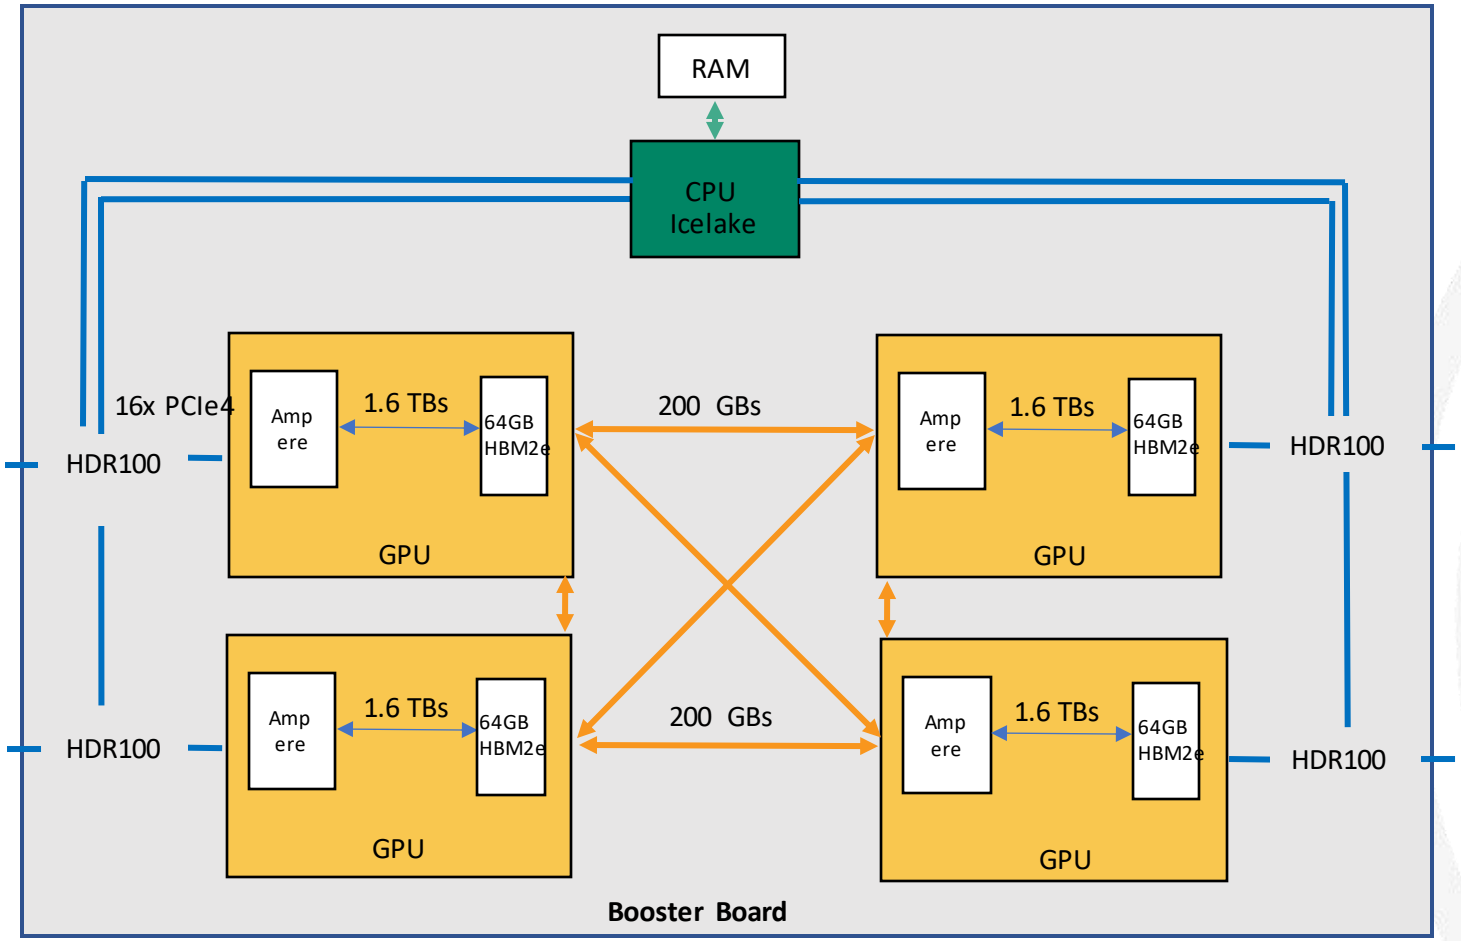

- ad hoc board from ATOS
- 4 NVidia Ampere GPUs (SXM)
- CPU-GPU connection via PCIe4 16x connection through HDR Connect6 HCA
  - PCI passthrough
  - 16 PCI links towards CPU, 16 links towards GPU
  - Bandwidth: 64 GBs duplex
- Full NVLink GPU-GPU connection
  - 200 GB/s bi-directional
- No PCI switch between host and external network
  - Low latency
- Out-of-band telemetry information
- GPUdirect

Data centric node

- Based on BullSequana X2610 compute blade
- 2x Intel Sapphire Rapids
  - >50 cores
  - 350 W
- 512 GB DC
  - DIMM 32GB DDR5 4800 MTS
  - Bandwidth 250 GB/s x socket
- 3.8 TB NVM

