Referaggio calcolo GR1 Esperimenti non-LHC 2025

b.Giacobbe

July 2024

PROCEDURE

- As usual thanks to all experiments and to all referees for useful interactions, information, ideas, solutions, etc…

- This year I created a shared file to minimize risk of mistakes (18 experiments to review !)
  - https://docs.google.com/spreadsheets/d/1LsJLSRgcHGgFLm2IpaqTjZI2vEN4NeSV6sltaworkYE/edit?pli=1&gid=820258402#gid=820258402

BELLE-II

Requests: 60 TB TAPE Tier1

Decision: final

CPU/DISK/TAPE in Italy vs requests/milestones

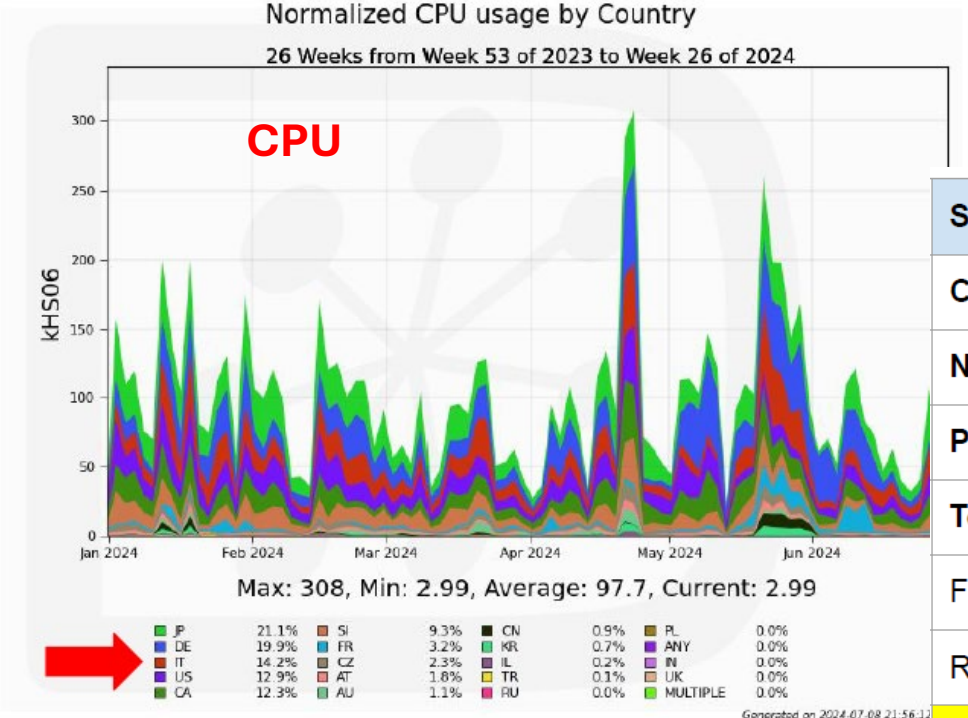

DISK - Italy

SITE	TB
CNAF	820
Napoli	590
Pisa	200
Torino	350
Frascati	11
Roma3	2
Total Pledge	**1.973 TB**
Gran Toal	**1.986 TB**

TAPE - Italy

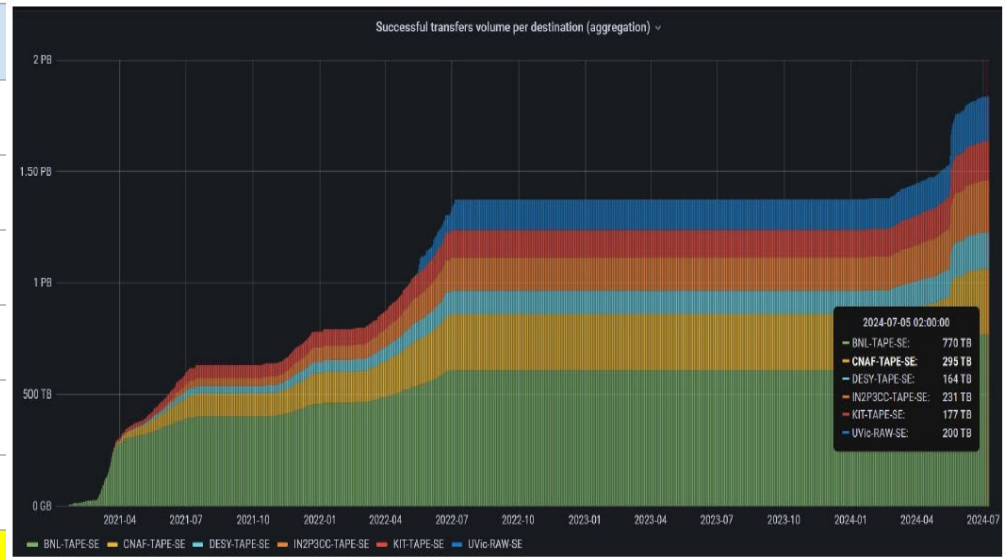

Italy 20% of total

Milestone 2024: 12%
Italy so far: 14.2%

Next challenges

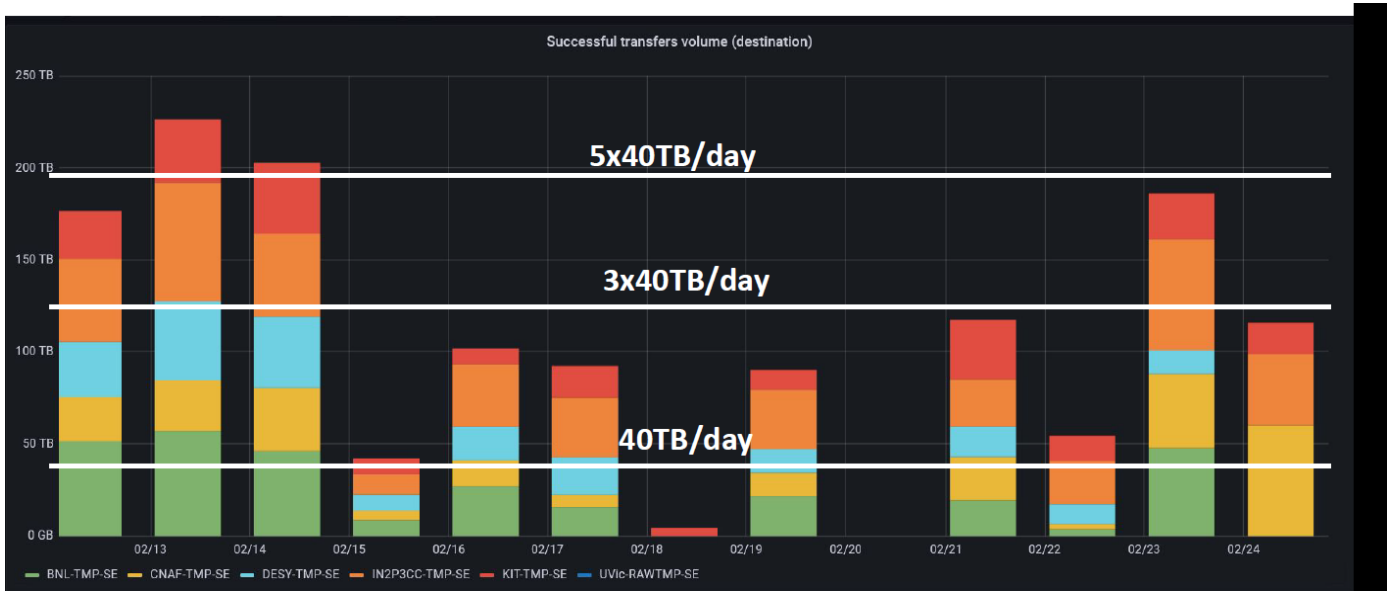

- Emulate data-transfer conditions in high-lumi scenario: 40 TB to be transferred/day to all sites (20% to CNAF)
- CentOs7 deprecation ➔ migration to AlmaLinux/RHEL 9
  - Consequences on distributed computing
- Jennifer Computing Workshop discussed ideas on:
  - Cloud integration in the computing model for user analysis

BELLE-II: Tier1 resourc

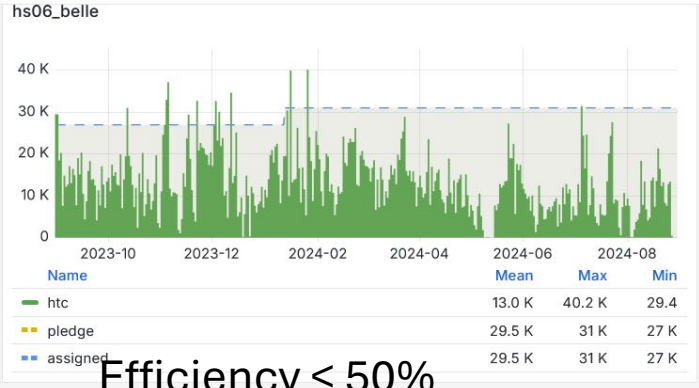

Efficiency < 50%

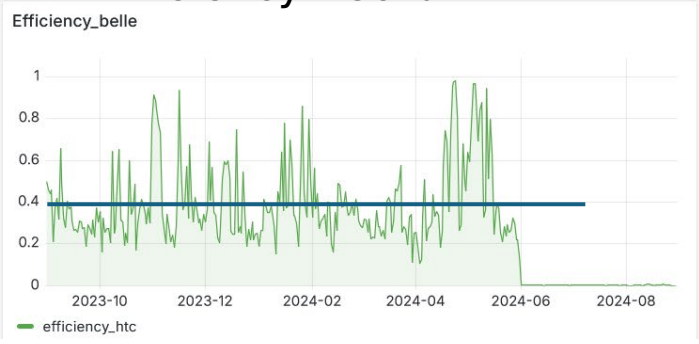

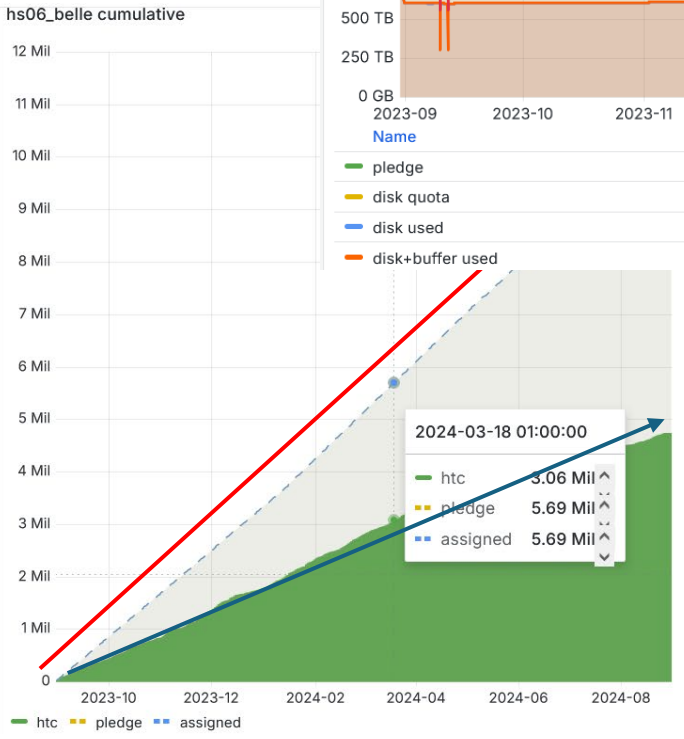

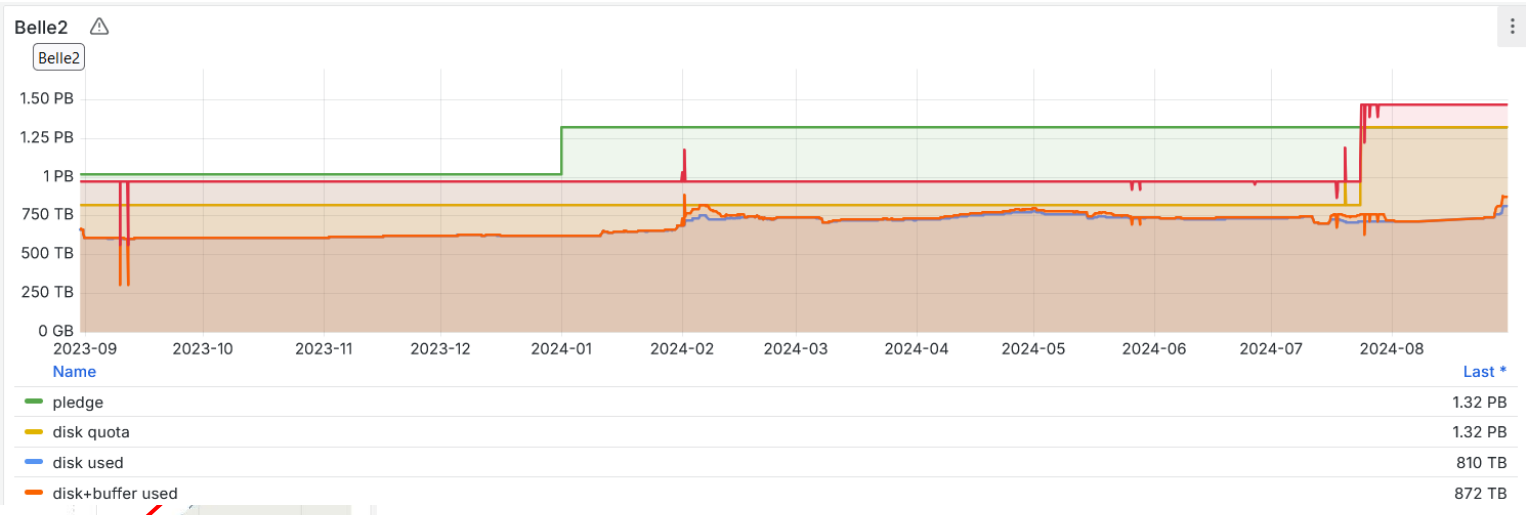

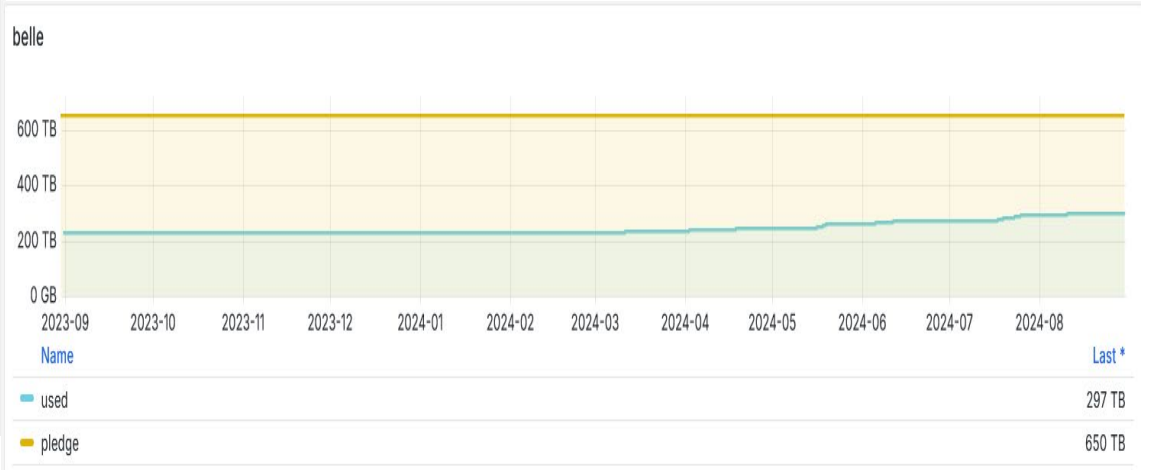

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHSP06	13	31	31	42%	42%	not optimal CPU usage/efficiency
DISK TB	810	1320	1320	62%	62%	
TAPE TB	300	650	650	45%	45%	low tape usage but not negligible

