Migrating the INFN-CNAF datacenter to the Bologna Tecnopolo

a status update

C. Pellegrino - CHEP 23 - May 9th 2023

INFN CNAF

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 centres for distributed computing in Italy
- It is currently located close to the Bologna city centre
- Co-located with the Physics department of the Bologna University

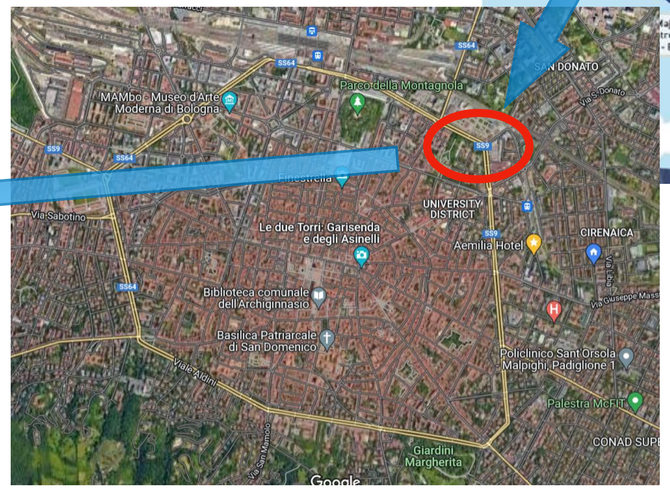

Scientific motivations

There are three main drivers for the transfer of INFN-CNAF data centre to a new location:

1. an expected huge increase in IT resources to be installed in the coming years, mainly driven by HI-Lumi LHC projections
2. infrastructural problems of the current location
3. the expansion opportunities offered by the new location:
    - Proximity to Leonardo
    - Higher electrical power availability (up to 10MW, currently ~1.6MW)
    - Larger room space (~3000 m^2, currently ~800 m^2)

