T1 highlights CdG

January 2025

D.Cesini – INFN-CNAF

Pledges T1 2024-2025

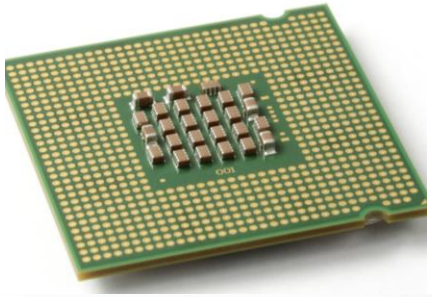

TOTAL	2024	2025
CPU Pledge (HS)	792.000	825.000
Disk Pledge (TBN)	82.949	101.000
Tape Pledge (TB)	193.581	233.000

Pledge LHC da 01/04, pledge non-LHC da 01/01
Assegnati pledge CPU 2025 non LHC

+ ExtraPledge: cpu=22kHS; disk=1620TBN; tape=19100TB (+9000 extra da 2023)

+ PledgeCloud Delta_CPU=1500HS; Delta_Disk=110TBN

+ HPC Bubbles

No GPU/FPGA pledged, yet some already available via cloud.
No ARM CPU pledged, yet already available via HTCondor

Stato Risorse Tecnopolo - Storage

- **Disco**
  - **AQ 2023-2024**
    - installato al tecnopolo il primo Appalto Specifico (64PB) su fondi ICSC
    - In fase di collaudo il secondo Appalto Specifico (16PB) su fondi ICSC
  - Migrati i sistemi con manutenzione attiva
    - DDN gara 2019 – 6PB
    - Lenovo-DDN gara 2022 – 14PB
    - Huawei os5k8 – 9PB

- **TAPE**
  - Nuova libreria IBM installata al tecnopolo
    - Andrà in produzione la prossima settimana
  - Vecchia libreria IBM portata al tecnopolo
  - Vecchia libreria Oracle in fase di repack portata al tecnopolo
    - Repack bloccato in attesa di nuovi nastri
    - Tornata in produzione per alcuni esperimenti
  - Gara 96PB nastri in fase di aggiudicazione
    - Aperte buste economiche – valutazione offerte anomale

Stato Risorse Tecnopolo - CPU

- **CPU**
  - 200 nodi di Leonardo in produzione da fine estate 2024 (572kHS)
  - Spenti i nodi di CINECA-Casalecchio per allagamento (270kHS)
    - In fase di migrazione verso CNAF@Tecnopolo
    - Da capire se e quanti ne riaccendiamo
  - 31/01 spegnimento vecchie CPU@b.pichat (110kHS)
  - 20kHS su ARM

- **HPC**
  - Spento il vecchio cluster (del 2012)
  - Migrato al tp il cluster dedicato ai teorici acceleratoristi del CERN
  - In fase di installazione le HPC Bubbles con fondi PNRR_Terabit

Stato trasferimento al Tecnopolo

- Da migrare 3 rack zona certificata ISO (EPIC)
- Da migrare 3 rack SSNN
  - Dipendenza da rete dedicata Tecnopolo-LNL (GARR)
- Decomissioning vecchia sede

Gara "HPC Bubbles"

- Accordo Quadro Nazionale
  - Lotto1
    - CPU, GPU, FPGA
  - Lotto2
    - Storage
  - Sedi Coinvolte: CNAF, BARI, MI-BI, PI, TO, LNGS, NA, RM1, PD/LNL

✱ **Nodi con fondi DARE – Terabit per Spoke8**

	Nodo CPU	Nodo GPU	Nodo FPGA Xilinx	Nodo FPGA Terasic	Nodo storage
BA_DARE	12	6	0	0	6
BA_TerabitS8	0?	0?	0	0	0?
CNAF_DARE	10	9	0	0	16
CNAF_TerabitS8	8?	8?	0	0	6

Quantità nodi con fondi Terabit-ICSC-DARE

	Nodo CPU	Nodo GPU	Nodo FPGA Xilinx	Nodo FPGA Terasic	Nodo storage
BA	24 ✱	6	0	0	32 ✱
CNAF	26 ✱	30 ✱	2	2	52 ✱
MIB	0	0	2	2	0
NA	18	1	2	0	8
PD	6	6	0	0	0
PI	20	0	0	0	0
RM1	12	0	0	0	0
TO	14	6	0	0	0
LNGS	0	6	0	0	12
CT	12	0	0	0	8
LNF	12	0	0	0	0
LNFESA	8	6	0	0	6
LNL	4	0	0	0	0
MI	4	0	0	0	0
TOTALE	160	61	6	4	118

Core: 30 kcore fisici
Circa 34 HS/core

GPU: 244 NVIDIA H100
40 FPGA
InfiniBAnd 400Gbs

45 PB RAW

HPC Bubbles

Nodo CPU	192 core fisici 1.5TB RAM DDR5 IB NDR 400G 20TBL (SSD) + dischi di sistema
Nodo GPU	Come CPU + 4x NVIDIA H100 SXM5 con minimo 80GB e memoria HBM2e
Nodo FPGA	32core RAM 768GB DDR5 IB NDR 440G 4 x XILINX U55C o 4 x TerasicP0701
Nodo Storage (CEPH Bricks)	64 core fisici 1TB RAM DDR5 384 TBL HDD + 25.6 TBL NVMe
Accessori	Switch IB, Switch ETH Cavi IB, Cavi ETH Transceiver vari Assistenza 3+2

Backup

What can the Tecnopolo host?

