

T1 highlights CdG

February 2024


D.Cesini – INFN-CNAF

Resources@T1 2023-2024



ALL VO No Cloud	2023	2024	Delta
Pledge CPU (HS06)	660000	792000	132000
Pledge disk (TBN)	69576	82949	13373
Pledge tape (TB)	158282	193581	35299

CPU farm



CPU in 2024 – Leonardo integration

- No direct CPU acquisition in 2024
- We will use up to 300 Leonardo-GP@CINECA nodes
 - Dual 56 cores sockets Intel Sapphire Rapids
 - **2800 HS06**

GP partition online

Attualmente usata da INFN CS4 con sottomissione diretta

In preparazione VM per lanciare i «CNAF WN»

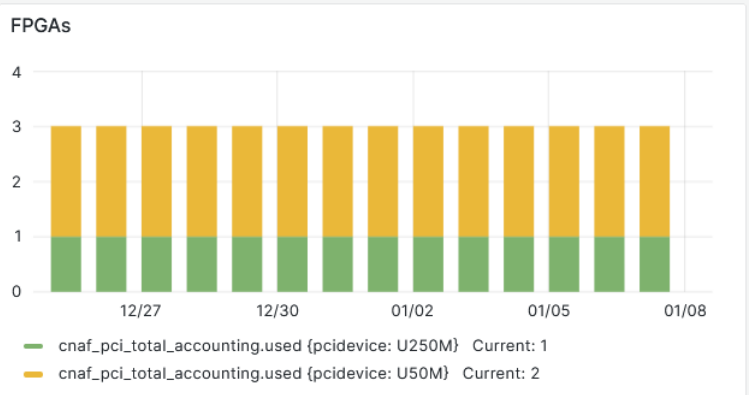


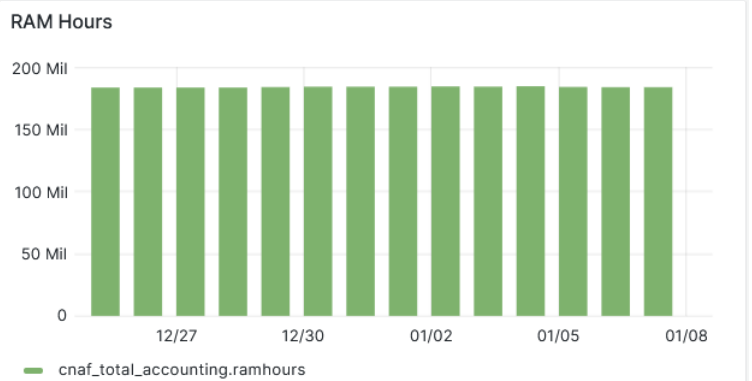
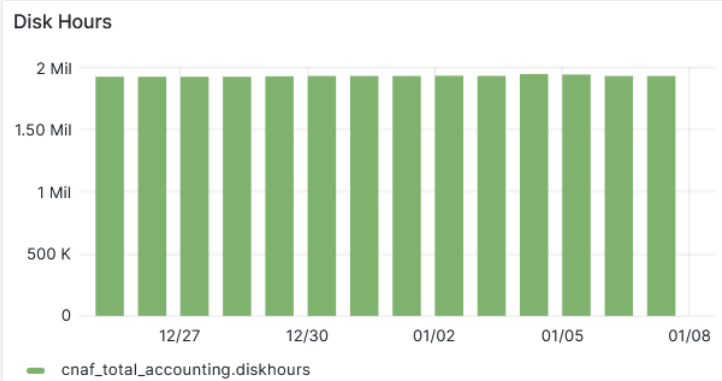
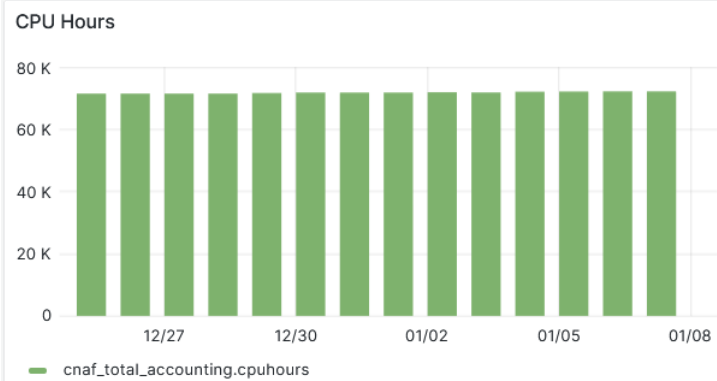

In progress cablaggio e configurazione degli skyway

- Integration Plan
 - “inifnite” SLURM jobs to launch VM containing “our” Condor WN
- PCI pass-through to see the IB cards on Leonardo
- Mellanox Skyway IB-ETH bridges to reach our LAN
 - 16 x 100Gbs



- > Standard 2U appliance
- > 1.6Tb/s solution
- > 8-port HDR/HDR100/EDR InfiniBand
- > 8-port 200/100Gb/s Ethernet

Cloud@CNAF



90 Tenant Di cui 11 per INFNCLOUD

tier1_internal

Home > Dashboards > Cloud@CNAF Accounting

2024-01-05 00:00:00

ctrl+k

Last 15 days

	min	max
cnaf_region_accounting.servers {region: sdds}	312	32
cnaf_region_accounting.servers {region: tier1}	392	42