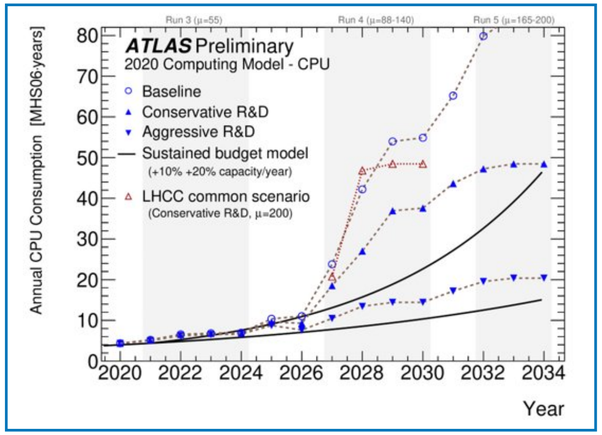

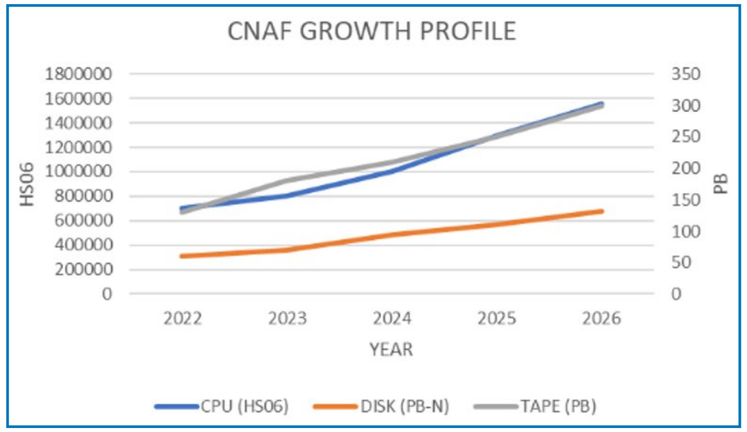

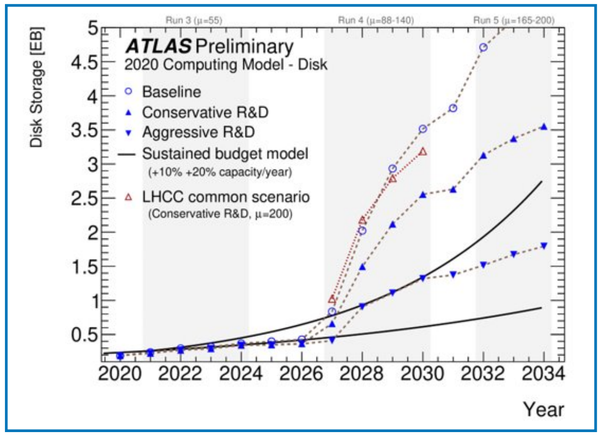

The new INFN Data Centre 1/5

Location

- New INFN Data Centre will be hosted in Hall **B5**
  - Infrastructure (power, cooling) shared with CINECA (**C2**)
  - Newest technologies:
    - DLC ready
    - Low PUE (1.08-1.2)
- Two phases depending on the total available power
- **Phase 1**: 2023-25: 13 MW
  - **INFN-Tier1**, Leonardo => Lisa, ICSC, Quantum, CNR Tier-1, INAF, ...
- **Phase 2**: 2026: up to 24 MW (up to 10 MW for INFN)
  - INFN-Tier1, ICSC, ready for post-exascale machine (CINECA)
    - Leonardo to be decommissioned in 2026

The new INFN Data Centre 2/5

Low-density zone

- 740 m², 132 racks, up to ~2.4 MW of power
  - 6 blocks of 16 kW racks
  - Cold-aisle containment
  - Plenum-air cooling

- It will host:
  - Disk-based Storage systems
    - 44 racks, ~400PB
  - Cloud infrastructure
  - ISO-certified Cloud
  - Other ICT services

Low-density zone

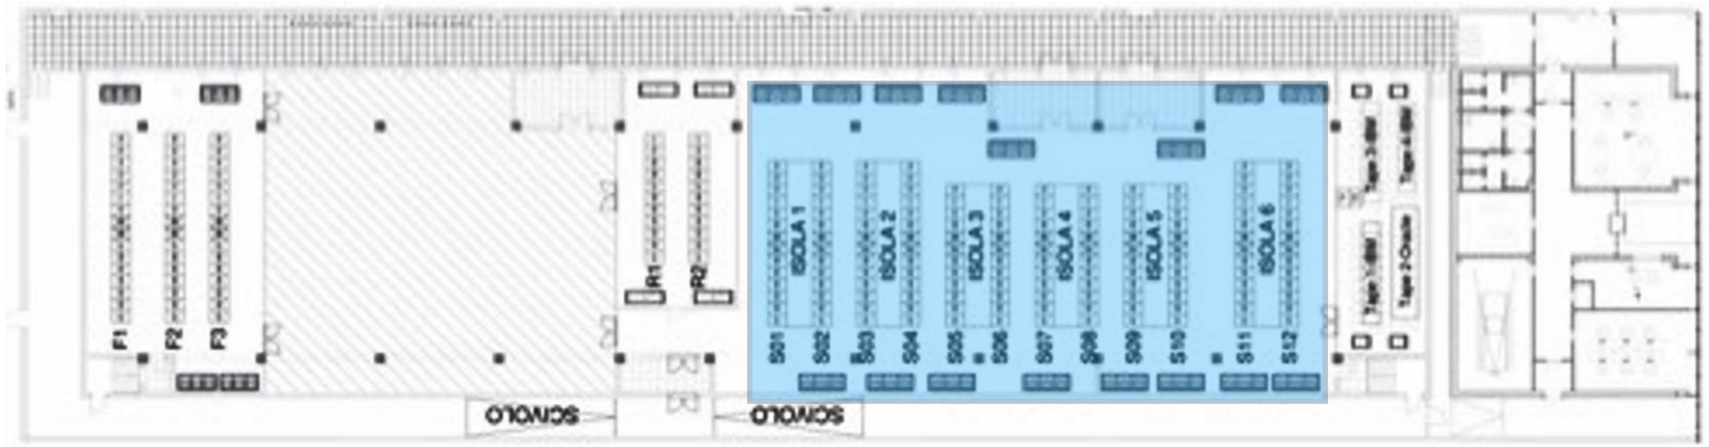

1 block: 20 or 24 racks + 4 for network

The new INFN Data Centre 3/5

High-density zone

- 200 m², 42 racks, up to ~3.4 MW of power and (at least) 4 MHS06
  - 3 rows each with 14 DLC racks (up to 80 kW/rack)
- From Phase 2
  - DLC technology to be decided by the end of 2025 (constraints from water temperature, due to Leonardo settings)

