



State of Storage

CdG 15 dicembre, 2023




Business as usual




Last month

All servers network traffic out (reading)

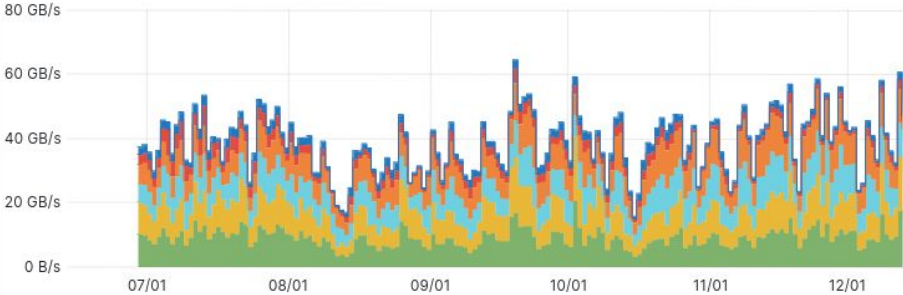


Gateway traffic out (non POSIX reading)




Last 6 months

All servers network traffic out (reading)



Gateway traffic out (non POSIX reading)



Disk storage in produzione

Installed: **53.64** PB, Pledge 2023: **69.6** PB, Used: **48.8** PB

Storage system	Model	Net capacity, TB	Experiment	End of support
ddn-10, ddn-11	DDN SFA12k	10120	ALICE, AMS	12/2022 (+10 spare hdd)
os6k8	Huawei OS6800v3	3400	GR2, Virgo	12/2023
md-1,md-2,md-3,md-4	Dell MD3860f	2308	DS, Virgo, Archive	05/2024
md-5, md-6 e md-7	Dell MD3820f	50	metadati, home, SW	11/2023 e 12/2024
os18k1, os18k2	Huawei OS18000v5	7800	LHCb	7/2024
os18k3, os18k5, os18k5	Huawei OS18000v5	11700	CMS	6/2024
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

Acquisti recenti e futuri

- Gara storage 2022 (14PB netti)
 - LENOVO DE6600: Collaudo non superato
 - Nuova proposta con apparati DDN SFA7990X
- AQ storage 2023-2024
 - Il vincitore è Huawei con sistemi OceanStore Micro 1500/1600
 - Richiesta fornitura di 64PB nel 2023
 - Installazione direttamente al Tecnopolo da inizio dicembre
- Gara Tape Library
 - Contratto presumibilmente ad ottobre / novembre
- Gare nastri
 - 53 PB: pledge + overpledge 2023
 - installato
 - Ulteriori 30 PB di nastro per repack dei dati dalla libreria Oracle
 - Gara in preparazione




Figure 6 - Structure of the 4U high density disk enclosure




Figure 7 - Rear view of the 4U high density disk enclosure


ALICE

Filesystem	Size	Used	Avail	Use%	Mounted on
gpfs_tsm_alice	574T	5.2T	569T	1%	/storage/gpfs_tsm_alice
gpfs_alice	7.7P	6.5P	1.2P	86%	/storage/gpfs_alice

gpfs_alice:

- ddn-10, ddn-11 fuori manutenzione da un anno
- Per coprire guasti abbiamo tolto dalla PROD 2 LU (2x64TB, 20HDD)
- In un anno abbiamo consumato 18 dischi su 20
- Per recuperare 10 dischi “spare” abbiamo tolto un altro LU da 64TB dallo FS gpfs_alice

gpfs_tsm_alice: no data written to tape since July???

