State of Storage

CdG 20 Giugno, 2025

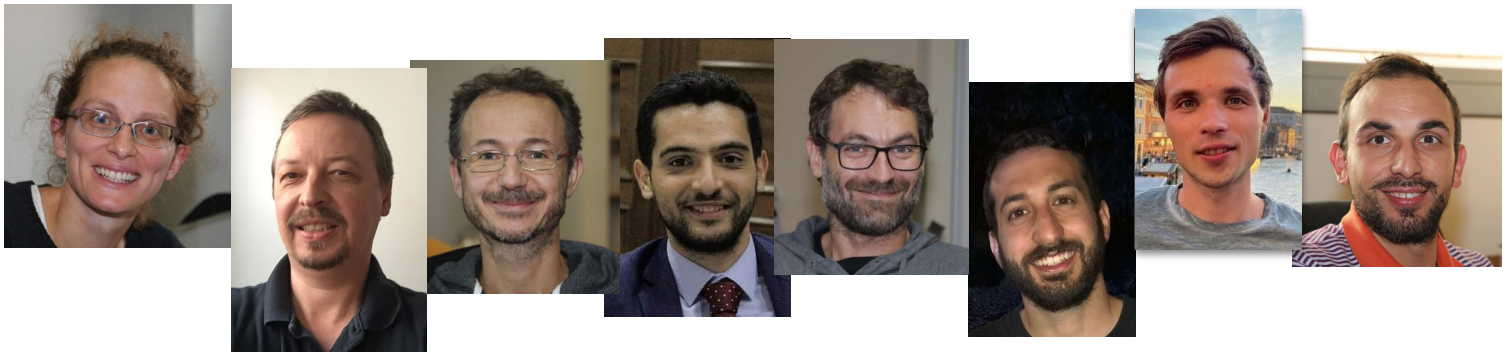

Back to Business as usual

Last month

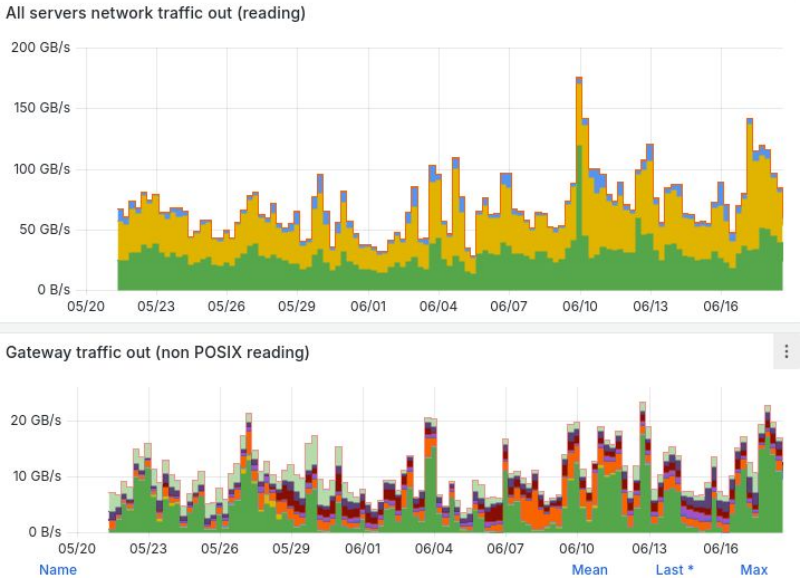

Last 6 months

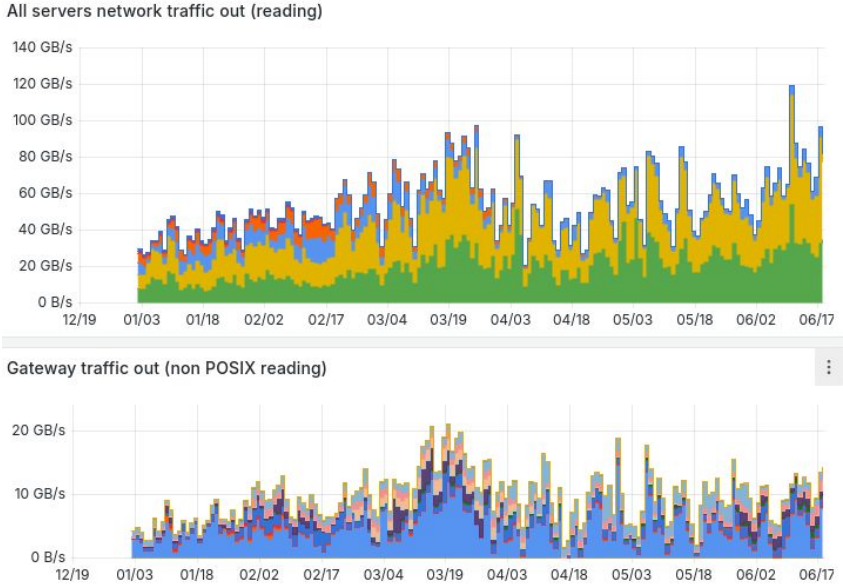

Disk storage in production **96.3**PB

Installed: **100.1**PB; Pledge 2025: **101**PB; Used: **71.2**PB

Storage system	Model	Net capacity, TB	Experiment	End of support
os5k8-1,os5k8-2	Huawei OS5800v5	8999	GR2	2027
ddn-12, ~~ddn-13~~ →ddn-12	DDN SFA 7990	~~5840~~ →4550	LHCb	2025
ddn-14, ddn-15	DDN SFA 2000NV (NVMe)	24		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

Some adjustments

- DDN-12 and DDN-13 out of support contract as of 06.2025 (and EoL 11/25)
  - 5 years old
  - Support provider cannot guarantee spare parts
  - Decided to rebuild one system out of two
    - Each system has 1 controller unit with 90 HDDs + 2 JBODs (90 HDDs each)
    - Rebuild (recabled) two systems into one with 4 JBODs keeping 1 controller unit and 90 HDDs as “cold” spare
    - To minimize on-site interventions for disks replacement we created “hot spare pool” of 20 HDDs
