



State of Storage

CdG 21 aprile, 2023





Business as usual

Last month



Last 6 months



Disk storage in produzione

Installed: **53.64** PB, Pledge 2023: **69.6** PB, Used: **45.8** PB

Storage system	Model	Net capacity, TB	Experiment	End of support
ddn-10, ddn-11	DDN SFA12k	10120	ALICE, AMS	12/2022 (20 spare hdd)
os6k8	Huawei OS6800v3	3400	GR2, Virgo	12/2023
md-1,md-2,md-3,md-4	Dell MD3860f	2308	DS, Virgo, Archive	11/2023
md-5, md-6 e md-7	Dell MD3820f	50	metadati, home, SW	11/2023 e 12/2024
os18k1, os18k2	Huawei OS18000v5	7800	LHCb	2023
os18k3, os18k5, os18k5	Huawei OS18000v5	11700	CMS	2024
ddn-12, ddn-13	DDN SFA 7990	5840	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

Storage allocation for LHC experiments

Experiment	FS name	1TB blocks	Used	Use%	Assigned, TB	Pledge, TB	Deficit, %
ALICE	gpfs_tsm_alice	574	484	85.00%	10843	11430	5.14%
	gpfs_alice	7669	6928	91.00%			
	ceph_alice	2600	570	22.00%			
ATLAS	gpfs_tsm_atlas	574	456	80.00%	10248	12240	16.27%
	gpfs_escape	36	1	2.00%			
	gpfs_atlas	9638	9064	95.00%			
CMS	gpfs_tsm_cms	10651	9375	89.00%	10651	12740	16.40%
LHCb	gpfs_lhcb	7984	7530	95.00%	7984	11561	30.94%

Storage allocation for non-LHC experiments

Experiment	FS name	1TB blocks	Used, TB	Use%	Assigned, TB	Pledge, TB	Deficit/overbooking
Darkside	gpfs_ds50	2300	1748	77.00%	2300	2340	1.71%
VIRGO	gpfs_virgo	757	685	91.00%	881	856	-2.92%
	gpfs_virgo4	124	94	77.00%			
KLOE	gpfs_kloe	1	1	5.00%	1	33	
AMS	gpfs_ams	2631	1569	60.00%	2631	2800	6.04%
Teorici	gpfs_cnaf	62	16	26.00%	62	?	
GR2/GR3	gpfs_data	8010	7718	97.00%	11909	11909	32.74%

Acquisti recenti e futuri

- Gara storage 2022 (14PB netti)
 - Il vincitore è Lenovo con ThinkSystem DE6600
(Derivato dal NetApp DE6600 con un enclosure con i dischi NVME)
 - Installazione dal 26 aprile
- AQ storage 2023-2024
 - 66PB nel 2023 + 14PB nel 2024
 - 5 offerte (3 Huawei, Lenovo e DDN), In fase di valutazione tecnica



Current SW in PROD

- GPFS 5.1.2-8 (to be updated soon to 5.1.2-10)
- StoRM BackEnd 1.11.21 (latest)
- StoRM FrontEnd 1.8.15 (latest)
- StoRM WebDAV 1.4.1 (latest, waiting for testing new rpm)
- StoRM globus gridftp 1.2.4
- XrootD 5.4.2-1
 - 5.3.1-1 on CMS redirectors (local and EU/IT/FR)
- Ceph 16.2.6 (Pacific)

Tickets and problems


- Ready to configure access with *tokens* for all the storage areas upon request
- ATLAS
 - Transfer failures from UK sites, due to a discrepancy in the root CA during TLS handshake; a new ca-policy-egi-core rpm was released to fix this (GGUS [160759](#))
 - Upgrade of CaNL libraries in StoRM WebDAV strongly suggested, a new rpm should come out soon
 - “Bring-online timeout has been exceeded”: a not-finished recall process on our backend had set a lock on recalls from a specific cartridge (GGUS [160730](#))
 - Failing transfers: the source of the failures was found and fixed at IN2P3-CC (GGUS [160587](#))
 - TPCs in push-mode to GOOGLE_EU fail for all StoRM sites (GGUS [158487](#))
 - Issue for StoRM developers, on hold

Tickets and problems


- CMS

- Gridftp still used, only for SAM tests :-)
- CMS suggests an XrootD upgrade to 5.5.4: following this CDG, we would do that
- [Issue](#) opened to developers on thread saturation, log files stuck, need restart
 - Recently we observed a remarkably high load on all 4 servers, but no stuck process and no big difference linked to async on/off
 - Given more than 20k files per hour were accessed on each server, we increased maxStatCache within GPFS from 10k to 50k

