



The new INFN Data Center at Bologna Tecnopolo

Luca dell'Agnello – INFN-CNAF; Tommaso Boccali – INFN-PISA;
Daniele Cesini – INFN-CNAF

ACAT 2022

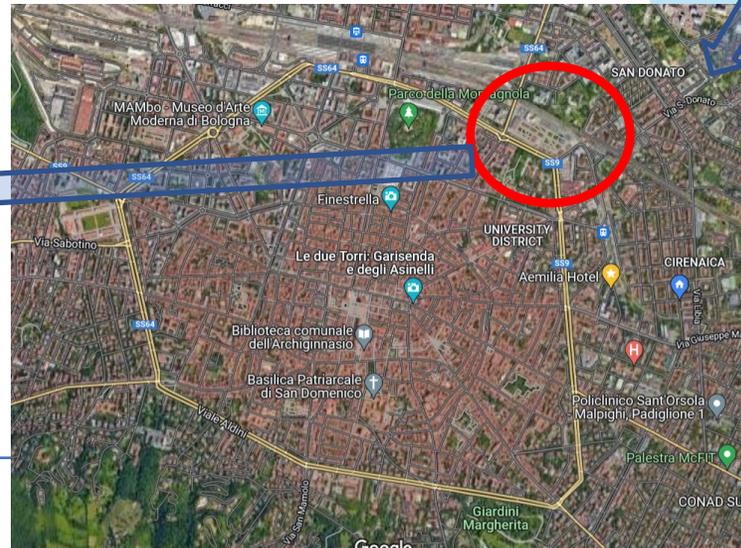
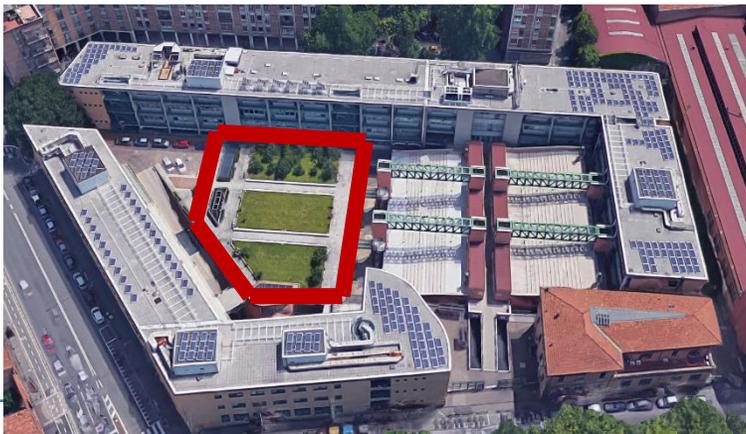
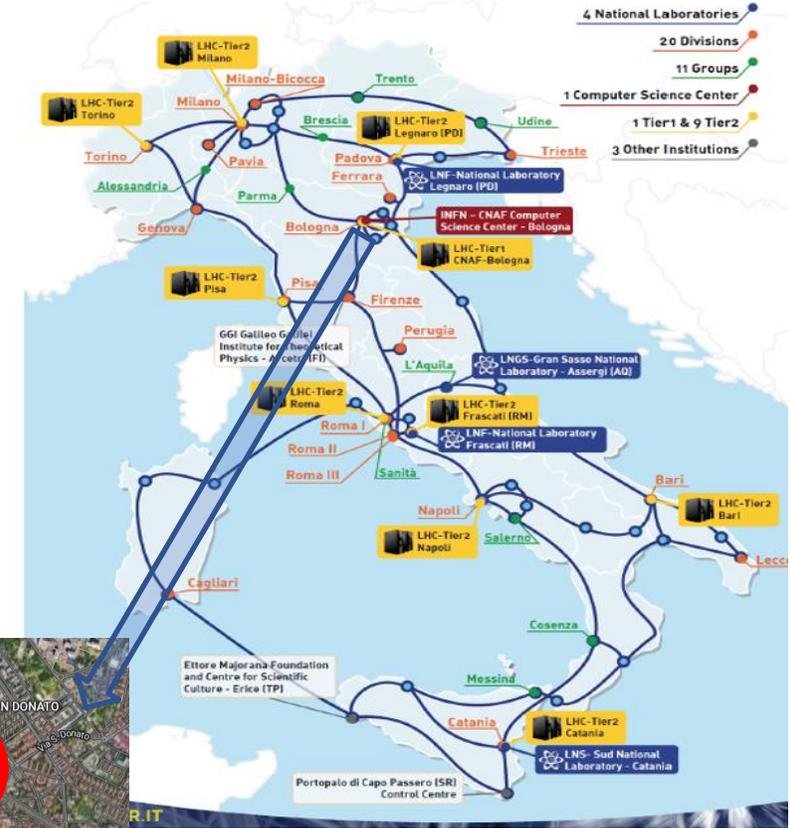
25 Oct 2022 – Bari (IT)



The INFN-CNAF



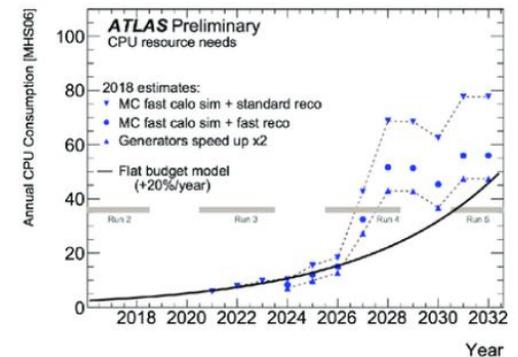
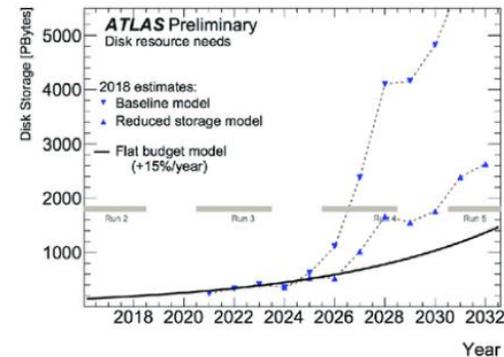
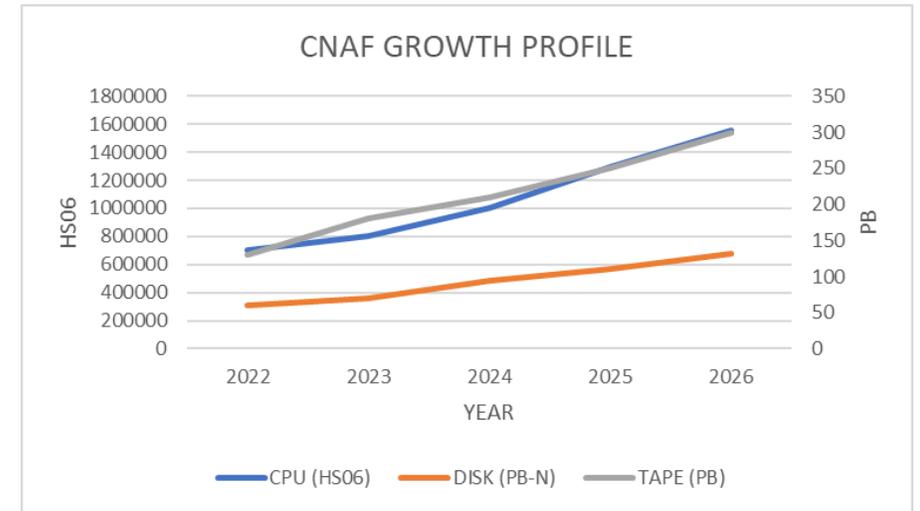
- INFN-CNAF: the national center of INFN (Italian Institute for Nuclear Physics) dedicated to Research and Development on Information and Communication Technologies
 - <https://www.cnaf.infn.it/en>
- CNAF hosts the Italian Tier-1 data center for the high-energy physics experiments at the Large Hadron Collider in Geneva
- CNAF is one of the most important centers for distributed computing in Italy
- It is currently located in close to the Bologna city center
 - Co-located with the Physics dept of the Bologna University



