

State of Storage

CdG 21 Marzo, 2025



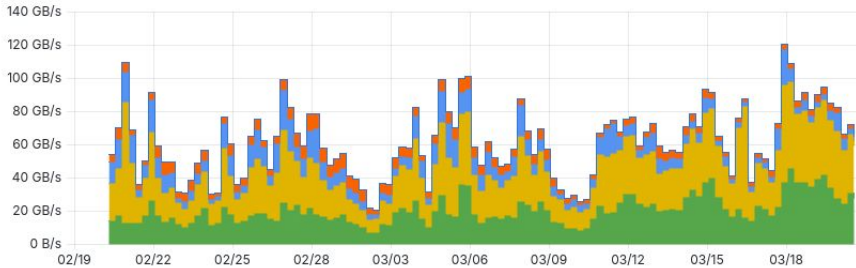
Business as usual + migration to TP




Last month

Last 6 months


All servers network traffic out (reading)




All servers network traffic out (reading)



Gateway traffic out (non POSIX reading)



Gateway traffic out (non POSIX reading)



Disk storage in produzione 77.5PB

Installed: **101.4PB**; Pledge 2025/01: **86.95PB**; Pledge 2025/04: **101PB**; Used: **60PB**

Storage system	Model	Net capacity, TB	Experiment	End of support
os5k8-1,os5k8-2	Huawei OS5800v5	8999	GR2	2027
ddn-12, ddn-13	DDN SFA 7990	5840	LHCb (commissioning)	2025
ddn-14, ddn-15	DDN SFA 2000NV (NVMe)	24	Metadata (commissioning)	2025
ddn-16	DDN SFA 2000NVX2 (NVMe)	96	LHCb metadata, hotdata	2031
ddn-17,ddn-18,ddn-19	DDN SFA 7990X	14000	LHCb	2031
od1k6-1,2,3,4,5,6	Huawei OD1600	60000(-10%!)	ALICE,ATLAS,CMS, GR2	2031
od1k6-7,8	Huawei OD1600	18000	CMS, GR2	2031
od1k5-1,2	Huawei OD1500 (NVMe)	400	Metadati, varie buffer	2031

Disk storage in dismissione: 32.2PB


Storage system	Model	Net capacity, TB	Experiment	End of support
os6k8	Huawei OS6800v3	3400		07/2024
ddn-10, ddn-11	DDN SFA12k	10120		12/2022
md-1,md-2,md-3,md-4	Dell MD3860f	2308		12/2024
md-5,md6,md-7	Dell MD3820f	50	SW (4TB)	11/2023 e 12/2024
os18k1, os18k2	Huawei OS18000v5	6520		7/2024
os18k3, os18k5, os18k5	Huawei OS18000v5	9780	ATLAS (560TB buffer tape)	6/2024

Traslochi recenti e futuri

- Trasloco della libreria IBM: completato
- Trasloco della libreria Oracle: completato
- Gara nastri: conclusa
 - 96PB di tape JF (50TB/cassetta) in arrivo entro 2 mesi (?)
- Spostamenti dati CNAF→Tecnopolo
 - AMS 2.6PB - fatto
 - HOME 2.6TB→20TB (NVMe) - fatto
 - Area SW 2TB→4TB (NVMe) - da fare
 - Buffer tape di ATLAS 560TB - in progress
- Spostamenti “interni” per ridurre lo spazio “assegnato” sui sistemi storage con ThinProvisioning - fatto
- 10 server per data management (con sostituzione schede di rete 100Gb→25Gb)
 - + varie recuperi (Dell MD7, Switch FC e Eth, etc)

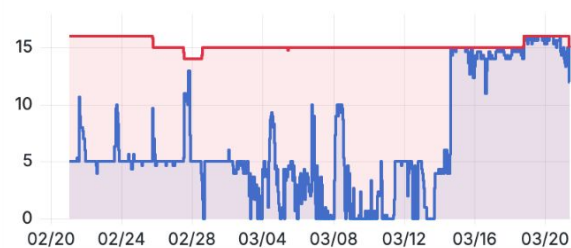
Stato tape (Last month)

MSS bytes in/out (per day)

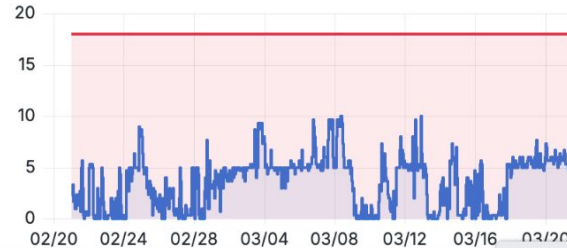


Name	Mean	Last *	Max	Min	Total
— out traffic (recalls)	79.8 TB	118 TB	204 TB	8.85 TB	2.31 PB
— in traffic (migrations)	179 TB	213 TB	253 TB	28.4 TB	5.19 PB