Project View

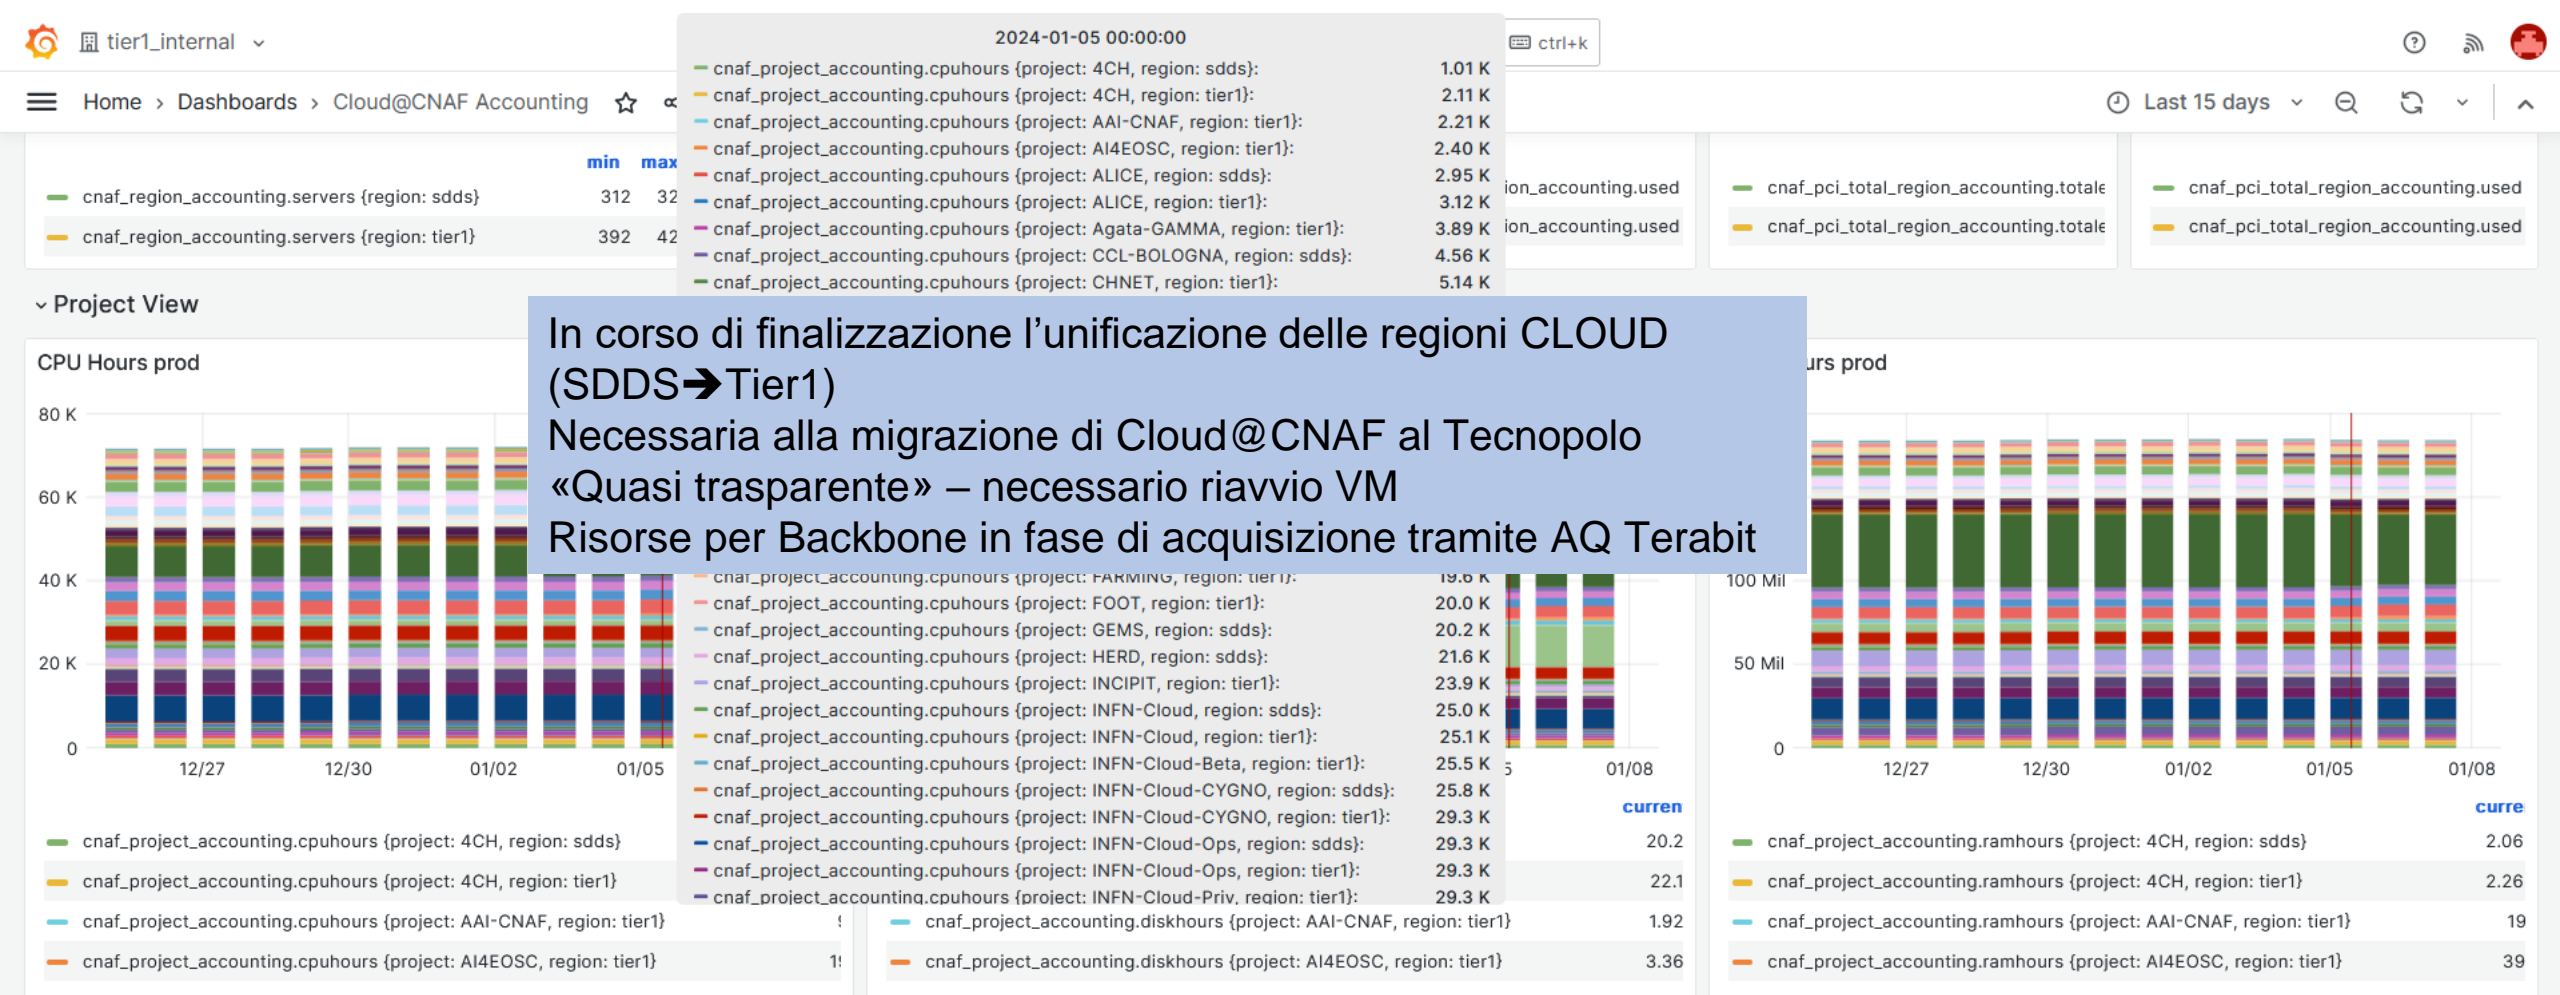
CPU Hours prod

RAM Hours prod

cnaf_project_accounting.cpuhours {project: 4CH, region: sdds}	1.01 K
cnaf_project_accounting.cpuhours {project: 4CH, region: tier1}	2.11 K
cnaf_project_accounting.cpuhours {project: AAI-CNAF, region: tier1}	2.21 K
cnaf_project_accounting.cpuhours {project: AI4EOSC, region: tier1}	2.40 K
cnaf_project_accounting.cpuhours {project: ALICE, region: sdds}	2.95 K
cnaf_project_accounting.cpuhours {project: ALICE, region: tier1}	3.12 K
cnaf_project_accounting.cpuhours {project: Agata-GAMMA, region: tier1}	3.89 K
cnaf_project_accounting.cpuhours {project: CCL-BOLOGNA, region: sdds}	4.56 K
cnaf_project_accounting.cpuhours {project: CHNET, region: tier1}	5.14 K
cnaf_project_accounting.cpuhours {project: CNAF, region: tier1}	5.18 K
cnaf_project_accounting.cpuhours {project: CNAF-CERN-WP3, region: sdds}	6.10 K
cnaf_project_accounting.cpuhours {project: DARKSIDE, region: sdds}	6.19 K
cnaf_project_accounting.cpuhours {project: DARKSIDE, region: tier1}	6.53 K
cnaf_project_accounting.cpuhours {project: DODAS, region: sdds}	12.8 K
cnaf_project_accounting.cpuhours {project: DODAS, region: tier1}	15.9 K
cnaf_project_accounting.cpuhours {project: EA-I-ENERGY, region: sdds}	19.0 K
cnaf_project_accounting.cpuhours {project: EEE, region: sdds}	19.5 K
cnaf_project_accounting.cpuhours {project: EEE, region: tier1}	19.6 K
cnaf_project_accounting.cpuhours {project: EGIOPS, region: tier1}	19.6 K