Scientific motivations

There are three main drivers for the transfer of CNAF resources to a new location:

- 1) an expected huge increase in IT resources to be installed in the coming years
 - Ready for HI-Luminosity LHC
- 2) the infrastructural problems at the current site
- 3) and the opportunity offered by the new location



A brand-new data center for CNAF

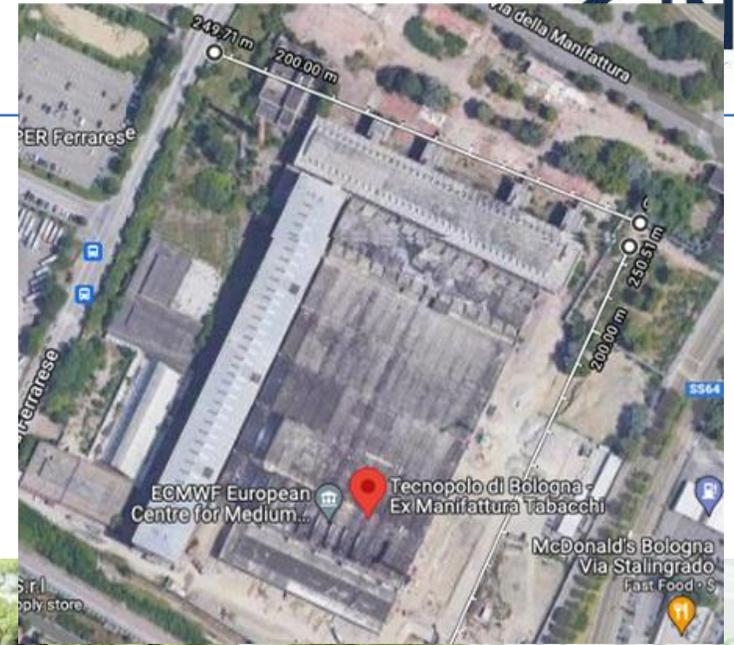


- Renew infrastructures to be ready for the HL-LHC era
 - up to ~ 2035 and beyond
- Use more compact computing
 - from today's ~ 20 kW/rack to 80 kW/rack DLC
 - Integration with CINECA-Leonardo Supercomputer
- Lower the PUE (*power usage effectiveness*)
 - Targeting 1.08-1.10
- Extend and expand networking for a future-proof infrastructure

The opportunities

- In **2017**, Bologna won a bid to host the datacenter of the “*European Centre for Medium-Range Weather Forecasts*” - ECMWF
- The Emilia Romagna region decided to repurpose the “*Manifattura Tabacchi*” area to host a technology district, hosting ECMWF and more