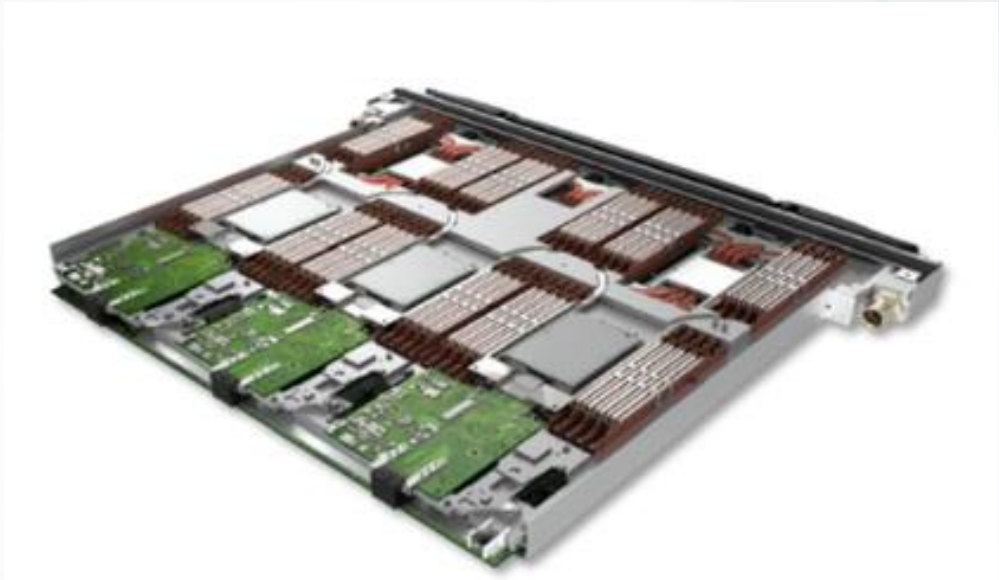

Network Topology

Based on Mellanox Networking Infiniband HDR hardware for switches, cables and NICs

Dragonfly+ topology

- All nodes are divided into cells
- Non-blocking, two-layer Fat Tree within the cells
- All to all connection between cells
- Mellanox QM8700 40-ports switches
- NIC Connect-X6

19 cells for Booster Module nodes

1 I/O cell

2 Data-Centric & General Purpose Cells

1 Hybrid cell, made of Booster and Data-Centric & General Purpose nodes

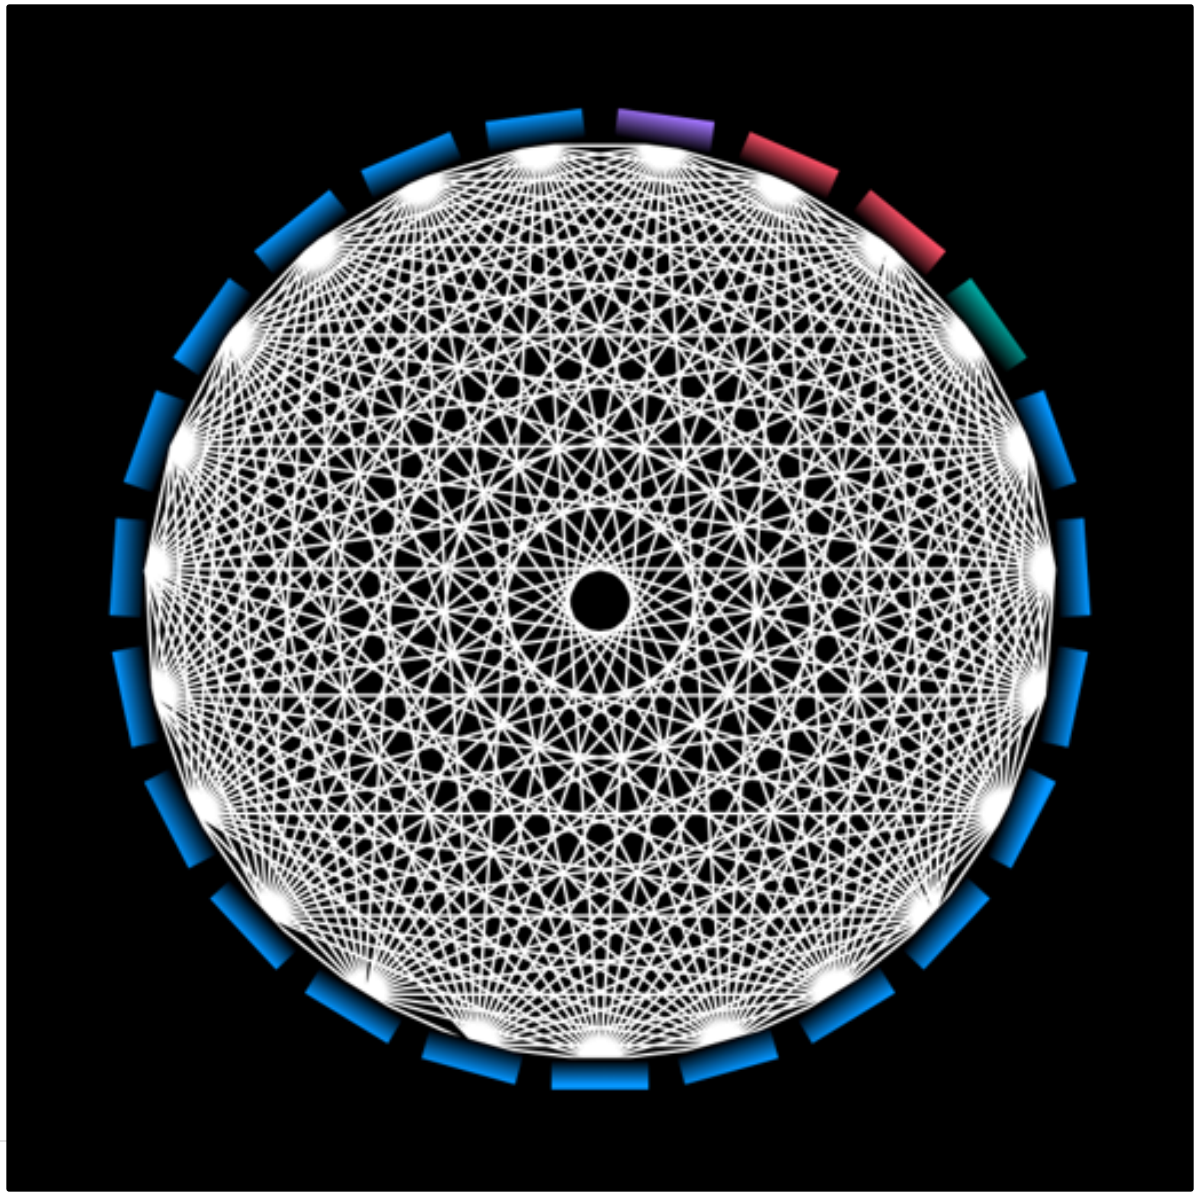

IO partition

- 31 x DDN Exascaler ES400NVX2 constitute I/O performance layer
  - 24 x 7,68 TB SSD NVMe with encryption support
  - 4 x InfiniBand HDR ports