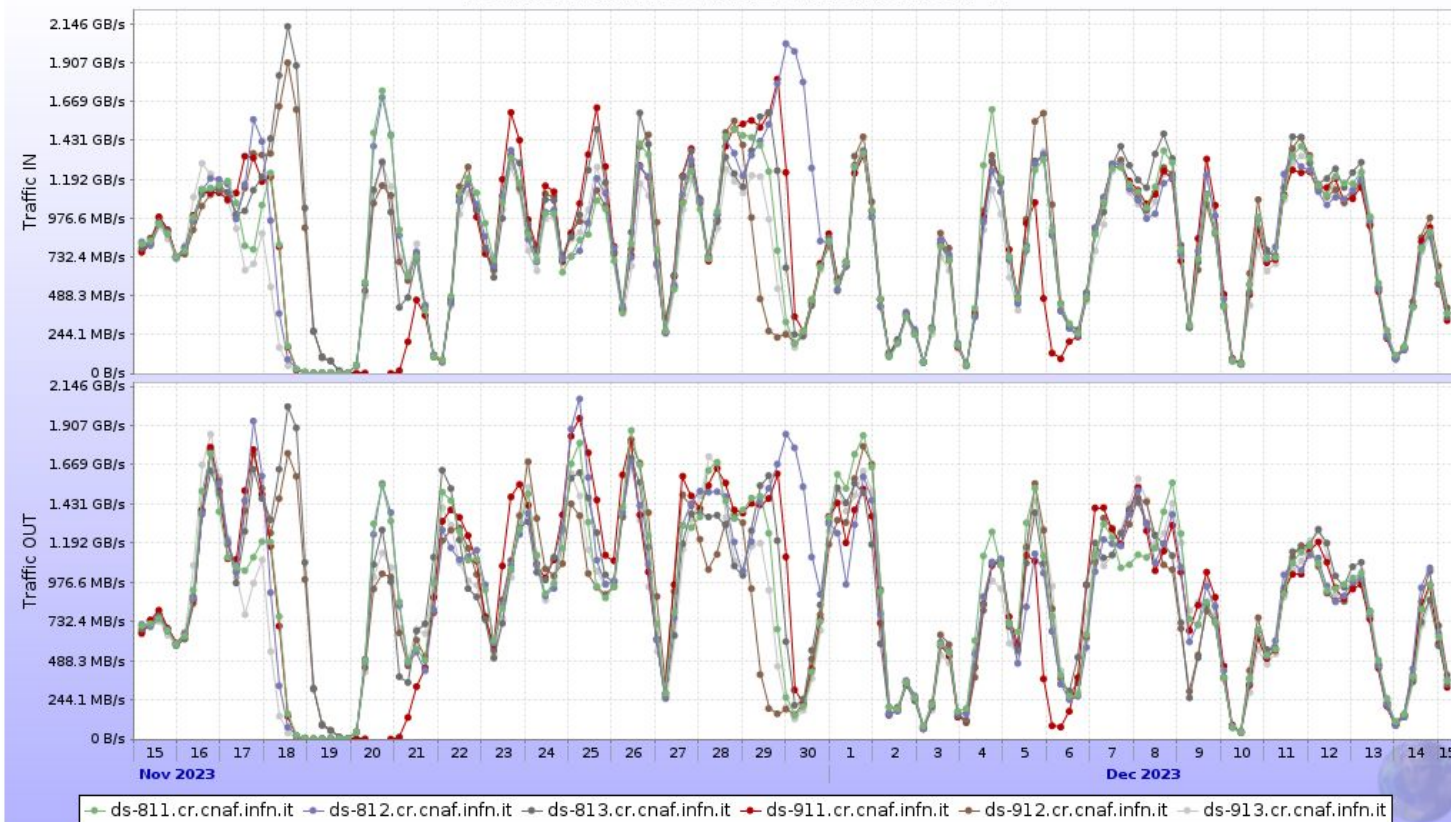


Current SW in PROD

- GPFS 5.1.2-11
- StoRM BackEnd 1.11.22 (latest)
- StoRM FrontEnd 1.8.15 (latest)
- StoRM WebDAV 1.4.2 (latest)
- StoRM globus gridftp 1.2.4
- XrootD 5.5.4-1
 - ALICE CEPH updated to 5.5.5-1.el8
- Ceph 16.2.6 (Pacific)