    - During “rebuild”, discovered source of frequent HW errors - damaged cables
  - Usable capacity reduced by 1290 TB: (from 5840TB to 4550TB)
  - In PROD for LHCB

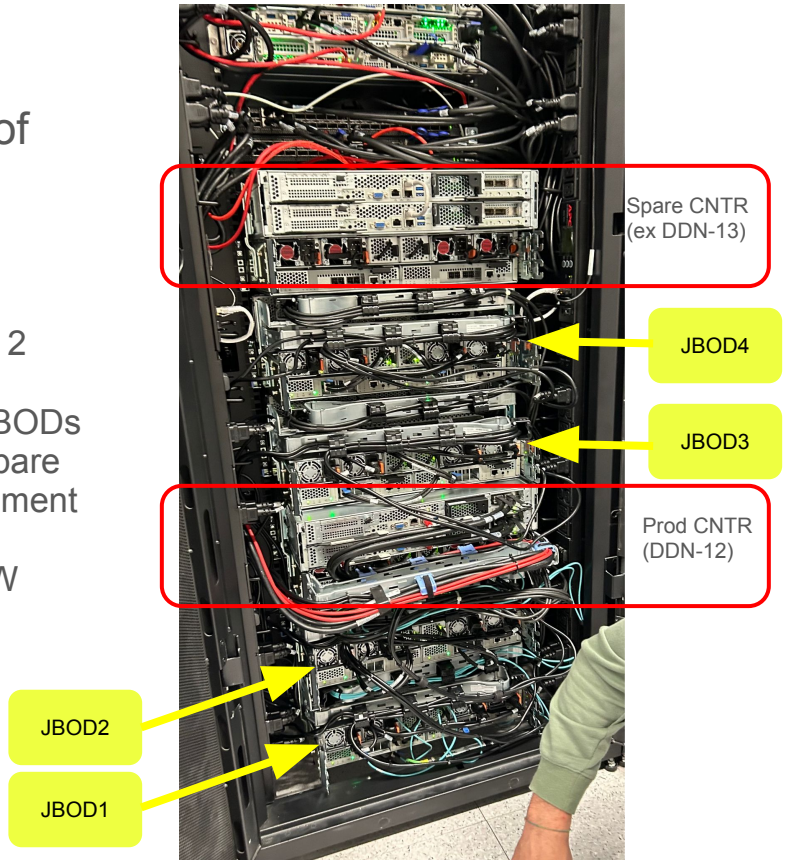

Use of NVMe disks as high-performance buffer

- Traditional (rotating) disks are not no longer suitable for some workflows
  - IO intensive data processing (for ex. “Data Sprucing” of LHCb)
  - IO to tapes (new tape drives are capable of doing 400MB/s)
- Allocated 80TB of NVMe storage for each LHC exp (on a shared storage system)
- Defined “placement policy” to write data going to tape first to NVMe pool
- Defined migration policy to NOT keep data on buffer after migration to tape
  - “Migrate” instead of “premigrate” in terms of HSM
  - This is not the case for ALICE, as they re-read data from the buffer to calculate checksums.
- Data recalled from tape are placed on HDD storage pool (as before)
- In case of LHCb, NVMe pool is shared between tape buffer and “hot” data
- Significant increase in migration rate

