


Referaggio risorse 2023 – Tier1

D. Cesini – INFN-CNAF


CPU - Farm

- Pledge 2022: 600kHS06
- Nuove risorse in produzione dal CINECA → + 232kHS06



4 nuovi rack in produzione al CINECA
800 HS06 per nodo
Rete resta per il momento SPOF
Manca un core switch
Sarà necessario un down completo
di CINECA (previsto dopo l'estate)

'Yearly' Graph (1 Day Average) Bit per sec su CINECA




Max In: 72.9 Gb/s (18.2%) Average In: 2628.2 Mb/s (0.7%) Current In: 4556.8 Mb/s (1.1%)
Max Out: 367.7 Gb/s (91.9%) Average Out: 70.0 Gb/s (17.5%) Current Out: 45.3 Gb/s (11.3%)

Non vediamo aumento che ci saremmo attesi (!?)

Composizione farm

- Siamo leggermente sopra pledge 2022: **630k HS06**

In corso overpledge
significativo CUORE
11kHS06



YEAR	CPU	DISK	TAPE
	kHS06	PB-N	PB
2023	820	78	172
2024	990	94	206
2025	1320	110	247

If ALL the CPU pledges (23-25) comes from Leonardo:

2023 - 11% (165 nodes) of Leonardo GP

2024 - 20% (300 nodes)

2025 - 36% (540 nodes)

Nuova gara CPU 2022

- 600k euro per 60k HS06
- **Pubblicata**, chiusura il 7 settembre
 - Base a 10eu/HS06
- Consegna non prima di dicembre (?)
 - Ci permetterebbe di sostituire gara 2015