T10kD drive in use vs. active



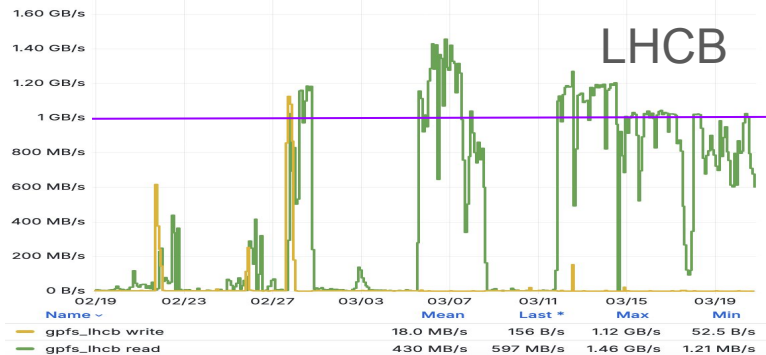

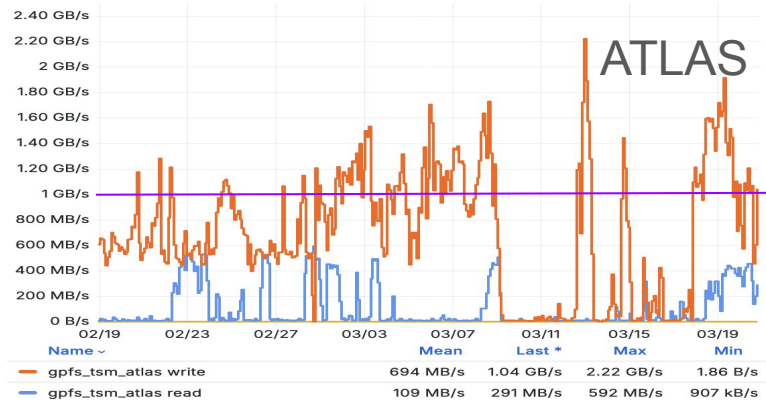

TS1160 drive in use vs. active



TS1170 drive in use vs. active



HSM activity by LHC experiments



Stato tape

- Spazio tape sulle 2 librerie IBM praticamente esaurito
- SL8500 back in business (per almeno 2 mesi)
- Liberi ~41 PB (cassette parzialmente piene o vuote)
- Usati 156 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	9439	46.2	30
TS4500 (IBM)	19*TS1160	400	6198	20	5474	99.7	14
TS4500-2(IBM)	18*TS1170	400	7844	50	163	5.4	3.5

Current SW in PROD

- GPFS 5.1.9-4, migrating to 5.1.9-7
- StoRM BackEnd 1.11.22 (latest, 1.12.0 in test)
- StoRM FrontEnd 1.8.15 (latest, 1.8.16 in test)
- StoRM WebDAV upgraded to 1.5.0 and 1.7.1 (latest)
 - [Scitags](#) enabled for ATLAS and to be enabled for CMS
- StoRM Tape 0.8.0 (latest)
- XrootD 5.6.9 for ALICE (disk), XrootD 5.7.3 for CMS, XrootD 5.5.5 for LHCb and XrootD 5.6.8 for no-LHC
 - Still 5.5.4 for ATLAS (needed?)

Tickets and more

- ALICE

- Installed and configured 4 EL9 servers for disk cluster (xs-606,xs-607,xs-706,xs-707)
 - Anything to be done to enable Scitags?
- Open action: finalize the configuration for the XrootD tape cluster (xs-204, xs-304)
 - Waiting for new servers @Tecnopolo to install and test rpm for interaction xrootd-tape

- ATLAS

- GGUS [#682733](#) (solved): transfer errors due to “spontaneous” reboot of one WebDAV server
- GGUS [#682434](#) (solved): "INFN-T1 staging failures"
 - 144 failed stage requests all at the same time due to a temporary network problem
- GGUS [#682211](#) and [#682109](#) (on hold): transfer failures and deletion timeout errors
 - Very small timeout (20s) for deletions (?)
 - The disk storage (gpfs_atlas) reordering is completed and there are several concurrent activities. Shall we change the status to “in progress” or “solved”?
 - Tape buffer storage (gpfs_tsm_atlas) movement is still ongoing

Tickets and more

- CMS

- GGUS #[682638](#) (solved): “Request for XRootD Upgrade and Network Packet Labeling Configuration”
- GGUS #[682009](#) (in progress): “CNAF archived transfers and failures requests”
 - Further problems: requests to tape-rest-API timed out
 - zombie processes hanging on the filesystem were reaped
 - It seems CMS is trying to overwrite files on buffer by resubmitting FTS requests
 - DM team involved: FTS reports error, but Rucio marks the transfer as good
 - It will be fixed by the end of the year
 - FTS also issues an /archiveinfo call to the StoRM Tape REST API for each file every 10 minutes and declares a transfer as failed after 7 days
 - But the pre-migration process is asynchronous!

Tickets and more

- LHCb
 - GGUS #[682253](#) and #[682260](#) (solved): failed data transfers
 - Problems with 2 StoRM WebDAV endpoints (out of 8 in xfer-lhcb)
 - Re-installed on e19
 - GGUS #[682244](#) (solved): a bunch of files stuck in staging
 - A couple of cartridges were in error state
 - GGUS #[682599](#) (old [169555](#), [681921](#)) (to be closed): “Corrupted files at CNAF v3”
 - Network issues on the 28th of Dec resulted in 4 corrupted files on disk
 - We checked all the files transferred between the 18:00 and the 20:00 (23k) and there are no other corrupted files

Tickets and more

- Gsiftp protocol switched off also for Xenon and CTA-LST
 - Replaced by StoRM WebDAV (<https/davs>)
- Belle II
 - Belle community is testing StoRM Tape REST API with fine grained AuthZ supporting more than one experiment
- Dune
 - GGUS [#682553](#) (in progress): “Enabling WLCG Token access for DUNE storage”
 - GGUS [#682396](#) (to be closed): “Compute and storage submissions failing to CNAF”
 - Errors with XrootD 5.5.5 from el7 clients, but no errors from el9 ones
 - Upgrading to 5.6.8 solved the issue
- KLOE
 - New disk StoRM WebDAV storage area, AuthN/Z with JWT of iam-t1-computing
 - <davs://xfer-archive.cr.cnaf.infn.it:8443/kloe>