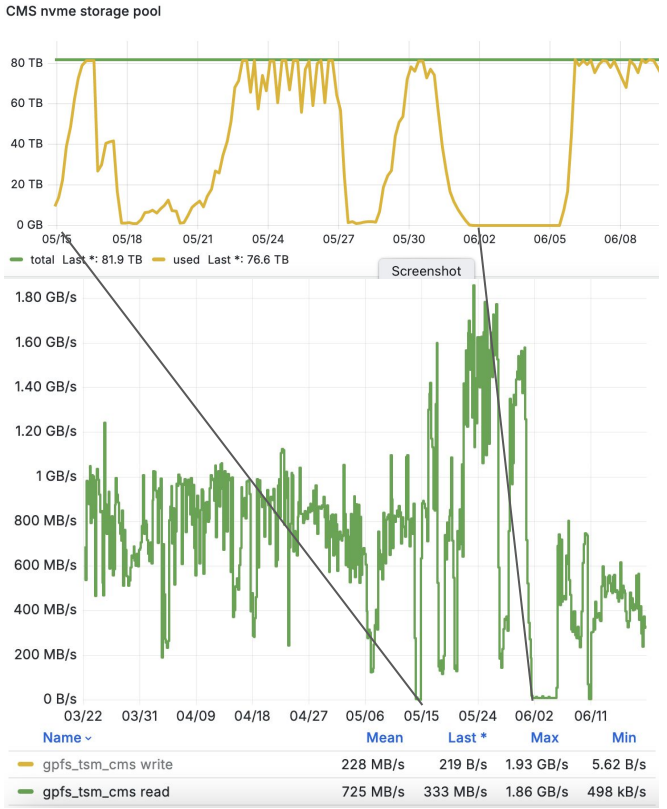

Stato tape (Last month)

