






State of Storage

CdG 21 giugno, 2024





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 (+10 spare hdd)
os6k8	Huawei OS6800v3	3400	GR2, Virgo	07/2024
md-1,md-2,md-3,md-4	Dell MD3860f	2308	DS, Virgo, Archive	12/2024
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)
 - Nuova proposta con apparati DDN SFA7990X
 - In attesa per la consegna entro giugno
- AQ storage 2023-2024
 - Huawei OceanStore Micro 1500/1600
 - 8 sistemi di 10PB + 40 server
 - Richiesta fornitura di 64PB nel 2023
 - Installazione e collaudo in corso (60PB)
- Tape Library
 - Installata, manca cablaggio FC e le cassette JF da 50TB - in arrivo (speriamo bene!) prossima settimana
- Gare nastri
 - Acquistati 14PB (JE e JF)
 - Nuova gara di acquisto tape JF in preparazione



Current SW in PROD

- GPFS 5.1.2-13 (in preparazione migrazione alla 5.1.9-4 che supporta RHEL9)
- 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)

Tickets and more

- ALICE
 - Finishing XrootD configuration restyling of GPFS cluster:
 - Manage configuration files with Puppet
 - Upgrade to latest version in production (5.5.4-1.el7)
 - Check on the status of the service has been included within sensu framework of check and remediation
 - Finalizing the configuration with the tape cluster (xs-204, xs-304)
 - Issues with the “tsm” RPM building
 - We need to migrate data on CEPH FS (~2PB) to GPFS at Tecnopolo. Two proposed solutions:
 - Bulk rsync transfer
 - Set an XrootD SE at Tecnopolo and transfer data via grid

Tickets and more

- ATLAS
 - StoRM Tape REST in production; three storage areas
 - ATLASGROUPTAPE, ATLASMCTAPE, ATLASDATATAPE
 - GGUS [166125](#) (closed): “INFN-T1 has low transfer efficiency as dst site : No space left on device”
 - Too high writing rate (recalls included) on the disk buffer
 - DDM Atlas team reduced the writing rate to 1.2GB/s
 - But only for the related campaign
 - SSR did not report the right free space
 - Despite the failures and the FS status, transfers were submitted
 - Implemented a script that automatically reports the right status of the FS in the report.json file used by Atlas monitoring

Tickets and more

- ATLAS
 - GGUS [166882](#) (open): “INFN-T1 has staging errors”
 - 30/05: concurrent activities on two tape SA filled the disk buffer again
 - We asked again to reduce the writing rate to 1.2GB/s
 - 17/06: 645 stage requests in only one minute failed with 502 error
 - 22k file stages failed
 - Nginx closes the connection to the StoRM Tape REST unit when there is a high load on the server and a consequent timeout error occurs
 - We are setting proper Nginx parameters to avoid this behaviour
 - 19/06: Atlas reported 56k file stages failed with 403 error in the last 2 days
 - Permission denied due to trailing slash
 - Same problem reported at LHCb GGUS [164634](#)
 - 20/06: changing StoRM tape REST configuration files to accept ending slash in the path solved the issue

Tickets and more

- CMS
 - StoRM Tape REST in production, a single storage area with two access points:
 - /cmstape, /cmstape/store
 - GridFTP switched off starting from April
 - Removed “srmhttpstests” tape data
- LHCb
 - GGUS [167045](#) (open): 4 servers handle both http requests, acting as StoRM WebDAV servers, and requests to the underlying file system, acting as GPFS NSD servers
LHCb has been overloading these servers in both ways
 - We planned to improve the situation with new servers already installed at the Tecnapolo area in order to separate NSD and StoRM WebDAV servers

Tickets and more

- Gsift protocol via StoRM backend is still available for two experiments
 - New StoRM release should allow to switch GridFTP off (Xenon, CTA-LST)
- Belle II
 - GGUS [166666](#) (closed): some transfers failed because the FS was full
 - GGUS [166715](#) (closed): an old gfal version used by Dirac did not allow to read the report.json via “srm+https”
 - Dirac upgrade solved the issue
- Cygno
 - Renaming files (55k files, 50TB) on tape in order to align the path to that one on disk
 - Rucio requirement
 - Manual intervention
 - From /storage/gpfs_archive/cygno to /storage/gpfs_archive/cygno/cygno-data

Tickets and more

- Dampe
 - GridFTP “plain” still used
 - TPCs between XrootD server at IHEP and CNAF are working well
 - Rucio+FTS (https) should replace the current gsiftp transfers (WP6-DataCloud)
- DUNE
 - Request to expose data in read mode also via XrootD
 - 4 XrootD servers shared among other experiments
 - Waiting for the configuration files from DUNE community


Tickets and more

- HyperK
 - GGUS [167154](#) (open): random errors contacting the StoRM WebDAV SA `davs://xfer-archive.cr.cnaf.infn.it:8443/hyperk`
 - One of the servers in the “xfer-archive” alias is being banned by ngs.ac.uk
 - Hence CRLs do not get updated
 - The endpoint has probably been banned following a false security alarm raised one month ago from CERN (<https://www.virustotal.com/gui/ip-address/131.154.128.183/community>)
- Newchim
 - Tape data restore - 150k file are now ready to be transferred from the disk buffer
 - Manual intervention
 - New tape StoRM WebDAV storage area configured for this purpose
 - `davs://xfer-archive.cr.cnaf.infn.it:8443/newchim-tape`
 - Group-based AuthN/Z with JWT issued by iam-t1-computing

Stato tape



MSS bytes in/out (per day)

Last 6 month




18.6 PB of new data written to tapes since 1/1/2024
 +20PB migrated from SL8500 (see next slide)

Typical usage of the tape drives




Tapes: Migration from Oracle to IBM library


Repack - data moved per week (reclamation vs migration)



Repack - Library Space Occupancy



Number of tape drives used for repack



Repack - number of free cartridges



Stato tape

- Liberi ~4.6 PB (Scratch tape sulla libreria IBM).
- Usati ~125 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	36	-
TS4500 (IBM)	19*TS1160	400	6198	20	5104	89	4.6

BACKUP slides

ATLAS DC24

Target: (min) 24.2 Gbit/s = 3.03GB/s

(max) 27.7 Gbit/s = 3.46GB/s

Archived week average (disk)= 5.31 GB/s

Total (disk+tape) = 7.21GB/s

Total network utilization (IN) - all hosts of selected VOs