Resources in Italy & requests for 2025

SITE	CPU	STORAGE	TAPE
CNAF	31kHS06	820TB	650 TB
Napoli+Cosenza	19.9kHS06	860TB	
Pisa	8kHS06	200TB	
Torino	6kHS06	350TB	
TOTALE OGGI	**64.9kHS06**	**2.230TB**	**650TB**

+500 TB pledge 2023-2024 now available

- No reprocessing foreseen in 2025 ➔ relaxed requests (will increase in 2026)

	2025	2026	2027	2028
Total tape (PB)	11.8	15	20.4	25.9
Total disk (PB)	19.4	25.3	30.3	37.7
Total CPU (kHS06)	247	492	547	643

	Pledge Italia 2024	Provided	Pledge Italia 2025	Needs
TAPE (TB)	490	650	710	+60
DISCO (TB)	2.522	2.230	2.460	No needs
CPU (kHS06/kHS23)	64,86	64,9	33,17	No needs

60 TB Tape on **Tier1**

OK

0.6 kEUR (GE)

AMBER/COMPASS

Requests: TIER1 (50 TB disk + 1 kHSP06) & Tier2 TO & LOCAL FARM TS

Decision: final

Richieste calcolo

- TIER1: prima richiesta calcolo
  - per simulazioni apparato
- TN (farm locale non finanziata da CSN1): CPU su farm per analisi finali
- TO TIER2
  - Ultime richieste 3 anni fa
  - Richieste per aggiornamenti CPU (acquistati prima del 2015)
- TS (farm locale finanziata anni fa da CSN1)
  - Ultime richieste 3 anni fa
  - Aggiornamento CPU; in particolare:

Nome	Modello	N. CPU	N. core	Desc CPU	RAM MB	Anno acquisto
farm073	Dell PowerEdge C6100	2	12	Intel Xeon L5640 2.26GHz	49152	2011
farm074	Dell PowerEdge C6100	2	12	Intel Xeon L5640 2.26GHz	49152	2011
farm075	Dell PowerEdge C6100	2	12	Intel Xeon L5640 2.26GHz	49152	2011
farm076	Dell PowerEdge C6100	2	12	Intel Xeon L5640 2.26GHz	49152	2011
farm061	Supermicro X9DRT	2	12	Intel Xeon E5-2630v2 2.60GHz	65536	2013
farm062	Supermicro X9DRT	2	12	Intel Xeon E5-2630v2 2.60GHz	65536	2013
farm063	Supermicro X9DRT	2	12	Intel Xeon E5-2630v2 2.60GHz	65536	2013
farm064	Supermicro X9DRT	2	12	Intel Xeon E5-2630v2 2.60GHz	65536	2013

AMBER/COMPASS : requests 2025

- Tier1: GE
  - 1 kHS06 for detector simulation ➔ 10 kEUR
  - 50 TB disk for simulation temporary storage ➔ 5 kEUR
- Tier2 TO: CSN1
  - 1 kHS06 10 kEUR for CPU substitution of old (<2015) ➔ 10 kEUR
- Local Farms: CSN1
  - TS: 15 kEUR for CPU substitution of old (<2013) ➔ 7.5 kEUR
    - (vedi pagina seguente)
  - ~~TN: 6 kEUR for CPU for analysis ➔ 6 kEUR ➔ NO (*)~~
  - ~~TN: contributo storage ➔ 5 kEUR ➔~~ NO (**)

- (*) = farm locale non finanziata da CSN1
- (**) = contributo a storage Teri1 non da CSN1

50 TB Disk on **Tier1**

5 kEUR

1 kHS06 on **Tier1**

10 kEUR

1 kHS06 CPU: Tier2 TO

10 kEUR

7.5 kEUR TS **local farm** CPU (obsolete)

15 kEUR

Commento su Farm TS

- Farm di Trieste è una farm di Sezione originariamente finanziata da CSN1 (ma utilizzata anche da altri utilizzatori ad es. Alice)
- CPU indubitabilmente obsolete
- La discussione tra referee e con la commissione calcolo ha fornito questo esito:
  1. Per non creare problemi nel breve periodo all'esperimento la richiesta viene parzialmente accolta, **MA**:
  2. In quanto farm di sezione si richiede il co-finanziamento al 50% da parte della Sezione e/o degli altri utilizzatori
  3. Non verranno finanziati in futuro sostituzioni di dischi o cpu che dovessero risultare fuori garanzia
  4. Si suggerisce quindi caldamente una transizione smooth verso il Tier2 di Torino (ad esempio) per adattarsi sul lungo periodo alla politica dell'INFN che non finanzia farm locali:
     1. Copia dei dati
     2. Training del calcolo sul Tier2 di TO
  5. Rimane inteso che questo parizale finanziamento NON constituisce un precedente per le richieste di finanziamento di FARM locali nuove (TS è farm esistente e finanziata anni fa da CSN1)

RD_FCC

No request for 2025 (except for recovery of last year CPU not recorded)

Decision: final

RD_FCC : Tier1 resources

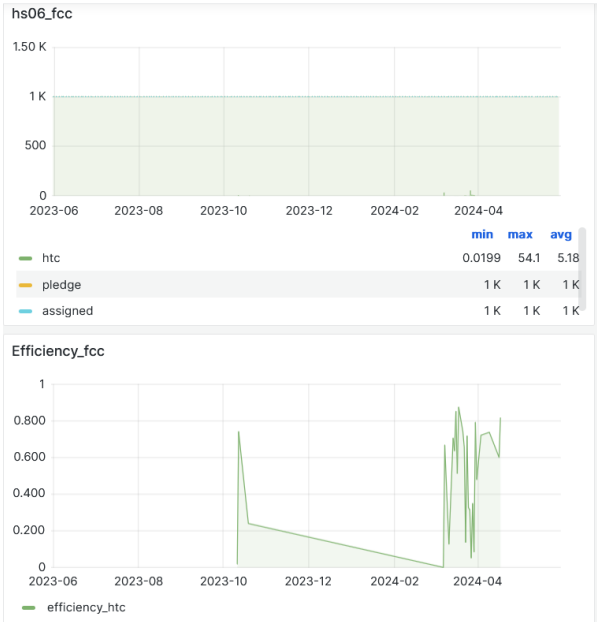

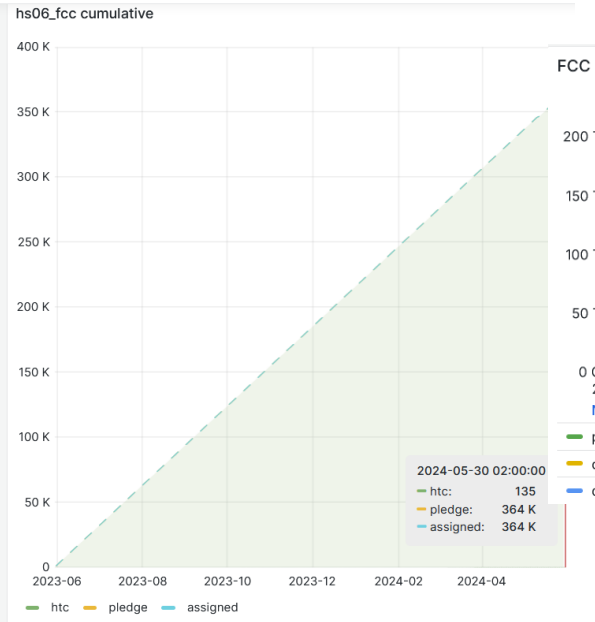

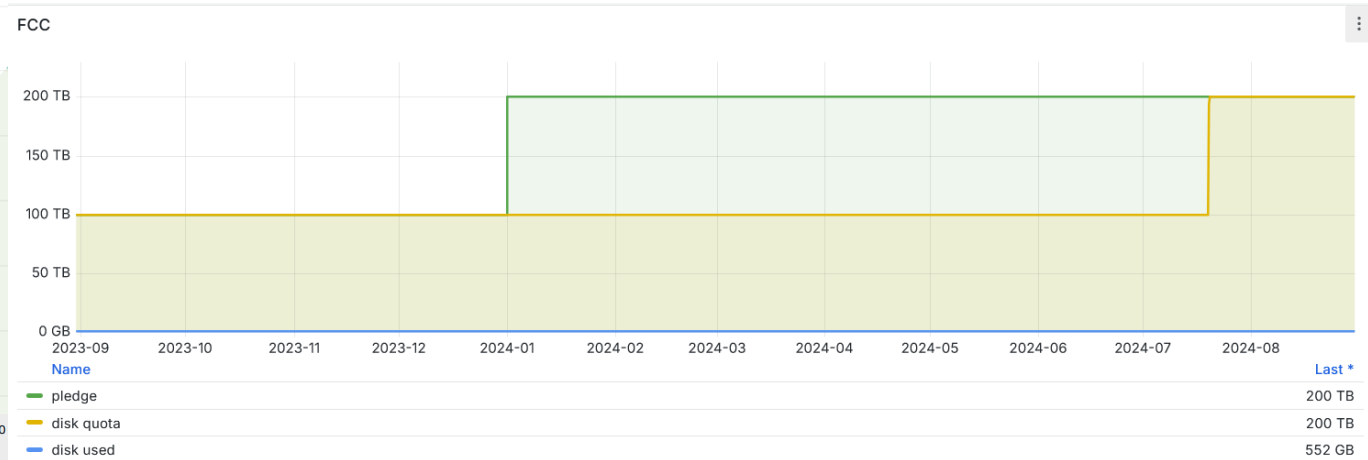

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHSP06	0	1	1	0%	0%	no use of CPU
DISK TB	0.5	200	200	0%	0%	no use of DISK

Should be 2 kHS06: See note at the end of RD_FCC

Current resources for FCC project overall

- VO fcc existing and several Italian users subscribed
- CERN
  - EOS volumes
    - 500 TB for central productions (157 TB free, stil used by some CDR files)
    - 200 TB for analysis, starts to be used
  - CPU: 9000 HS06 on lxbatch, used for FCC analyses by the Italian community
  - Integrated in iLCDirac, already tested with test jobs at CNAF and INFN Bari
    - F. Fanzago (INFN Padova) provides support for testing procedures, users
    - Steps (from test to production):
      run distributed MC production and FCC analysis chain with DIRAC and Italian resources
- Other sites integrated
  - INFN T2 BARI
    - - 10 TB mostly used for testbeam data with drift chamber (about 1TB)
    - 250 CPU for hyperparameter optimization of NN-based cluster counting algo
  - INFN T1 CNAF 200 TB, several users with an account at CNAF for interactive, batch and grid-based processing → limited use so far, testbeam data from DCH and DR communities archived there
  - Glasgow (storage only) and DESY & MIT (work in progress)
- Some GPU resources
  - CERN, EuroHPC, INFN Bari
- FCC interactive resources: already exploited and used

❑ Storage and computing requirements for FCC overall

Table 4 RAW data estimates for FCC-ee

Run	$\sqrt{s}$ (GeV)	Statistics	RAW data
Z	91.2	3×10^{12} Z decays (visible)	3–6 EB
WW	160	10^8 W^+W^- events	0.1–0.2 PB
ZH	240	10^6 ZH events	1–2 TB
$t\bar{t}$	350, 365	10^6 $t\bar{t}$ events	1–2 TB

≈HL-LHC

Table 5 AOD data estimates for FCC-ee

Run	$\sqrt{s}$ (GeV)	Statistics	AOD data
Z	91.2	3×10^{12} Z decays (visible)	15–30 PB
WW	160	10^8 W^+W^- events	0.5–1 TB
ZH	240	10^6 ZH events	5–10 GB
$t\bar{t}$	350, 365	10^6 $t\bar{t}$ events	5–10 GB

Table 6 Time estimated to generate $q\bar{q}$, $\tau^+\tau^-$ and $\mu^+\mu^-$ events at the Z peak

Generator	Process	100k/core (s)	Z sample/core	Z sample/9000 HS06 (days)
Pythia8	$q\bar{q}$	148	4×10^9 s = ~126 y	50–75
KKMCee	$q\bar{q}$	151	4×10^9 s = ~126 y	50–75
KKMCee	$\tau^+\tau^-$	195	0.25×10^9 s = ~8 y	3–4.5
KKMCee	$\mu^+\mu^-$	334	0.44×10^9 s = ~14 y	5–7.7

Table 7 Time estimated to simulate $q\bar{q}$ events at the Z peak

Process	1k/core	Z sample/core	Z sample/9000 HS06
$q\bar{q}$	20k s = 5 h 33 min	6×10^{13} s = ~ 1.9×10^6 y	$2.1–3.2 \times 10^3$ y

2000-3000 years!