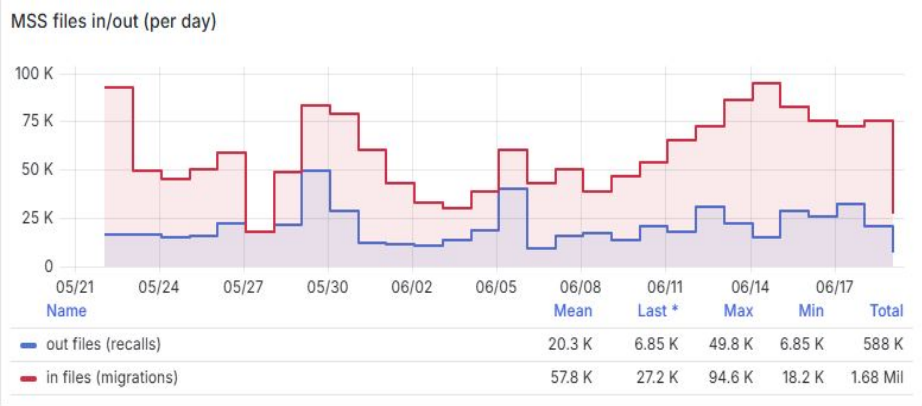

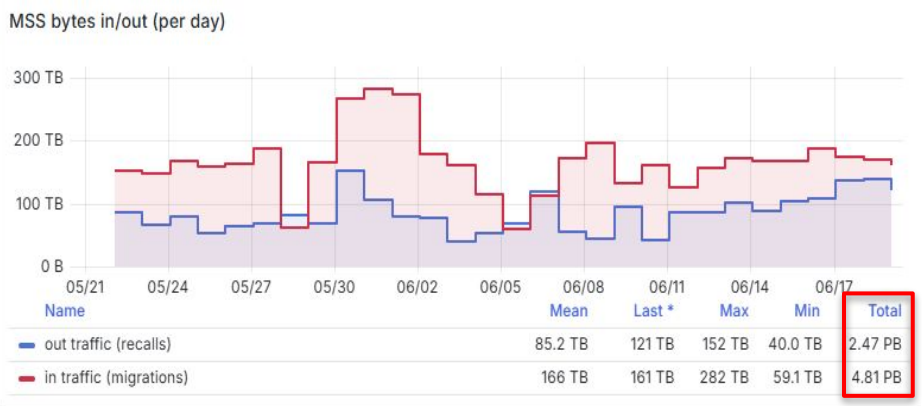

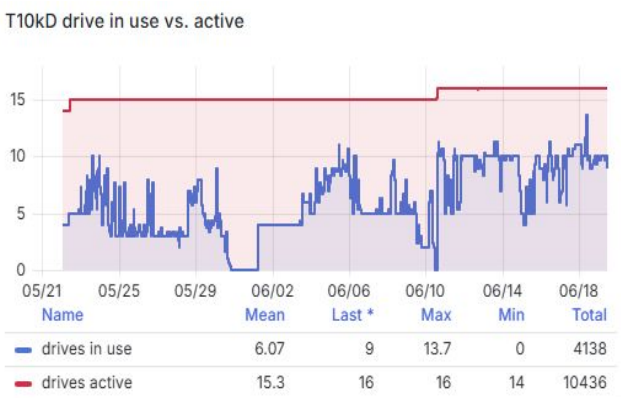

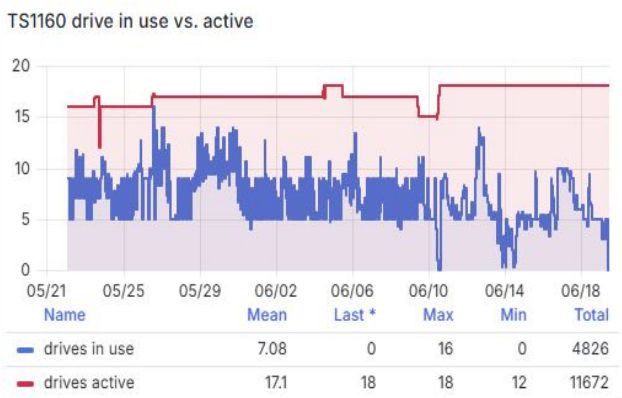

HSM activity by LHC experiments

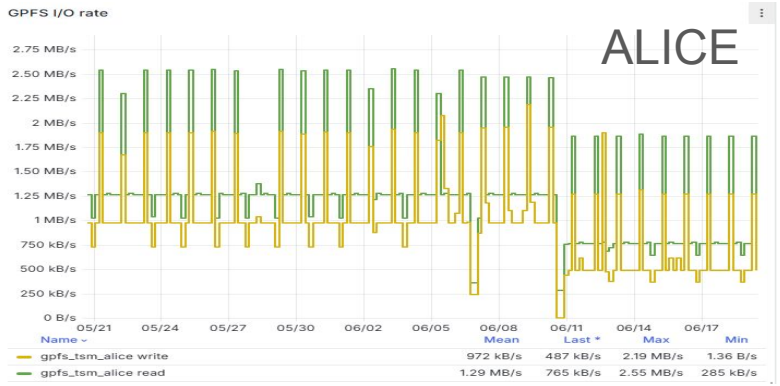

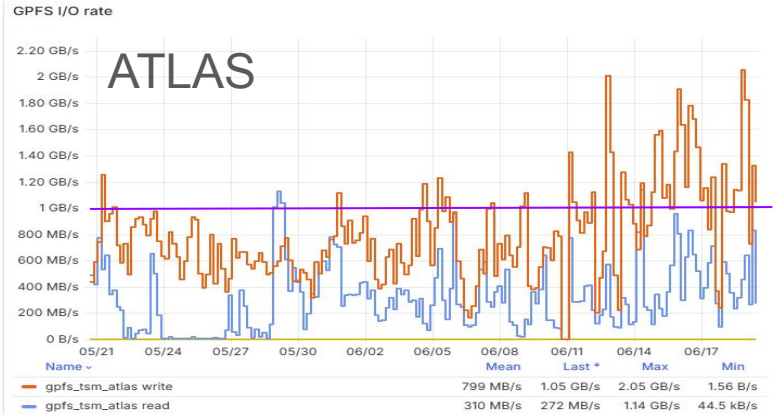

Stato tape

- TS4550-2 installazione acquisto 96 PB in corso
- SL8500 ripartenza delle operazioni di migrazione dati su IBM
- Stima della durata del repack: ulteriori 16 mesi con 5 processi in parallelo