cnaf_project_accounting.cpuhours {project: FARMING, region: tier1}	19.6 K
cnaf_project_accounting.cpuhours {project: FOOT, region: tier1}	20.0 K
cnaf_project_accounting.cpuhours {project: GEMS, region: sdds}	20.2 K
cnaf_project_accounting.cpuhours {project: HERD, region: sdds}	21.6 K
cnaf_project_accounting.cpuhours {project: INCIPIIT, region: tier1}	23.9 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud, region: sdds}	25.0 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-Priv, region: tier1}	25.1 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-Beta, region: tier1}	25.5 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-CYGNO, region: sdds}	25.8 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-CYGNO, region: tier1}	29.3 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-Ops, region: sdds}	29.3 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-Ops, region: tier1}	29.3 K
cnaf_project_accounting.cpuhours {project: INFN-Cloud-Priv, region: tier1}	29.3 K
cnaf_project_accounting.diskhours {project: AAI-CNAF, region: tier1}	1.92
cnaf_project_accounting.diskhours {project: AI4EOSC, region: tier1}	3.36

cnaf_pci_total_region_accounting.totale		
cnaf_pci_total_region_accounting.used		

cnaf_project_accounting.ramhours {project: 4CH, region: sdds}	2.06
cnaf_project_accounting.ramhours {project: 4CH, region: tier1}	2.26
cnaf_project_accounting.ramhours {project: AAI-CNAF, region: tier1}	19
cnaf_project_accounting.ramhours {project: AI4EOSC, region: tier1}	39