Table 8 Time estimated to simulate $q\bar{q}$ events at the Z peak with DELPHES

Process	100k/core	Z sample/core
$q\bar{q}$	212 s	6.4×10^9 s = ~ 202 y

- ❑ We would host a fraction of 15 % of data at INFN in the next few years
- ❑ MC event generator can be challenging for full-scale production
- ❑ Full simulation times: Geant4 is up to a factor of 2 faster for ATLAS
- ❑ Fast simulation techniques might help
- ❑ Other option: selected simulation, i.e. non simulating what is never touched
- ❑ Reconstruction:Between 10% and (ALEPH) - 30% (BELLE) of simulation
  - ❑ Could really benefit from using heterogenous resources

- ❑ The FCC-ee resource needs are of the same order of HL-LHC
- ❑ The resources currently available are O(1000) off for full simulation/reconstruction
- ❑ Numbers should be multiplied by the number of detector variations and analysis

Plans and computing request for RD_FCC

Steps towards a definition of computing requests:

- import 100 TB of MC samples for physics analysis at CNAF
- exercise the FCC analysis at CNAF and T2 Bari in remote
- test the workflow of the hyperparameter optimization at CNAF
- enforce use of interactive analyses resources , GPUs

Requests:

NO REQUESTS for 2025

- given the limited use of current CNAF resources we don't request additional ones and we plan to use more efficiently the existing ones in 2025
- given the:
  - status of workflow for hyperparameter of NN-based algos for DCH data processing
  - the starting of Higgs HZ analyses and the need of CPU for processing
  - the need of pushing for interactive analyses

  we would need 50TB, 50 CPU and a 2-3 GPU at INFN Bari in the next couple of years so could postpone this request to next year

Commenti su assegnazioni per il 2024

- Le richieste fatte per il 2024 a Settembre 2023 erano:
  - 100 TB disco Tier1 aggiuntivi rispetto ai 100 TB esistenti: ok
  - 1 kHS06 aggiuntivo rispetto a quanto già a pledge: forse per misunderstanding non sono stati assegnati ➔ si propone assegnazione in questa riunione.

kLOE

Requests: Maintenance libraries & Disk

Decision: final

KLOE: data preservation & computing model

- Data reconstruction and preservation:
  - RAW data on TAPE for preservation
  - ROOT files on disk for analysis
    - KLOE-II : Data & MC fully reconstructed
    - KLOE-I : started production of data & MC for 2001-2006 data taking
- Computing Infrastructure
  - RAW & DST files on Tape libraries (Old+new) ➔ data migration: next slide
    - TAPE available also at TIER1 but low usage (old data only)
  - ROOT files on LNF-KLOE server disk for storage & analysis
  - CPU on LNF-KLOE server for analysis
- Due to problems of LNF-KLOE server (see after) LNF provided 650 TB of disk space to backup all relevant areas (DST, ROOT, user/group areas)
  - Possible request by LNF to free disk sooner or later ➔ request of 300 TB SJ (+250 approved in 2024) at the LNF computing center

Problems with infrastructure

- Problems:
  - KLOE server: some disks damaged after controlled shutdown (Temperature alarms)
    - 10% of data lost
    - Available 650 TB from LNF for temporary backup ➔ backup done
  - OLD library:
    - First hardware problems (robotic arm, some units)
      - Company called for maintenance not able to solve
      - F. Fortugno, only expert, managed to solve problem with home-made procedure
    - New problem with dust and mixture of water and oil on robotic arms
      - Not possible to solve yet
  - No computing expert anymore
- New library OK but copy of data hanging due to OLD library unavailability
  - 90% of RAW copied
  - Only 30% of reconstructed DST copied
- LNF director promised AdR for KLOE computing ➔ effective if AdR interacts with CNAF ! Anyway difficult to identify person ➔ HELP REQUESTED !

Referees concerns and proposed strategy

- Infrastructures show criticality & weakness
  - Person-power dedicated is now absent
- Data must be preserved and analyses completed
- Libraries strategy:
  - D. Cesini provided companies hopefully able to fix the problem with OLD library
    - same companies ensuring maintenance of TIER1 libraries ➔ funding requested
    - Possible alternative: use robotic arm of another library of the same type ➔ to be investigated
  - If fixed, finalize the copy into the NEW library ➔ contract for NEW lib maintenance requested
- Longer-term strategy that makes KLOE computing model less critical discussed with the experiment and the referees:
  1. **Gradually move to Tier1:**
     - Copy of all ROOTUPLES to CNAF in order to perform all final analyses @ T1 (allows easy access also to external groups (Liverpool) ➔ rootuple analysis possible on Tier1 processors (not reconstruction of old data)
     - Safety backup of all data into Tier1 TAPE (3.1 PB already available)
  2. keep the NEW library (complete data transfer from OLD lib)
  3. One year to organize next steps depending on OLD/NEW lib evolution

KLOE: Tier1 resources

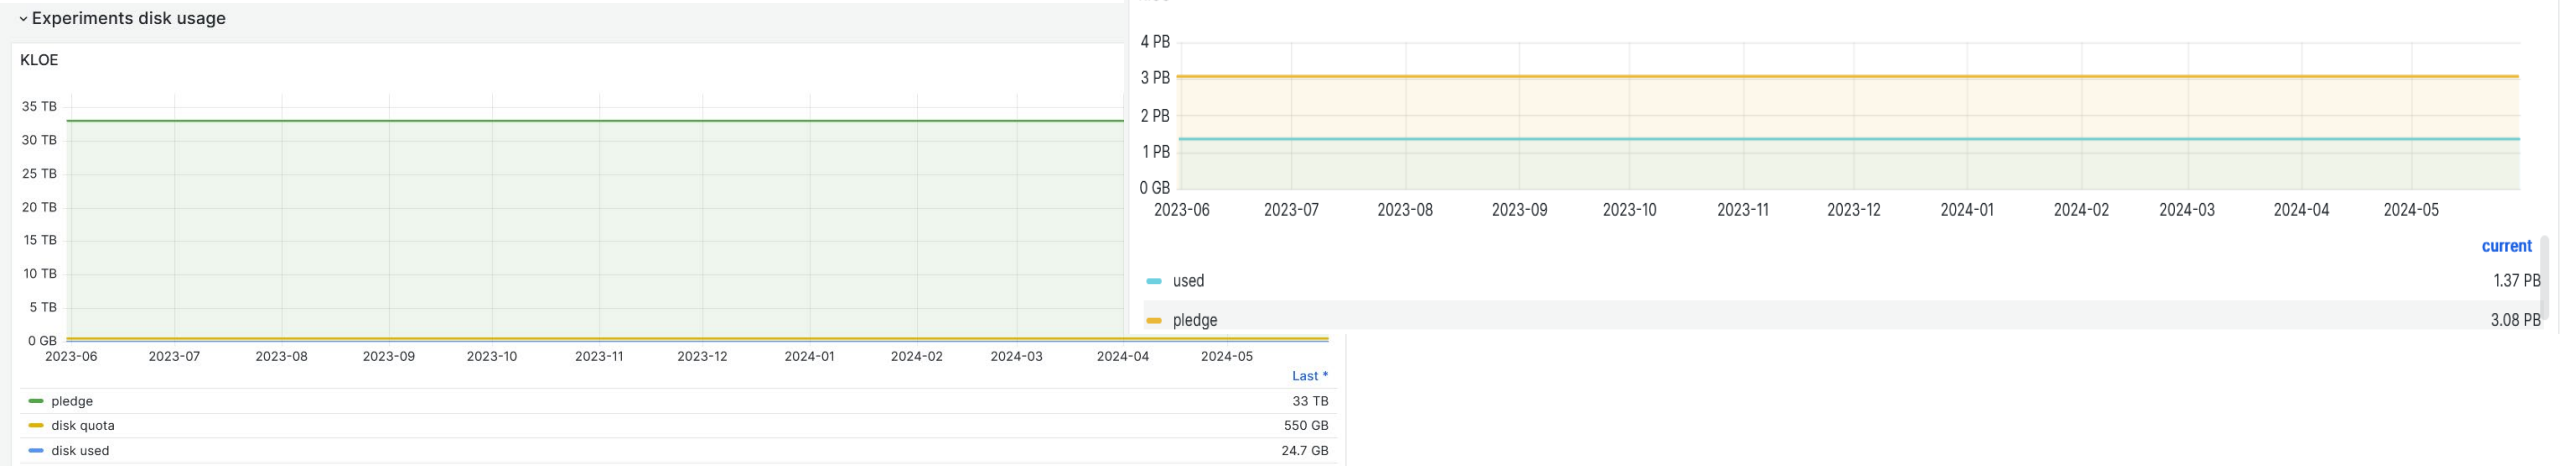

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
DISK TB	0.02		33		0%	no use of disk
TAPE TB	1370	3100	3100	45%	45%	old data presumably

1.7 PB Tape available

33 TB disk already available

KLOE requests for 2025 & referee proposal

- Maintenance of the NEW library ➔ 26 kEUR - OK CSN1
  - Not the 3 kEUR SJ
- Maintenance of the OLD library ➔ 5 kEUR SJ – OK CSN1

CSN1
26 kEUR New lib maintenance
5 kEUR (SJ) for OLD library
OK
26 + 5 (SJ) kEUR

- ~~300 TB @ LNF computing center~~
  - ~~quoted 61 kEUR by LNF computing center, double than @ TIER1~~ ➔ NO
- BUT:
  - 300 TB disk @ TIer1 ➔ 30 kEUR (Giunta Esecutiva)
    - Note: 33 TB already available
  - 1 kHS06 @ Tier1 ➔ 10 kEUR (Giunta Esecutiva)
    - To perform data analysis of ROOTUPLES ➔ incentive to actively pursue the new strategy !

1 kHS06 CPU: Tier1
ok
10 kEUR

300 TB disk @ Tier1
OK
30 kEUR

KLOE situation from 2025 @ T1

- Full “chain” for data analysis & preservation @ T1
  - 330 TB disk for ROOTUPLE for final data analysis
  - 3.1 PB Tape for RAW/DST/ROOTUPLE storage
  - 1 kHS06 for final data analysis
- KLOE server to be used only for finalization of data reconstruction & ROOTUPLE production (impossible @ T1)
- OLD library hopefully fixed and then dismissed after copy finalization
- NEW library will remain at least for the time being

LHCf

Requests: TIER1 6 kHS06 + 50 TB TAPE

Decision: final

LHCf : Tier1 resources

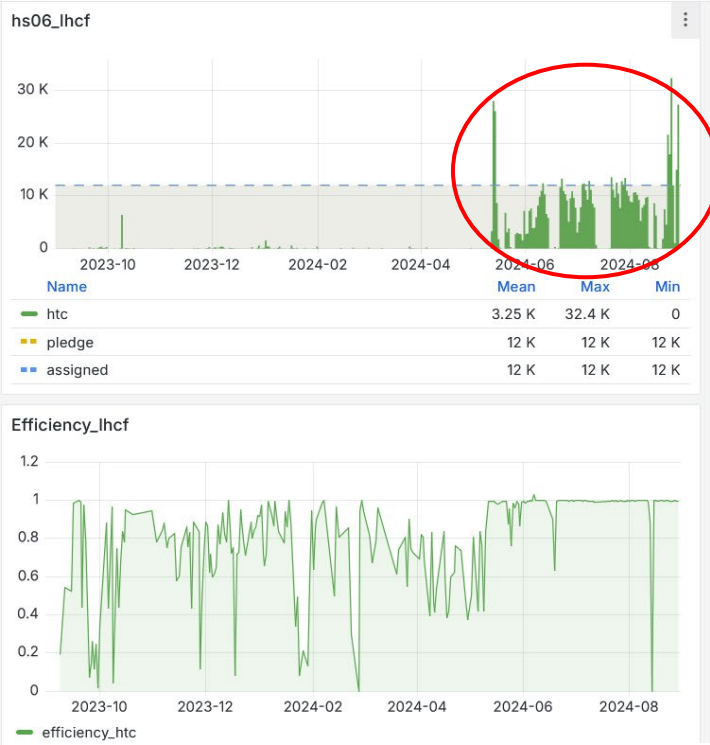

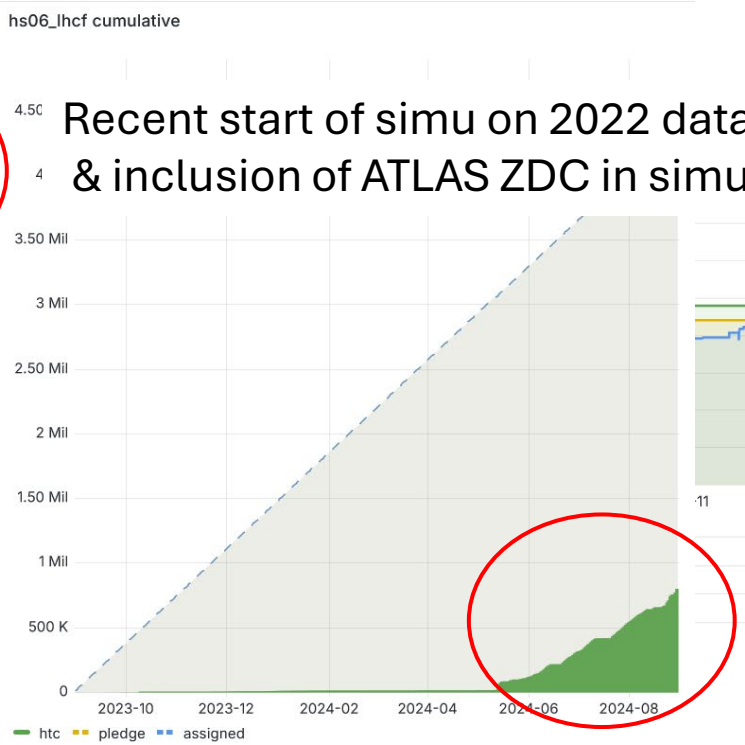

Recent start of simu on 2022 data
& inclusion of ATLAS ZDC in simulation

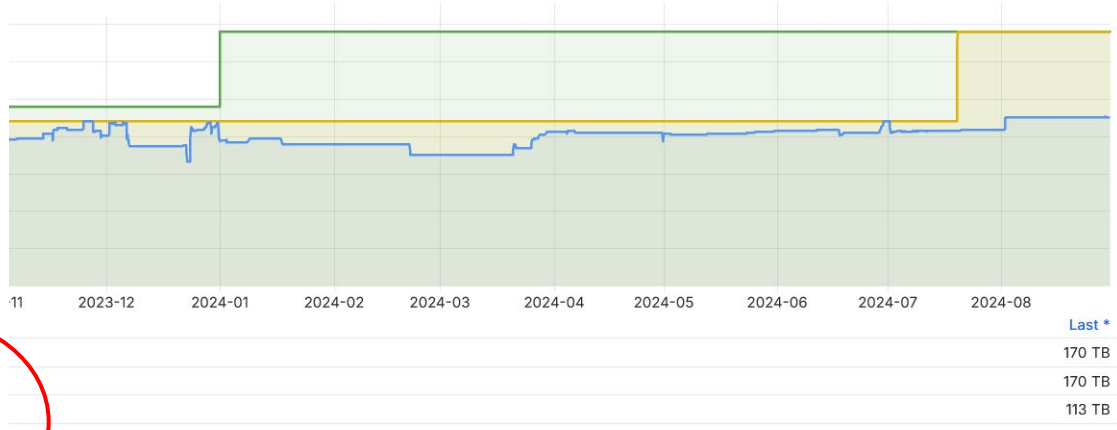

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHSP06	3.2	12	12	26%	26%	recent start of CPU heavy use
DISK TB	113	170	170	67%	67%	

plans

- Foreseen simulations:
  - 2022 13 TeV in 2 different detector positions
    - A) 2.5x10^8 collisions with EPOS-LHC, QGSJetII-04 and SIBYLL 2.3
    - B) 2.0x10^8 π0 events
    - C) 2.5x10^7 η0 events
    - D) 2.5x10^7 K0