High-density zone

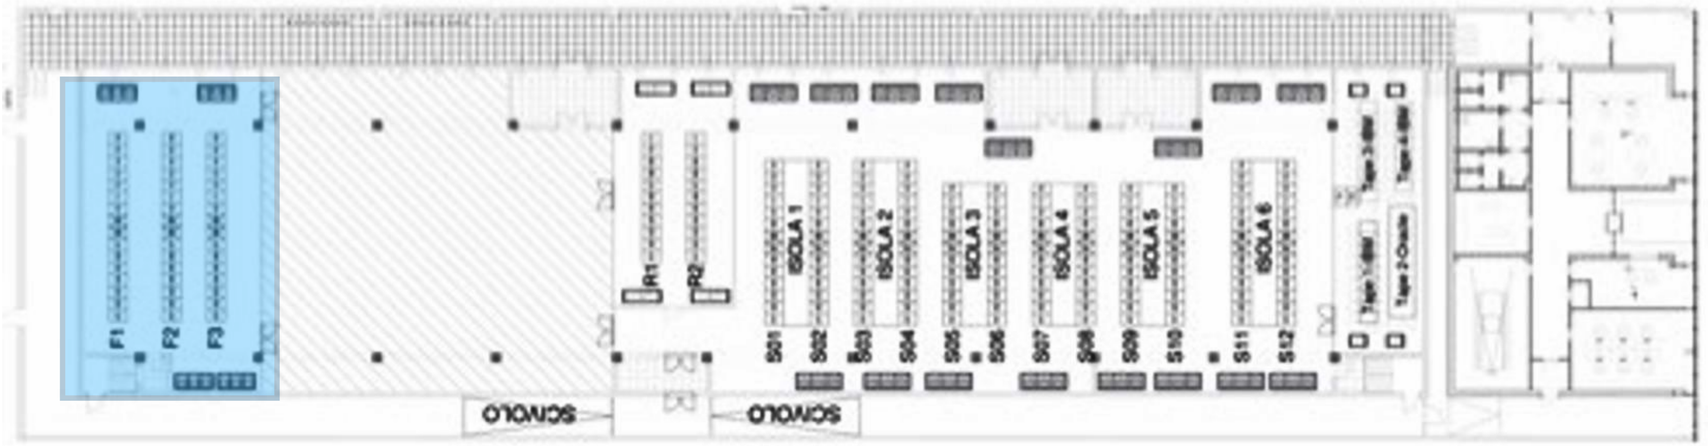

The new INFN Data Centre 4/5

Other zones

- Up to 4 tape libraries
- Network zone
  - GARR GP and LHC OPN/ONE uplinks
  - 1.2 Tb/s LAN extension with the current CNAF
  - 1.6 Tb/s LAN extension with Leonardo (local)
  - Tb/s DCI CERN (R&D)