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	9384	**47.8**	**32**
TS4500 (IBM)	19*TS1160	400	6198	20	5470	**97.6**	**11,4**
TS4500-2(IBM)	18*TS1170	400	7844	50	1363	**15.7**	**47.1**

Stato Repack Oracle –> IBM

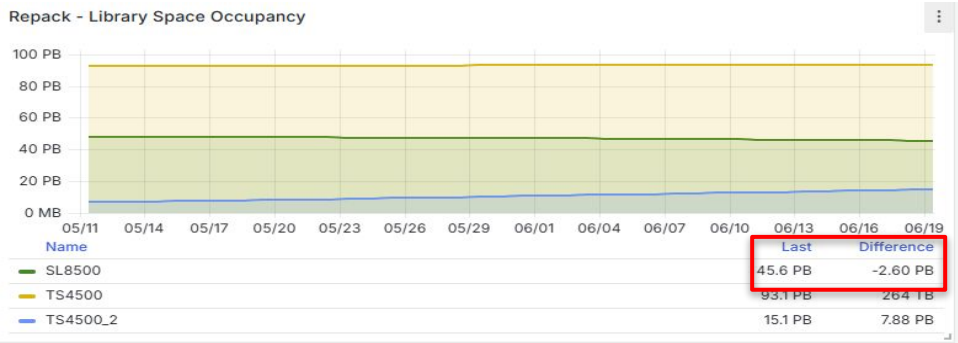

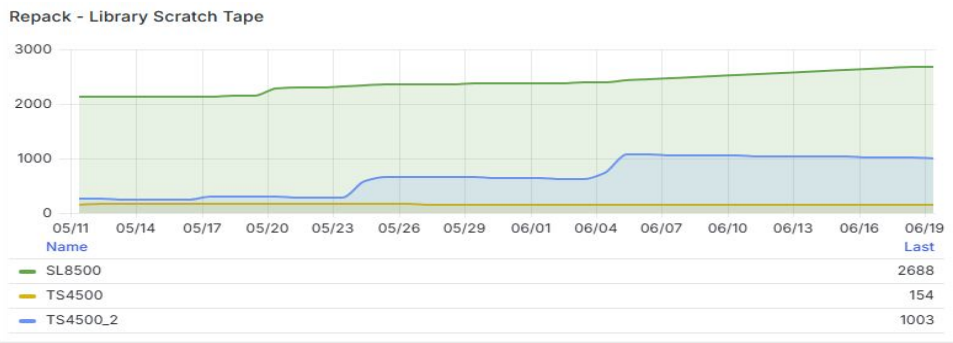

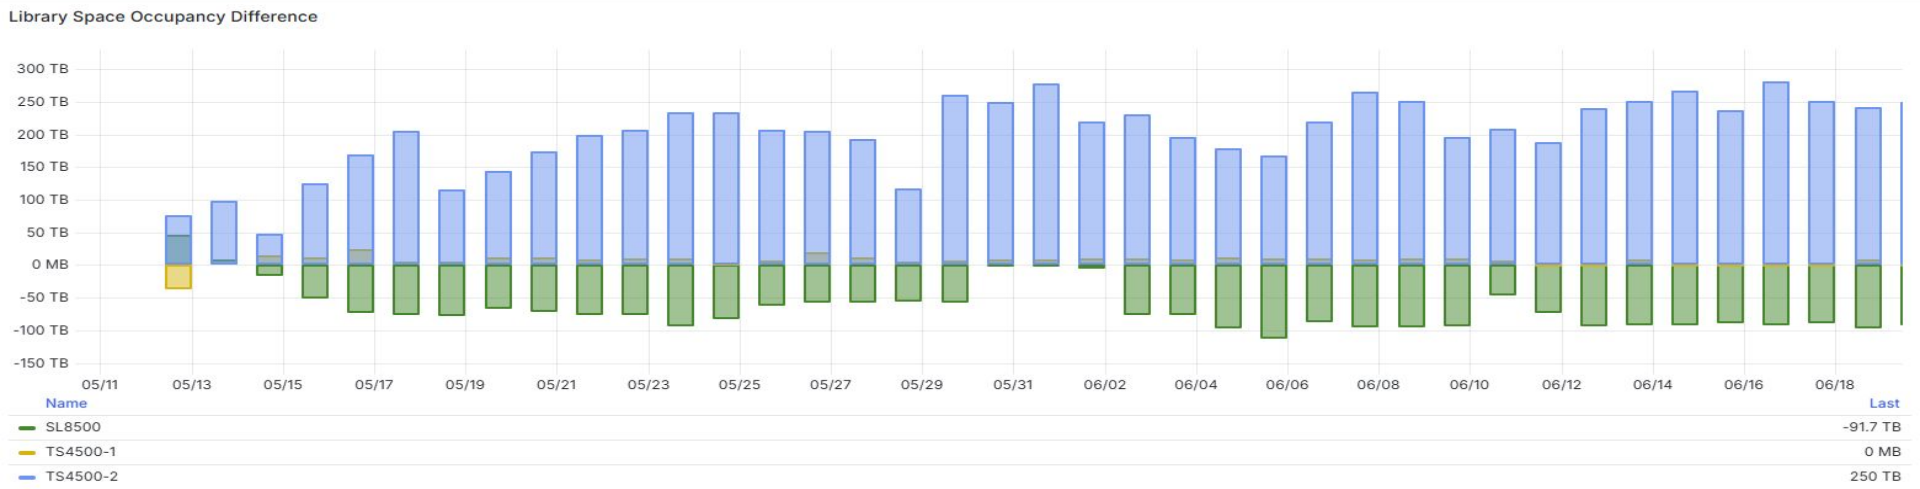

Current SW in PROD

- GPFS 5.1.9-9
- StoRM BackEnd 1.11.22 and StoRM FrontEnd 1.8.15 (latest, testing upgrade)
- StoRM WebDAV upgraded to 1.9.0 (latest, in beta)
  - Spring Boot 3, scitags, new monitoring metrics and log format, tape storage areas, nginx support
- StoRM Tape 0.9.0 (latest) and a new StoRM Tape Authz 1.1.0 (latest)
  - Telemetry via OpenTelemetry, new authorization policies (similar to the StoRM WebDAV fine-grained authorization), new deployment model
- 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
  - Two servers dedicated to xrootd for LHCb following discussion in GGUS #683513
  - One dedicated xrootd server for ATLAS (needed?)

Current SW in PROD (kernel upgrade)

- Since January, we observed several kernel crashes and “spontaneous reboots” of the new gateway servers installed at Tecnopolo
- Finally, we collected info and opened a case to Red Hat, which suspected a known bug in the ice driver with a MTU of 9000 that causes unexplained memory corruption and crashes (https://access.redhat.com/solutions/7093261)
- They recommended to upgrade to kernel-5.14.0-503.38.1.el9_5 (was 5.14.0-427.37.1.el9_4)
- During last week, we upgraded the kernel of all gateway servers in a transparent way