    } 10^9 x 2 collisions
  - 2025 p-O 9.9 Tev (presumably):
    - Reasonable initial value (to prepare the run): **10^8** collisions to be increase to 10^9 in the next years
- With present resources (12 kHS06) by end-2024 p-O initial set + small part of 2022 can be simulated.
  - Note: 1 collision takes about 3 minute to be simulated including ZDC+LHCf
- With present resources full data set needs 1.5 years to be simulated
  - Increase by 50% CPU would correspongingly decrease needed time and speed up analyses (see next)

Foreseen analyses & referee comments

- Foreseen analyses (in 2 detector positions):
  - K0s: 2022
  - pi0: 2025, 2022, 2016 [+ 2015 (colpevolmente in ritardo)]
  - eta: 2025, 2022, 2016
  - gamma: 2025, 2022, 2016
  - neutroni: 2025, 2022, 2016
- Comments:
  - Plan of analysis very demanding with short person-power
  - Simulation approach non really optimized (long simu time / event) BUT investment on optimization out of reach for small community and would delay analysis and publications

Requests 2025

- Requests on TIER1: GE
  - 50 TB TAPE ➔ 0.5 kEUR
    - in line with disk availability
    - Will allow better use of disk by storing Level2 and Level3-OLD data
  - 6 kHS06 CPU ➔ 60 kEUR
    - Important request but justified by demanding simulation program
    - No requests in recent times (not optmized scheduling…)
    - Available CPU will eventually go out of maintenance and will have to be replaced. Granting new resources should allow to terminate simulation before that and prevent future purchase to substitute old resources

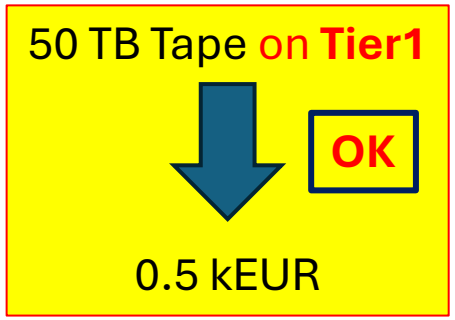

RD_MuColl

Requests: Cloud Veneto 150 TB disk

Decision: final

RD_MuCOLL : Tier1 resources

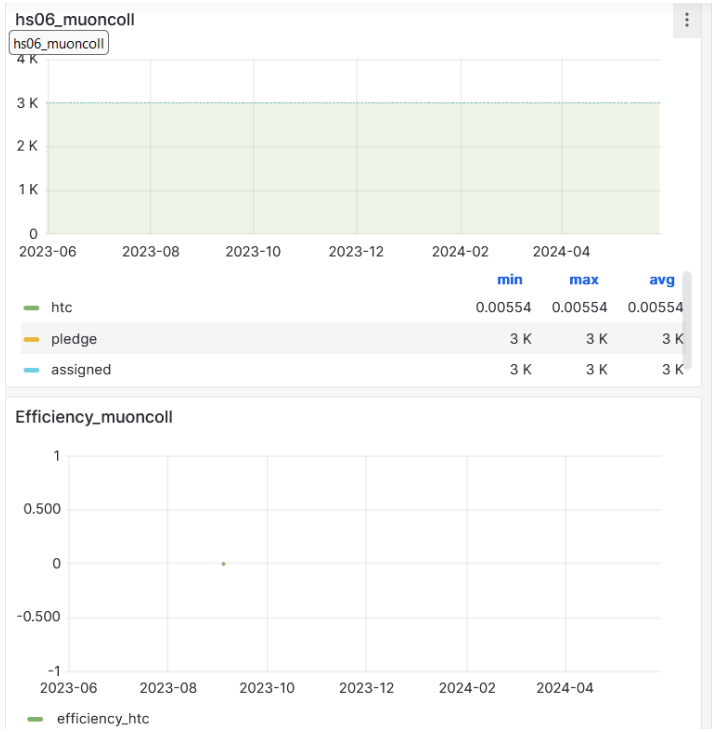

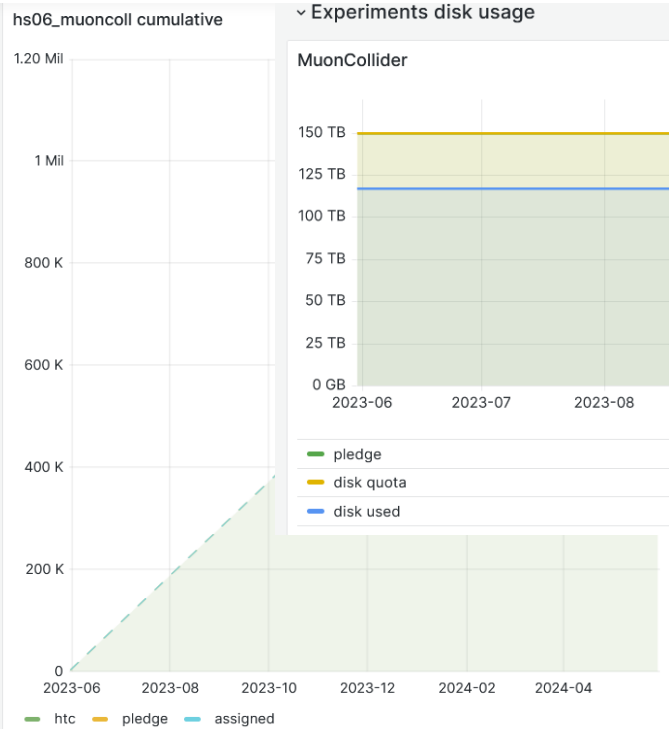

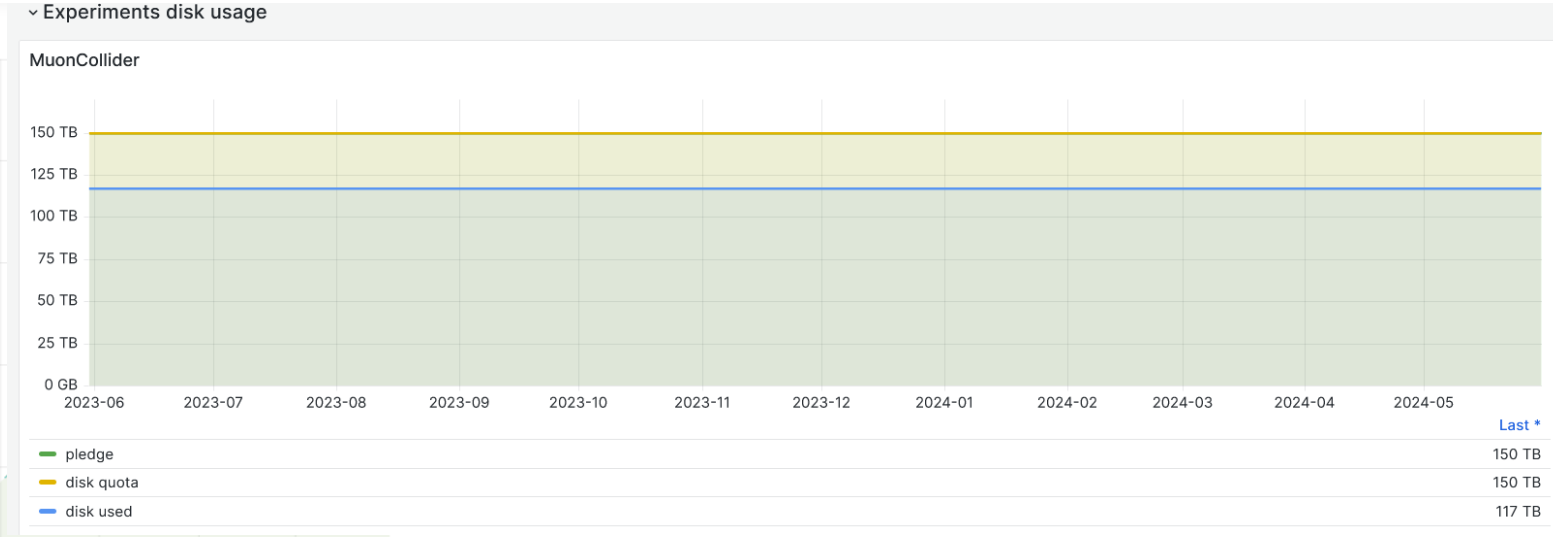

Cloud Veneto disk used @ \sim 90%
Note that in september the 150 TB requested in 2022 will be available

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHS06	0	3	3	0%	0%	no use of CPU
DISK TB	117	150	150	78%	78%	

Risorse calcolo RD_MUCOL

- **Cloud-Veneto**: 200 VCPU, 740 GB di RAM, ~100 TB di storage
- **CNAF**: batch system basato su HTCondor, 150 TB di storage, 6 CE
- **IBISCO-Bari**: risorse condivise con altri progetti allocate al momento della richiesta
- **CERN**: batch system basato su HTCondor, 300 TB di storage su CERN EOS
- **Risorse locali**: Farm Trieste (modalità opportunistica), Pavia etc.
- **Richieste 2023-2024 (ancora da acquisire)**: 150 TB storage e 512 GB di RAM su Terabit/Cloud-INFN

In arrivo

- **Richieste 2025**: altri 150 TB di storage su Cloud-Veneto

Cloud-Veneto

Report sull'utilizzo per periodo	2023-07-24	2024-07-24				
ID progetto	1d4bbed70b794917acaaaa69990873fb					
Istanze attive:	16					
Utilizzo totale VCPU (Ore):	1130997	4				
RAM totale attiva (MB):	753664					
Utilizzo totale memoria (Ore):	4829081160	46				
Dimensione totale disco (GB):	406					
Utilizzo totale disco (Ore):	2060912	63				
Nome Istanza	VCPUs	RAM (MB)	Disco (GB)	Utilizzo (Ore)	Age (Seconds)	Stato
MuonC_V02_07_A	8	32768	25	3059,15	11012944	Attivo
MuonC_v02_07_B	8	32768	25	5829,09	20984713	Attivo
MuonC_32	32	32768	28	3064,66	11032771	Attivo
MuonC_Lorenzo	8	16384	25	880,94	3171368	Attivo
Whizard	32	32768	28	2904,32	10455560	Attivo
MuonC_UI	8	8192	25	3059,32	11013560	Attivo
MuonCServer	2	4096	25	882,22	3176000	Attivo
MuonC_Users01	8	32768	25	8793,11	35850301	Attivo
MuonC_Monster	32	348160	25	788,22	2837602	Attivo
MuonC_Up01	8	32768	25	883,91	3182071	Attivo
MuonC_Users02	8	32768	25	882,36	3176493	Attivo
ILC_Nazar	8	32768	25	2202,08	7927503	Attivo
TestInst	8	32768	25	214,82	773337	Attivo
MuonC_Stream	8	32768	25	3058,83	11011787	Attivo
TestSpack	8	32768	25	47,12	169629	Attivo
MDI	8	16384	25	3059,08	11012684	Attivo

- Risorse in condivisione tra INFN e Università
- **Accesso via INFN IdP, necessita account centralizzato**
- Grazie alla flessibilità della Cloud (ad es. nessun limite ad allocazione RAM o al tempo massimo di esecuzione), riusciamo a girare qualunque tipo di job (generazione, simulazione, ricostruzione etc.)
- Spesso le risorse Cloud-Veneto sono al limite, sia per l'esaurimento dello storage che per le CPU occupate

```
10.64.3.12:/mnt/home   59T  53T  2.4T  96% /users/muoncollider
10.64.3.12:/mnt/data   30T  25T  3.3T  89% /muoncdata
```

CNAF

- **Accesso via VOMS,** istruzioni https://confluence.infn.it/display/muoncollider/Storage+Element
- **In questo momento usato solo come storage**
- Riscontrati problemi dovuti alla durata dei jobs a alla scadenza del proxy, soprattutto per jobs di ricostruzione
- Attuale occupazione Disco: 117 TB/150 TB

IBISCO-Bari

- Le risorse di IBISCO-Bari sono state acquisite da RD_MUCOL nel 2022: 7k HS06 e 300 TB
- Destinata alla produzione campioni di b, c e light jets (full simulation + ricostruzione con BIB) per studiare algoritmi di ricostruzione/identificazione basati su machine learning, e a campioni per lo sviluppo di HCAL
- **Le risorse sono attualmente condivise con gli altri progetti, vengono allocate al momento della loro richiesta**
- L'accesso avviene tramite account a Bari, può essere aperto anche da utenti di altre sedi

+ CERN

Richieste 2025

- **Richieste 2025:** 150 TB di storage su Cloud-Veneto
- Questa richiesta è specifica per Cloud-Veneto: viene usata intensivamente e lo storage è al limite
- Se possibile, **chiederemmo di acquisirli in anticipo nel 2024**, visto che la deadline per la European Strategy è marzo 2025. **Nel 2024 non abbiamo avuto nuove risorse!**
- **Non ci sarebbero problemi con i tempi di acquisto**: Cloud-Veneto può anticipare lo storage su server già presenti a Padova, e acquistarne nuovi alla fine dell'anno