XrootD ALICE CEPH

Network traffic on ALICE::CNAF::CEPH



Tickets and more

- ALICE
 - Planning XrootD configuration restyling of GPFS cluster. Goals:
 - Manage configuration files with Puppet
 - Revert to original configuration with all servers acting as managers as well
 - Upgrade to latest version in production (5.5.4-1.el7)
 - ds-802.cr.cnaf.infn.it was removed from production and upgraded using puppet (22/11)
 - Check on the status of the service has been included within sensu framework of check and remediation
 - Waiting for Francesco and Mario feedback to proceed with the other endpoints

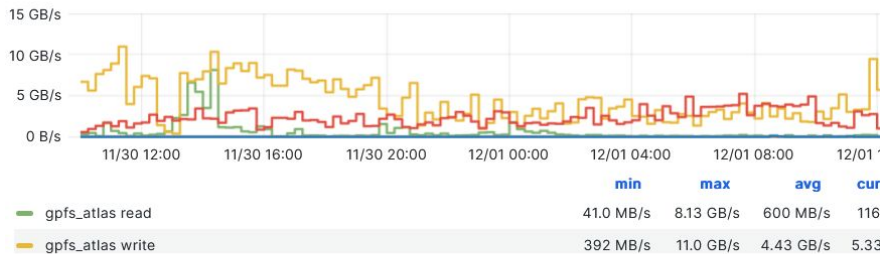
Tickets and more

- ATLAS
 - Low transfer efficiency (GGUS [164631](#))
 - StoRM WebDAV was not restarted after certificate renewal, fixed
 - StoRM Tape REST installed and configured; no tests yet (ATLASGROUPTAPE)
- CMS
 - GridFTP still used, only for SAM tests :-(
 - StoRM Tape REST installed and configured; no tests yet and it did not help in getting rid of GridFTP :-(
 - CMS tape (4796 files) with problems at mount time
 - Waiting for feedback about data recovery to remove volume
 - Problems when archiving files, resulting in 0 checksum files archived on tape (GGUS [162945](#))
 - Either a StoRM or a FTS issue; waiting for reply
 - We reported 28 additional files without checksum which have been migrated to tape


Tickets and more

- 30/11: in view of DC24, ATLAS and CMS simultaneously transferred data from T0 to all T1 for 12 hours
- No saturation for us, but we would like to know more about failures observed
- According to our monitoring, CMS injection led to much lower error rate.


GPFS speed




Total network utilization (IN)



GPFS speed



Total network utilization (IN)



Tickets and more

- LHCb
 - Tests with StoRM Tape REST started, and here is a ticket: GGUS [164634](#)
 - GGUS [164502](#): Data transfer problems, due to misconfiguration of restart remediators for StoRM WebDAV
 - GGUS [164032](#): User fts transfers from INFN are failing
 - Still in progress, but FTS and StoRM WebDAV developers have been involved
 - StoRM WebDAV provides storage tokens (macaroons) via the oauth/token endpoint, but FTS retrieves macaroons from the resource path. This should not be allowed by StoRM WebDAV ([STOR-1602](#)), no matter permissions of the storage area.

Tickets and more

- Gsiftp protocol via StoRM backend is still available for a few experiments
 - Tests srm+https and feedback very **welcome** (Belle, Xenon)
 - Goal: remove gsiftp protocol and switch off GridFTP
- CTA-LST
 - GridFTP switched on for CTA-LST storage areas to allow Third-party copies from PIC (gsiftp)
- Dampe
 - GridFTP “plain” still used
 - Testing XrootD server at IHEP to perform the transfers to CNAF (WP6-Datacloud)
- Dune
 - New storage area with scope-based authN/Z (WLCG profile) and VOMS roles
 - Errors due to old LSC file: unfortunately we are not notified of changes for FNAL voms servers