Risorse in arrivo con Terabit al CNAF



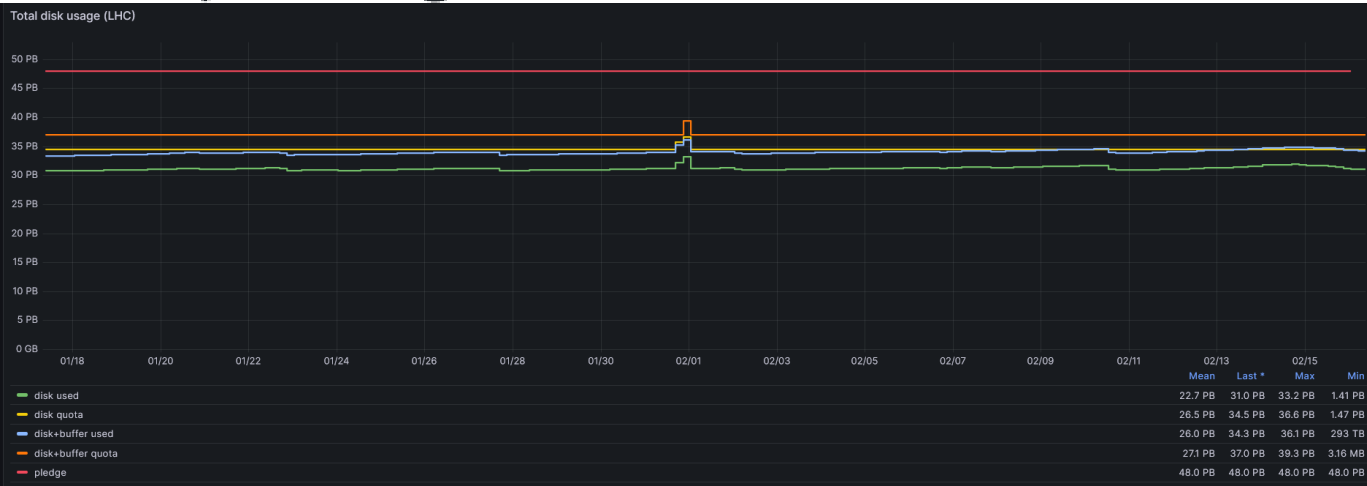
Fondi Terabit+DARE

Nodi CPU	Nodi GPU	NODI FPGA	NODI STORAGE	SW ETH	SW IB
26 (16+19)	30 (21+9)	4 (2+2)	52 (36+16) 20PB raw	12	6

Ieri in delibera approvazione atti gara «HPC Bubbles»
2 mesi per avere il contratto
Da avviare preparazione Appalti Specifici

Disk Usage @T1 ALL VOs – no cloud


Disk space usage LHC



DISK

- about 15 PBN underpledge
- Missing 2022 and 2023 resources
- Extending maintenance contracts for older systems that should be dismissed
- First 64PBN of the 23-24 AQ ready to be delivered and installed at Tecnopolo – only the contract is missing

Disk space usage NO-LHC



eri in delibera l'Appalto Specifico su AQ storage 23-24
Avvio anticipato per iniziare subito l'installazione
Il materiale per i primi 4PB è già al tecnopolo
(mancano solo gli switch, in arrivo la prossima settimana)
La gara 2022 non siamo ancora riusciti a sbloccarla

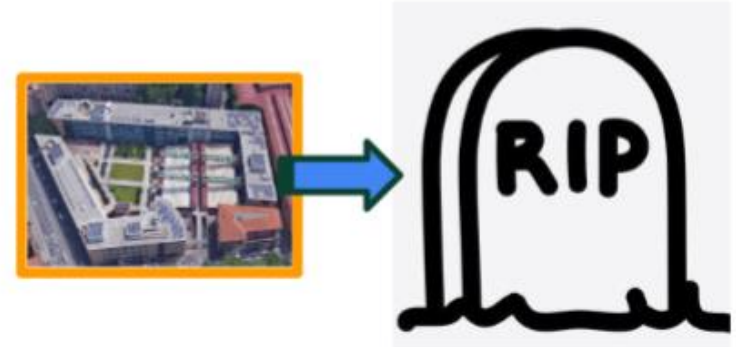
CNAF Tape Libraries and Drives

- **1 x Oracle SL8500**
 - **1 tape library with 16 tape drives T10000D** (8.5TB/cartridge)
 - 80PB installed, 64PB USED
 - Repack on the other libraries needed
 - After completion of repack this library will be dismissed
- **2 x IBM TS4500**
 - **1 tape library with 19 tape drives TS1160** (20TB/cartridge)
 - 102 PB Installed, 50PB USED
 - cannot be further extended due to physical constraints in the current room
 - This library will be moved to the new data center
 - **1 tape library with 18 tape drives TS1170** (50TB/cartridge) acquired and will be installed at new data center Q1 2024
Tender completed and approved by the GE – the contract is missing




CNAF Tape Libraries and Drives

- **1 x Oracle SL8500**
 - **1 tape library with 16 tape drives T10000D** (8.5TB/cartridge)
 - 80PB installed, 64PB USED
 - Repack on the other libraries needed
 - After completion of repack this library will be dismissed
- **2 x IBM TS4500**
 - **1 tape library with 19 tape drives TS1160** (20TB/cartridge)
 - 102 PB Installed, 50PB USED
 - cannot be further extended due to physical constraints in the current room
 - This library will be moved to the new data center
 - **1 tape library** **Gara Assegnata e atti deliberati – siamo in attesa del contratto – sono emersi problemi durante i controlli di AC sull’Operatore Economico**