150 TB disk on **Cloud Veneto**

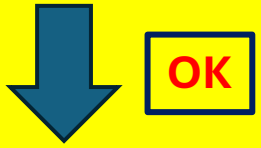

OK

15 kEUR (CSN1)

DOVREBBERO ARRIVARE GIA' NEL 2024

CPU al CNAF

- 3 kHS06 inutilizzati al CNAF
  - Mu_coll già lo scorso anno aveva proposto di renderli disponibili. In tabella riassuntiva ho segnato quindi –3 kHS06 su mu_coll

NA62

No requests

Decision: final

NA62 : Tier1 resources

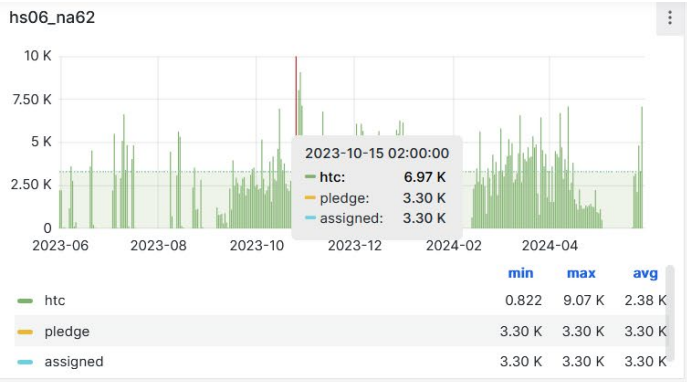

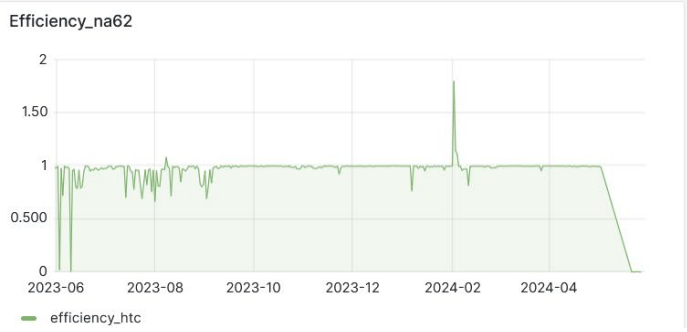

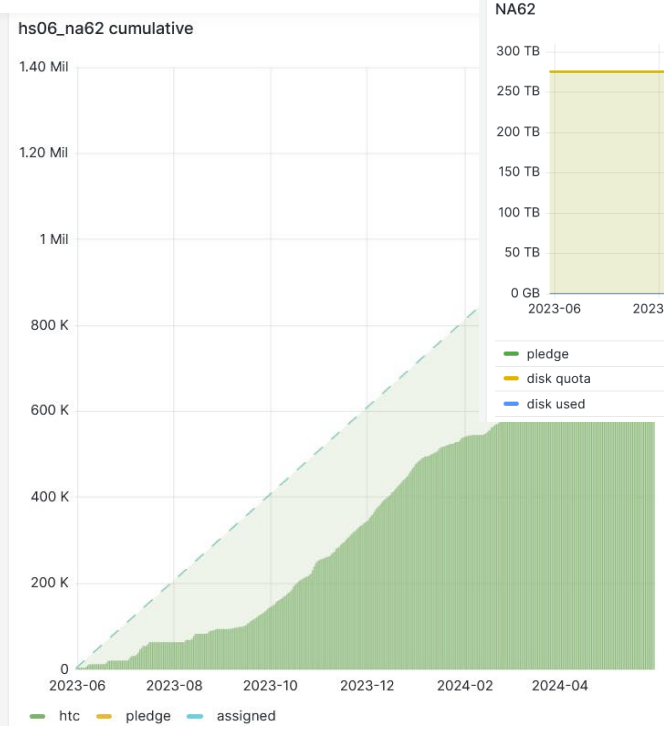

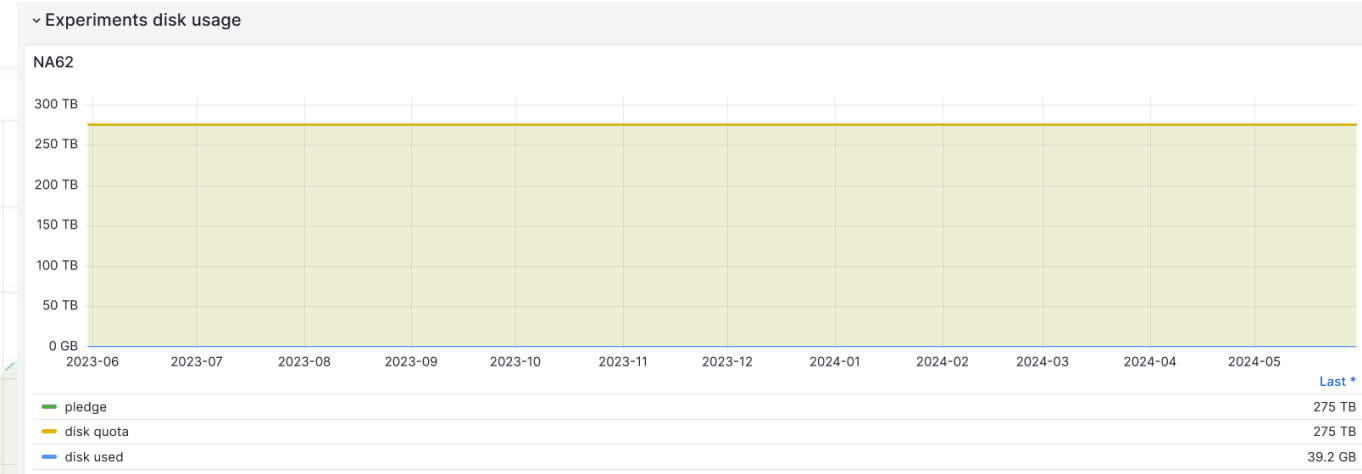

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHS06	2.4	3.3	3.3	85%	85%	OK
DISK TB	0.04	275	275	0%	0%	No disk usage

PADME

Requests: TIER2 LNF (100 TB disk SJ + 6 kHSP06 end maintenance)

Decision: final

PADME : Tier1 resources

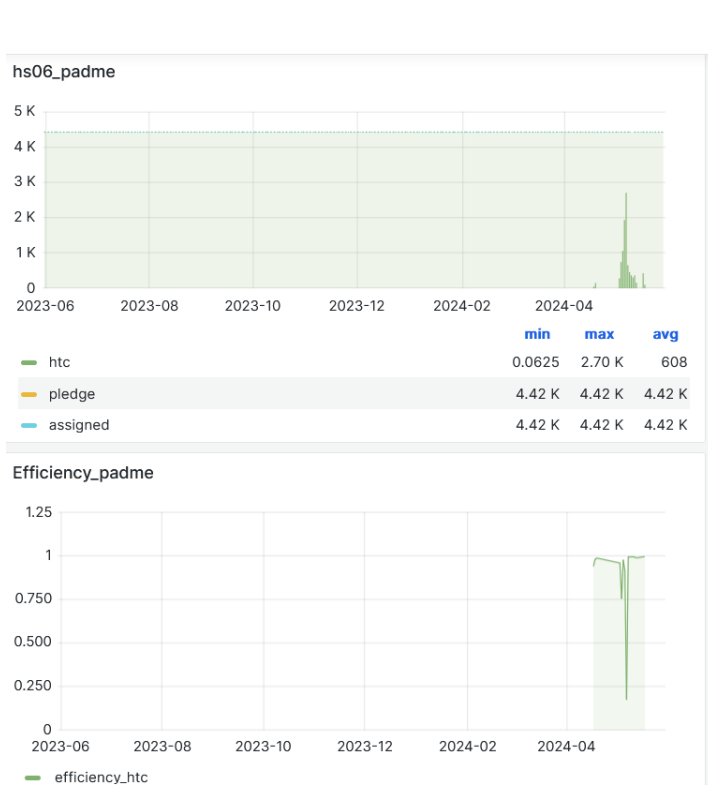

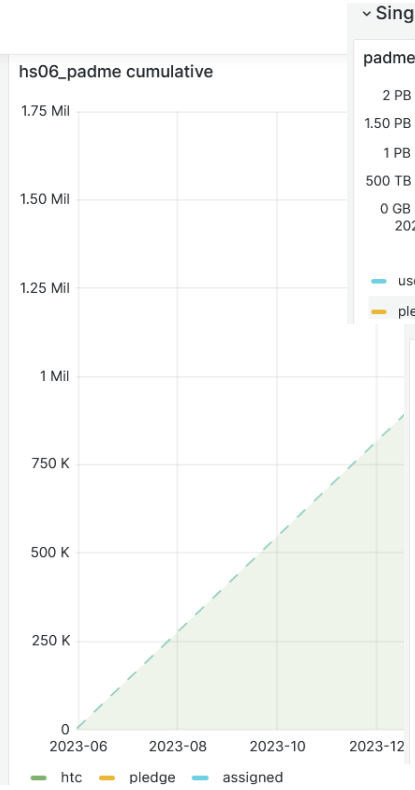

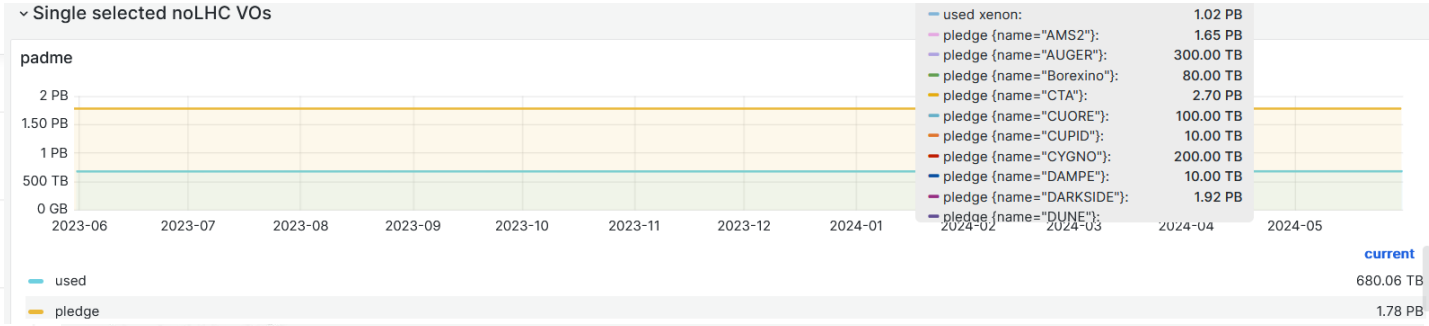

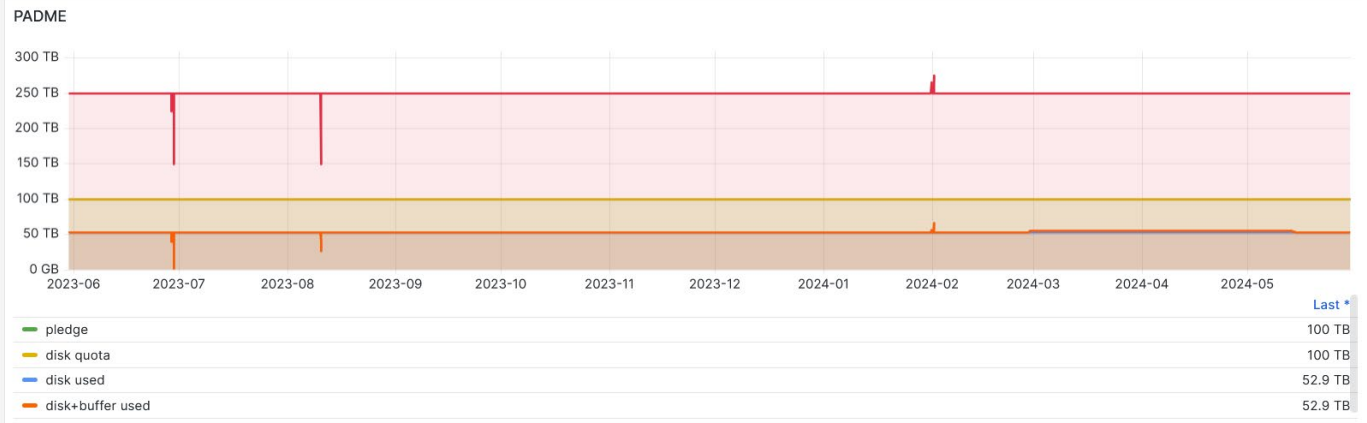

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHS06	0.6	4.4	4.4	15%	15%	low use of CPU.
DISK TB	53	100	100	53%	53%	OK
TAPE TB	680	1800	1800	40%	40%	low TAPE usage

Requests 2025

- Use of Tier2 LNF CPU was extensively used with GRID
  - Monte Carlo production of background with equivalent statistics to data
  - Multiple reconstruction campaigns of full Run-III data set + MC bkg + signal
- 6 kHS06 off maintenance (2018-2019). Replacement mandatory:
  - Further reprocessing campaigns possible on Run-III
  - Reconstruction/reprocessing of Run-IV data (up to x4 wrt Run-III)
  - MC background campaign exceeding Run-IV statistics (more than x4 wrt Run-III)
- Disk: new tracker with 0 suppression requires 100TB new disk (additional to 280TB available)
  - Uncertainty in 0-suppression real impact on new tracker data size. Can be better estimated with Run-IV data. Therefore for the moment request SJ

- Request Tier2 LNF: CSN1
  - 6 kHS06 for old CPU replacement ➔ OK 60 kEUR
  - 100 TB disk space (SJ to estimation of 0-suppression) ➔ worse case scenario. If 0-suppression is better than this, will request only a fraction of the 10 kEUR ➔ partial/total request of assignment MUST be justified in detail ➔ OK 10 kEUR SJ

