


State of Storage


CdG 17 novembre, 2023



Business as usual


Last month

All servers network traffic out (reading)




Last 6 months


All servers network traffic out (reading)



Gateway traffic out (non POSIX 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 (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	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	1/2024
os18k3, os18k5, os18k5	Huawei OS18000v5	11700	CMS	12/2023
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)
 - Collaudo non superato
 - Vista l'emergenza, tutti gli esperimenti sono invitati a cancellare i dati superflui
- 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
 - Consegnato e in fase di installazione nella libreria IBM
 - 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




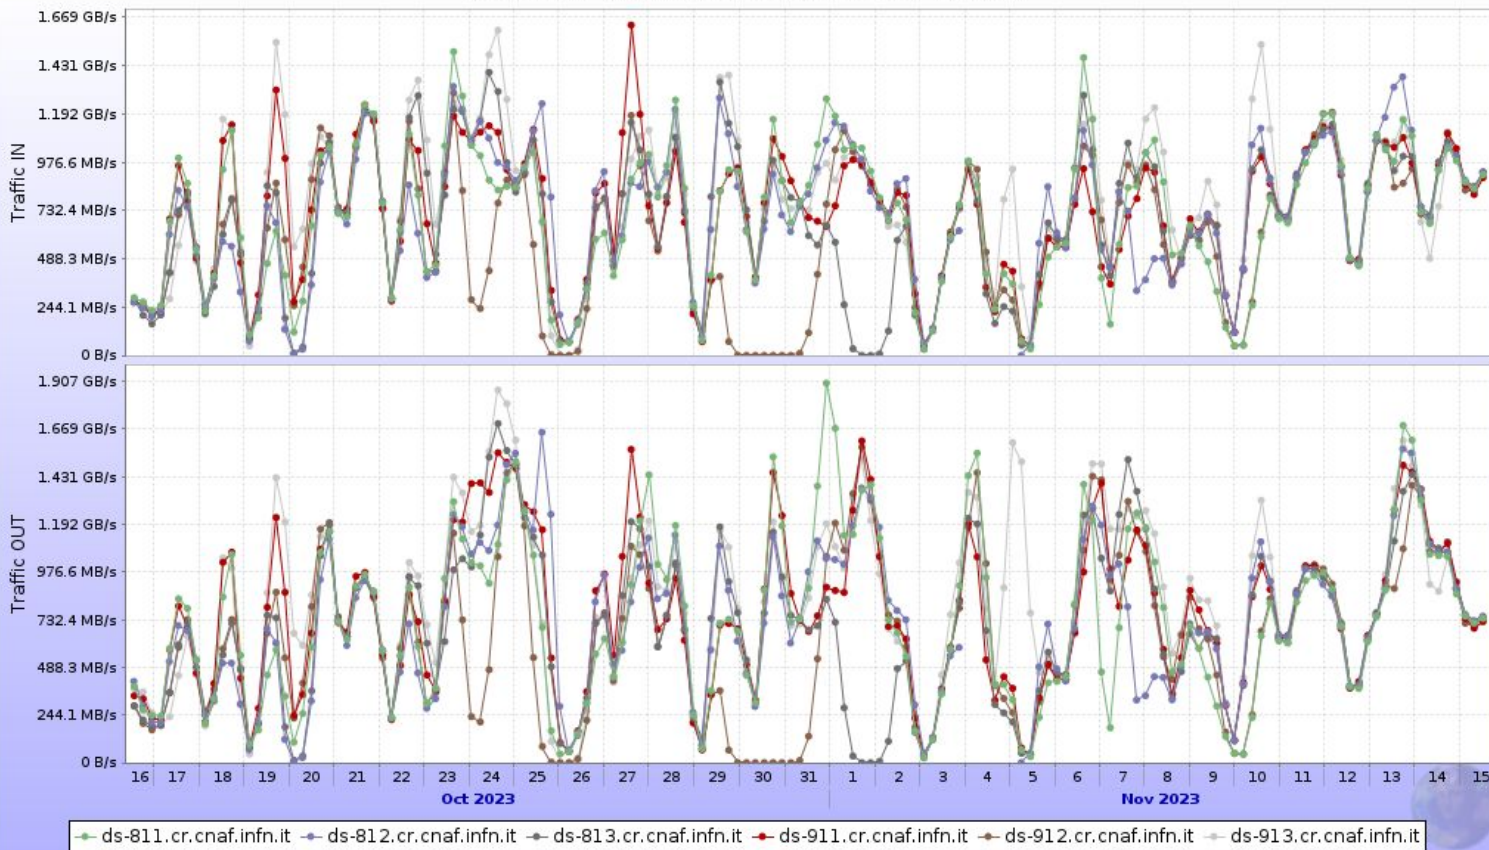
Figure 7 - Rear view of the 4U high density disk enclosure