Tickets and more

- ALICE
  - 6 EL9 servers installed and configured for disk cluster (xs-606, xs-607, xs-706, xs-707, xsod44, xsod45)
    - Anything to be done to enable Scitags?
  - Finalized the configuration for the XrootD tape cluster (dm-12-14-01, dm-12-14-03)
    - Problems reported by Francesco on June 9th, due to a “spontaneous reboot”

Tickets and more

- ATLAS
  - GGUS #683563 (in progress): staging errors, 415 Unsupported Media Type
    - Currently debugged by StoRM developers
  - GGUS #683487 (in progress): file transfers timed out
    - Switching to storm webdav behind nginx. This should improve performances.
  - GGUS #683273 (solved): 18 corrupted files
    - All transferred via the same server within the same hour, preceding a “spontaneous reboot”
    - Found errors in GPFS log files (problems in the communication among storage servers)
    - We ran a checksum consistency analysis on all the 2000 files transferred in some 30 minutes-time ranges around the time of all these errors and found 99 corrupted files
  - GGUS #683043 (solved): drops in transfer efficiency, due to a bug with Spring Boot 3
  - GGUS #683132 (solved): missing fs storage dumps
  - GGUS #683174 (solved): staging failures due to tape in error state
  - GGUS #682830 and #682829 (solved): low staging efficiency, fixed changing nginx configuration in the StoRM tape REST API with a proper subrequest_output_buffer_size directive, to accomodate for the (eventually large) response of STATUS requests