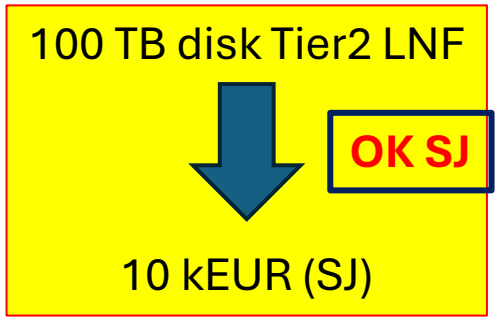

DUNE

Requests: TIER1 TAPE 500 TB

Decision: final

DUNE : Tier1 CPU

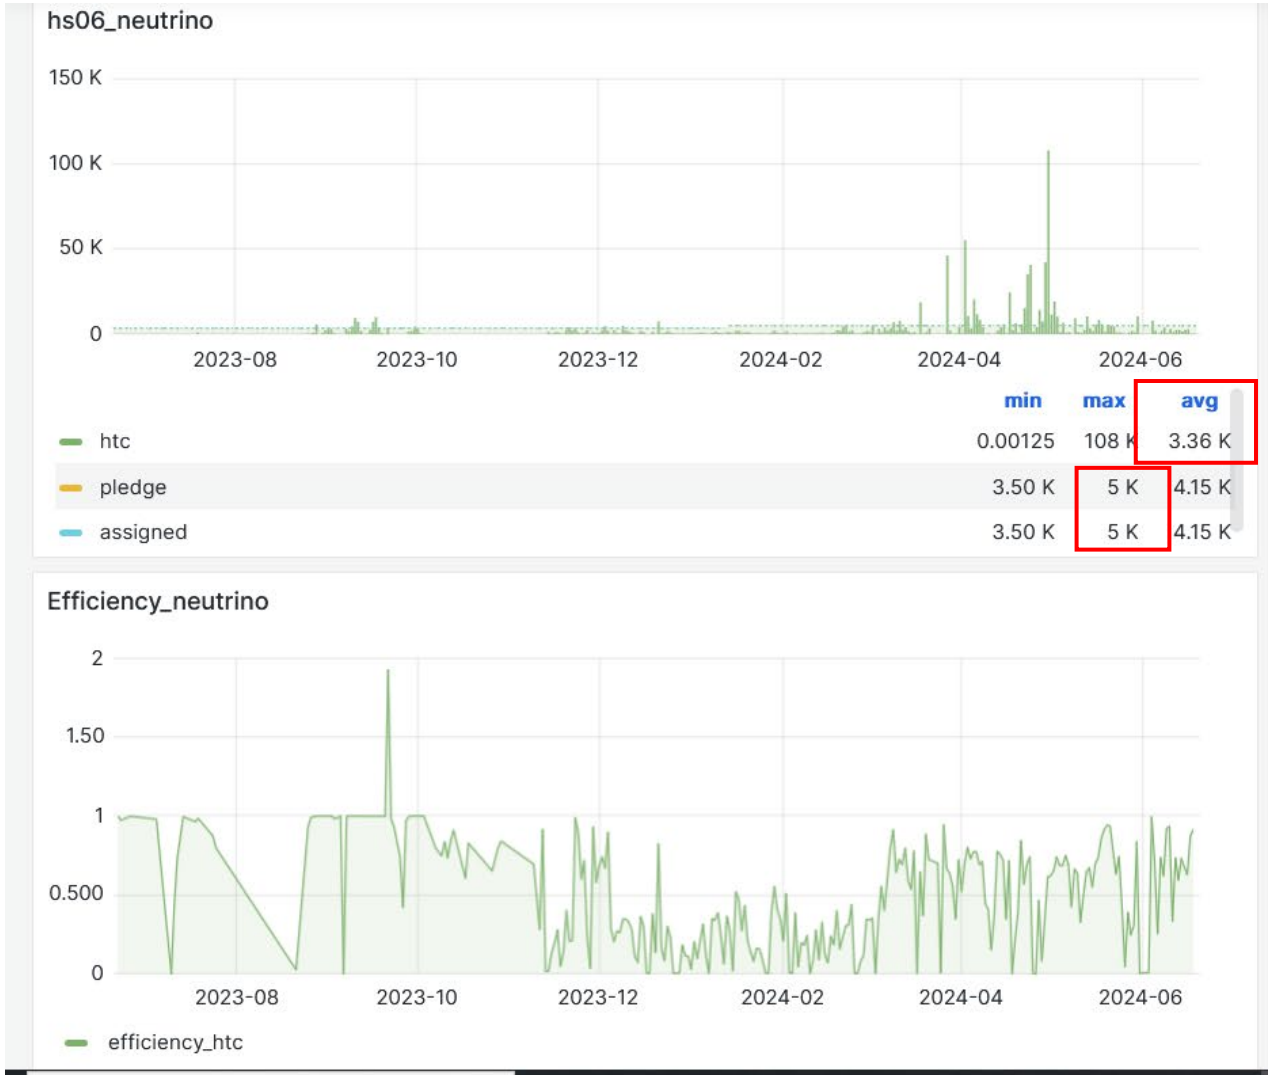

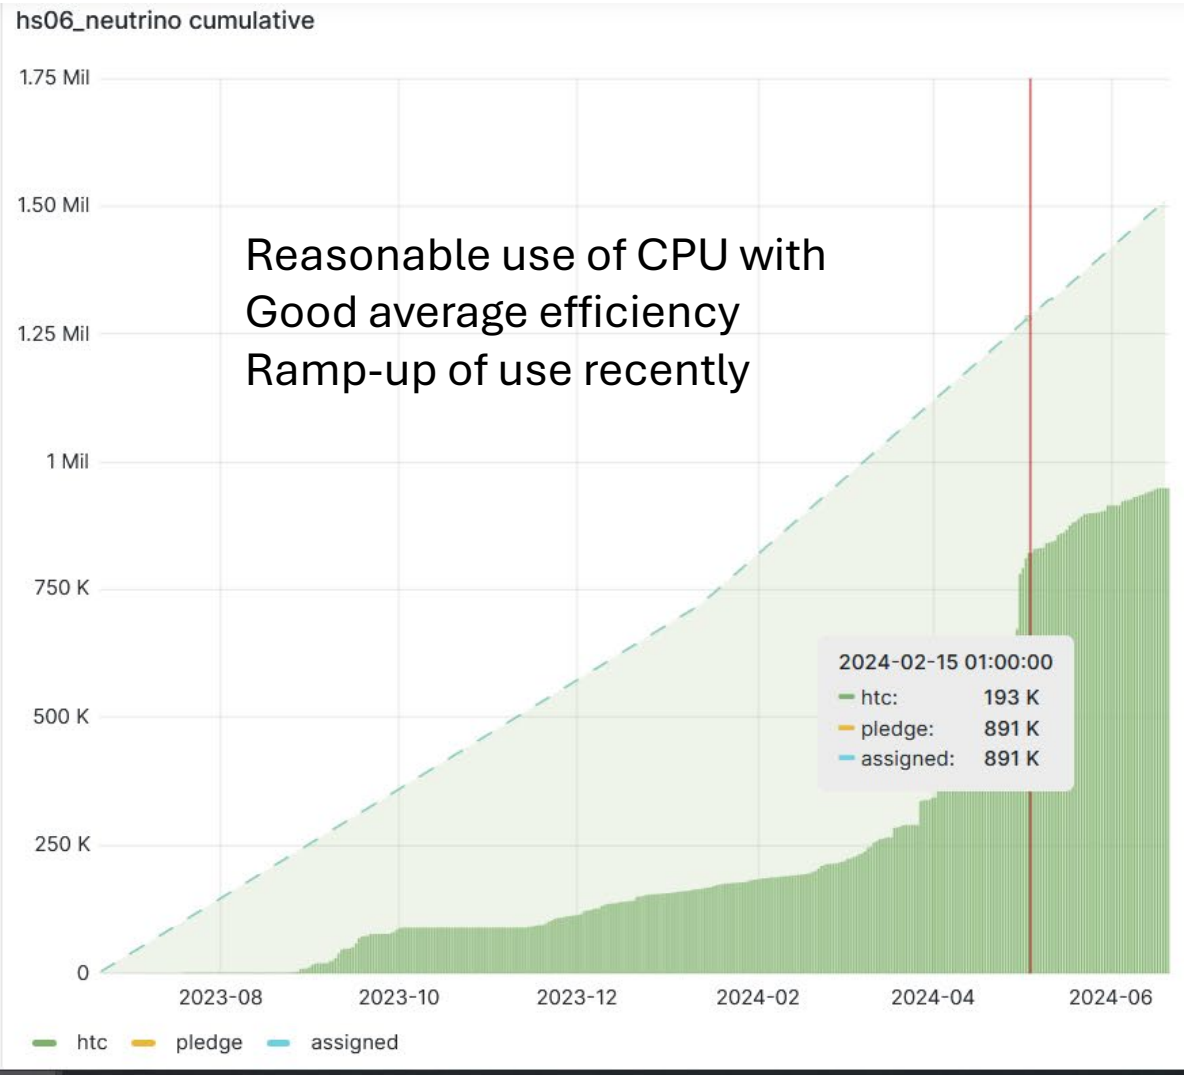

Reasonable use of CPU with
Good average efficiency
Ramp-up of use recently

DUNE : Tier1 resources

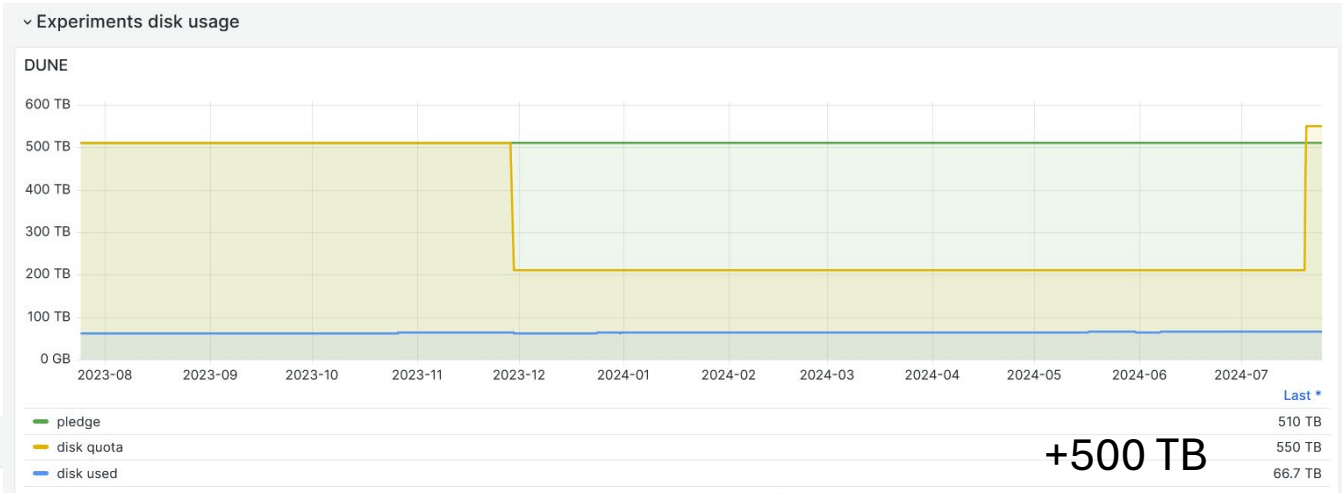

+500 TB

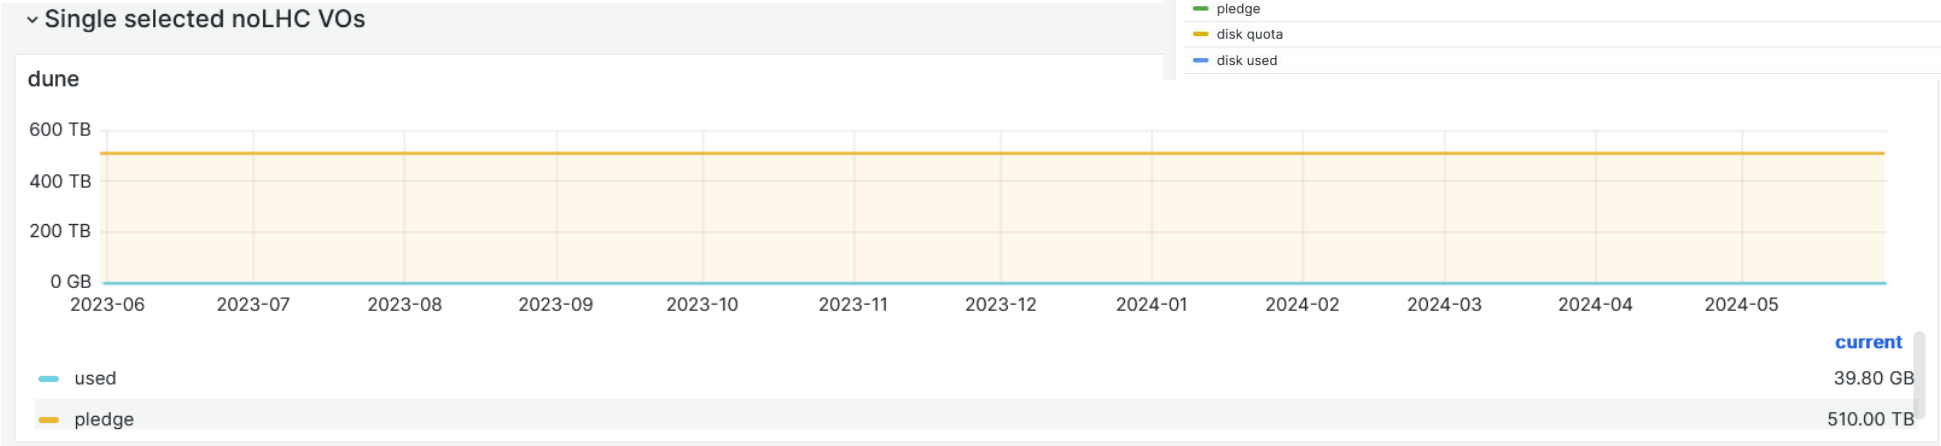

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU	3.4	5	5	68%	68%	OK
DISK TB	66	1100	1100	6%	6%	little use of DISK
TAPE TB	0.04	510	510	0%	0%	no use of TAPE