Roughly
250x250 m²

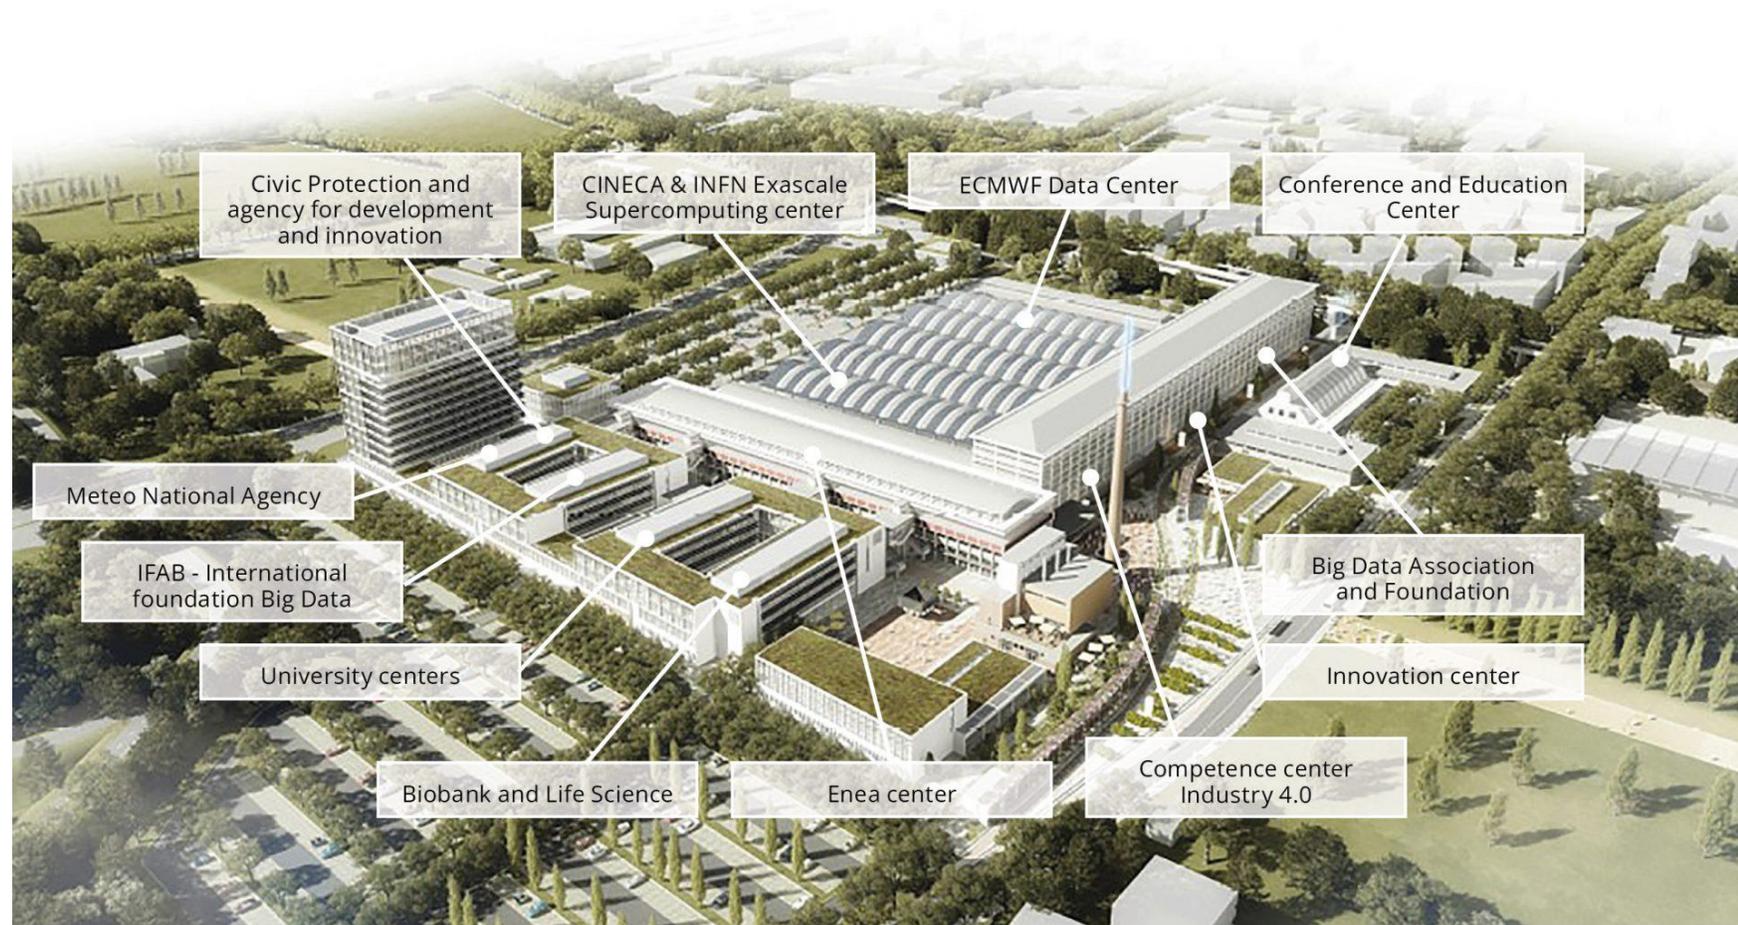


How it will be



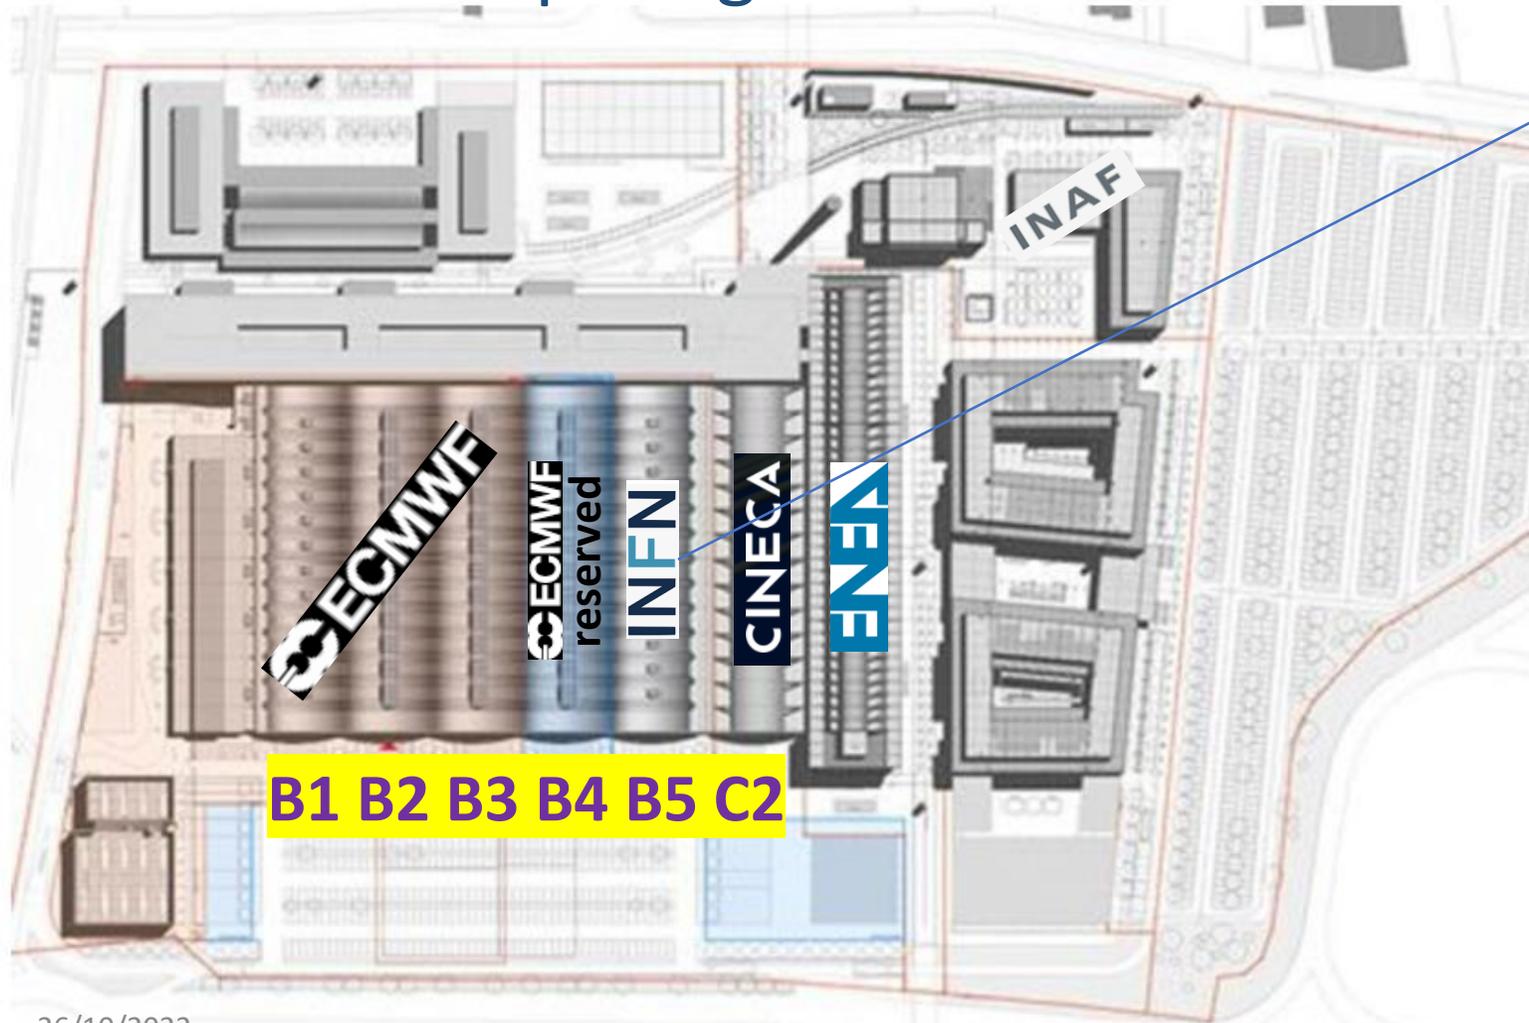
What can the Tecnopolo host?