Storage Usage @T1 ALL VOs – no cloud

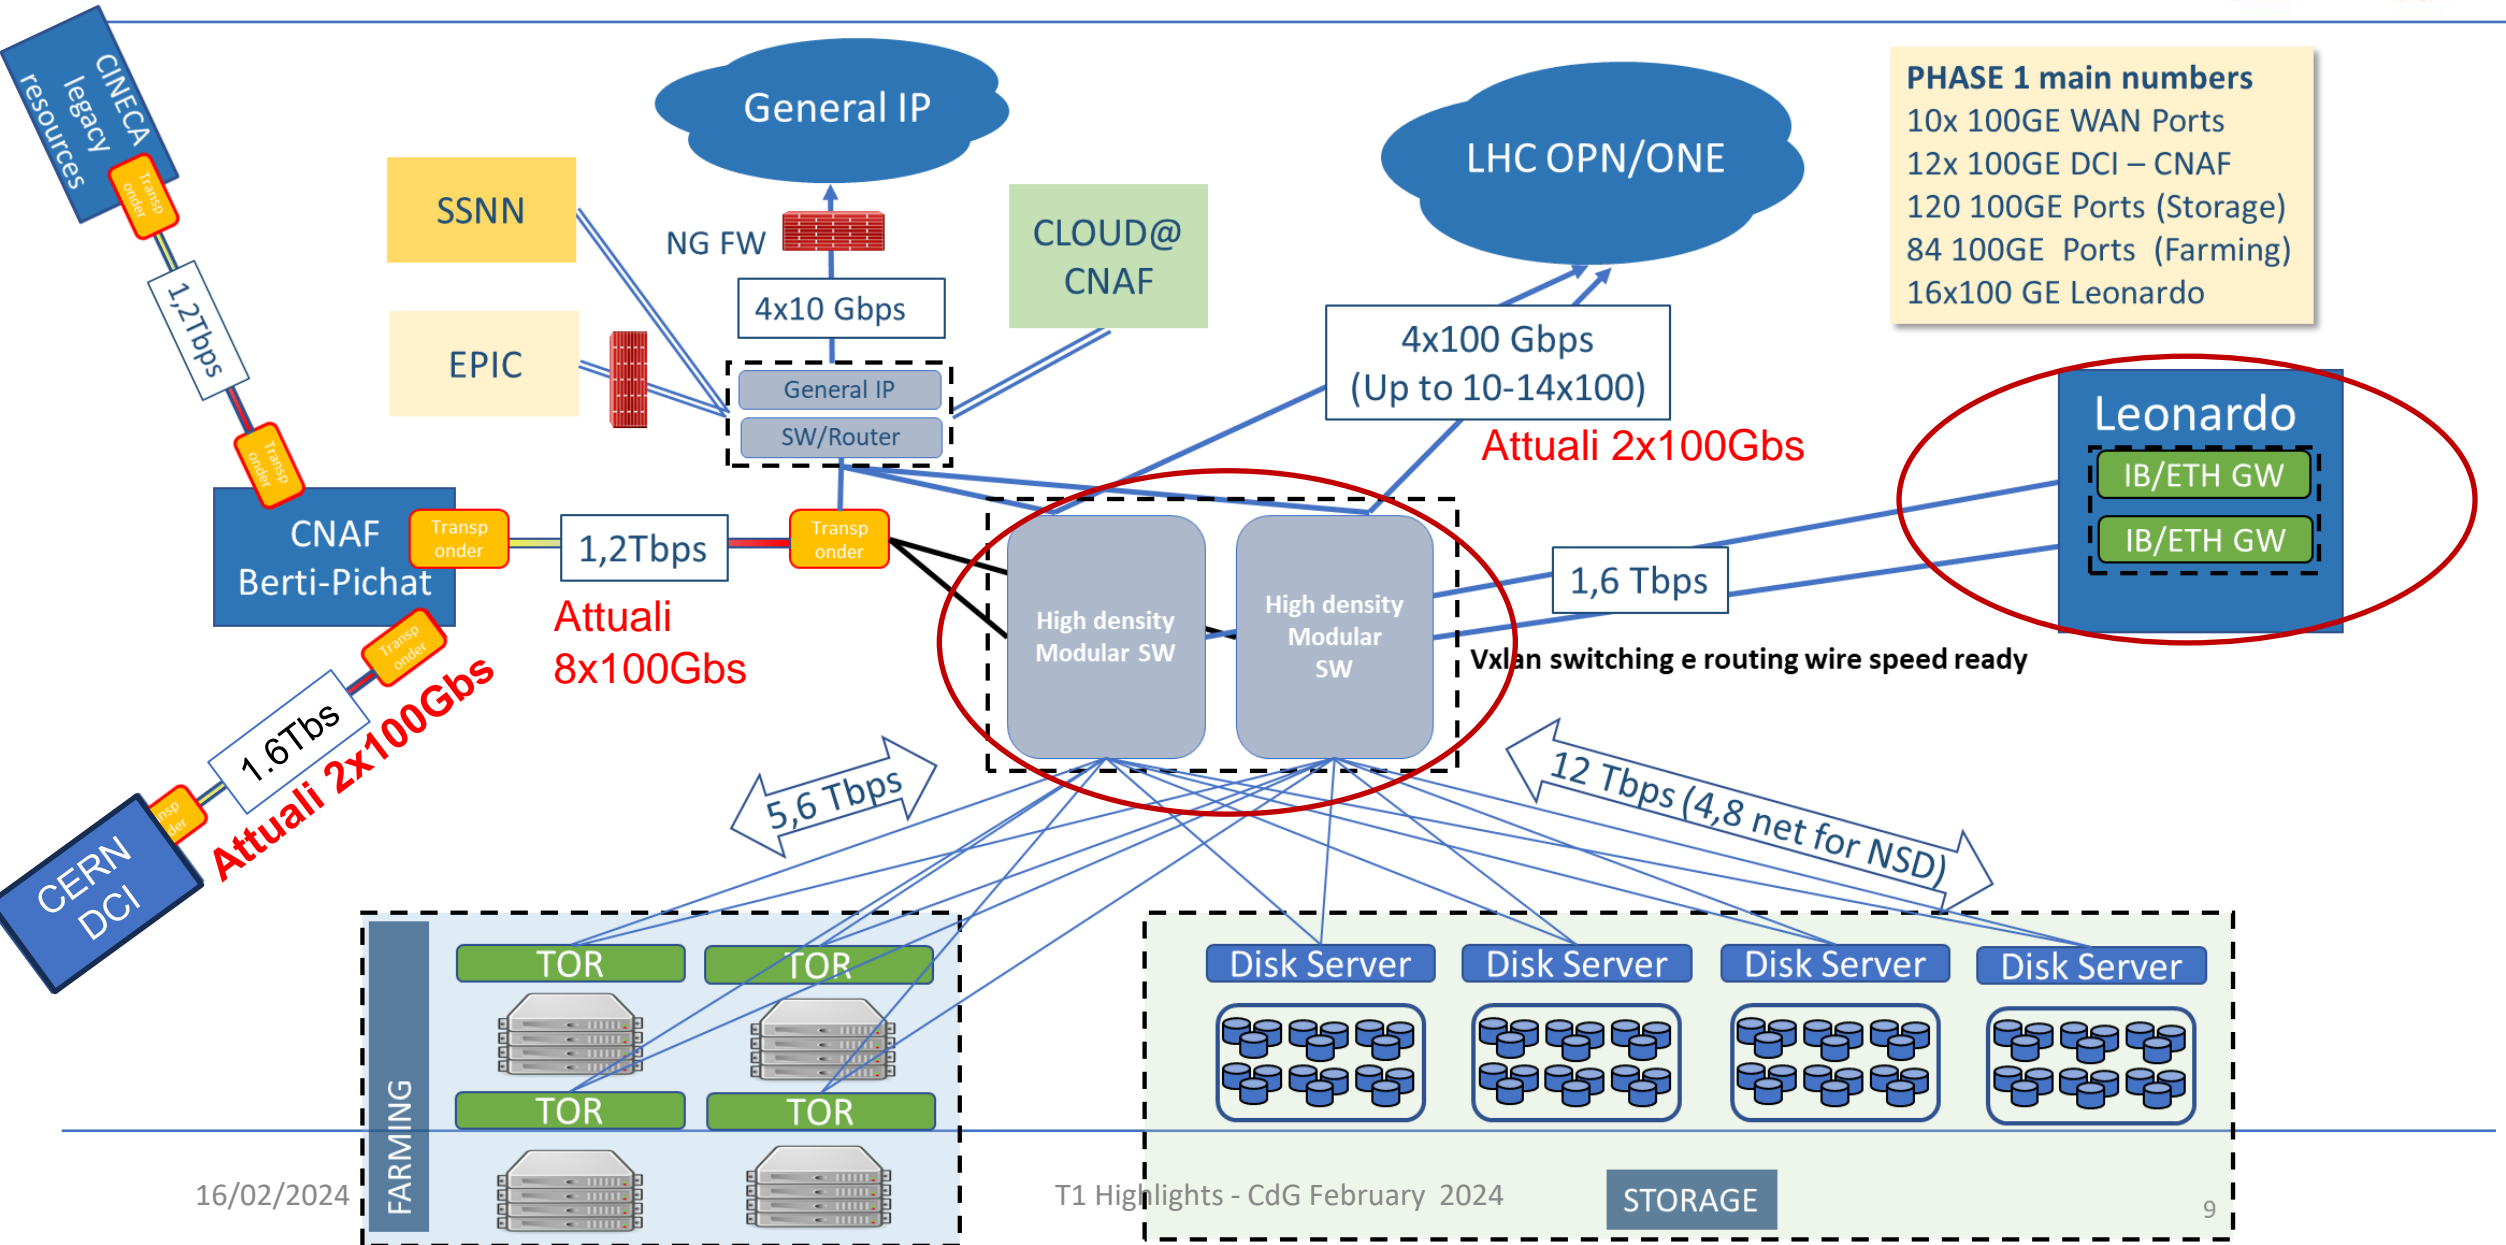


TAPE