Requests 2025

- Tier1: GE
  - 500 TB Tape ➔ 5 kEUR

Follow-up with referees

HyperK

Requests: TIER1 300 TB TAPE

Action: riunione preliminare il 25/7 poi 4/9

HyperK : Tier1 resources

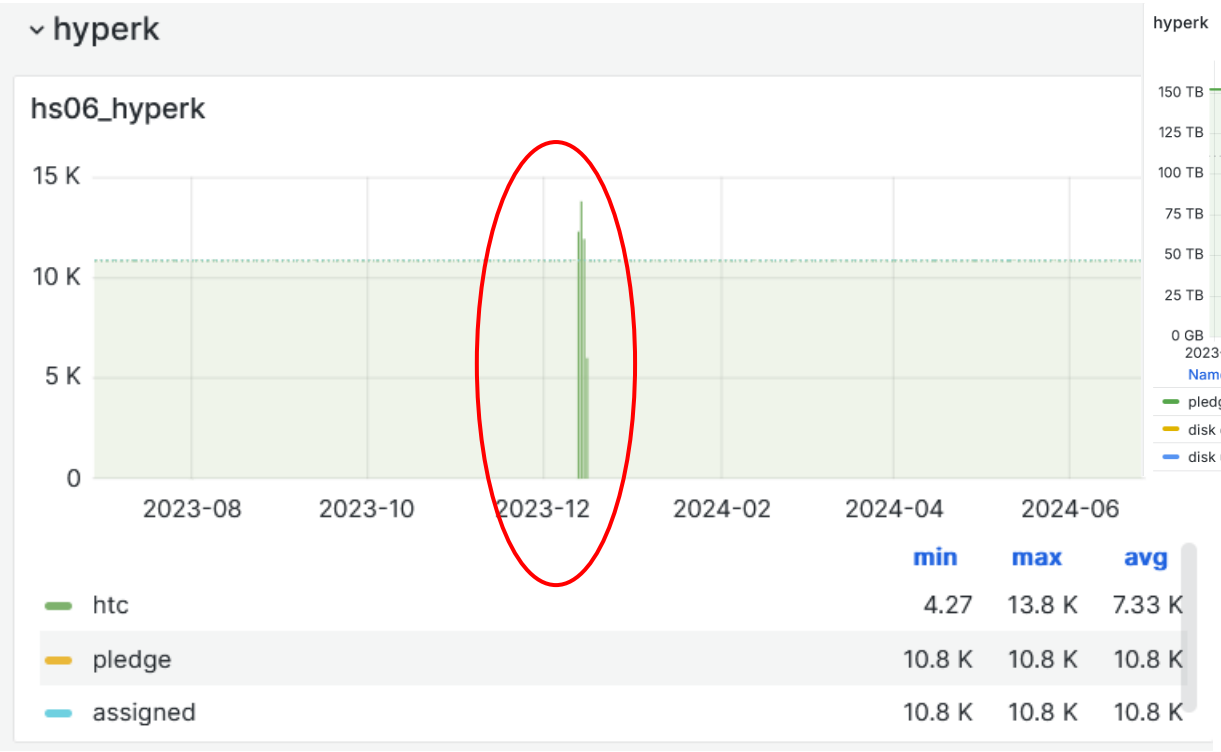

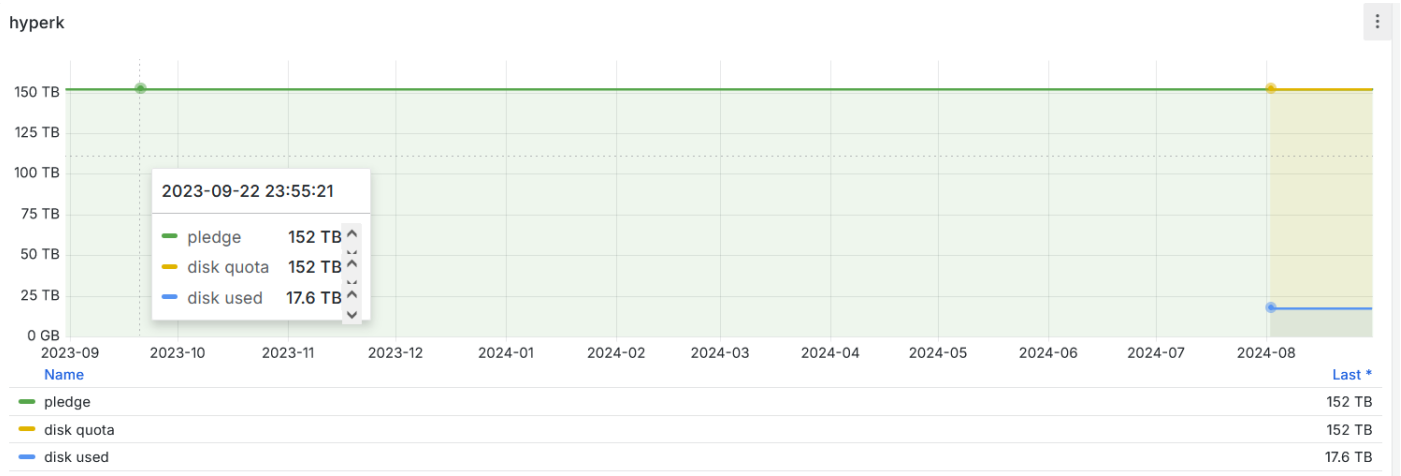

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU	7.3	10.8	10.8	70%	70%	
DISK TB	18	152	152	12%	12%	little use of disk
TAPE TB						

Requests 2025

- Tier1: 300 TB Tape ➔ 3 kEUR

- To be discussed with referees. Apparently no TAPE pledge for HyperK (?)

BES-III

No requests

Decision: final

G-2

No requests

Decision: final

G-2: requests for 2025

- No requests for 2025
  - Disk, CPU and Tape requested for 2024 OK (tape to be configured)

- shutdown of Fermilab started 26/8 but
  - Computing ok (limited support but ok)
  - Anyway storage available at cnaf and working

MEG-II

Requests: “MOF-like” contribution 20 KEUR

Decision: final

MEG-2 : status & requests for 2025

- Official decision by PSI to migrate computing to HPC-center in Zurich
- From this summer possible to:
  - Install reconstruction Software
  - Test PSI-Zurich connection for online-data transfer starting with old data and then really online
- AS LONG AS DATA TRANSFER IS NOT 100% STABLE AND RELIABLE no change in model
- IF data-transfer will not work, need to purchase new disks

- Last year decided to assign 20 kEUR / year as
  - «MOF»-like contribution if migration to Zurich
  - To update infrastructure otherwise

OK 20 kEUR CSN1

Mu2e

Requests: TIER1 (50 TB Tape) & TIER2 Pisa (50 TB disk + 6 kHS06)

Decision: final

INFN contribution to Mu2e computing

INFN contribution to Mu2e computing can be a win-win solution:

lu2e computig requests to DOE will be more robust → See problem of ICARUS with DOE !

oung people trained in Italy can provide precious uncosted man power he collaboration

NFN will consolidate his active role in simulation, calibration and data alysis

lost of the contribution can be accounted for as in-kind contribution to mmon Funds → RECOGNIZED AS COMMON FUNDS: 25 kEUR

er 2 in Pisa has already experience on setting up grid calculus for g-2 Fermilab. The computing power cost will be just the one of the new rvers: ~5K per Million of CPU hours

025 is the right moment to start this contribution

INFN contribution to Mu2e computing power

Current Mu2e Plan for CPU resources

	2025	2026	2027
HTC Mhours (kSH06)	**15 (22,5)**	15 (22,5)	40 (60)
HPC Mhours (kSH06)	**5 (7,5)**	5 (7,5)	5 (7,5)

CPU

Proposal for INFN contribution (20% of total):

	2025	*2026*	*2027*
HTC Mhours (kSH06)	**3 (4,5)**	*3 (4,5)*	*8 (12)*
HPC Mhours (kSH06)	**1 (1,5)**	*1 (1,5)*	*1 (1,5)*

2025 will be used to test the infrastructure, requests for following years will profit of the learned lessons

INFN contribution to Mu2e disk space and storage

In 2025 the main computing activity will be:

- Cosmic Ray simulation (currently done at Argonne using HPC nodes) for the detectors commissioning
- Data processing to test calibration procedures
- simulation of beam pileup and creation of signal+pileup datasets to prepare the Run 1 analysis

A typical simulation campaign foresees **10 live seconds of beam simulation** corresponding to 5 M CPU hours and **12 TB** of disk space. This disk space must be **doubled** to consider the storage of the pileup mixed datasets.

In order to have this contribution evaluated as inkind contribution to teh Common Funds, the created datasets must be available to the whole collaboration.

A backup on tape at CNAF will also be a good test to evaluate the opportunity to have a storage on tape of the data collected during Run 1.

Requests 2025

- Tier1: GE
  - 50 TB Tape ➔ 0.5 kEUR
- Tier2 PI: CSN1
  - Disk: 50 TB ➔ 5 kEUR
  - 3 server HTC + 1 HPC (corresponding to 6 kHS06) ➔ 20 kEur

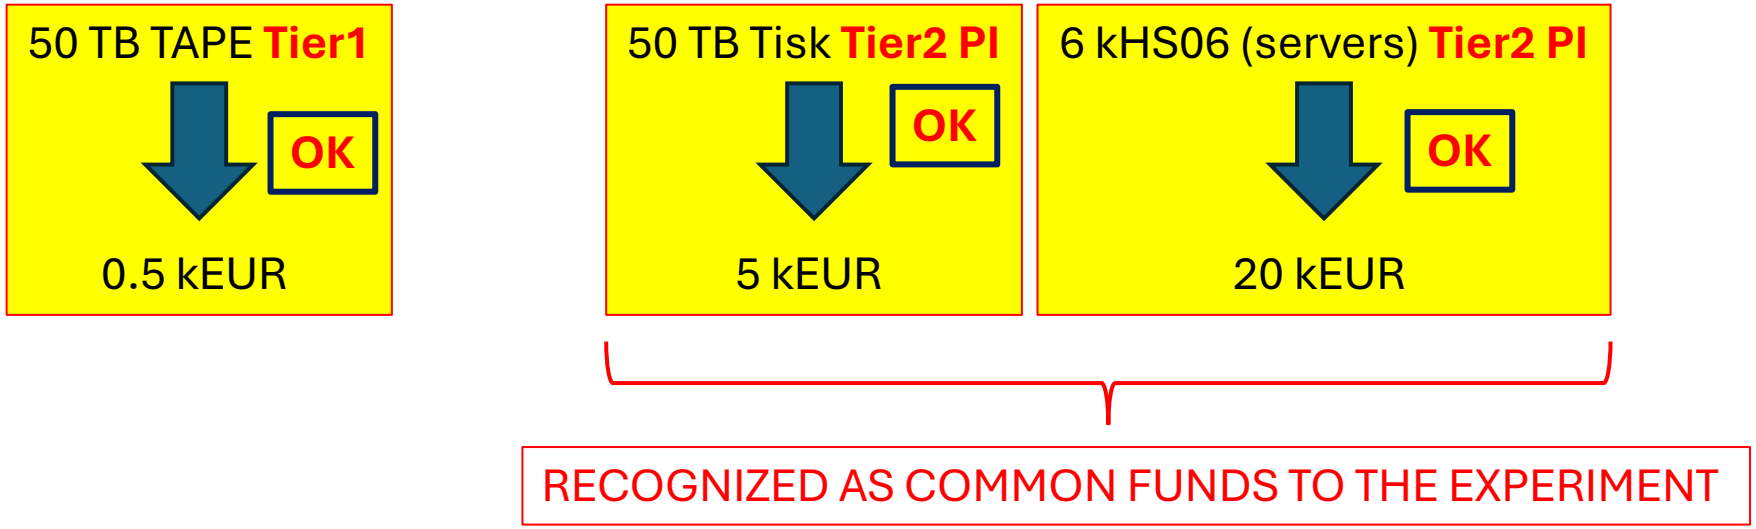

MuonE

No requests

Decision: final

MUONE : Tier1 resources

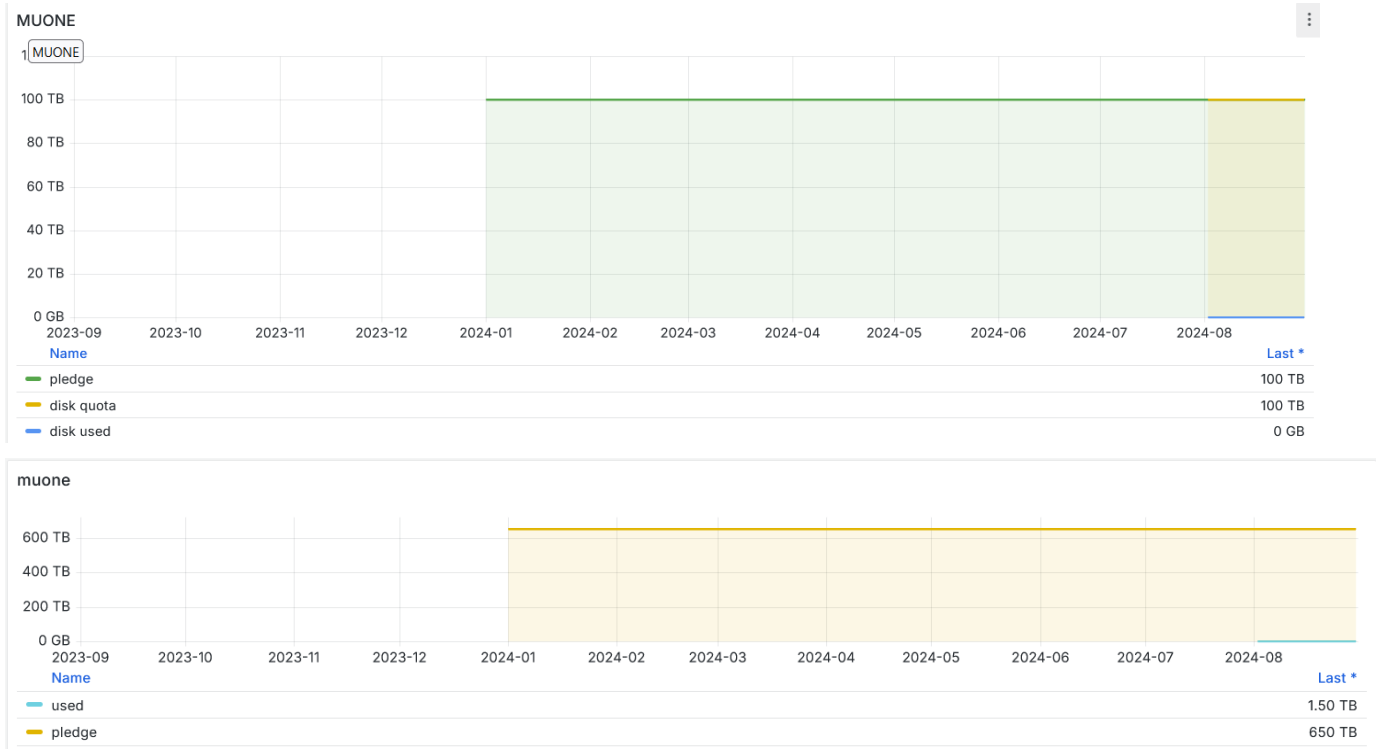

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHS06						missing 1 kHS06 from 2024 request
DISK TB	0	100	100			
TAPE TB	1.5		650	0%	0%	little use so far