- Expansion area ~ 450 m^2
- Meeting room and control room

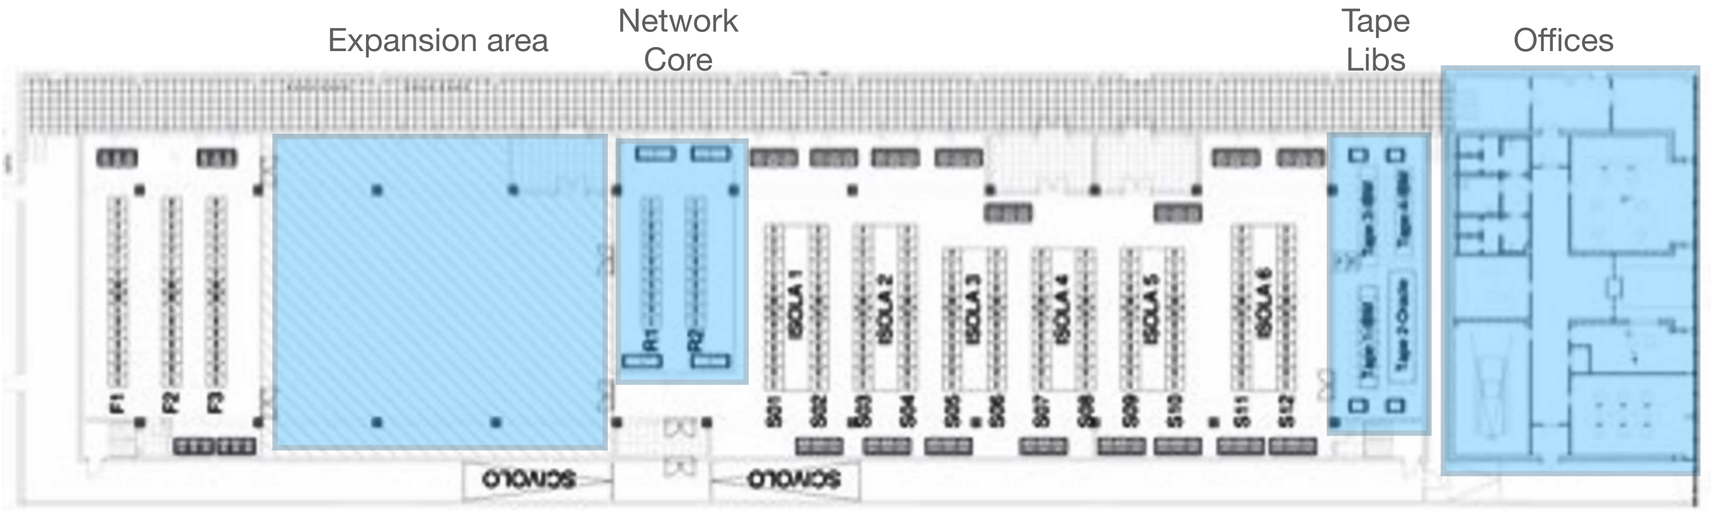

High density zone

DLC pipes

With the floating floor on

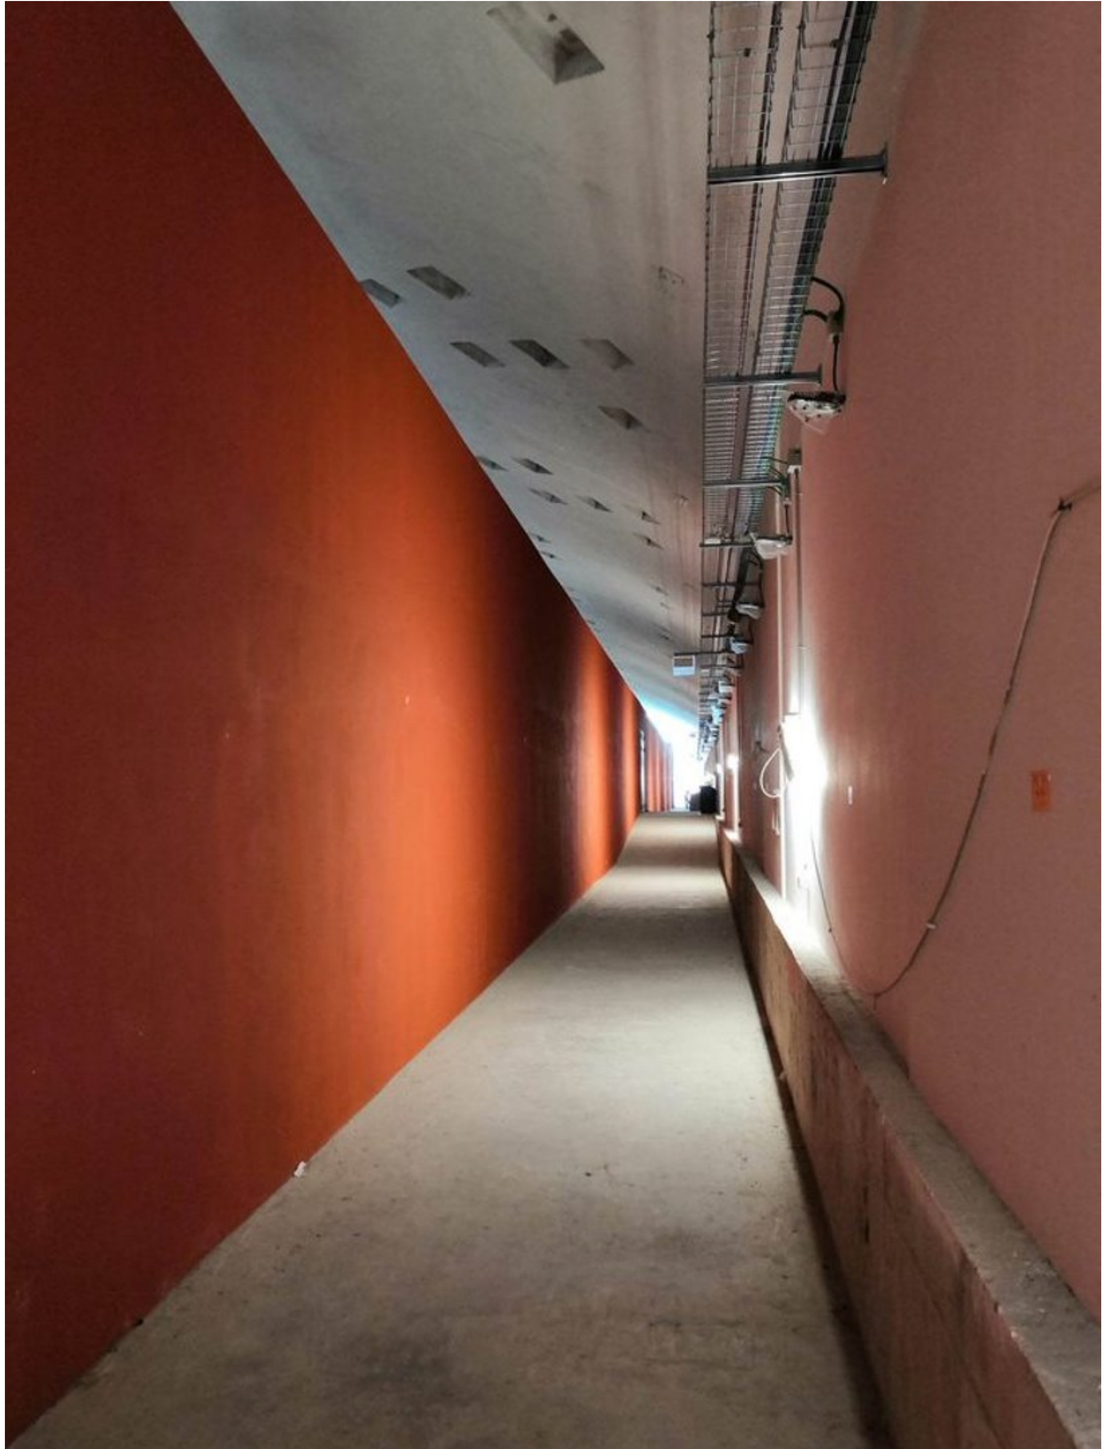