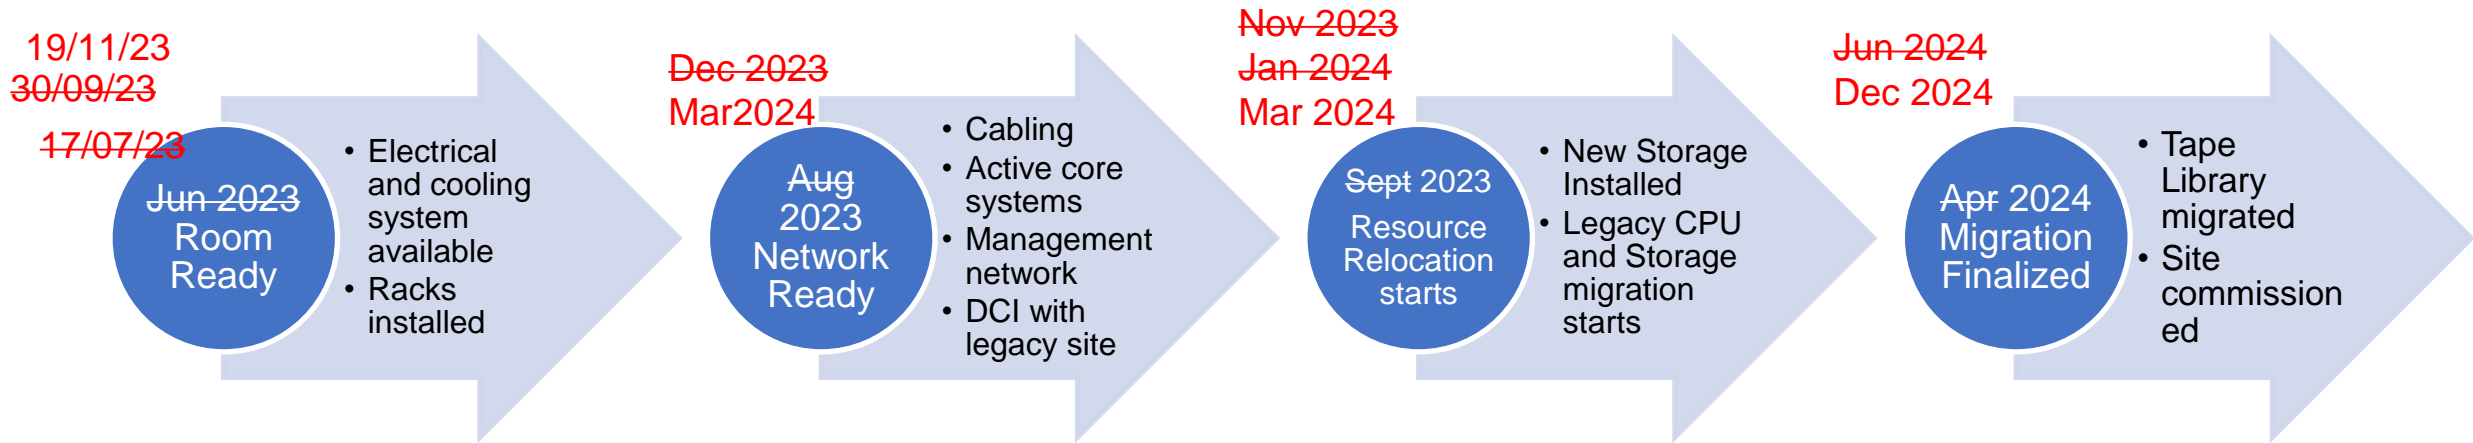
- Installed greater than the pledge by 30PB
- due to Overpledge2023 + REPACK needs