Capacity Tier

- 31 x DDN EXAScaler SFA799X appliances for HDD storage:
- Controller node: 82 x 18 TB HDD SAS 7200 rpm and 4 x HDR100 ports
- 2 x JBOD expansion per controller, each with 82 x 18 TB HDD SAS 7200 rpm

4 x DDN EXAScaler SFA400NVX appliances for metadata

- 21 x 7,68 TB SSD NVMe with encryption support
- 8 x InfiniBand HDR100 port

Storage Fast Tier	
Net capacity	5.4 PB
Disk technology	full flash (NVMe and SSD)
Bandwidth:	Aggregated: 1400 GB/s r/w io500[6]: 676

Storage Capacity Tier	
Net capacity	106 PB
Disk technology	NVMe and HDD
Bandwidth:	Aggregated: read performance of 744GB/s and write performance of 620GB/s io500[8]: 197 GiB/s

Status

Leonardo roadmap

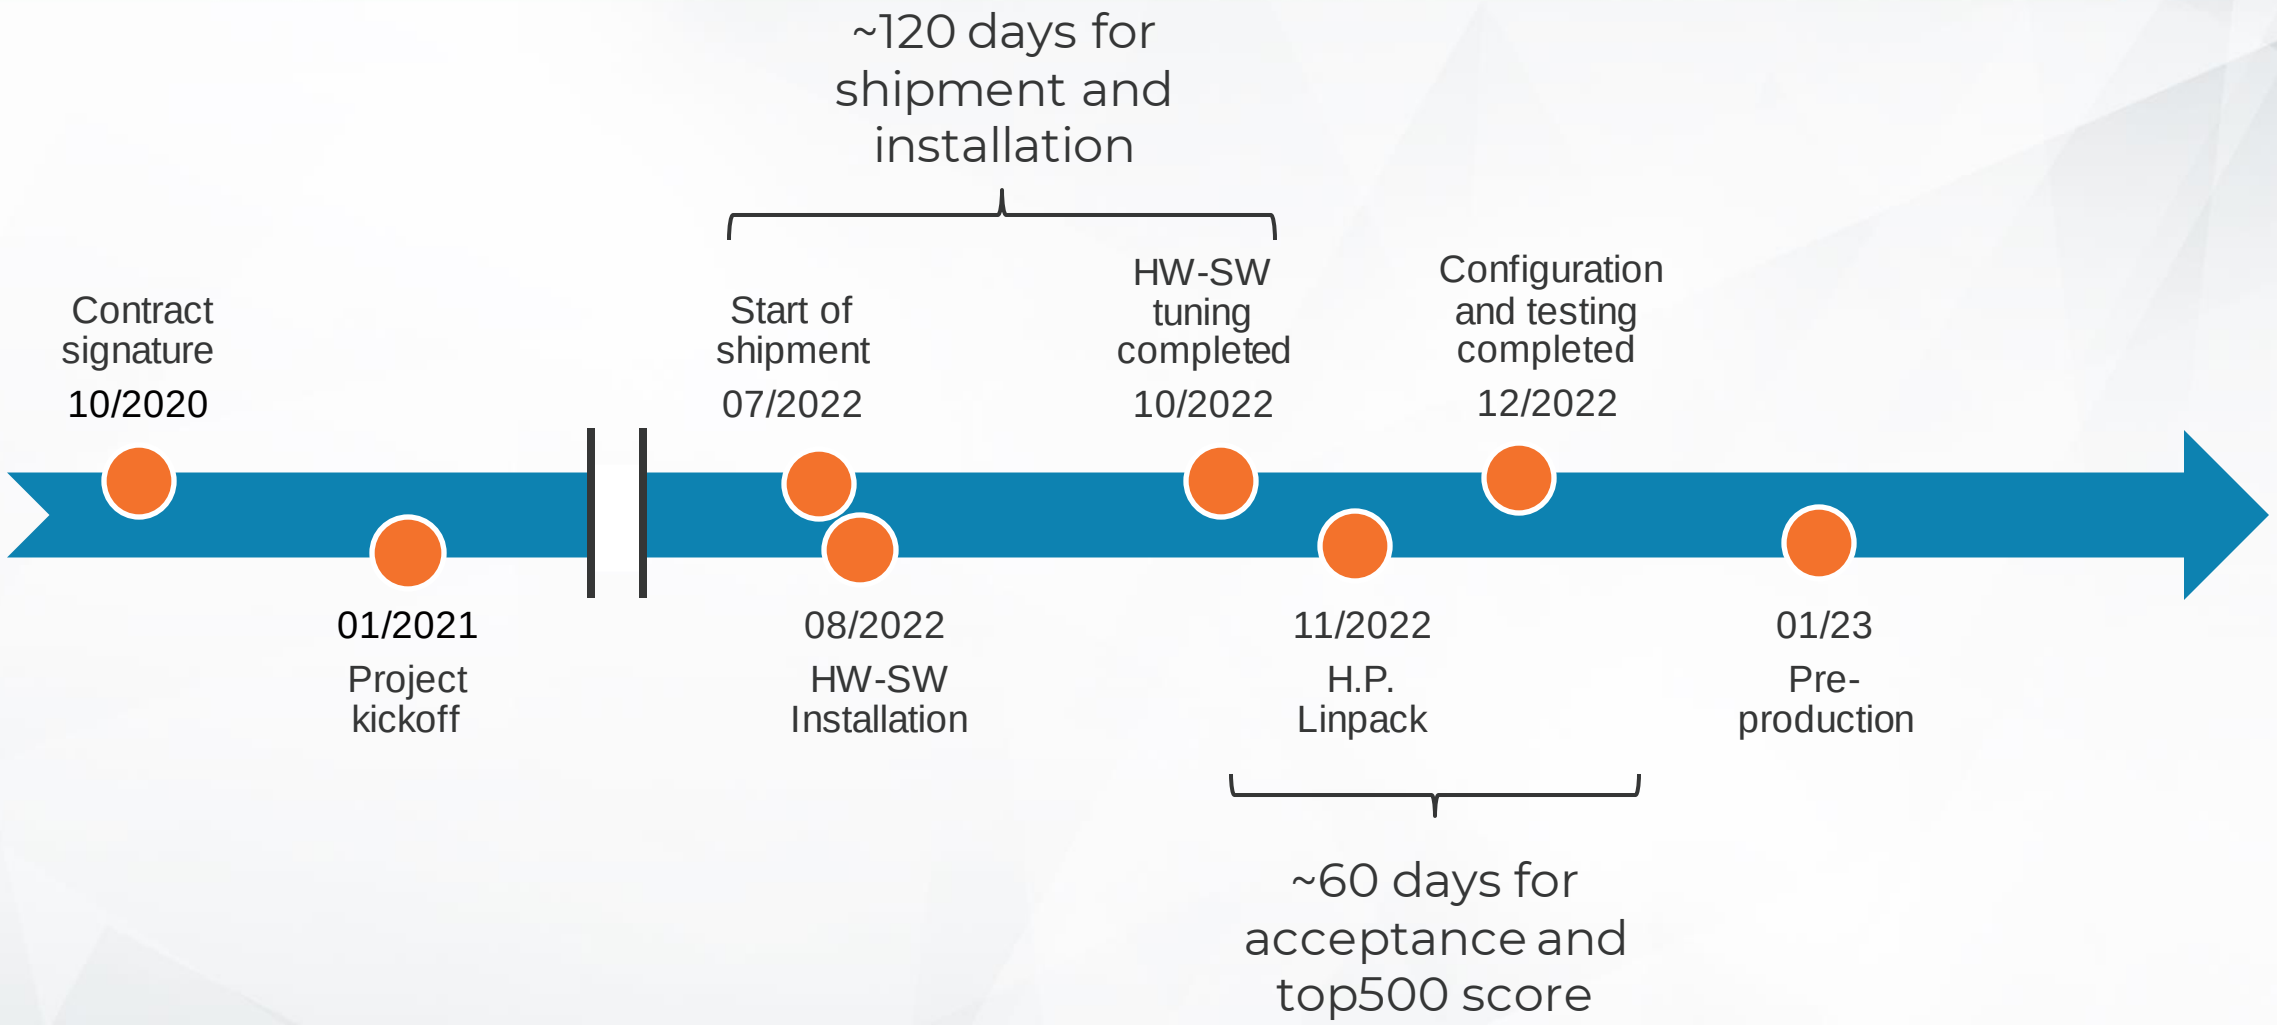

TECNOPOLO

CINECA / INFN

Capannone Miscela C2 - LEONARDO

Botte B5 - INFN Data center

Ballette building

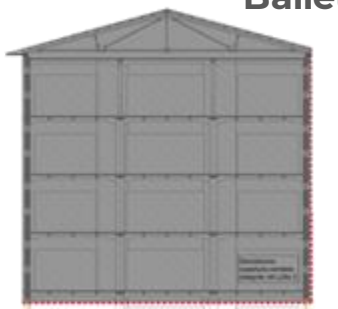

