




# State of Storage

CdG 21 ottobre, 2022





# Business as usual

## Last month



## Last 6 months



# Disk storage in produzione

Installed: **50.07 PB**, Pledge 2022: **59.1 PB**, Used: **41.4 PB**

Sistema	Modello	Capacita' netta, TB	Esperimenti	Scadenza
ddn-10, ddn-11	DDN SFA12k	10752	ALICE, AMS	03/2021 → <del>06/2023</del>
os6k8	Huawei OS6800v3	3400	GR2, Virgo	06/2022 → 12/2023
md-1,md-2,md-3,md-4	Dell MD3860f	2308	DS, Virgo, Archive	11/2021 → 12/2022
md-5, md-6, md-7	Dell MD3820f	28	metadati, home, SW	12/2022 → 12/2023 ?
os18k1, os18k2	Huawei OS18000v5	7800	LHCb	2023
os18k3, os18k5, os18k5	Huawei OS18000v5	11700	CMS	2024
<b>ddn-12, ddn-13</b>	<b>DDN SFA 7990</b>	<b>5060+870</b>	<b>GR2,GR3</b>	<b>2025</b>
ddn-14, ddn-15	DDN SFA 2000NV	24	metadati	2025
os5k8-1,os5k8-2	Huawei OS5800v5	8999	<b>ATLAS</b>	2027
Cluster CEPH	12xSupermicro SS6029	3400	<b>ALICE, cloud, etc</b>	2027

# Prossimi acquisti

- Gara storage 2022 (14PB netti)
  - Gara pubblicata il 13/06/2022
  - Il vincitore (provvisorio) e' **Lenovo con ThinkSystem DE6600**  
(Devivato dal NetApp DE6600 con un enclosure con i dischi NVME)
  - Installazione e messa in PROD verso fine dell'anno
- AQ storage 2023-2024
  - La documentazione e' stata sottomessa alla GE



# Scheduled Storage Maintenance

Huawei OS5800-1 and OS5800-2 (**ATLAS**)

- firmware upgrade
- **At Risk** 27/10/2022 08:00-18:00

# Current SW in PROD

- GPFS 5.0.5-13 (to be updated soon to 5.1.2-7)
- StoRM BackEnd 1.11.21 (latest)
- StoRM FrontEnd 1.8.15 (latest)
- StoRM WebDAV 1.4.1 (latest)
- StoRM globus gridftp 1.2.4
- XrootD 4.11.2
  - updated to 4.12.4 in the 4 CMS servers
  - 5.3.1-1 on CMS redirectors (local and EU/IT/FR)
- Ceph 16.2.6 (Pacific)

# Tickets and problems - July/October

- ATLAS

- TPCs in push-mode to GOOGLE\_EU fail for all StoRM sites (GGUS [158487](#))
  - Issue for StoRM developers, on hold
- Transfer failures from SARA and TRIUMF; GARR involved (GGUS [158067](#))

- CMS

- Delete tape files to allow Rucio retry the transfer (GGUS [159175](#), also [158898](#), [158515](#), [158179](#), and [157983](#))
- Failing SAM tests in tape storage area, cleaned by garbage collector (GGUS [158953](#), [158751](#))
- Tape rule approval, agreed that A.Pascolini takes care of this (GGUS [158536](#), also [158192](#))


# Tickets and problems - July/October

- Belle
  - Storm-fe-archive timeout in bringonline (GGUS [158492](#))
    - Fixed issue with stuck recalls
    - The experiment Increased 12 hours-timeout to 3 days-timeout
- ALICE
  - Need to upgrade kernel (reboot) for ds-801, ALICE XrootD redirector (currently a single point of failure)
    - Switch to redirector in each server mode on hold
      - Already in production in the xrootd-ceph cluster
- Borexino
  - 30 TB exposed via StoRM WebDAV (xfer-archive) to be downloaded with tokens (dedicated group in iam-t1-computing) by 3 (authorized) Russian researchers



# Traffic stops when XrootD threads saturate - CMS

- Threads saturated during unusual amount of read requests from JINR
- Log files stuck
- All details in the [issue](#) opened to developers, who are suggesting to double max\_threads



# Stato tape

20 Aug 2022 - 19 Oct 2022

MSS bytes in/out (per day)



	min	max	avg	current	total
— out traffic (recalls)	821 GB	240 TB	48.6 TB	96.1 TB	2.92 PB
— in traffic (migrations)	153 GB	113 TB	38.7 TB	33.5 TB	2.32 PB

# Stato tape

- Liberi 30 PB (su cassette vuote, complessivamente sulle 2 librerie).  
Usati 93 PB.
  - Pledge 2022: 130.5 PB
  - Installato attuale: 130.5 PB
  - Alcune recenti cancellazioni (CMS, AMS)
  - ATLAS ha 600 TB oltre pledge
    - Sono previste cancellazioni a breve (non solo ATLAS)?

Library	Tape drives	Max data rate/drive, MB/s	Max slots	Max tape capacity, TB	Installed cartridges	Used capacity, PB
SL8500 (Oracle)	16*T10KD	250	10000	8.4	~10000	61
TS4500 (IBM)	19*TS1160	400	6198	20	2450	32

# Repack

- Libreria Oracle da dismettere il prima possibile per costi di manutenzione in aumento
- Da spostare su IBM: **61 PB**
- Rate repack agosto/settembre: 4 PB / mese con massimo 11 drive
- Possibilità di aumento rate del 15-20% con seconda CPU al TSM server
- Tempo stimato repack completo: minimo 14 mesi
- E' previsto che resti spazio inutilizzato per qualche esperimento a fine anno?

Repack - data moved per day (all tasks)