- Not only research infrastructures and supercomputers
- Areas for
 - Technological institutions
 - university
 - innovation hubs
 - technology transfers
 - Industry 4.0
- Restaurants

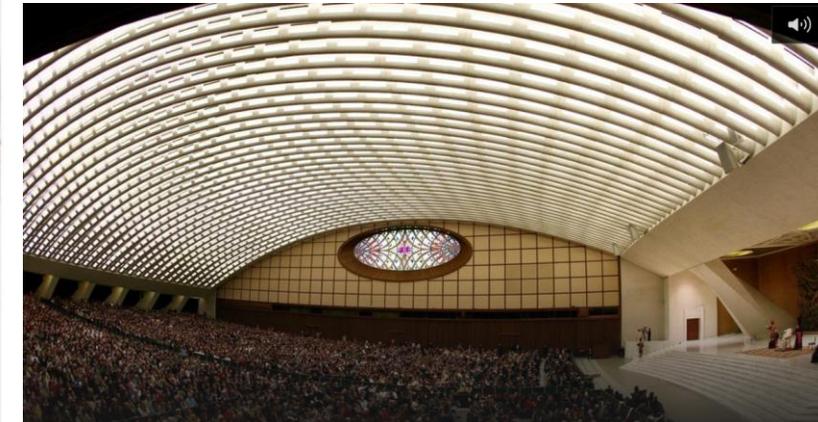


What can the Tecnopolo host?

The computing infrastructures



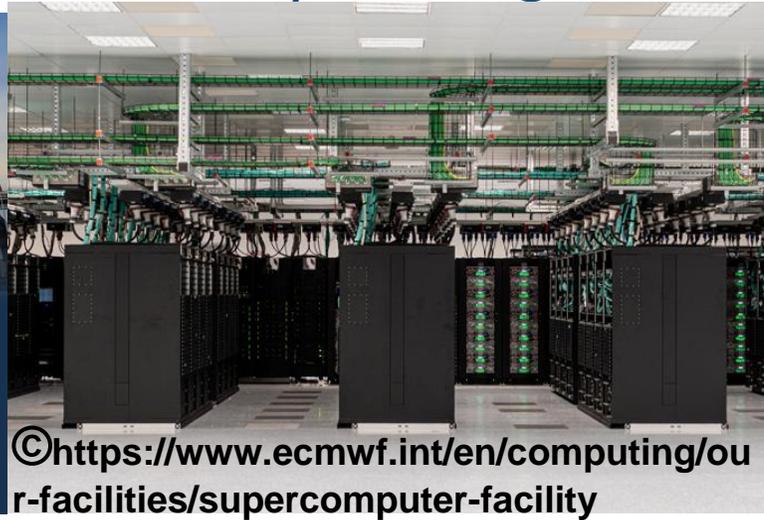
Each of the 6 “botti” (barrels) is
~5000m² of usable IT space



Same architect and design of the
“Sala Nervi” in the Vatican

The INFN+CINECA project

- The ECMWF is already running!



- CNAF “B5” Barrel expected to be ready by mid 2023
- Two phases expected
 - **Phase-1 (2023-2025)**
 - Leonardo + T1-CNAF → 13 MW
 - **Phase-2 (2025+)**
 - infrastructure up to 25 MW ready for post-exascale and for HL_LHC

- CINECA Leonardo is being commissioned in October 2022



Current status.....

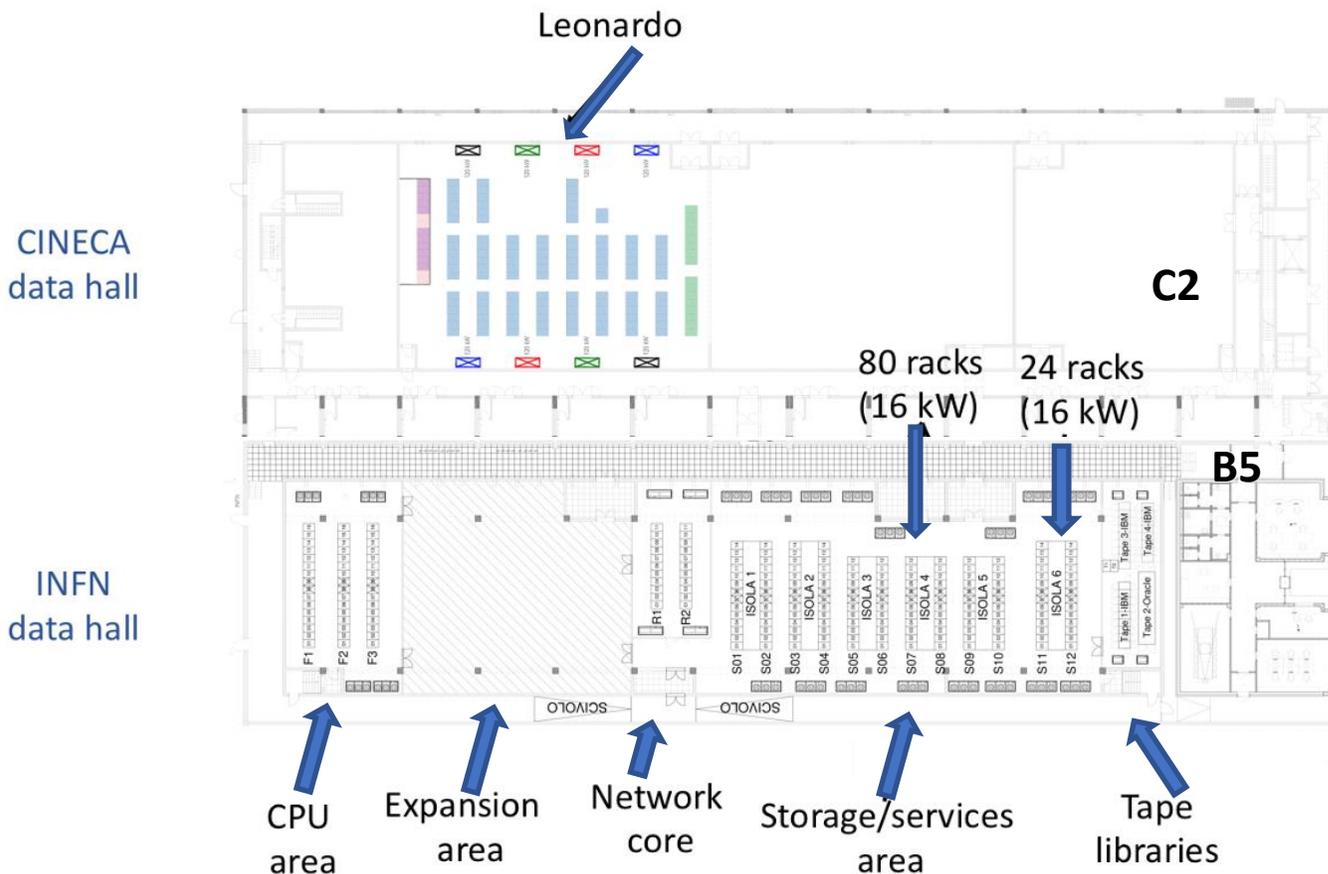


CNAF Barrel

Current status...