Rimpiazzi – 198kHS06

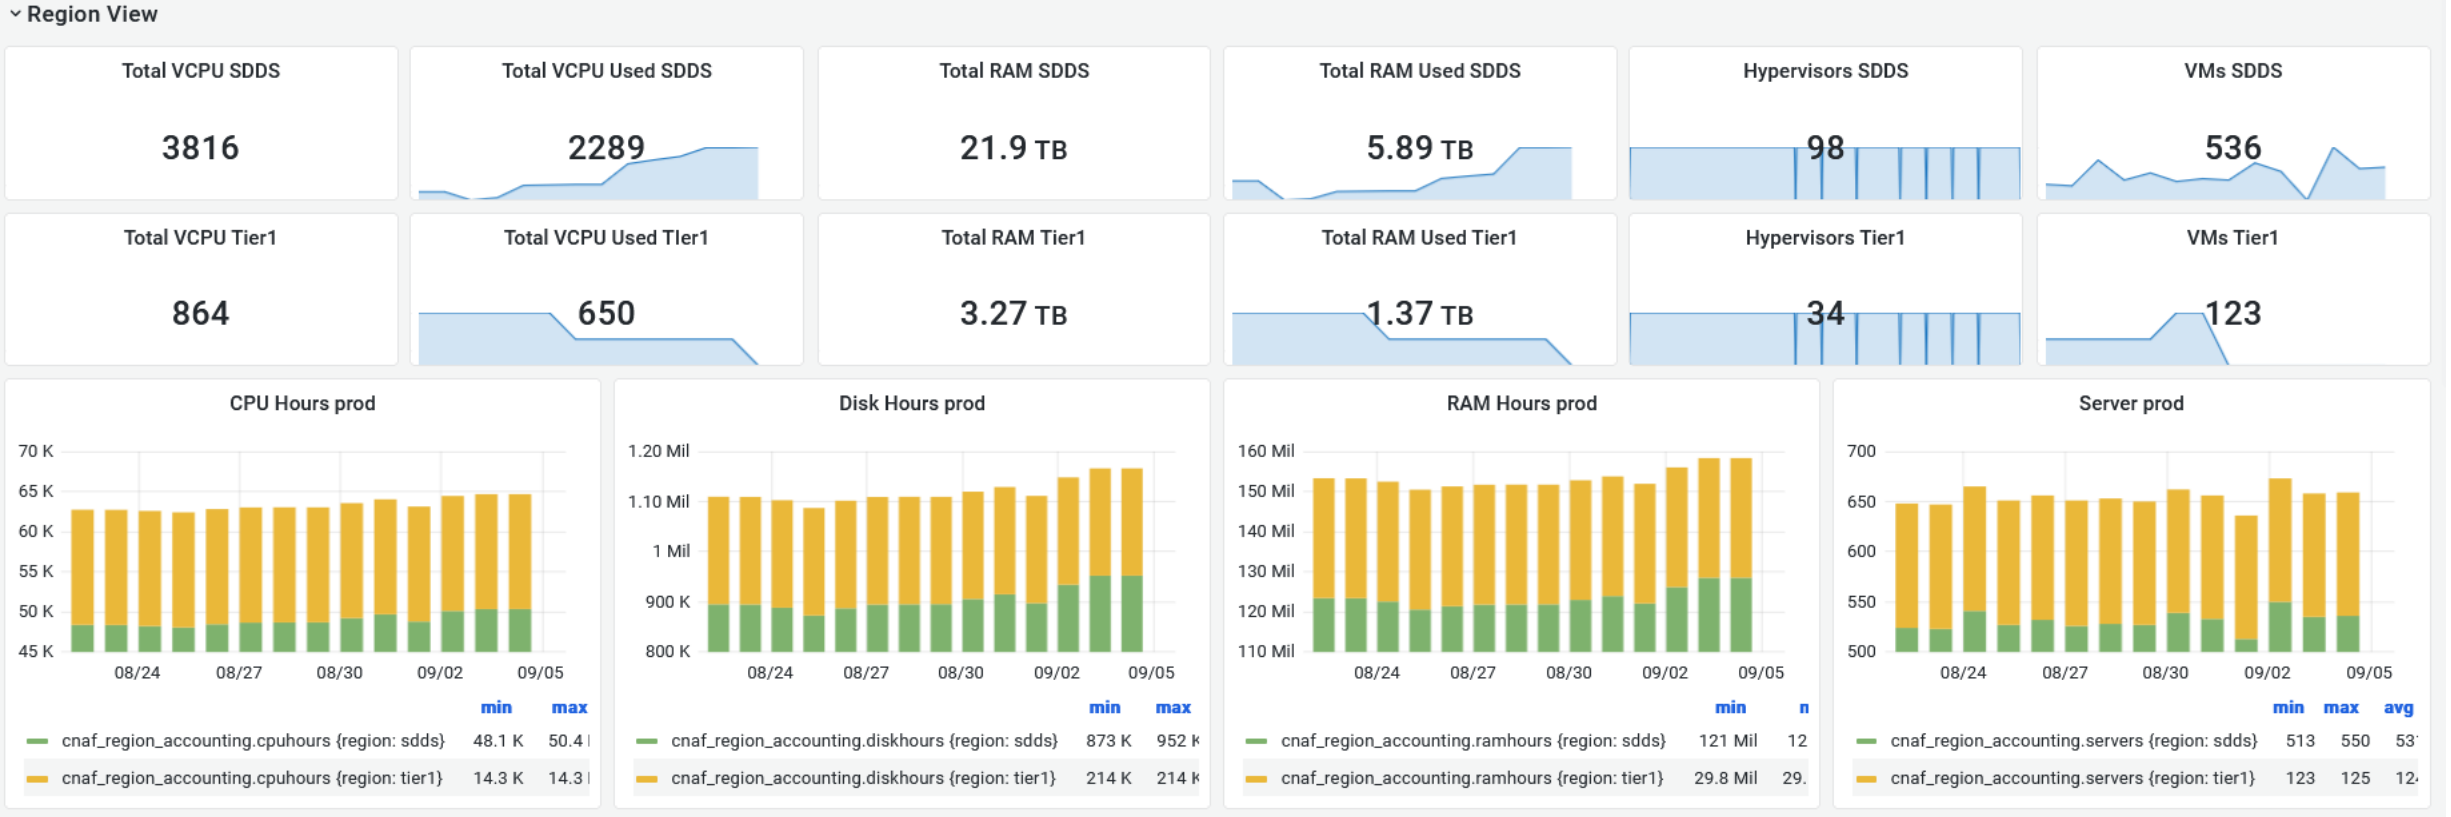
- Gara 2015: 54kHS06
- Gara 2016: 102kHS06
- Gara 2017: 42kHS06