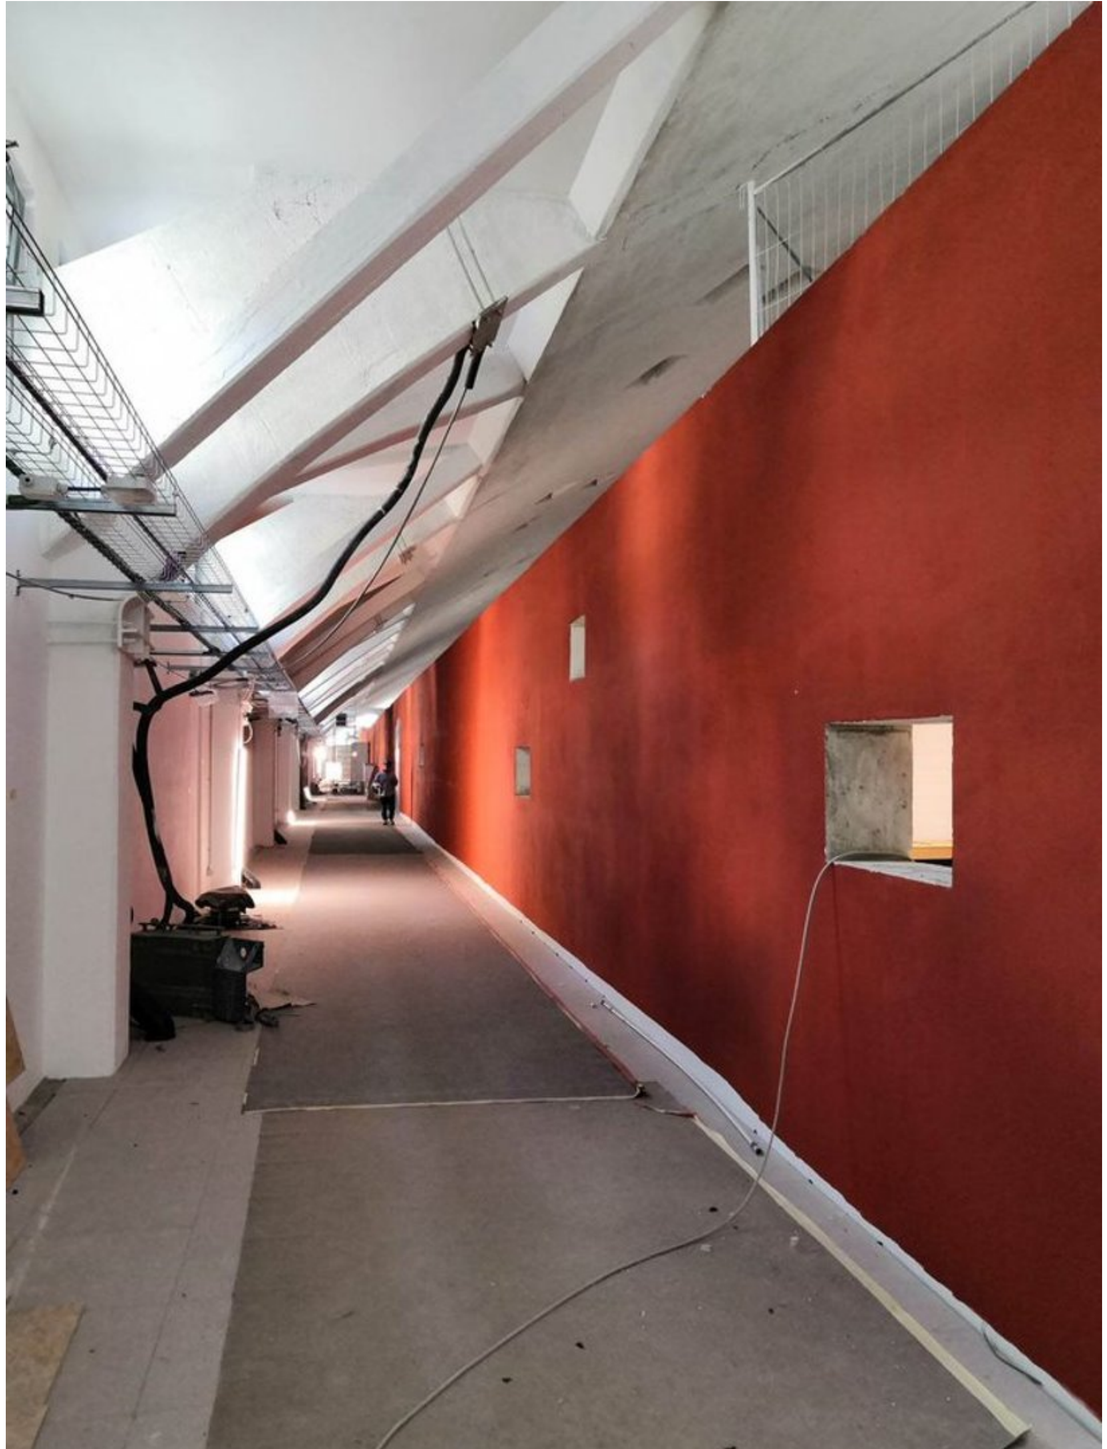

Low Density zone

Low Density & expansion zones with floating floor on

Technical challenges

- Minimise downtimes of services
  - mainly the live copy of data through DCI
- Network connections:
  - LAN
  - to current DC
  - to LHC OPN/ONE
  - to Leonardo (CINECA)
  - to Marconi (CINECA, Casalecchio)
  - to GARR