Technological center G1

Data Hall Leonardo

Capannone Miscela C2

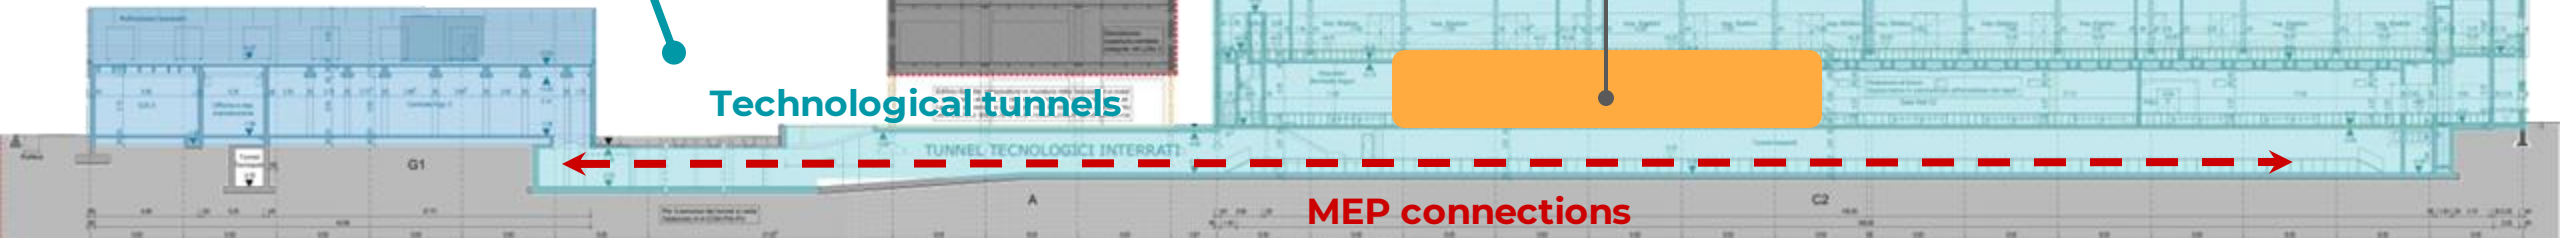

Technological tunnels

MEP connections

C2
Section

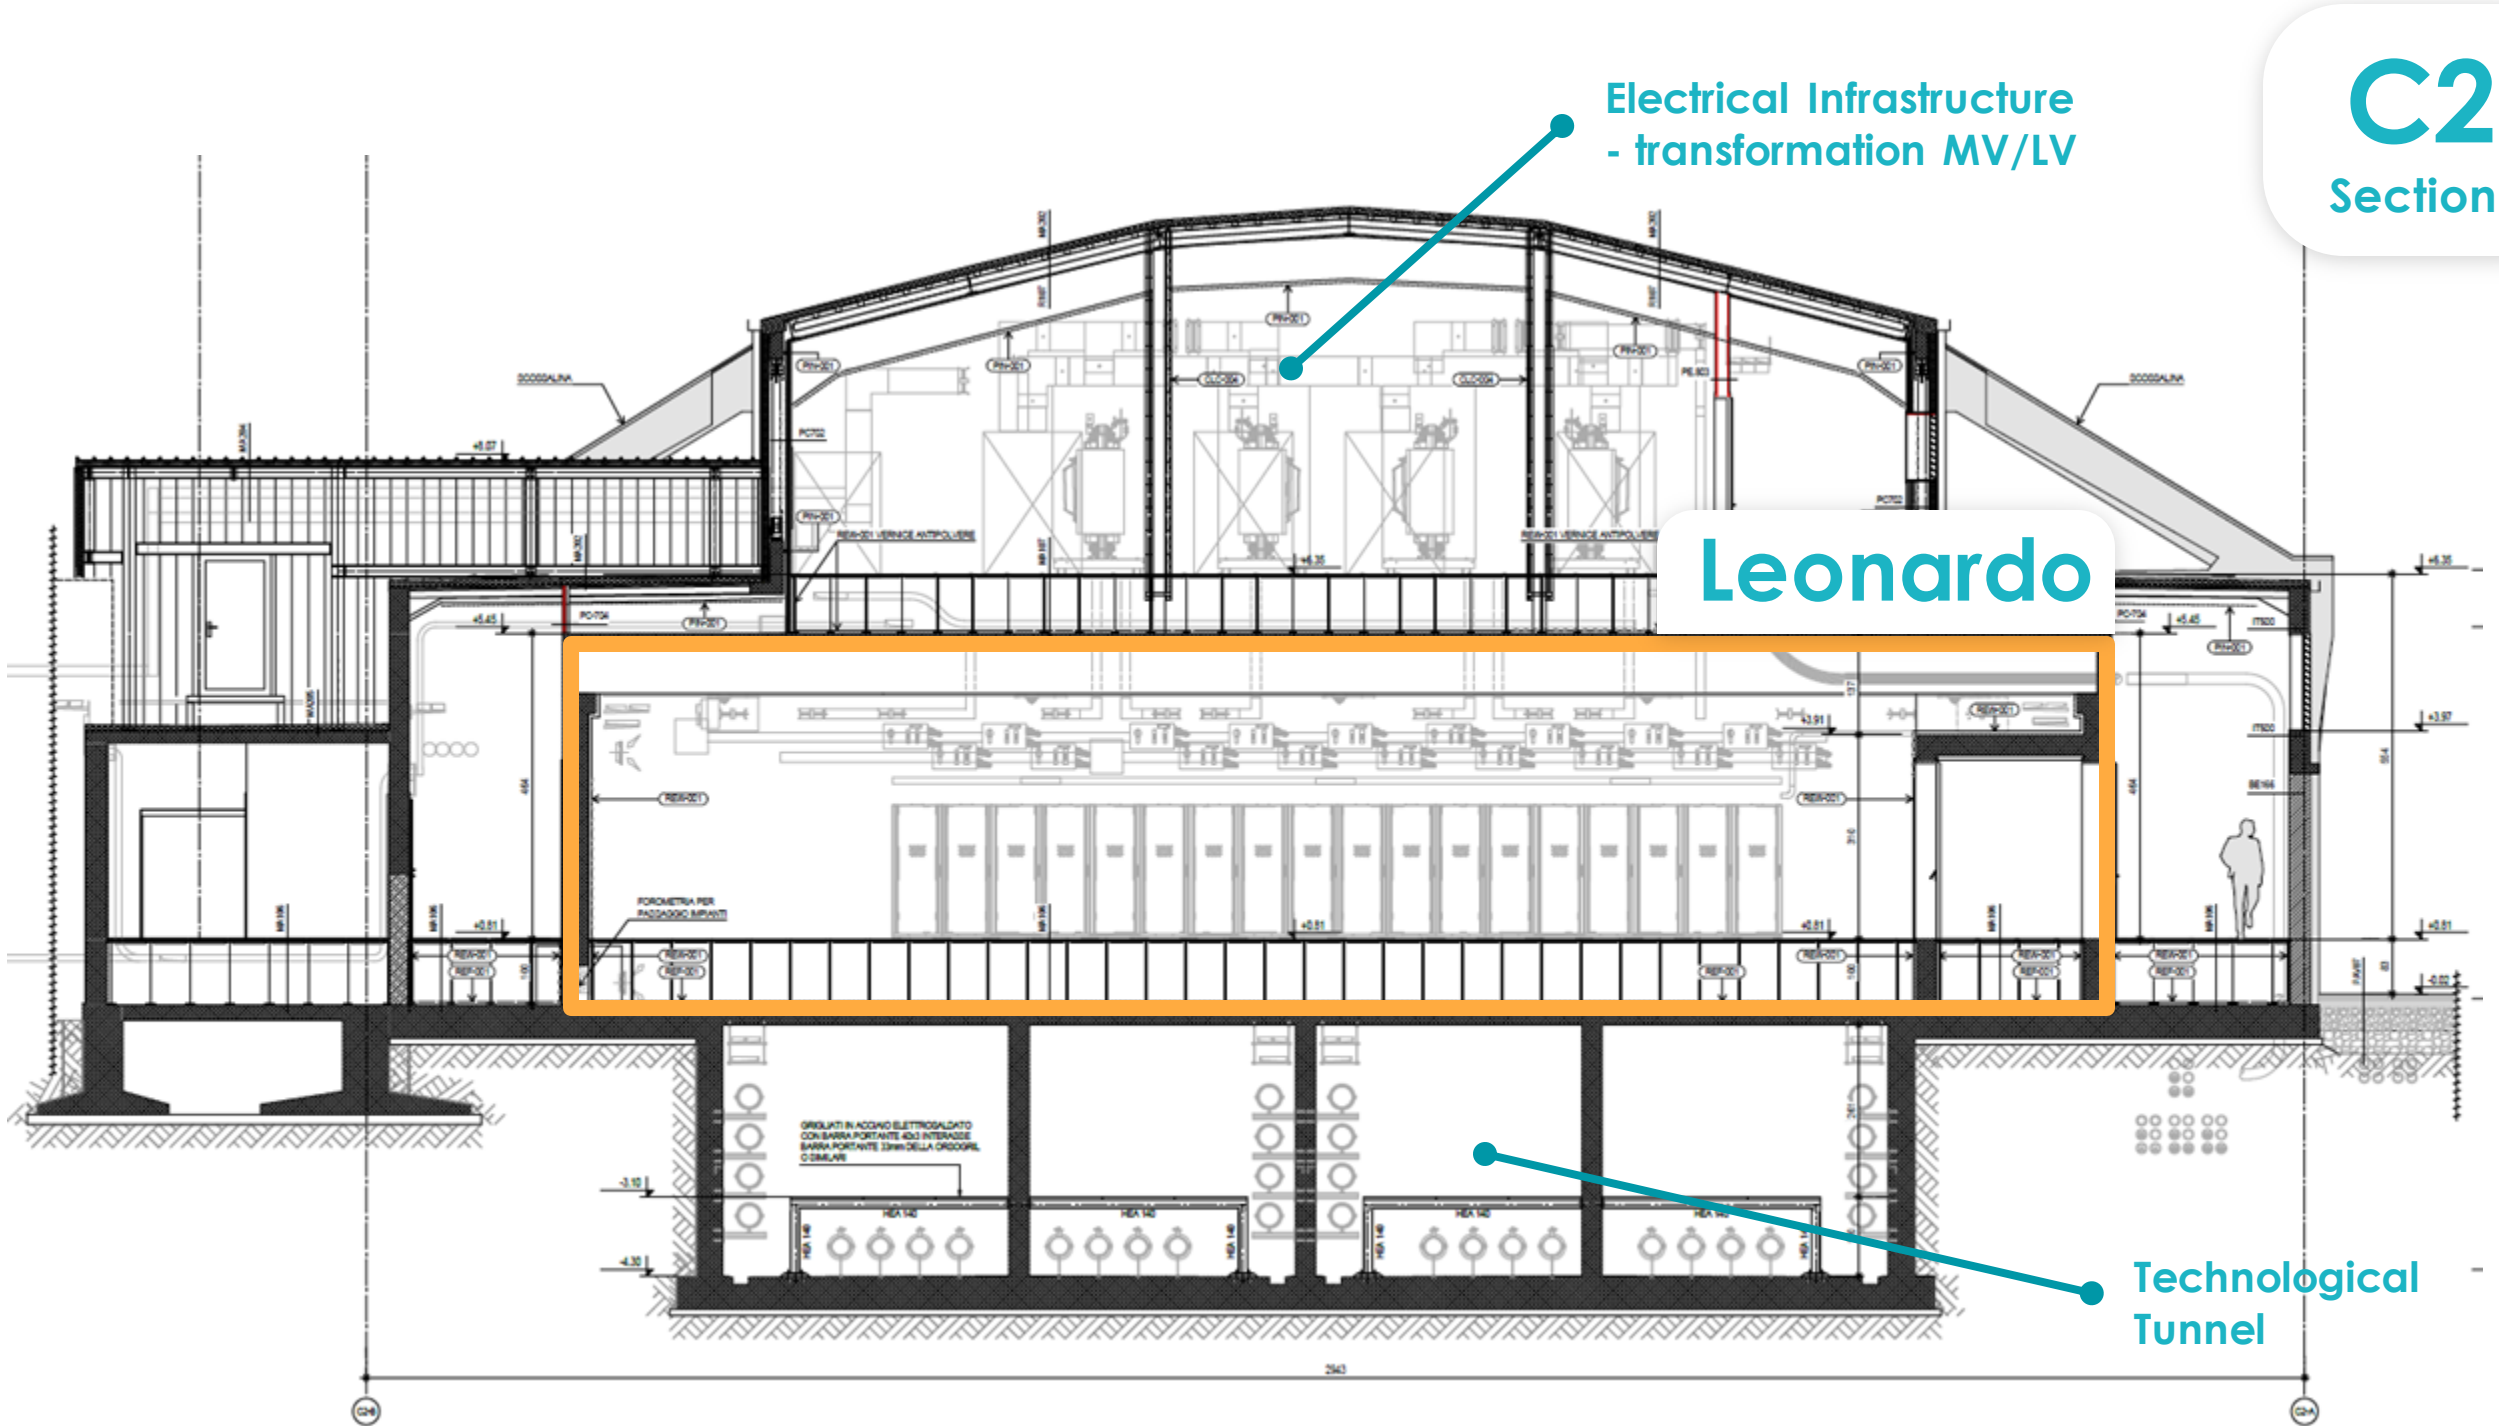
Electrical Infrastructure
- transformation MV/LV
Leonardo
Technological
Tunnel