CNAF and CINECA data halls



- The new CNAF Datacenter will feature the following main areas
 - High Density – 2-3 rows for 80kW racks
 - Low density – 80+24 16kW racks
 - Expansion area
 - Tape libraries areas
 - Up to 4 libraries
- The CPU area can host up to 3MW of CPUs via 42 DLC high density racks
- The low-density area will be used to host
 - Storage systems
 - CNAF Cloud Infrastructures
 - ISO certified Cloud racks
- Cooling
 - Air cooled Cold Corridor aisles
 - Direct Liquid in High Density
- 3+1 redundancy in all the infrastructure facilities

DLC 80kW



26/10/2022



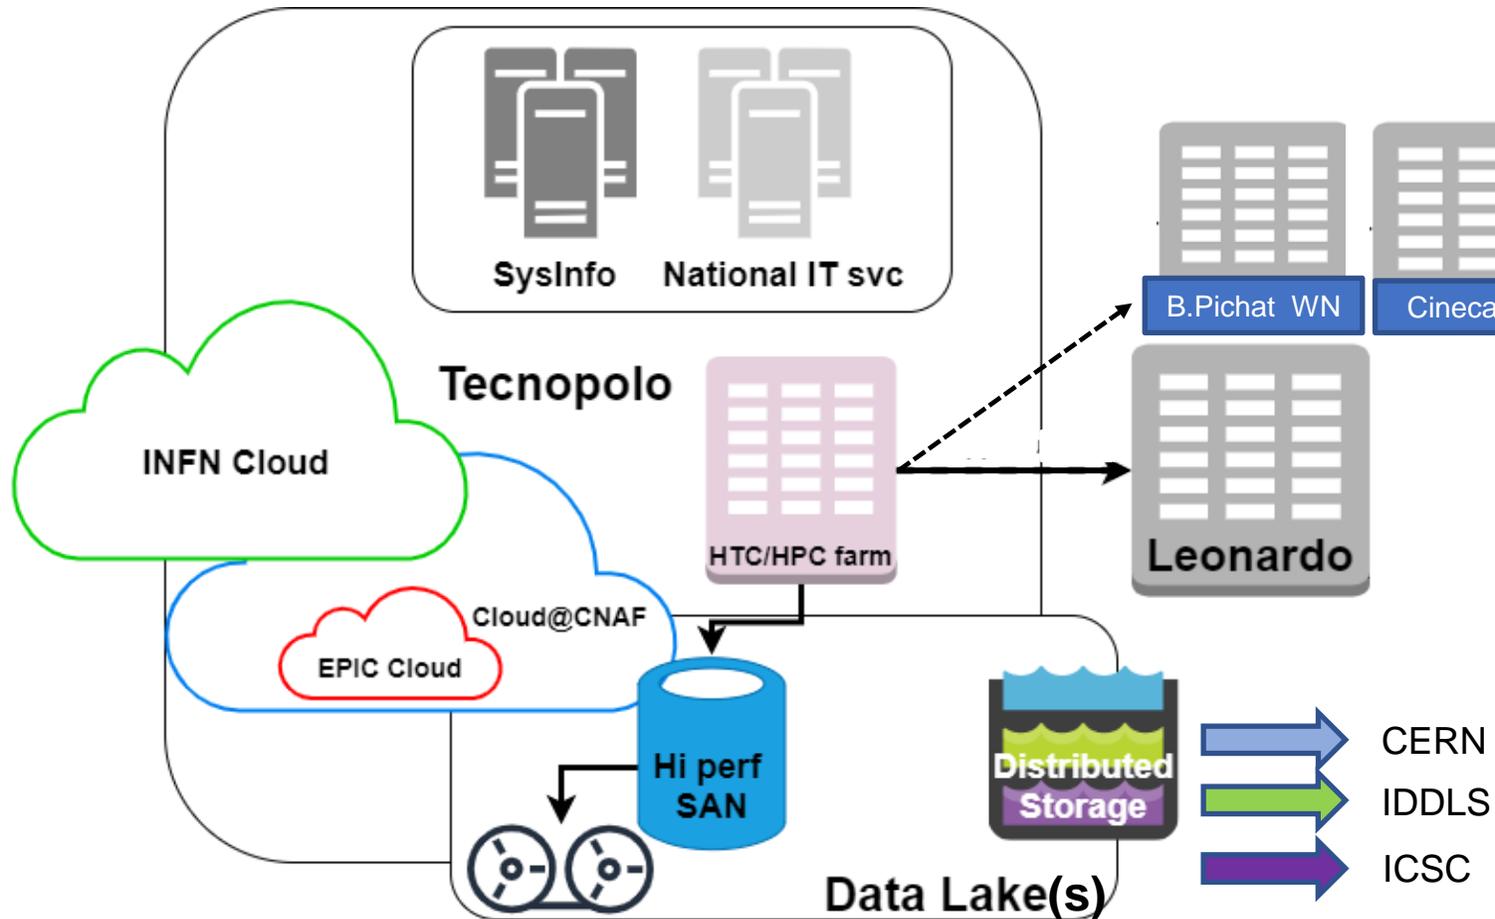
The cooling system and the PUE

- 4 central refrigerator Units
 - 3+1 redundancy
- Chilled water 19-26 °C for the low density air cooled racks
 - 2 MW Chillers
 - Total/partial free cooling is possible
- Warm water 37-47 °C for DLC racks
 - 2,25 MW Chillers
- To be doubled in the second phase
- **High Density CPU Area**
 - 4 CRAH - 200 kW each (3+1)
- **Network Area**
 - 4 CRAH - 75 kW each (3+1)
- **STORAGE Area**
 - 16 CRAH - 200 kW each (12+4)
 - Cold corridor aisles
- **TAPE Area**
 - 4 CRAH - 25 kW each (3+1)