Tickets and more

- CMS
  - Following discussion between Katy, Francesco, and Andrea at the WLCG WS, we dedicated one of the CMS StoRM WebDAV endpoints to test the deployment of StoRM WebDAV behind an external reverse proxy NGINX, delegating GET requests but also VOMS proxies parsing to NGINX
    - Two weeks of tests between May and June, “using” production activity to fix configuration and improve logging and monitoring
    - We will coordinate with Katy to test performances (and also run mini data challenge)
  - We are ready to install a proxy cache for Leonardo. However, it is possible that authz should be revised due to the discontinuation of myproxy server
  - GGUS #683408 (solved): SAM tests failing
    - Failures affected only one of the 5 endpoints in the alias, the one with nginx enabled
    - StoRM developers provided a fixed nginx-module-http-voms

Tickets and more

- CMS: spikes patterns starting in June 8th
  - Very high requests rates on single endpoints for very short periods of time.
  - Daniele identified a cleaning agent removing stuff from unmerged/ and failing with 412 status (DELETE a non-empty folder): hopefully this was fixed by CMS.

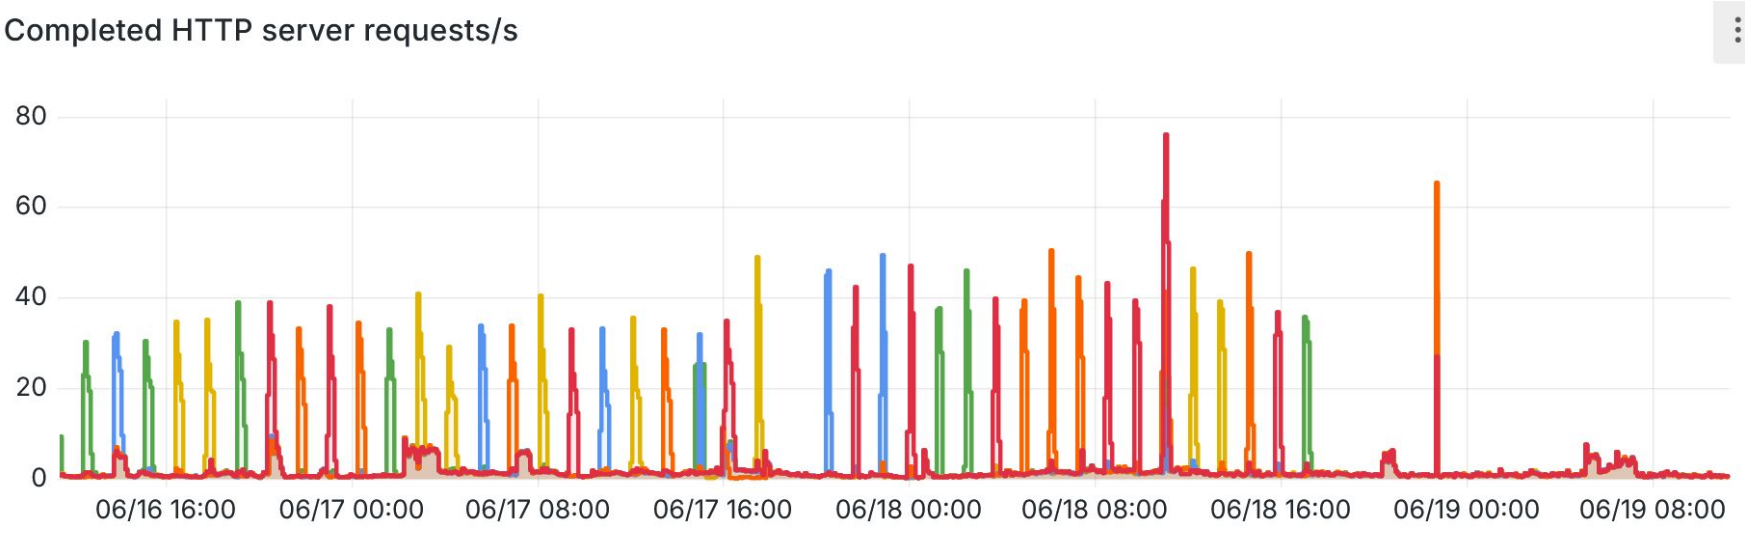

Tickets and more

- LHCb
  - GGUS #683529, #683523, #683489, #683290, #683444: failing transfers and/or upload from farm
    - One StoRM WebDAV server gets stuck and needs restart
    - Switching to storm webdav behind nginx (already in production in 1 out of 5 servers)
  - GGUS #683548, #683122 (solved): user problems contacting srm
    - Please, move away from srm :-)
  - GGUS #683513 (solved): about xrootd usage (to retrieve root-based urls instead of srm) and actual need for data transfer (by jobs running outside CNAF)
    - Dedicated two xrootd endpoints, on different servers wrt StoRM WebDAV
  - GGUS #683295 , #683221 (solved): files stuck in staging