Load average



Number of threads per Xrootd



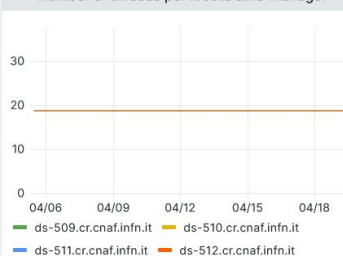
Tickets and problems

- Hyperk
 - Failing uploads: inodes used up. Inodes multiplied by a factor 10 (GGUS [161522](#))
 - We are working to add inodes metrics to our monitoring and alarming system
- Virgo
 - Stash-cache upgraded from OSG 3.5 to OSG 3.6
 - The request of publishing virgo/ligo data with unauthenticated access has been postponed
- Belle
 - Tests srm+https on hold (undergoing DIRAC upgrade)
- AMS02
 - Access to storage enabled for VOMS proxies obtained with iam-ams
 - Discussing the switch from srm + gsftp -> srm + https
 - Need to modify the access point for the tape storage area (/ams -> /ams-cnaf) to avoid nested storage areas, unsupported by StoRM WebDAV
- Configured StoRM WebDAV storage areas for Muoncoll and KM3net

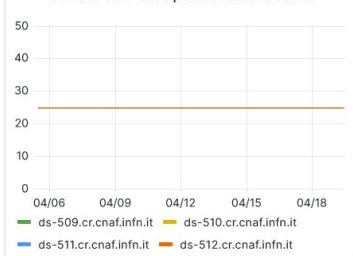
Tickets and problems

- All the no-LHC xrootd instances have been moved to 4 servers (ds-509, ds-510, ds-511, ds-512) with the alias xrootd-archive/xrootd-ams.
 - Dampe is the main user by far
 - Unauthenticated access (!) from the farm WNs
 - Only one out of 4 servers is used (all of them subscribe to the redir in Bari)
 - Currently trying to identify a contact person for these activity within Dampe

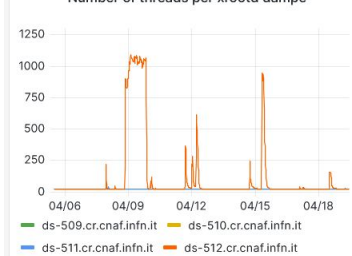
Number of threads per xrootd ams-manager



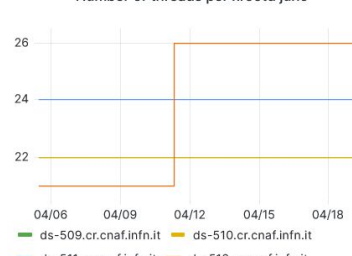
Number of threads per xrootd ams-server



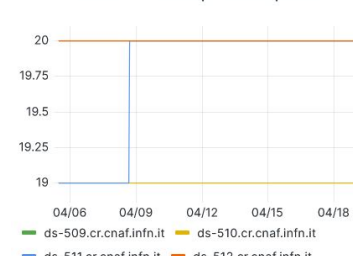
Number of threads per xrootd dampe



Number of threads per xrootd junio



Number of threads per xrootd padme



Stato tape

20 Feb - 20 Apr 2023

MSS bytes in/out (per day)



	min	max	avg	current	total
— out traffic (recalls)	24.5 GB	341 TB	72.7 TB	148 TB	4.36 PB
— in traffic (migrations)	1.10 TB	216 TB	50.7 TB	14.4 TB	3.04 PB

Stato tape

- Liberi 25 PB (su cassette vuote, complessivamente sulle 2 librerie).
Usati 98 PB.
 - Pledge 2022/23: 134.3 PB (130.5 PB pledge 2022)
 - Installato attuale: 130.5 PB
 - Gara da 53 PB pubblicata, per arrivare a 183.5 PB (pledge 2023)

Library	Tape drives	Max data rate/drive, MB/s	Max slots	Max tape capacity, TB	Installed cartridges	Used space, PB	Free space, PB
SL8500 (Oracle)	16*T10KD	250	10000	8.4	~10000	55	23.2
TS4500 (IBM)	19*TS1160	400	6198	20	2450	43	2.3

Evoluzione tape a medio termine

- Gara per acquisto nuova libreria pubblicata
 - Libreria con supporto a drive IBM e LTO
 - In uscita nuova tecnologia tape drive e cassette
 - Da installare al Tecnopolo
- Repack dati da Oracle a IBM (55 PB)
 - Senza spostamento della libreria Oracle al Tecnopolo
 - Ulteriore acquisto di 30 PB di nastro nel 2023 e il resto nel 2024
 - Repack presumibilmente ultimato entro il 2024

Repack Oracle -> IBM

Repack - data moved per day (all tasks)



- Oltre 100 TB / giorno negli ultimi 40 giorni
- Fermo in attesa dei nuovi nastri (giugno)