TECNOPOLO

Buildings

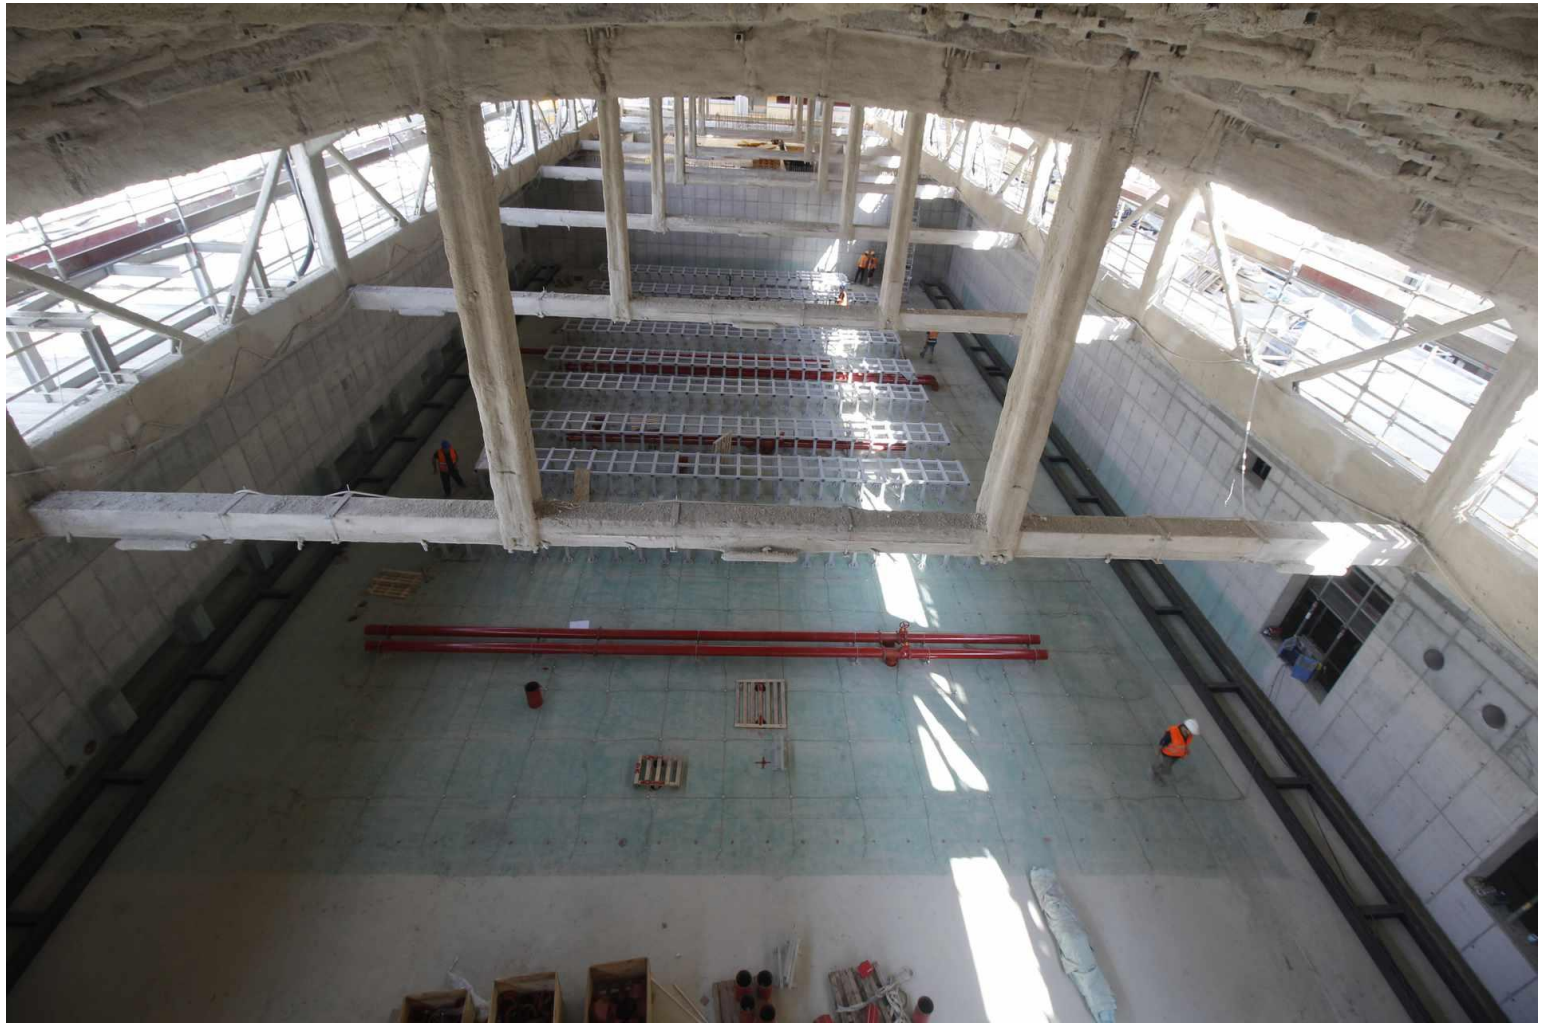

CINECA

MUR
Ministero dell'Università e della Ricerca

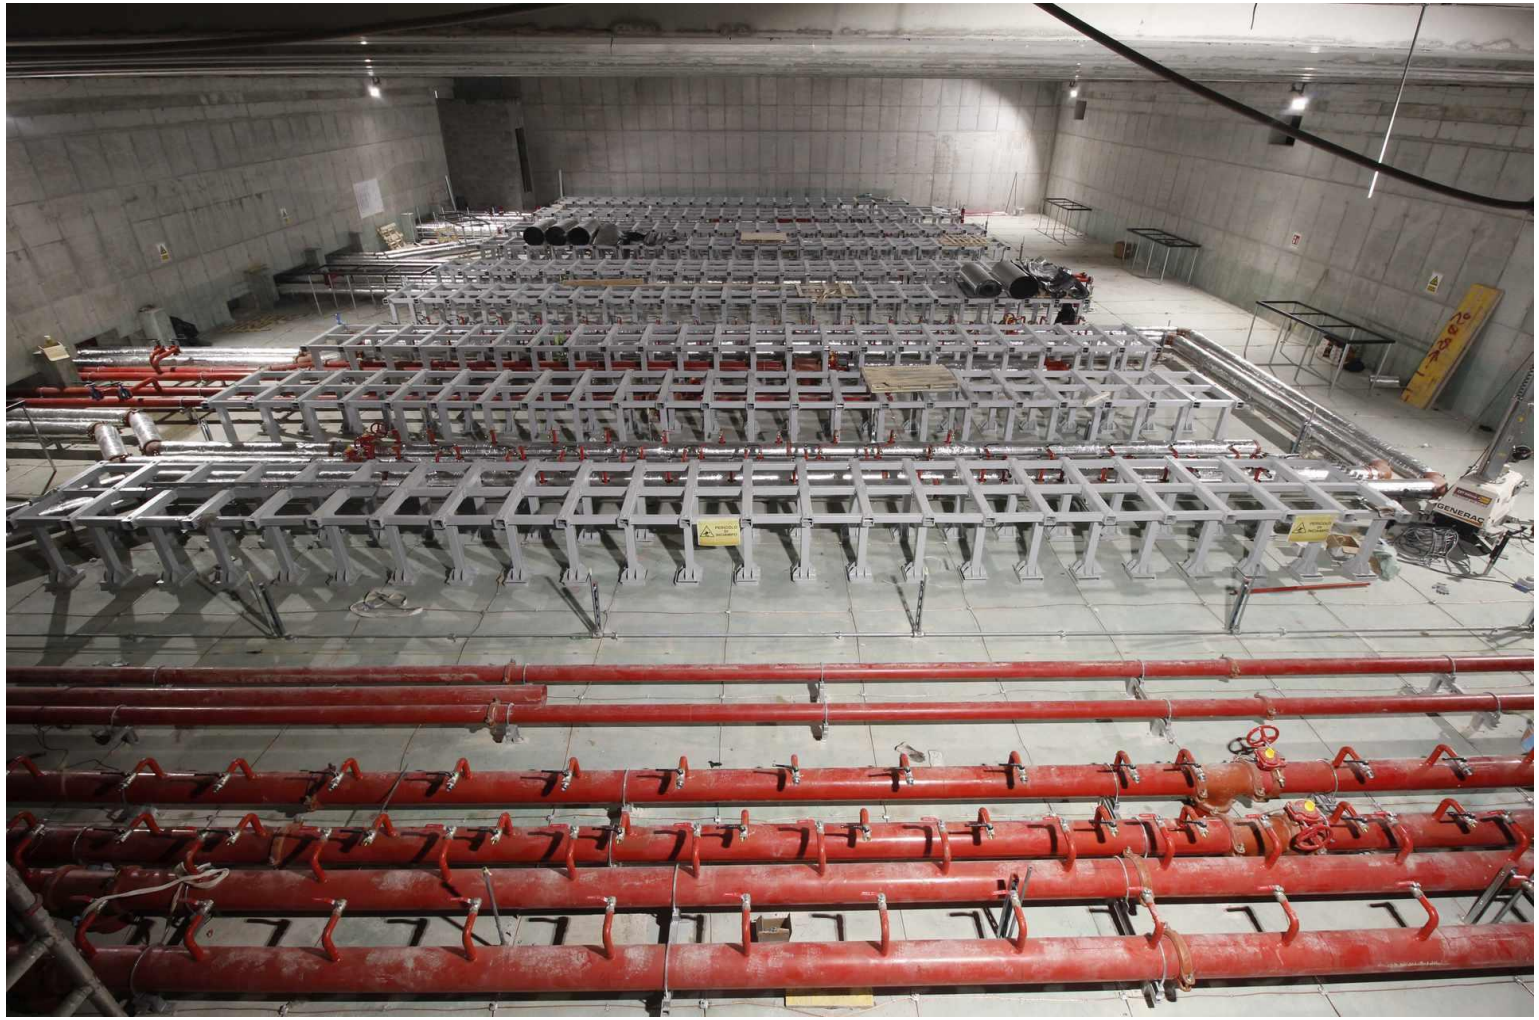

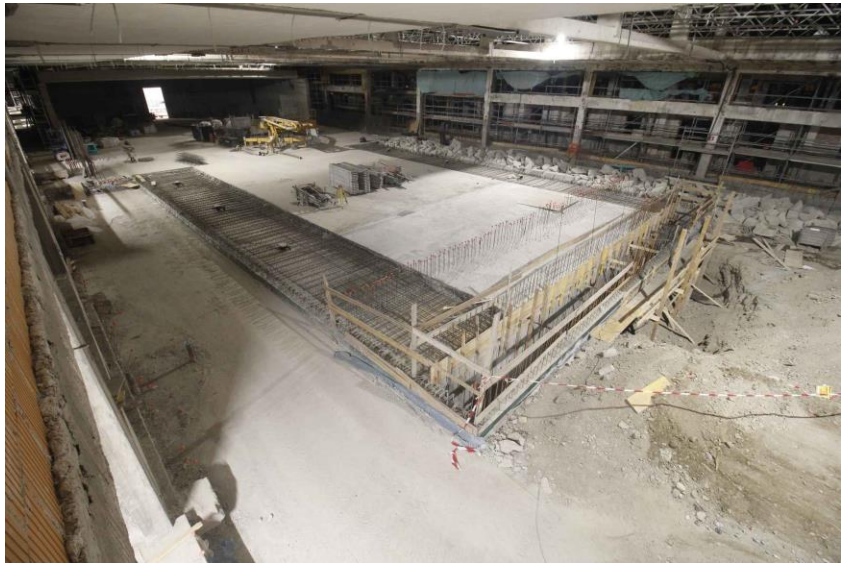

CINECA

MUR
Ministero dell'Università e della Ricerca


Atos
LEONARDO
CINECA

Bologna is the place...
we like all kind of high performance

			
Clear user interaction	yes	yes	yes
More is better	yes	yes	yes
State-of-the-art engine	yes	yes	yes
Can become really hot	yes	yes	yes
Can you enjoy it alone	yes	yes, but ask!	yes

EuroHPC Joint Undertaking

CINECA

MUR
Ministero dell'Università e della Ricerca

INFN
Istituto Nazionale di Fisica Nucleare

SISSA

MINISTRY OF EDUCATION, SCIENCE, RESEARCH AND SPORT OF THE SLOVAK REPUBLIC

Federal Ministry
Republic of Austria
Education, Science and Research

MINISTRY OF INNOVATION AND TECHNOLOGY

REPUBLIKA SLOVENIJA
REPUBLIC OF SLOVENIA
Ministrstvo za izobraževanje, znanost in šport
Ministry of education, science and sport

HELLENIC REPUBLIC
Ministry of Digital Governance

Thanks for coming to Bologna