The computing infrastructures

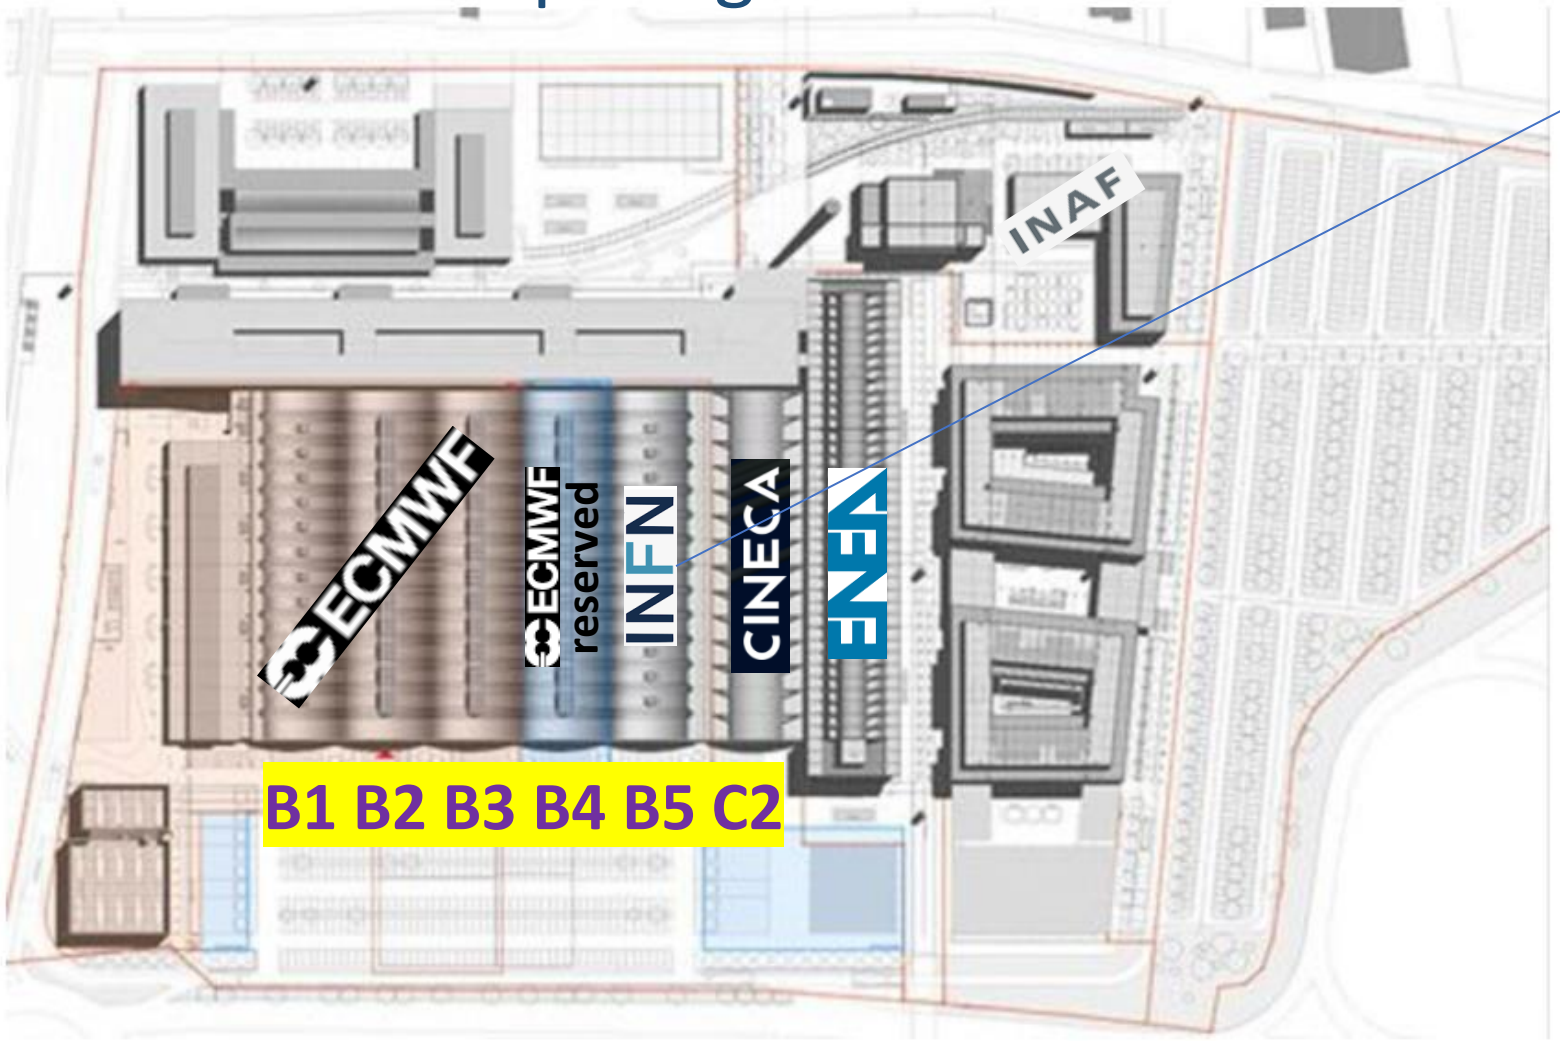

Each of the 6 “botti” (barrels) is ~5000m^2 of usable IT space

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

CNAF and CINECA data halls

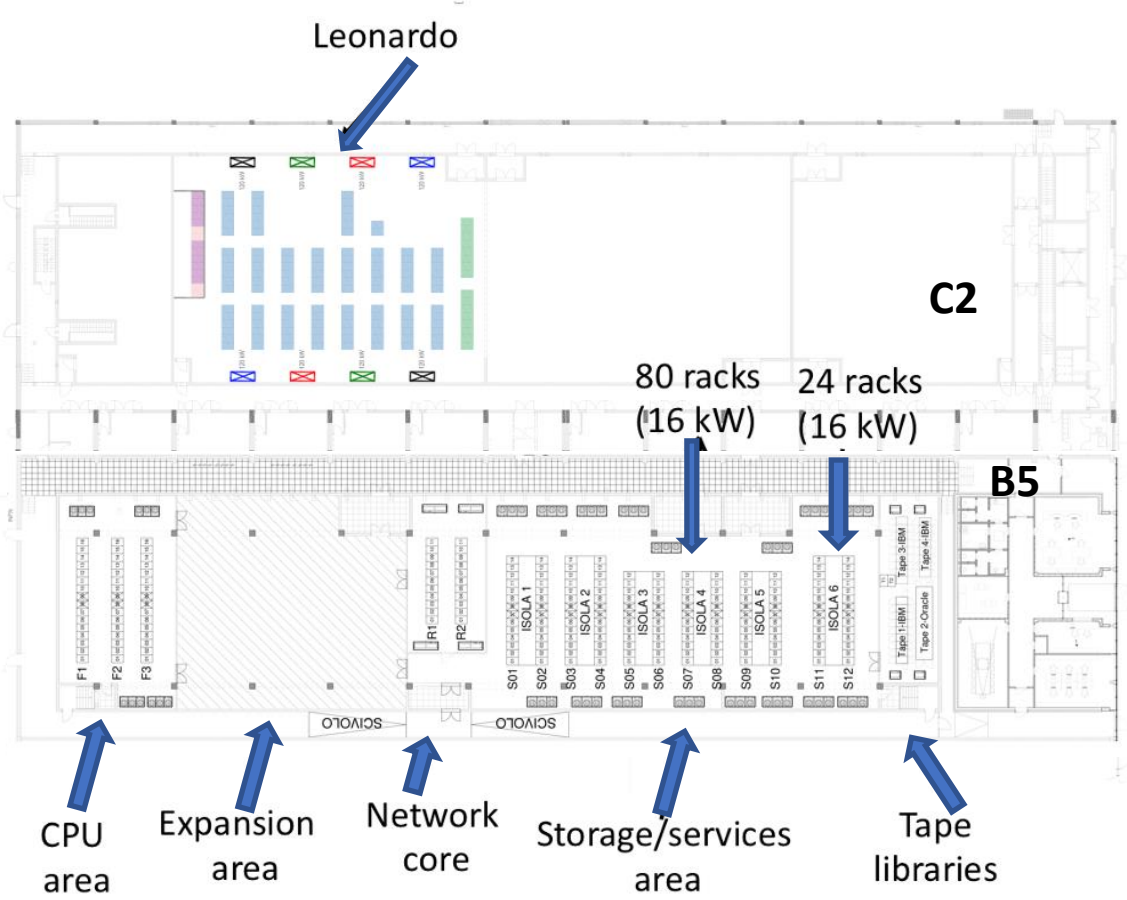

DLC 80kW

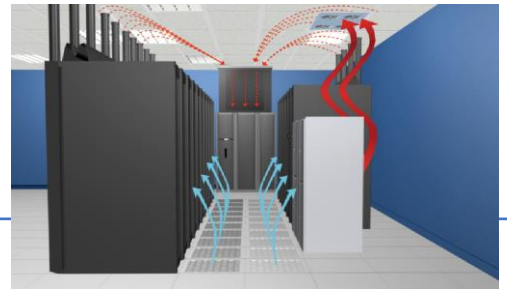

- The new CNAF Datacenter will feature the following main areas
  - High Density – 2-3 rows for 80kW racks
  - Low density – 80+24 16kW racks
  - Expansion area
  - Tape libraries areas
    - Up to 4 libraries
- The CPU area can host up to 3MW of CPUs via 42 DLC high density racks
- The low-density area 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

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)

$$PUE_{DLC} \approx 1.08$$

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

Networking Infrastructure