


State of Storage


CdG 17 Ottobre, 2025



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 production **96.3PB**

Installed: **100.1PB**;


Pledge 2025: **101PB**;

Used: **79PB**

Storage system	Model	Net capacity, TB	Experiment	End of support
os5k8-1,os5k8-2	Huawei OS5800v5	8999	GR2	2027
ddn-12	DDN SFA 7990	4550	LHCb	2025
ddn-14, ddn-15	DDN SFA 2000NV (NVMe)	24	LHCb	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

State of tape (last 3 months)


MSS files in/out (per day)




MSS bytes in/out (per day)




T10kD drive in use vs. active



TS1160 drive in use vs. active



TS1170 drive in use vs. active




State of tape

- Installed 96 PB (TS4500-2) in June
- SL8500 repack paused: → TS4500-2
 - 39.4 PB: still to migrate
- ~50 PB total available on IBM libraries
 - or only ~10 PB free considering space needed for repack


Library	Tape drives	Max data rate/drive, MB/s	Max slots	Max tape capacity, TB	Installed cartridges	Used Tape space, PB	Scratch Tape space, PB
SL8500 (Oracle)	16*T10KD	250	10000	8.4	9387	50.5	28.2
TS4500 (IBM)	19*TS1160	400	6198	20	5472	106	2.6
TS4500-2(IBM)	18*TS1170	400	7844	50	2082	53.4	50.5

Repack Oracle → IBM


Repack - Library Space Occupancy



Repack - Library Scratch Tape



Library Space Occupancy Difference



Current SW in PROD

- Updating GPFS up to 5.1.9-12 on el9 (5.1.9-4 on el7)
- StoRM BackEnd 1.11.22 and StoRM FrontEnd 1.8.15 (latest, testing upgrade)
- StoRM WebDAV upgraded to 1.11.0 (latest, in beta)
 - Spring Boot 3, scitags, new monitoring metrics and log format, tape storage areas, nginx support, virtual threads
- StoRM Tape 0.9.0 (latest) and a new StoRM Tape Authz 1.1.0 (latest)
 - Telemetry via OpenTelemetry, new authorization policies (similar to StoRM WebDAV fine-grained authorization)
 - CMS has a different deployment with dedicated endpoint for polling requests
- XrootD 5.6.9 for ALICE (disk and tape), XrootD 5.7.3 for CMS, ATLAS, and no LHC, XrootD 5.5.5 for LHCb
 - One dedicated xrootd server for ATLAS

Tickets and more

- ALICE
 - No AuthZ on 2 out of 6 XrootD disk servers: AuthZ files have been fixed
 - Anything to be done to enable Scitags?
- ATLAS
 - GGUS-Ticket-ID: #[1000134](#) (on hold) “Transfer and staging errors”:
 - Several different issues and solutions in one ticket, but the debug of the low efficiency of transfers from CNAF tape to CNAF disk is still ongoing together with StoRM devels
 - Kind of duplicates of this one: GGUS-Ticket-ID: #[1000072](#), GGUS-Ticket-ID: #[683563](#)

Tickets and more

- CMS
 - Xrootd proxy for Leonardo installed and ready for tests
 - Robot certificate of A. Pascolini installed
 - GGUS-Ticket-ID: #[1000548](#) (in progress) “Invalidating corrupted files to avoid FTS failures (making T1 CNAF unusable)”
 - We did provide a list of files without checksum in GGUS-Ticket-ID: #[682009](#) ; these were not invalidated but they are being discussed in a jira issue

Tickets and more

- CMS (long-standing issues with tape archival)
 - GGUS-Ticket-ID: #[682009](#) (in progress) “CNAF archived transfers and failure requests”
 - Several different issues with data archival: polling, when to declare successful/failed transfer to tape, missing 'transfer-done' messages in Rucio, and ultimately disagreement in monitored space and rates between site and CMS
 - Following two productive meetings in July and August with CMS Ops:
 - dedicated endpoint for /archiveinfo requests within the StoRM Tape REST API lead to strong decrease of 502 errors when polling files (solved)
 - generation of lists available in /info storage area containing all the files in tape storage area (CMS consistency check to be started soon)
 - generation of list with information on missing/present checksum for such files
 - changed a configuration parameter in GEMSS so that only files with the checksum extended attribute are staged out to tape asap, the other files being staged out to tape anyway after 5 days

Tickets and more

- LHCb
 - GGUS-Ticket-ID: #[1000616](#) (in progress) “More corrupted file on disk”
 - Found only one corrupted file
 - Investigated the problematic transfer and the corrupted file together with the StoRM developers, but we could not find any explanation apart from generic instabilities in the local network. We do not manage to reproduce the problem, and we cannot find any other evidence of data corruption

Tickets and more

- Belle

- GGUS-Ticket-ID: #[1000052](#) (solved): permission error when accessing data
 - Again KEK LSC file (<https://github.com/italiangrid/voms/issues/141>), this time with nginx in front of StoRM WebDAV
- “srm” protocol is still used. Can we switch off the SEs?

- KLOE

- 355k files from KLOE1 (210 TB) were repacked onto the new IBM library
- Checksums were verified, and only three files did not have them
- We are applying the same procedure to all the other files (1.17PB)

- VIRGO

- Upgraded Pelican OSDF cache ([OSG support ticket](#)) to 7.20, as requested by OSG
 - Throttling configured, as requested, to improve concurrency handling