Stato tape

Last month

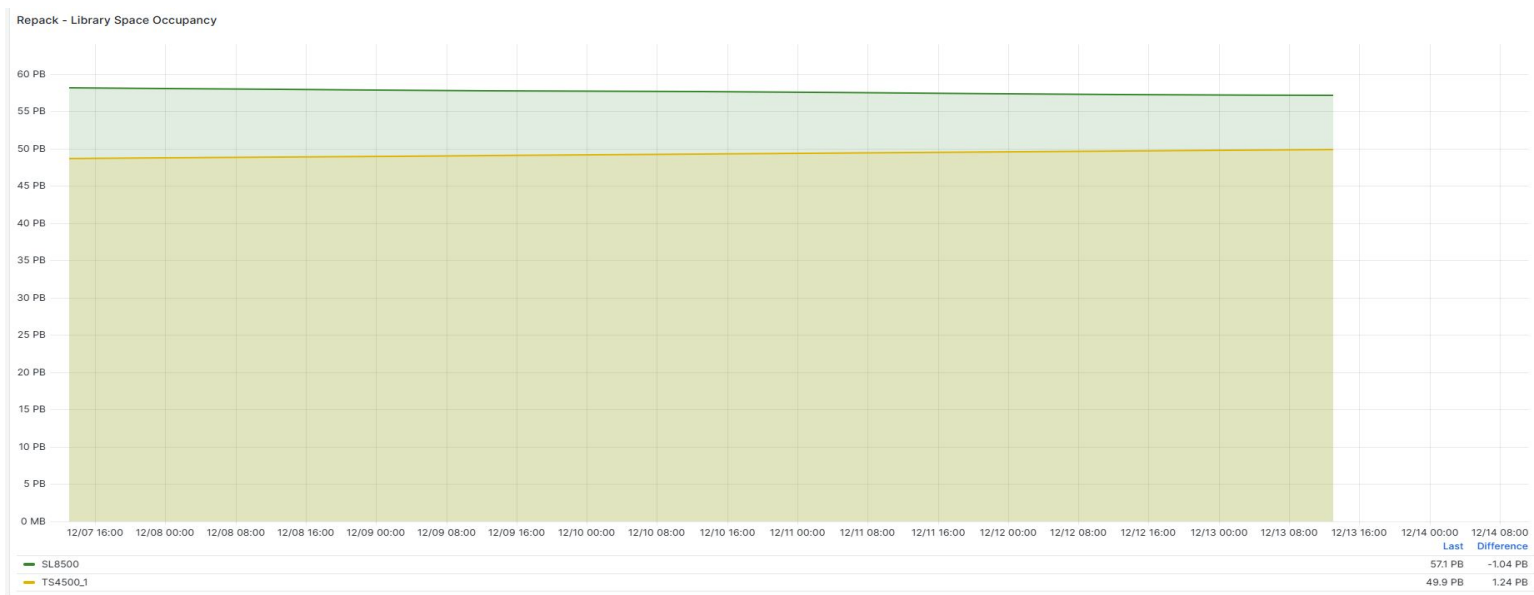
MSS bytes in/out (per day)



	min	max	avg	current	total
— out traffic (recalls)	4.12 TB	245 TB	83.7 TB	43.0 TB	2.43 PB
— in traffic (migrations)	8.99 TB	114 TB	58.5 TB	8.99 TB	1.70 PB