$$PUE_{DLC} \approx 1.08$$

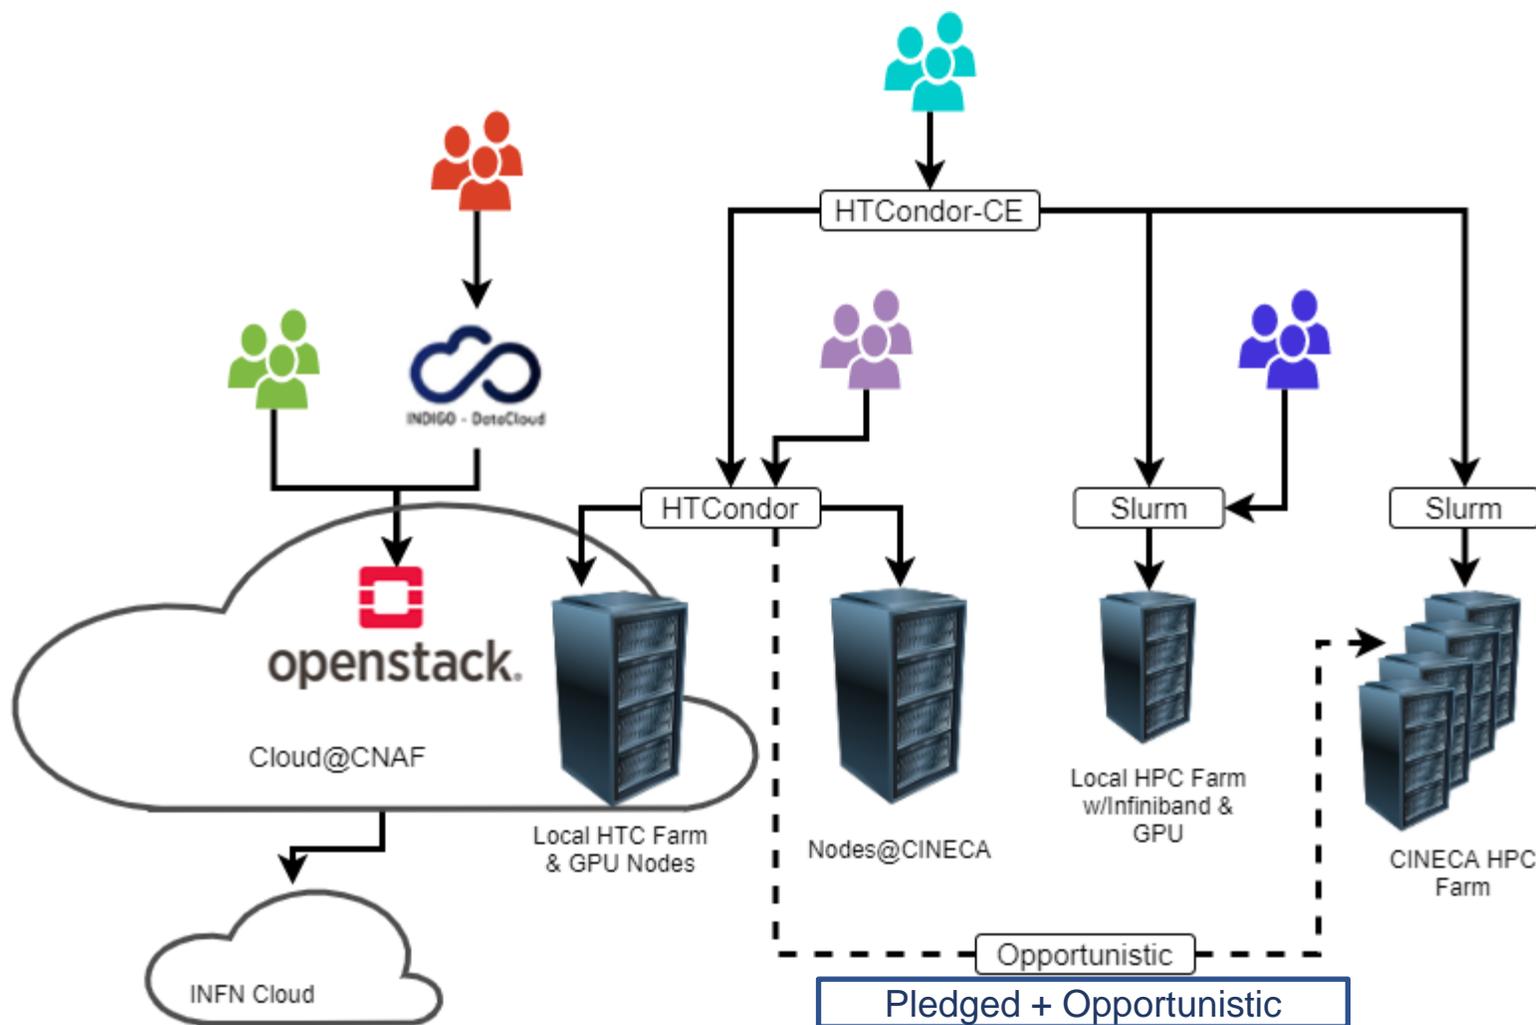
$$PUE_{Tot} \approx 1.2 - 1.3$$

A “distributed” datacenter



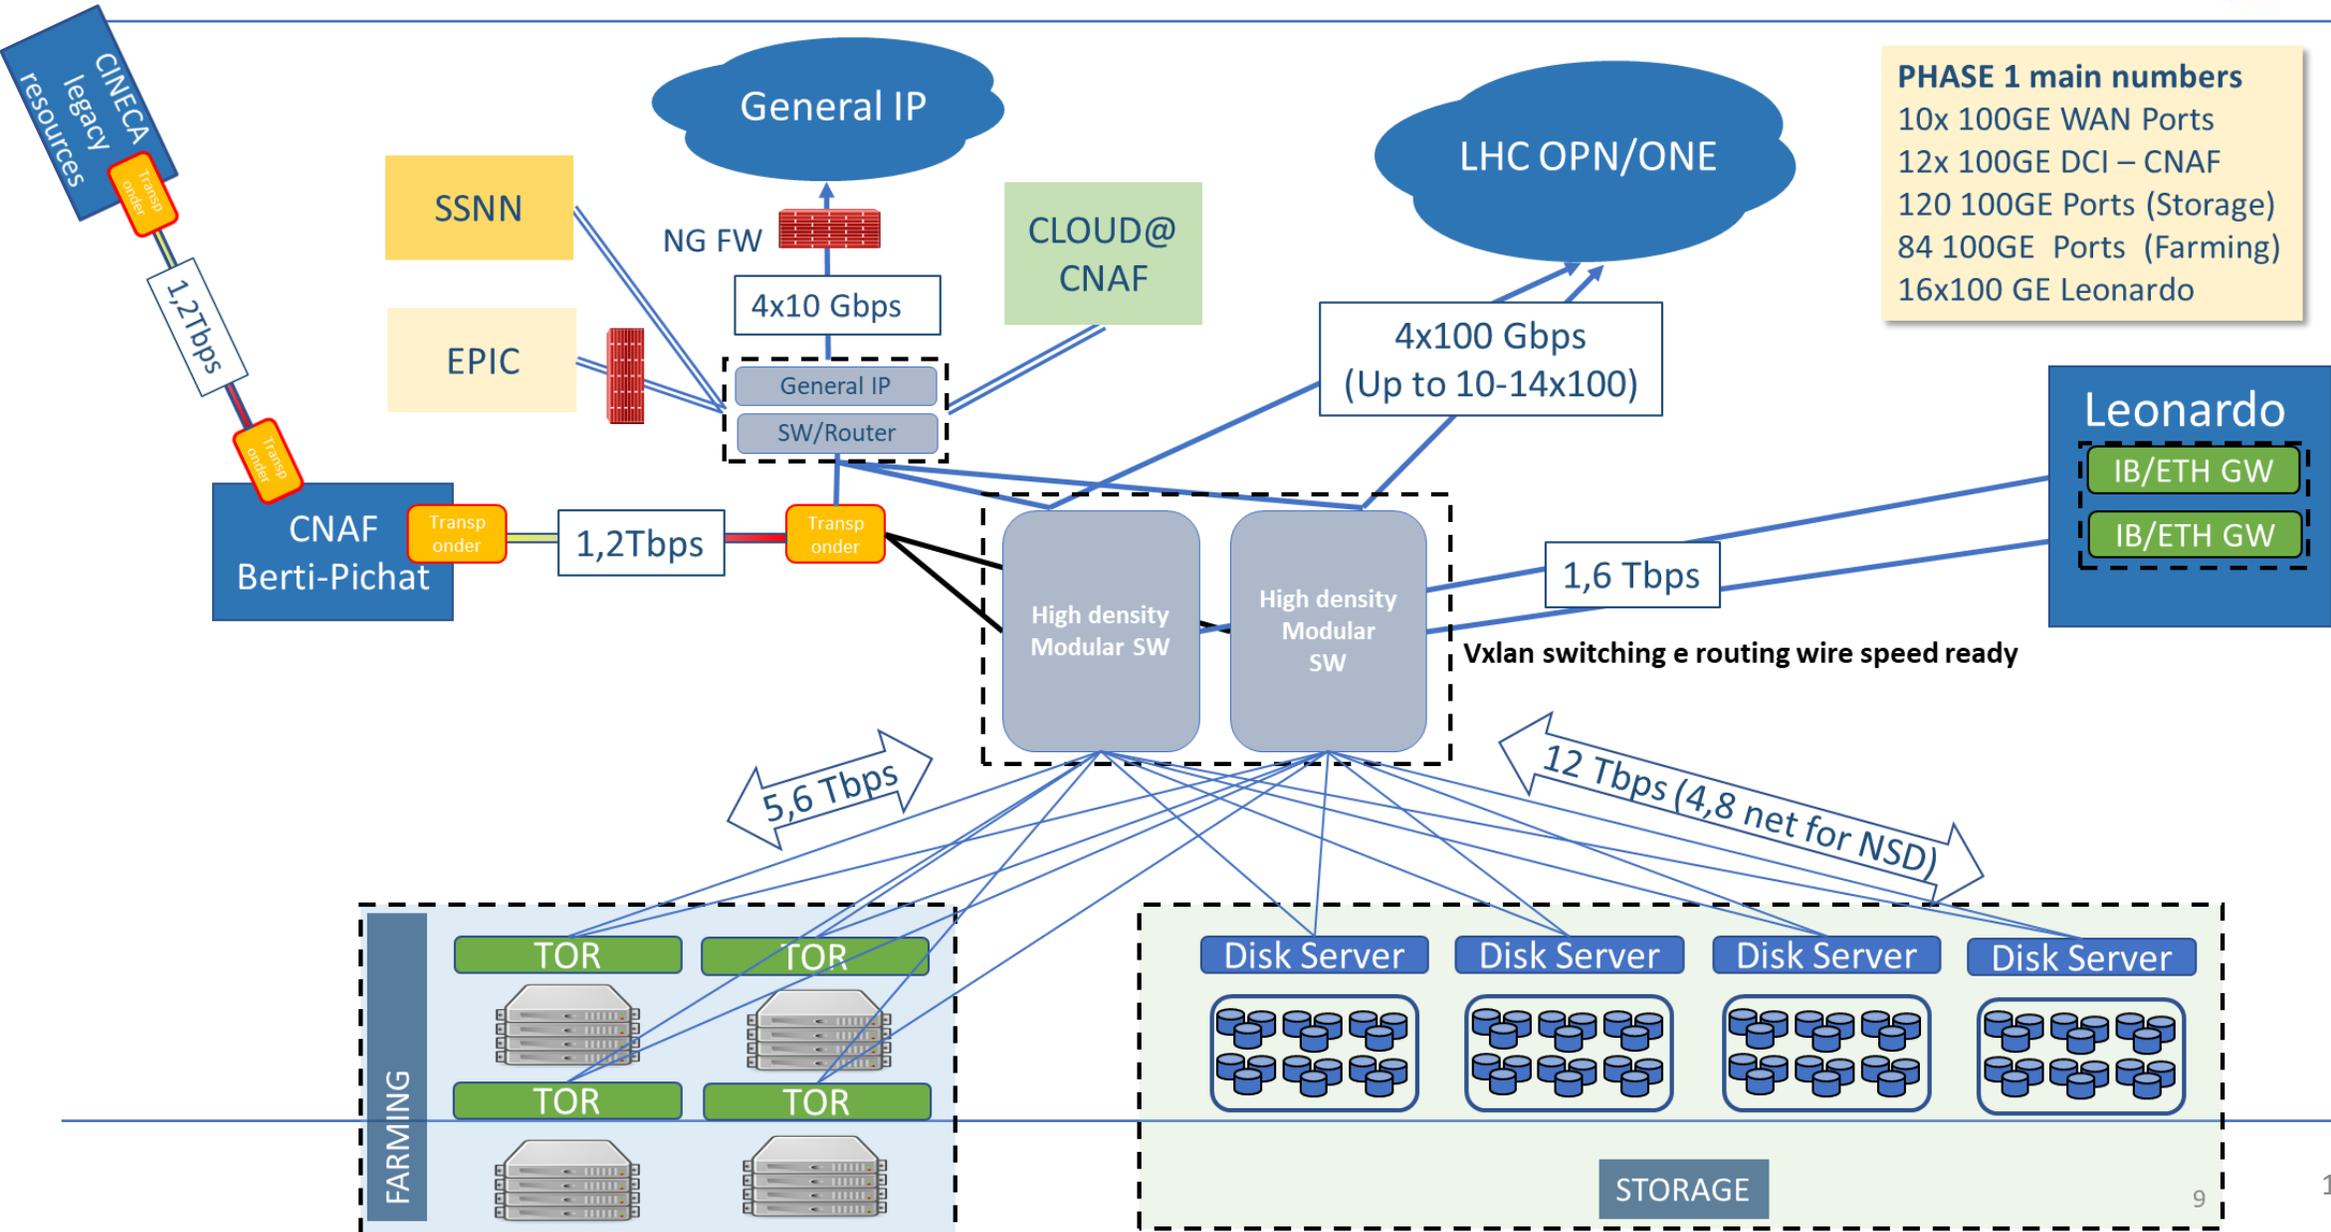
- Multiple “locations”
 - CNAF Technopole
 - CINECA Leonardo CPU access
 - INFN-CLOUD federated cloud
 - Data-lake(s)
 - DCI with INFN sites
 - DCI with CERN
 - New national data lake for the ICSC project
 - The ICSC headquarter will be at the Technopole