Leonardo@CINECA

Inseriti in DB preventivi quantificando secondo lo share avg(2015-2017) per CSN e scorporati per LHC – 15eu/HS06

Rimpiazzi CPU		
VO	AVG share HS06 (2015-2017)	keu
CMS	0.211	626
ATLAS	0.207	614
LHCb	0.122	362
ALICE	0.124	370
CSN1_NO_LHC	0.122	350
CSN2_NOVIRGO	0.130	387
VIRGO	0.086	256
CSN3_NOALICE	0.002	6.5

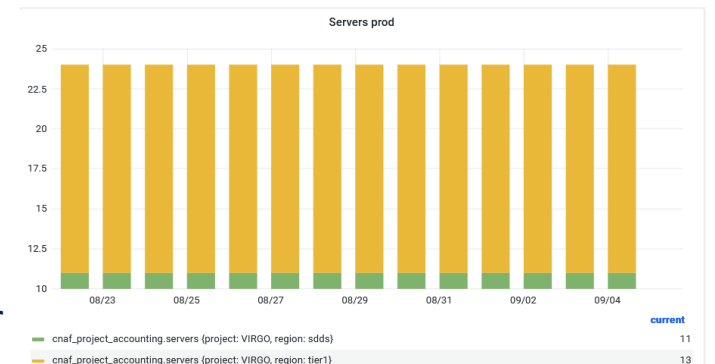
Stato Cloud@CNAF



Circa 100 tenant configurati, alcuni per esperimenti con pledge@Tier1

- AGATA-GAMMA
 - 2 istanze 8core 16GB RAM, 160GB HDD
 - 1 shared volume 2TB
- FOOT
 - 2 istanze 8core 16GB RAM, 160GB HDD
- NTOF
 - 5 istanze 8core, 16GB RAM, 160GB HDD
 - 1 shared volume 1TB
- ASFIN
 - 1 istanza 32core, 64GB RAM, 280GB HDD

- VIRGO Low Latency Cluster on K8s



Manteniamo stesso modello delle pledge standard? Ok, ma attenzione a Overlap Factor

Disk storage in produzione

- Installed: 50.07 PB, Pledge 2022: 59.1 PB, Used: 41.4 PB

Sistema	Modello	Capacita' netta, TB	Esperimenti	Scadenza
2015 ddn-10, ddn-11	DDN SFA12k	10752	ALICE, AMS	03/2021 → 06/2023
2016 os6k8	Huawei OS6800v3	3400	GR2, Virgo	06/2022
md-1,md-2,md-3,md-4	Dell MD3860f	2308	DS, Virgo, Archive	11/2021 → 12/2022
md-5, md-6, md-7	Dell MD3820f	28	metadati, home, SW	12/2022
2017 os18k1, os18k2	Huawei OS18000v5	7800	LHCb	2023
2018 os18k3, os18k5, os18k5	Huawei OS18000v5	11700	CMS	NOV 2023
ddn-12, ddn-13	DDN SFA 7990	5060+870	GR2,GR3	2025
ddn-14, ddn-15	DDN SFA 2000NV	24	metadati	2025
os5k8-1,os5k8-2	Huawei OS5800v5	8999	ATLAS	2027
Cluster CEPH	12xSupermicro SS6029	3400	ALICE, cloud, etc	2027

Rimpiazzati DISK 33PB_N		
VO	AVG share HS06 (2015-2018)	keu
CMS	0.201	927.5
ATLAS	0.193	892..5
LHCb	0.154	710.0
ALICE	0.172	794.5
CSN1_NO_LHC	0.035	162.5
CSN2_NOVIRGO	0.221	1021.5
VIRGO	0.024	110.0
CSN3_NOALICE	0.001	3.0