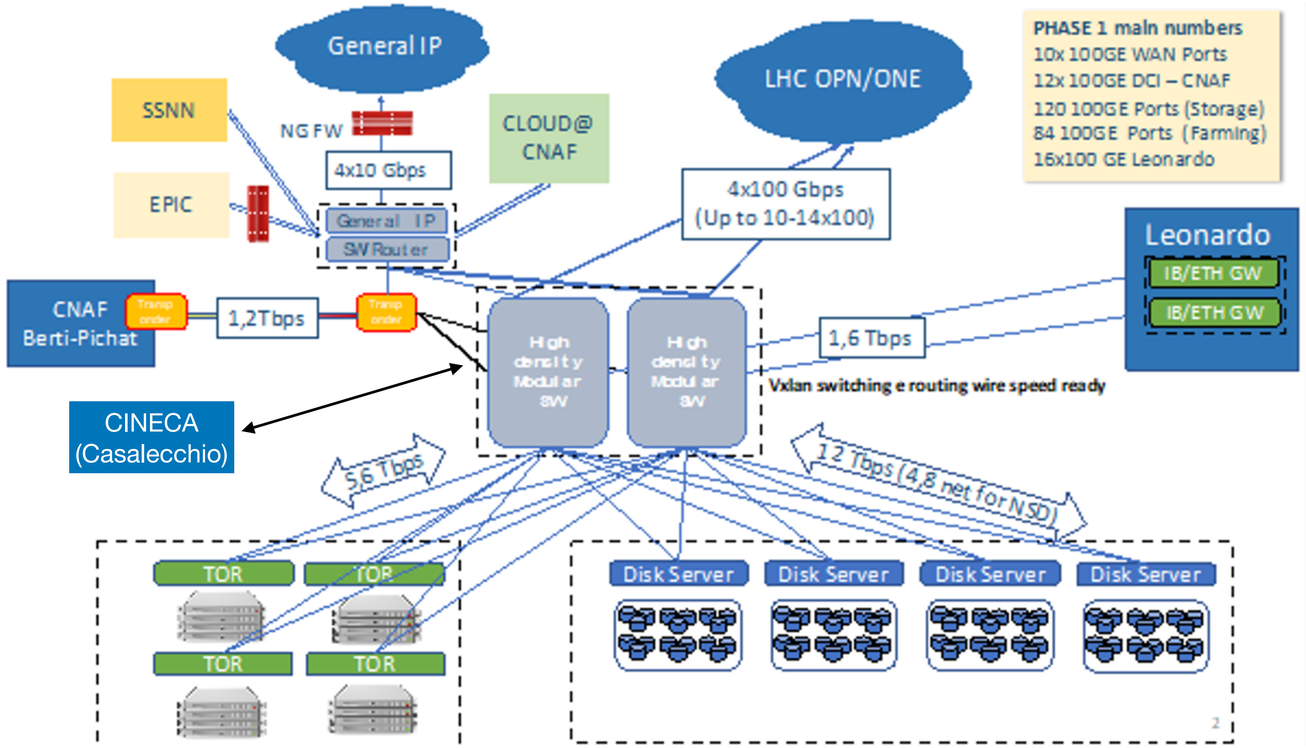

Network connection towards Leonardo

- (Still an) Open issue with Leonardo having IB connection only
  - Access from CNAF through 2 IB-Eth gateways (Mellanox Skyway) with aggregated bandwidth of 1.6 Tbps
  - Several known limitations:
    - No VLAN tagging
    - No IPv6
    - No support for large Eth frames (Skyway supports MTU=4092 while on CNAF LAN MTU=9000)
  - To avoid changing configuration on all disk-servers we configured the Path MTU discovery
  - Simulation of the scenario with a Skyway: a rack of WNs configured with MTU=4000 accessing production disk-servers having MTU=9000

Same performance, no errors

9000/9000: baseline

9000/4000: Skyway-like setup

Status and Conclusions

- Possibly we will interconnect to Leonardo before the actual start of the migration
  - we successfully tested the access to Leonardo from current CNAF
- After the actual delivery of Hall B5 (**July 2023**)
  - Network cabling installation (**July 2023**)
  - Delivery of network and storage systems (**September 2023**)
  - Core switch and core routers installation and validation (**October 2023**)
  - Storage installation and validation (**October 2023**)
- After network and storage in production: start of the migration
- The real challenge: minimise downtime for services
  - Install services on new hardware at the new data center
  - Copy the data and then move the storage devices
- Only move hardware with at least ~2 years of service life left
  - Part of storage will remain at the current DC until phase-out
  - To be migrated: storage systems (newer than 2019), cloud resources, ICT services, last year CPU tender
- The actual migration will be completed in 9 months, with minimal impact on the services
- The tape library will be moved later on

Electrical and cooling system available
Racks installed

July 23
Room ready

Cabling
Active core systems
Management network
DCI with legacy site

Aug 23
Network ready

New Storage Installed
Legacy CPU and Storage migration starts

Sep 23
HW relocation

Migration finalised
Tape Library migrated
Site commissioned

Apr 24
Finalisation

Thank you for your attention

Backup

The new INFN Data Centre 4/5

The cooling system

- 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 Phase 2

- 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)