Farm, Cloud, HPC integration

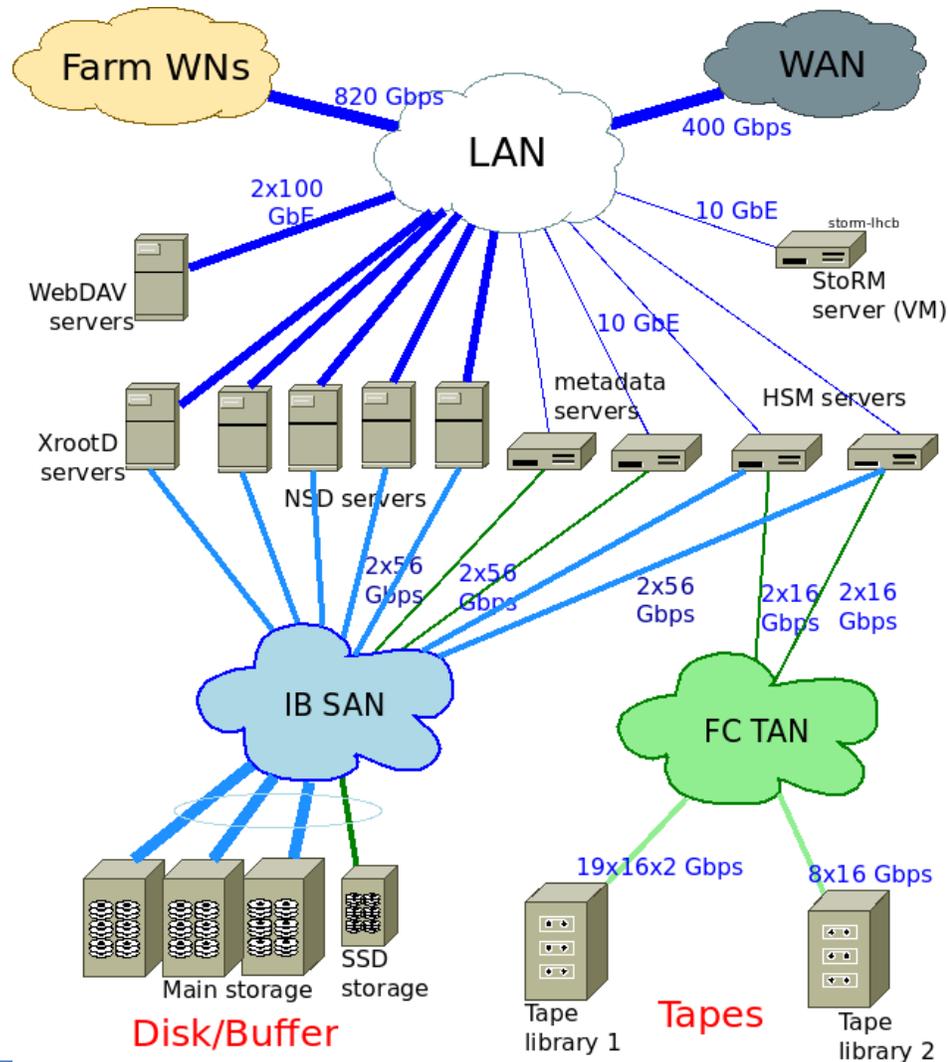


- Integration of access points for several different infrastructures
 - Cloud@CNAF
 - Cloud@INFN
 - Local HTC farm
 - Local HPC farms
 - CNAF WN@CINECA
 - Leonardo usage
- Based on
 - HTCondor
 - SLURM
 - HTC-SLURM connectors
 - INDIGO PaaS + Openstack

A Complex Networking Infrastructure



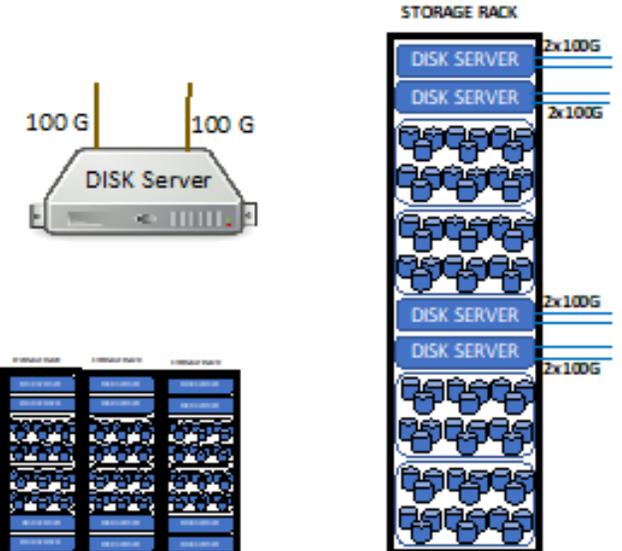
The storage architecture



Hypothesized network sizing for storage

STORAGE (DISCO) 120 PB

- About 60 Disk Servers connected at 2x100 Gbps
- Hypothesized density: 8 PB per Rack
- 15 Rack for the cluster Filesystem
- 3 Rack per Data Transfer R&D e metadata
- TOT 18 Rack



- To maintain 5MB/s for each TB we need to deploy one IO server for every 2PB of usable storage, i.e. 60 IO servers for 120PB of disk storage
- Data Mover servers: 32 ports 2x25Gbps
- Metadata servers 32ports x25Gbps

Live Relocation Timeline



- **Live Migration**
 - Legacy site “extended” through a DCI channel 1.2Tbit/s
 - Data moved to a new storage
 - CPUs moved in chunks
- Down only for tape libraries
 - Need dismantle and re-assembling

Communication



But since I was curious, I asked: what can you actually do with these supercomputers?

Data Valley:

<https://www.youtube.com/watch?v=96TfXHCWxf8>



They answered: everything you can think of... and other things you can't even imagine.

- CNAF is moving to a new location to cope with the increasing resources demand
 - From HEP communities, but not only
- It will be collocated with supercomputing infrastructure from CINECA and ECMWF
 - Strengthening the long-standing collaboration with CINECA
 - Mainly Theoretical Physics simulations, but increasing usage from WLCG
- A greener datacenter with a much lower PUE
- Targetting a «live» migration
 - With few exceptions, i.e. TAPE libraries