ICARUS

Requests: TIER1 (400 TB disk, 1 kHS06, 1 PB Tape)

Decision: final

ICARUS : Tier1 resources

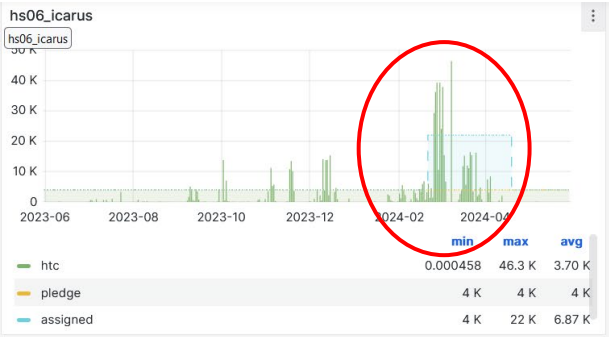

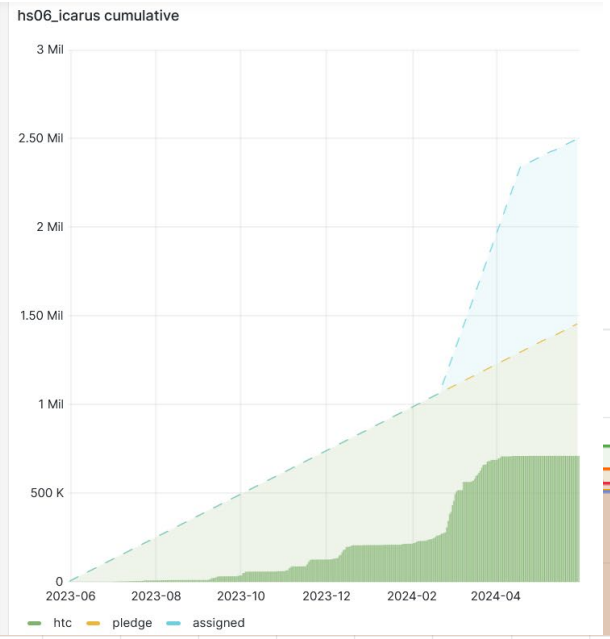

Disk restrictions limited the production campaigns at CNAF, concentrating the CPU usage in a short period of time

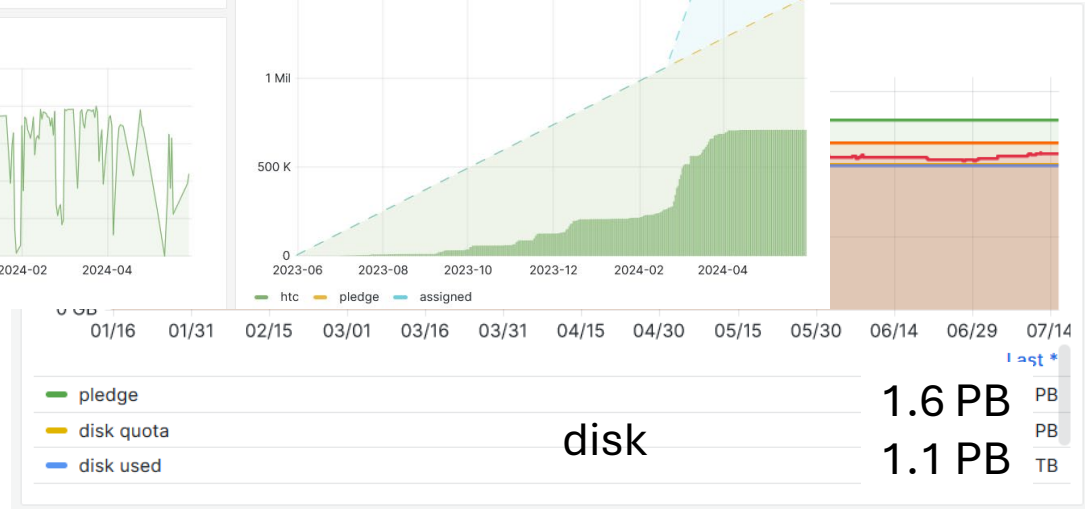

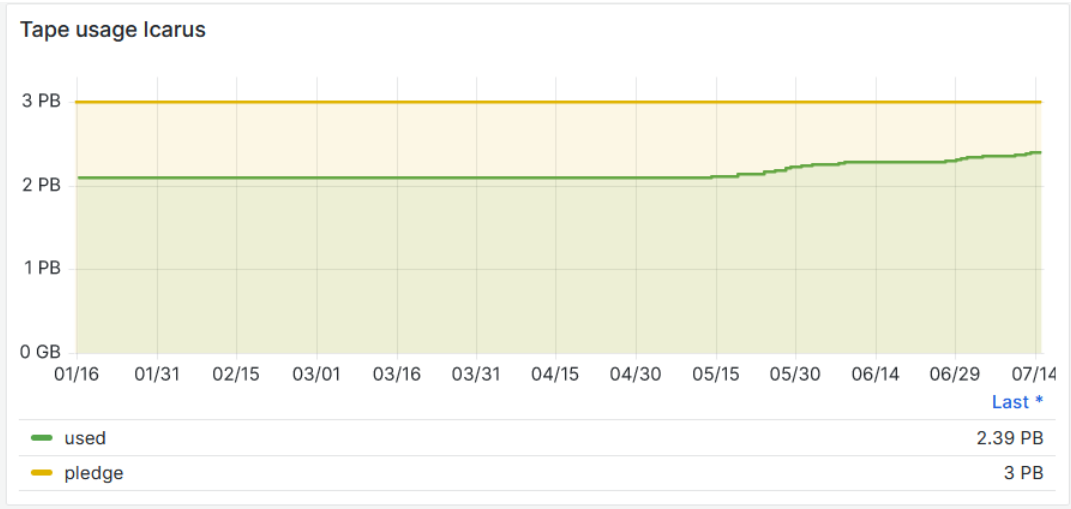

	Ave use	assigned	pledge	% /assigned	% / pledge	comment
CPU kHS06	3.7	4	4	92%	92%	discontinuous use of CPU
DISK TB	1100	1600	1600	70%	70%	
TAPE TB	2500		3000	80%	80%	

Status of computing at CNAF

- Run-3 and Run-4 (start October 2024 until 2026-27) will more than double the total statistics, requiring additional disk space.
- Disk limitation were handled thanks to a temporary increase in CPU power granted by CNAF and with a manual optimization by the production team, continuously monitoring the disk occupancy and freeing disk by removing intermediate files as the production progressed.
- The possibility to easily access the processed files and run the analysis code with the CNAF resources is extremely effective for the INFN analysis groups.
- Raw data are copied to tape at FNAL and replicated on tape at CNAF. The tape available at CNAF is 3 PB, occupied by 2/3. 1 PB additional tape space is required to save next year raw data and preserve MC productions for further reprocessing

Data handling – the recognized CNAF role in the experiment

- Collected data are transferred to FNAL central storage disks (190 MB/s) and saved on tape at FNAL after running the full event reconstruction chain:
  - Small size files (CAFs: Common Analysis Files) with the summary of the event reconstruction are maintained on disk, to be used for events selection and analysis;
  - Additional time overhead in case of reprocessing of data restored from tape to disk;
- Beam-on raw data and CAFs are transferred to CNAF with RUCIO package:
  - ~2 PB of raw data already transferred to CNAF (up to 200 MB/s) so far; additional disk (+600 TB) has been approved for 2023 and 2024 but not yet available.
- ICARUS resources at CNAF, including 3000 HS06, are exploited by a INFN team coordinated by M. Tenti. The ICARUS BNB analyses presented at the Summer conferences rely on CNAF activities:
  - Run2 event reprocessing with state-of-the-art reconstruction code (2.05 10^{20} PoT);
  - Generate/reconstruct a large (~2.8 10^{20} PoT) MC sample of BNB ν + cosmics events;
  - Subsequent analysis of the processed data, enabling fast evolution of the high-level analysis, tune the event selection and make detailed Data/MC comparisons.
- Due to budget restrictions FNAL didn't purchase additional disks for FY24, with negative impact on ICARUS computing, situation will worsen with the start of SBND data taking:

additional 400 TB of disk, 1 PB of tape and 1000 HS06 are required for the ICARUS data processing at CNAF

Requests 2025 & comments

- Tier1: GE
  - CPU: 1 kHS06 ➔ 10 kEUR
  - DISK: 400 TB ➔ 40 kEUR
  - TAPE: 1 PB ➔ 10 kEUR

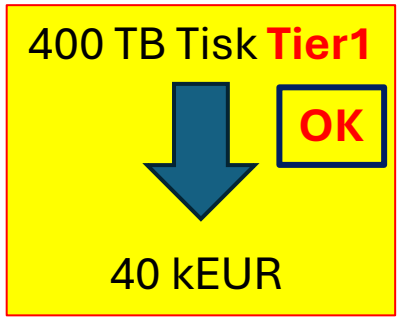

- Comments by referees:
  - ICARUS is a running experiment and Run-IV is expected to produce about double data than previous
    - Effort to reduce data-size to be actively pursued
  - CNAF resources heavily used and lack of disk is a problem for the efficient use of CPUs
  - Copy of RAW to CNAF tape is effective

BUT

- THERE IS A BIG CONCERN ABOUT FNAL ATTITUDE TO LOWER BOTH COMPUTING RESOURCES (no purchase of disk despite new run) AND PERSON-POWER IN USA
- INFN CANNOT REPLACE THIS LACK OF SUPPORT ON THE LONG TERM
- Worry about the implications for other experiments @ FNAL (Mu2e, g-2, DUNE) ?

Summary of Requests 2025

Totale richieste Tier1 - risorse

	CPU (kHS06)	DISK (TB)	TAPE (TB)	commenti
AMBER	1	50		OK
BELLE-II			60	OK
DUNE			500	OK
ICARUS	1	400	1000	OK
HYPERK			300	to be refereed
KLOE	1	300		not reported in calc1_tier1 db so far !
LHCf	6		50	OK
MEG-2				-
Mu2e			50	OK
PADME				-
RD_MUCOLL	**-3**			**“restituzione” 3 kHS06 (nota segno “meno”)**
RD_FCC	**1**			**non assegnate lo scorso anno (errore?) non in DB**

Totale richieste Tier1 - kEUR

1 HS06 = 10 Eur
1 TB disk = 100 Eur
1 TB tape = 10 Eur

	CPU (kEur)	DISK (kEUR)	TAPE (kEUR)	altro
AMBER	10	5		OK
BELLE-II			0.6	OK
DUNE			5	OK
ICARUS	10	40	10	OK
HYPERK			3	to be refereed
KLOE	10	30		not reported in calc1_tier1 db so far !
LHCf	60		0.5	OK
MEG-2				-
Mu2e			0.5	OK
PADME				-
RD_MUCOLL	**-30**			**“restituzione” 3 kHS06 (nota segno “meno”)**
RD_FCC	**10**			**non assgnate lo scorso anno (errore?) – non in DB**

TOTAL TIER1: 194.6 keur (non considerando la “restituzione” mucoll)

Totale richieste Tier2 – risorse (CSN1)

	CPU (kHS06)	DISK (TB)	TAPE (TB)	altro	note
AMBER	1 (T2 TO)				obsolete CPUs T2 TO – **PNNR** OK
BELLE-II					-
DUNE					-
ICARUS					-
HYPERK					-
KLOE					-
LHCf					-
MEG-2					-
Mu2e	6 (T2 PI)	50 (T2 PI)			**PNNR** Note: recognized as Common Funds by exp. OK
PADME	6 (T2 LNF)	100 (T2 LNF) - (SJ)			**PNNR** CPU out of maintenance (2018). OK
RD_MUCOLL					-

Totale richieste Tier2 - kEUR (CSN1)

	CPU (kEUR)	DISK (kEUR)	TAPE (kEUR)	altro (KEur)	note
AMBER	10 (T2 TO)				obsolete CPUs – T2 TO **PNNR** OK
BELLE-II					-
DUNE					-
ICARUS					-
HYPERK					-
KLOE					-
LHCf					-
MEG-2					-
Mu2e	20	5			**PNNR** Note: recognized as Common Funds by exp. OK
PADME	60 (LNF)	10 (LNF) - (SJ)			**PNNR** CPU out of maintenance (2018). OK
RD_MUCOLL					-

TOTAL TIER2: 95 kEUR + 10 kEUR SJ
PNRR: ???

Totale richieste other - kEUR (CSN1)

	CPU (kEUR)	disk (kEUR)	TAPE (kEUR)	altro (KEur)	note
AMBER	7.5 kEUR (TS)				TS: local farm originally from CSN1 (substitution of obsolete CPU) OK the remaining 7.5 should come from Sezione TS and/or other users
BELLE-II					-
DUNE					-
ICARUS					-
HYPERK					-
KLOE				26 + 5 (SJ)	26: maintenance of new lib (reject 3 kEUR SJ) OK 5 SJ: maintenance of the old lib. OK
LHCf					-
MEG-2				20	Contribution to PSI computing (decided 2024) OK
Mu2e					-
PADME					-
RD_MUCOLL		15			**PNRR ?** disk cloud Veneto OK

TOTAL other requests: 68.5 keur + 5 SJ

CSN1 & GE

- Total CSN1:
  - 163.5 kEUR (PNRR ??)
  - 15 SJ

- Total GE:
  - 194.6 kEUR
    - Not considering restitution by mucoll (3 kHS06)

STATUS OF PAST REQUESTS

Requests for 2024: status (missing) - I

RD_FCC	1 KHS06 T1	assigned for 2024 (September 2023)	not found in dashboard probably misunderstanding ? requested this year