

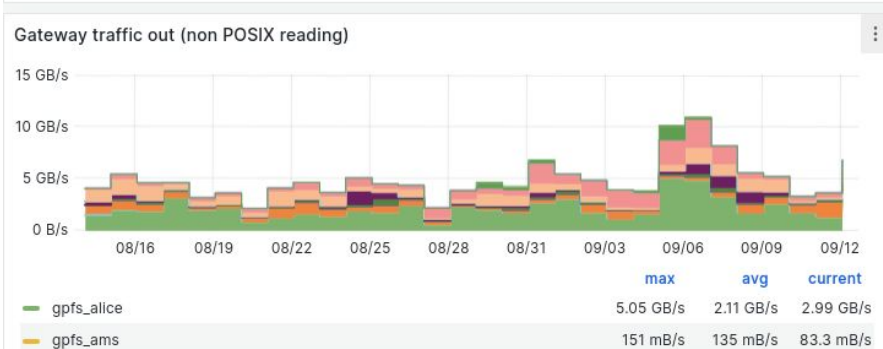
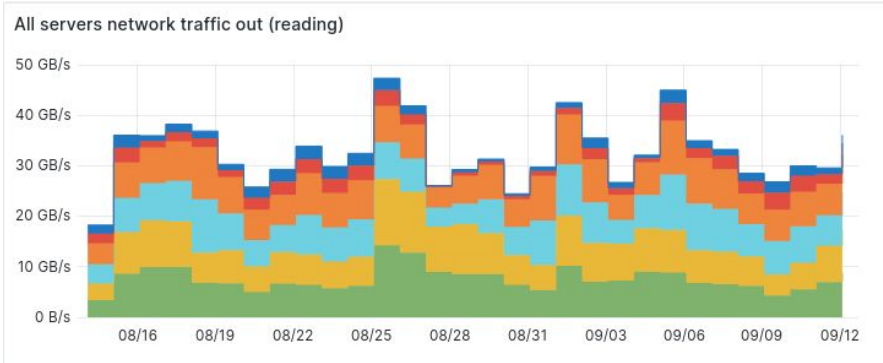
State of Storage

CdG 15 settembre, 2023

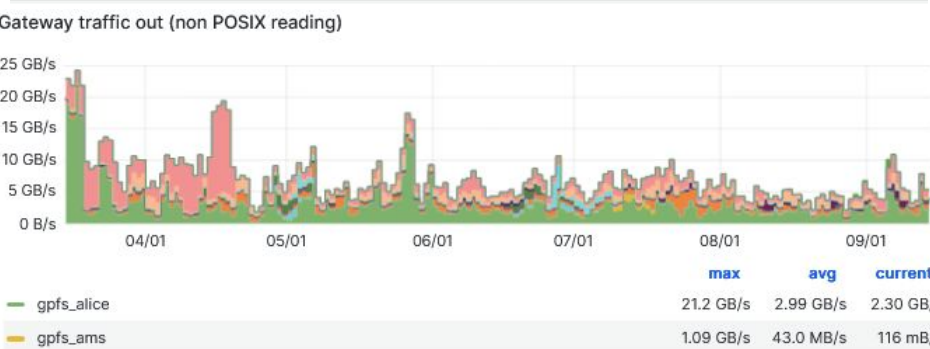



Business as usual

Last month



Last 6 months



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	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	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
- **Gara Tape Library**
 - Contratto presumibilmente ad ottobre / novembre
- **Gare nastri**
 - 53 PB: pledge + overpledge 2023
 - consegna prevista fine settembre / inizio ottobre
 - Ulteriori 50 PB di nastro per repack dei dati dalla libreria Oracle
 - da acquistare in autunno




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

- XrootD ALICE su CEPH
 - Più stabile del passato
 - I servizi (CEPH) che gestiscono i metadati vanno in sofferenza in caso di alto carico in lettura
 - Indagine in corso

AliEn SEs availability for reading



Tickets and more

- ATLAS

- StoRM Tape REST installed and configured; we'll contact Petr for tests (ATLASGROUPTAPE)
- 0% Transfer efficiency to tape (GGUS [162735](#)): space freed, efficiency back to normal
- TPCs in push-mode to GOOGLE_EU fail for all StoRM sites (GGUS [158487](#))
 - Issue for StoRM developers, under investigation

- CMS

- **Gridftp still used**, only for SAM tests :-)
- SAM tests failing due to network issues (GGUS [163170](#))
- Problems when archiving files, resulting in 0 checksum files archived on tape (GGUS [162945](#))
 - Either a StoRM or a FTS issue; waiting for reply

- LHCb

- Planning a tape challenge in Oct, hopefully also with StoRM Tape REST (ready, to be tested)
- Wondering where site-report info are published (GGUS [162826](#))
 - Correctly in <https://xfer-lhcb.cr.cnaf.infn.it:8443/info/report.json>
- Migration within filesystem done (GGUS [162740](#)), to avoid srm calls to deduce physical paths

Tickets and more

- **Https protocol now available for all storage areas exposed via StoRM backend**
 - Tests srm+https and feedback very **welcome** (Agata-Gamma, Auger, Belle, CTA, Dampe, DarkSide, Glast, Icarus, ILDG, NA62, Padme, Pamela, Virgo, Xenon)
 - Goal: remove gsiftp protocol and switch off gridftp
- **Asfin**
 - Data exposed via StoRM WebDAV (token-based auth) in a dedicated storage area
- **Cygno**
 - Discussing tape data access within WP6-Datacloud framework
- **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

Tickets and more

- To cope with space shortage within `gpfs_data`, we decided to move 200 TB from `gpfs_ams` to `gpfs_data` (ongoing)
- Darkside
 - Support users in the process of moving files out of `gpfs_ds50` (downloading them via a dedicated StoRM WebDAV storage area) to free space and in cleaning directory permissions
- Gamma
 - Data exposed via StoRM WebDAV (token-based auth) in a dedicated storage area
- Gerda
 - Request to backup data on tape, freeing disk: two copies, one on each tape library
- HyperK
 - Failing jobs because no space available (GGUS [163049](#)), quota increased from 10 to 20 TB
- Muone
 - Upon request, configured a StoRM WebDAV tape storage area (only write, token-based auth)
 - After a correction of the estimate of data amount, the experiment decided not to use it


Tickets and more

- Nucleus
 - Data exposed via StoRM WebDAV (token-based auth) in a dedicated storage area
- SWGO
 - Dedicated space for software configured
 - Moving data from /storage/gpfs_data/ctalocal/ctaplus to /storage/gpfs_data/swgo (ongoing)
- Virgo
 - Call to discuss transfer inefficiency from Cascina
 - Waiting for feedback from the experiment, namely the precise time periods where the problem was spot, and tolerance to short network saturation periods

Stato tape

17 Jul - 14 Sep 2023

MSS bytes in/out (per day)



— out traffic (recalls)

— in traffic (migrations)

min	max	avg	current	total
162 GB	167 TB	43.5 TB	3.90 TB	2.61 PB
946 GB	171 TB	68.0 TB	133 TB	4.08 PB

Stato tape

- Liberi ~16 PB (su cassette vuote, complessivamente sulle 2 librerie).
Usati 107 PB.
 - Installato attuale: 130.5 PB
 - Spazio quasi esaurito sulla IBM -> maggior parte delle scritture sulla Oracle in attesa dei nuovi nastri

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	61.2	15.5
TS4500 (IBM)	19*TS1160	400	6198	20	2450	45.8	0.3