Gara storage 2021-2022

- 14PB netti
 - 120eu/TB_N
- Chiusa il 15/7
- Aperte buste amministrative
 - 3 offerte
- Nominata commissione
- Da definire data per aperture buste tecnica ed economica

AQ Storage 23-24

- Fino a 68PB_N
 - Da installare al tecnopolo
 - In 2 Appalti Specifici: 54+14PB_N
- In approvazione GE fine Settembre

Stato Tape

- Liberi 12 PB (su cassette vuote, complessivamente sulle 2 librerie). Usati 90 PB.
 - Pledge 2022: 130.5 PB
 - Installato attuale: 116.5 PB
 - 14PB di media in consegna OGGI
 - Recente cancellazione CMS: 8.4 PB
- REPACK
 - Lo spazio liberato ci consente di far partire repack su nuova libreria
 - Dobbiamo dismettere quanto prima libreria Oracle per problemi sui contratti di manutenzione
 - **Media per completamento repack chiesto in WG DATCLOUD (67PB con nuova tecnologia a 7eu/TB)**
- ATLAS non ha spazio libero su tape

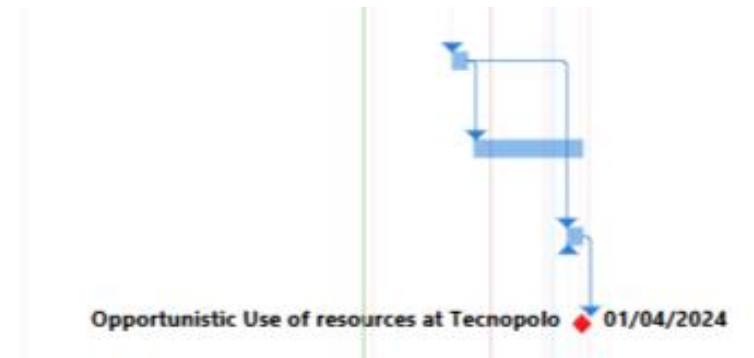
Library	Tape drives	Max data rate/drive, MB/s	Max slots	Max tape capacity, TB	Installed cartridges	Used capacity, PB
SL8500 (Oracle)	16*T10KD	250	10000	8.4	~10000	67.3
TS4500 (IBM)	19*TS1160	400	6198	20	1750	22.8

HPC

- Non abbiamo risorse “pledged” HPC
 - Piccolo cluster InfiniBand general purpose ormai obsoleto
 - 4 x NVIDIA GPU
 - Altre GPU recenti su cloud via progetti e ML_INFN
 - Dovrebbero arrivare le HPC bubbles con Terabit – tempistica incerta
- Uso partizione booster Leonardo
 - Da capire i tempi richiesti per risorse pledge 2023

Task per HTC-HPC integration su Leonardo in CNAF Reloaded


182	2.4.5	CNAF-Leonardo integration: duplicate M100 infrastructure (T)	T	A	Fri 07/04/23	Thu 18/05/23	3.6	5		0%
183	2.4.6	Production Phase accessing storage @ Berti (T)	T	A	Wed 07/06/23	Mon 01/04/24	6	182		0%
184	2.4.7	Opportunistic Production fully at Tecnopolo	T	A	Tue 20/02/24	Mon 01/04/24	5.4	182;215		0%
185	WP2-M20	Opportunistic Use of resources at Tecnopolo	T	M	Mon 01/04/24	Mon 01/04/24		184		0%