In preparazione un tender da circa 100PB per nastri sia in nuova che in vecchia tecnologia
In preparazione anche un acquisto “rapido” per tamponare in caso di tempi lunghi della gara

Networking Infrastructure



Live Relocation Timeline



- **Consegna sala rinviata innumerevoli volte per vari problemi tra cui amianto rinvenuto in un cantiere vicino (edificio ENEA)**
- **Il 04/12/2023 iniziati i lavori per il cablaggio di rete**
 - **Completamento cablaggio finito entro il 15/03**
- Appareti di rete attivi (core switch) installati – circa un mese per completarne la configurazione, incluso DCI con il CNAF attuale
- Da metà marzo 2024 iniziamo a trasportare il materiale dalla sede attuale
- Storage da AQ 2023-2024 inizio installazione dalla prossima settimana (primi 4 PB)





The new INFN Data Center at Bologna Tecnopolo











Sala
Espansione
07/02/24

16/02/20



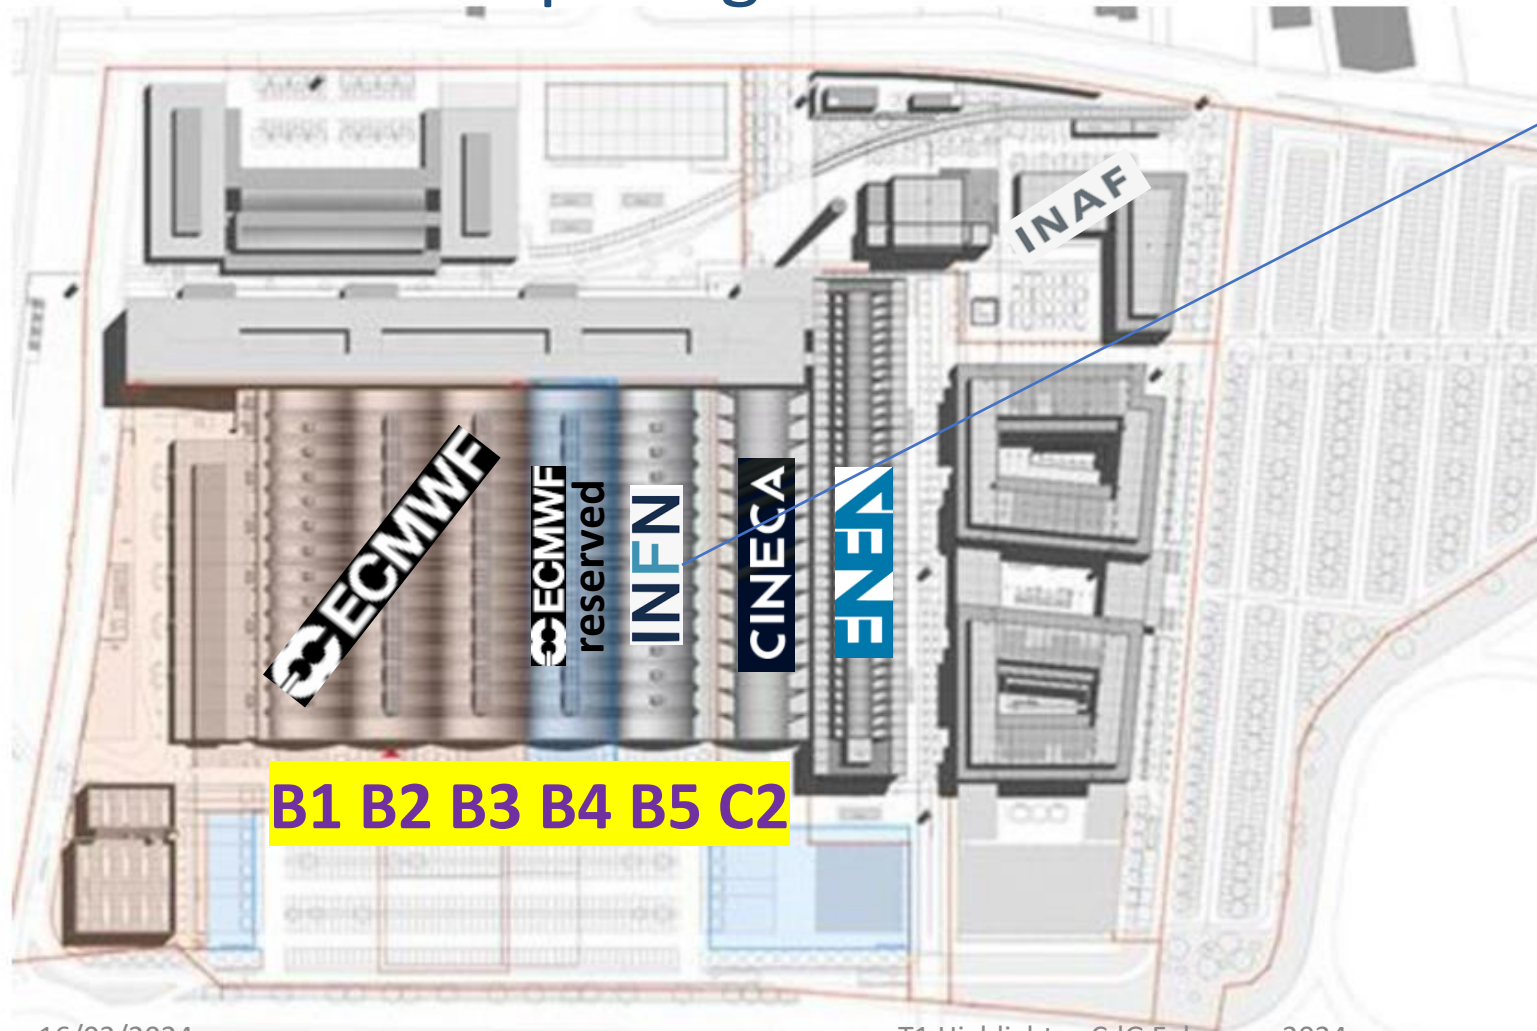
Sala Alta densità
07/02/24



Backup

What can the Tecnopolo host?

The computing infrastructures

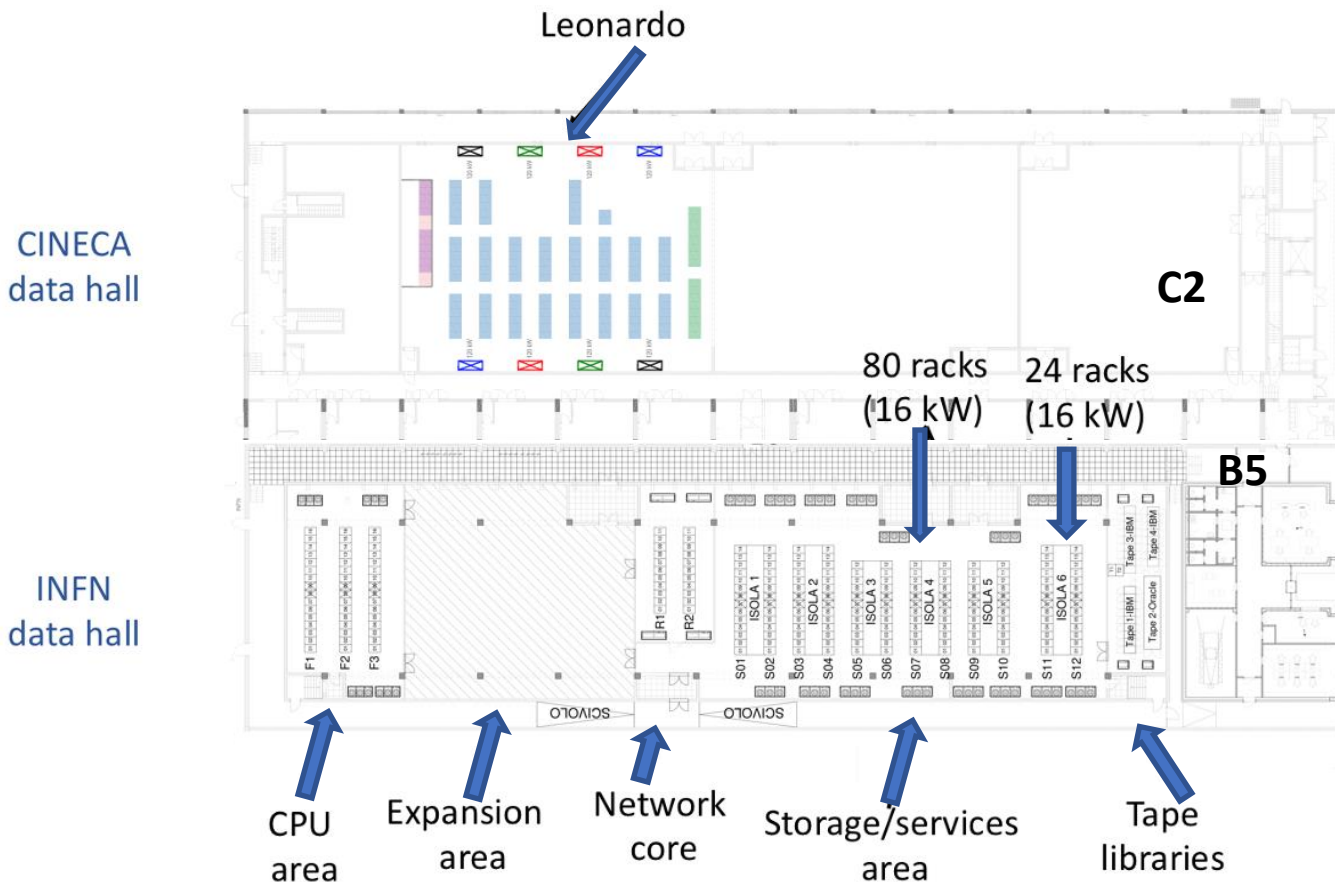


Each of the 6 “botti” (barrels) is
~5000m² of usable IT space



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

CNAF and CINECA data halls



DLC 80kW



16/02/2024



T1 Highlights - CdG February 2024

- 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)

$$\text{PUE}_{\text{DLC}} \approx 1.08$$

$$\text{PUE}_{\text{Tot}} \approx 1.2 - 1.3$$