

# State of Storage

CdG 17-Gen-2020

# Disk storage in manutenzione

Pledge 2019: 38.721 PB

Used: **35.44 PB**

Sistema	modello	Capacita', TB	esperimenti	scadenza
ddn-10, ddn-11	DDN SFA12k	10752	Atlas, Alice, AMS	2021
os6k8	Huawei OS6800v3	3400	GR2, etc	2022
md-1,md-2,md-3,md-4	Dell MD3860f	2308	DS, Virgo, Archive	05/2020
md-7	Dell MD3820f	20	Metadati, home, SW	05/2020
md-5, md-6	Dell MD3820f	8	metadati	2020-2022
os18k1, os18k2	Huawei OS18000v5	7800	LHCb, CMS	2023
os18k3, os18k5, os18k5	Huawei OS18000v5	11700	LHCb, ALICE	2024

# Ceph (Test) instance @CNAF

Test instance installation on the way

- HW: 4 server + 4 JBOD x 60 HDD 8TB = 1.9PB raw
  - Erasure coding 3 + 1 (1.44 PB usable)
  - Replicated pool with 2 copies (0.96 PB usable)
- Network (server side): 2x10Gbps x server => 10 GB/s theoretical limit (~10 MB/s/TB) but part of this BW will be used for replication traffic
- Purchased new hardware (4 server) in order to share a JBOD among two server instead of one
- After the installation the cluster can be considered pre-production

# New data management services

- Two dedicated hosts (xfer.cr.cnaf.infn.it) with latest version of StoRM WebDAV for the DOMA TPC testbed. Currently re-installing them CentOS7 and testing StoRM WebDAV installation via puppet.
  - The ESCAPE VO is also configured on the xfer.cr.cnaf.infn.it nodes (works with proxy generated by IAM).
- Xcache: configured on the same host for both Virgo (stash cache) and CMS (XCacheOnDemand).
  - Collecting monitoring metrics for our InfluxDB and sending some of them to DODAS server in Perugia every 5 minutes
- Apache: several experiments configured in the same host ds-814.cr.cnaf.infn.it: ASFIN, FAZIA, NEWSDM, NTOF. They share the same IAM instance, group-based authorization with tokens
  - Configuration with browsability to be introduced soon

# Migrazione a CentOS 7

- StoRM: puppet will be used to install and configure the services on CentOS 7
  - StoRM WebDAV: the migration can start (xfer-2.cr.cnaf.infn.it already migrated)
  - Storm GridFTP: logging each transfer info into separated file no longer supported, issue opened to globus (<https://github.com/gridcf/gct/issues/110>) but no plans to fix that
    - The migration can start, without one of the two log files :-(
  - StoRM frontend and backend: spring 2020
- XrootD
  - XCache currently works on CentOS 7
  - We could start. Any concern?
- IB and OFED
  - Resolved compatibility issue with Huawei and OFED 4.7-1
- HSM server
  - See next slide

# Tape server upgrade

- To use the new IBM tape library
  - It is needed a software upgrade of TSM and HSM servers
  - We scheduled also the upgrade to CentOS 7 for December 2019
    - Only software upgraded (GPFS, HSM e GEMSS) of all HSM, OS is still SL6 (except HSM ALICE that runs CentOS 7)
      - Because of problems with tape IBM and infiniband drivers, now solved
  - To upgrade TSM server we have to schedule a downtime (partial) of 20 hours for all the experiments. From midnight to 8 p.m. **January 23 (Thursday) or February 5 (Wednesday)?**
  - Later (after 2 weeks) we will update HSM servers (atlas, cms, lhcb, archive) to CentOS 7
    - An 8 hours (partial) downtime has to be scheduled for each HSM

# Some comments

- FS of Darkside is at 99% (almost full)
- CMS: disk buffer takes ~33% (2.6PB) of all available disk space
- ALICE: disk buffer in over pledge still needed? For how long?