BACKUP

CPU Usage CSN2


HS06_per_group



HS06_per_group cumulative




Global Disk Usage LHC



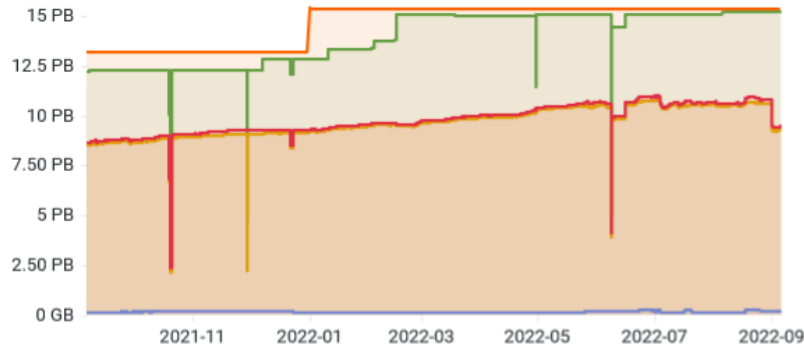
Global Disk Usage NOLHC

Total disk usage - Gruppo 1




disk quota	1.45 PB
used (no tape buffer)	806 TB
buffer used	4.04 TB
pledge	1.59 PB
disk+buffer used	810 TB

Total disk usage - Gruppo 2




disk quota	15.2 PB
used (no tape buffer)	9.26 PB
buffer used	187 TB
pledge	15.3 PB
disk+buffer used	9.44 PB

Total disk usage - Gruppo 3


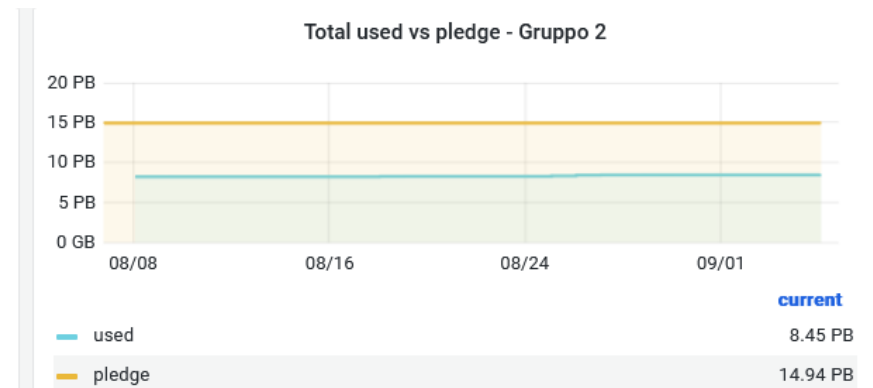



disk quota	565 TB
used (no tape buffer)	51.9 TB
buffer used	0.0318 GB
pledge	615 TB
disk+buffer used	51.9 TB

Tape Usage LHC

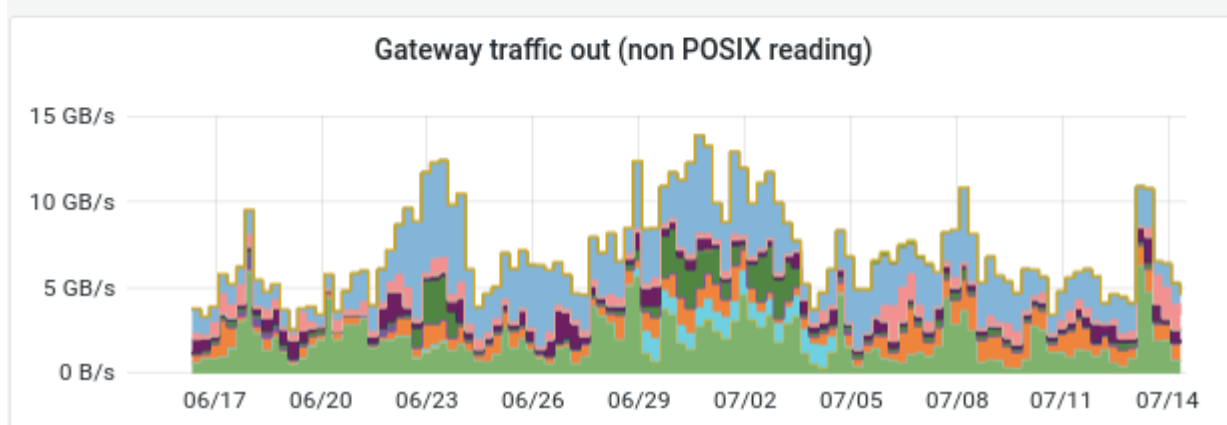



Tape Usage NoLHC



Global Disk servers traffic

Jun-Jul



Last 6 months