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

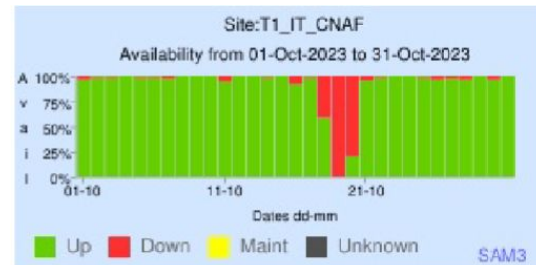


Tickets and more

- 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)

Tickets and more

- ATLAS
 - StoRM Tape REST installed and configured; no tests yet (ATLASGROUPTAPE)
- CMS
 - **Gridftp still used**, only for SAM tests, Daniele tried to remove it on 20/10 apparently causing site availability to fall :-(
 ■ srm and gsiftp being the same protocol for CMS, the only way to get rid of gridftp is move to the TAPE Rest API
 - Problems when archiving files, resulting in 0 checksum files archived on tape (GGUS [162945](#))
 ■ Either a StoRM or a FTS issue; waiting for reply
 ■ 28 files without checksum migrated on tape
 - StoRM Tape REST installed and configured; no tests yet



T1_IT_CNAF Avail: 92% Unkn: 0%

Tickets and more

- LHCb
 - ~~Planning a tape challenge in Oct~~, hopefully also with StoRM Tape REST (ready, to be tested)
 - Postponed to 2024
 - GGUS [163911](#): Failed uploads to CNAF SE
 - GGUS [164032](#): User fts transfers from INFN are failing
 - Still in progress, but FTS developers involved
 - GGUS [164150](#): Data transfer problem to INFN-T1 (solved)
 - Local uploads from Lancaster and Glasgow failed
 - Open tickets to source sites ([164193](#), [164172](#))

Tickets and more

- To cope with space shortage within `gpfs_data`, we moved 200 TB from `gpfs_ams` to `gpfs_data`
- Gsiftp protocol via StoRM backend is still available for a few experiments
 - Tests `srm+https` and feedback very **welcome** (Belle, CTA-LST, Icarus, Xenon)
 - Goal: remove gsiftp protocol and switch off GridFTP
- Belle II
 - Configured IAM at KEK as trusted token issuer for StoRM WebDAV storage areas
- CTA-LST
 - Third-party copies from PIC (gsiftp) to CNAF (`srm+https`) failed
 - Switching GridFTP on for CTA-LST storage areas solved the problem


Tickets and more

- Dampe
 - No feedback from Bari on the issue of XrootD redirector always contacting one out of 4 available XrootD servers at CNAF
 - GridFTP “plain” still used
 - XrootD server at IHEP to perform the transfers to CNAF
 - Waiting for tests in the WP6-Datacloud framework
- Dune
 - New storage area with scope-based authN/Z which follows the WLCG profile
- Limadou
 - Changed owner (from mergelimadou to ffollega) and access control list for all the files in `/storage/gpfs_data/limadou/data`
- SWGO
 - Data moved from `/storage/gpfs_data/ctalocal/ctaplus` to `/storage/gpfs_data/swgo`

Stato tape

15 Sep - 16 nov 2023

MSS bytes in/out (per day)



— out traffic (recalls)
— in traffic (migrations)

min	max	avg	current	total
14.9 GB	245 TB	62.7 TB	245 TB	3.95 PB
3.35 TB	154 TB	62.9 TB	45.7 TB	3.97 PB

Stato tape

- Liberi ~52 PB (su cassette vuote. Sulla libreria IBM).
- Usati ~108 PB.
 - Nuove cassette sulla libreria IBM - ripartite le scritte
 - Ripresa anche la migrazione dei dati da Oracle (in dismissione) a IBM
 - 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	62	-
TS4500 (IBM)	19*TS1160	400	6198	20	2450	45.8	52