Tickets and more

- All no-LHC experiments can move away from srm and use https instead, given the StoRM Tape REST API is ready: tape-archive.cr.cnaf.infn.it
  - A few pioneers: Belle, HyperK, LHCf, MIBLAT/QCDLAT, Newchim, PADME

- Belle
  - GGUS #682836 (solved): failing transfers upon LSC change to have a working StoRM Tape Rest API
    - C/C++ libraries can properly manage the shape of the file, but Java ones no
    - Waiting for KEK new LSC file (https://github.com/italiangrid/voms/issues/141)
- Darkside
  - Authorization policies revisited and corrected with Valerio

Tickets and more

- Dune
  - GGUS #683349 (solved): problems accessing storage (xrootd) from SL7 client
  - GGUS #682553 (solved): enabling WLCG Token access for DUNE storage
- HyperK
  - Tape storage area configured (pledge: 905 TB) using VOMS authz
- Virgo
  - Stashcache service installed and configured in xsod14.cr.cnaf.infn.it
    - However, the service does not start due to wrong registration in topology (commit)
  - Authz configured to allow Rucio rw tests from Torino group in a dedicated area

CNAF storage entered in top io500 ranking

- 16th (GPFS) and 22nd (CephFS) places among 24 10-node production HPC clusters
- We run io500 benchmarks before decomissioning Huawei OceanStor 18800v5 (was in production since 2018)
- https://io500.org/list/isc25/ten-production

	INFORMATION						IO500			
# ↑	BOF	INSTITUTION	SYSTEM	STORAGE VENDOR	FILE SYSTEM TYPE	CLIENT NODES	TOTAL CLIENT PROC.	SCORE ↑	BW (GIB/S)	MD (KIOP/S)	REPRO.
1	SC23	Argonne National Laboratory	Aurora	Intel	DAOS	10	2,080	2,885.57	734.50	11,336.27	
2	ISC23	LRZ	SuperMUC-NG-Phase2-EC-10	Lenovo	DAOS	10	1,120	1,008.81	218.38	4,660.23	
3	ISC25	Hudson River Trading	HRT	DDN	EXAScaler	10	1,600	348.08	136.05	890.51	
4	ISC24	Zuse Institute Berlin	Lise	Megware	DAOS	10	960	324.54	65.01	1,620.13	
5	SC24	Danish Centre for AI innovation AS	GEFION	DDN	EXAScaler	10	1,280	314.03	154.70	637.43	
6	SC24	SoftBank Corp	CHIE-2	DDN	EXAScaler	10	1,120	299.32	159.93	560.19	
7	SC24	University of Florida	HiPerGator AI	DDN	EXAScaler	10	640	243.61	124.89	475.20	
8	ISC25	Joint Center for Advanced High Performance Computing	Miyabi-G	DDN	Lustre	10	640	188.27	77.38	458.06	
9	ISC25	SAKURA Internet Inc and Prunus Solutions Inc	SAKURAONE	DDN	EXAScaler	10	1,280	181.91	133.03	248.74	
10	ISC24	NHN Cloud Corporation	NHN CLOUD GWANGJU AI	DDN	EXAScaler	10	640	176.57	62.58	498.22	
11	ISC24	University of Virginia	Afton	Dell	WEKA	10	850	105.94	33.28	337.29	
12	SC23	Japan Agency for Marine-Earth Science and Technology	Earth Simulator 4	DDN	EXAScaler	10	320	101.88	48.19	215.38	
13	SC24	Poznan Supercomputing and Networking Center	Proxima	xFiusion	Lustre	10	300	77.76	17.94	337.05	
14	SC23	Center for Research Informatics at University of Chicago	Randi	IBM	Spectrum Scale	10	160	60.88	31.05	119.36	
15	ISC24	Poznan Supercomputing and Networking Center	Altair	Huawei/ xFiusion	Lustre	10	200	54.03	11.39	256.36	
16	ISC25	Istituto Nazionale di Fisica Nucleare	Tier1-CNAF	Huawei	Spectrum Scale	10	100	37.26	11.30	122.93	