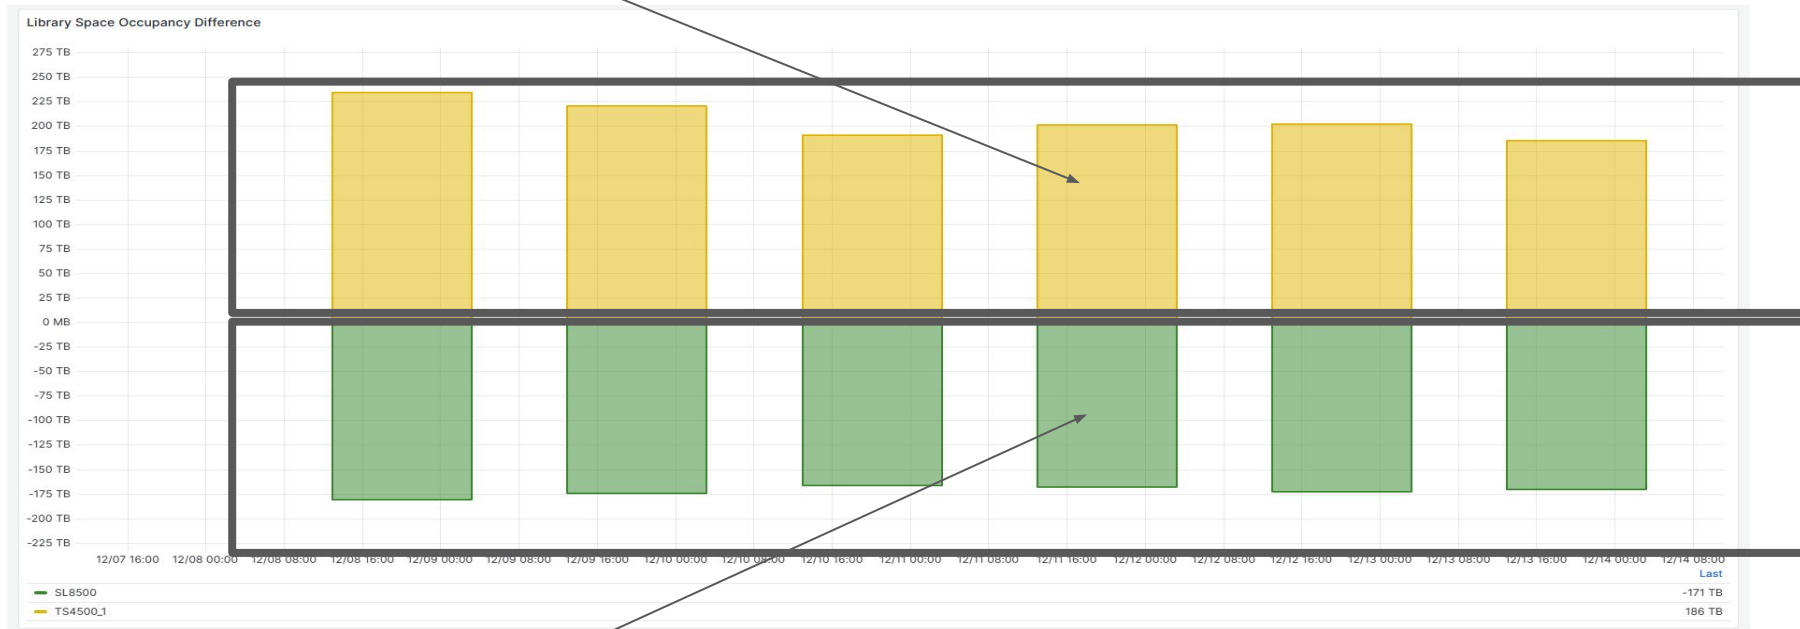
Migration Oracle -> IBM

- 1 anno di tempo stimato per terminare la migrazione alle condizioni attuali
(sono previste cancellazioni?)



Migration Oracle -> IBM

- Rate di migrazione attuale ~2.1 GB/s con 9 processi paralleli
- Rate riempimento IBM → Rate migrazione + Rate scrittura



- Svuotamento Oracle

Stato tape

- Liberi ~46 PB (Scratch tape sulla libreria IBM).
- Usati ~112 PB.
 - In preparazione gara per altri 30 PB

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	57.1	-
TS4500 (IBM)	19*TS1160	400	6198	20	5104	49.9	46