CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

CC-IN2P3

Tape Storage at CC-IN2P3

Status and Plans

Fabio Hernandez
Philippe Gaillardon
IN2P3 Computing Centre
Lyon (France)
fabio@in2p3.fr

LHC Computing Review
CERN, November 14th 2005

Status (1/3)

- Hardware
  - 3 automated tape libraries
  - STK 9310
    - 35 drives STK 9840A/B (20 GB)
    - 35 drives STK 9940B (200 GB)
    - 6 silos: ~36000 slots, 70% full
  - STK 9710
    - 10 DLT drives
    - ~400 slots, 80% full
    - To be phased out by the end of this year
  - STK L20
    - 2 LTO-2 drives
    - 20 slots, 100% full
  - DLT and LTO used mainly for data import/export

Status (2/3)

- Software
  - HPSS
    - Serving data for ~35 experiments
    - Uses 60 (out of 70) tape drives in the main library
      - 9840 cartridges: used for small files (< 64 MB) typical of astrophysics experiments
    - Uses 54% of 9840 and 94% of 9940 cartridges
    - Data accessed through xrootd, RFIO and dCache
    - Currently manages 1.2 PB
      - Volume of data managed by HPSS doubles every year

Status (3/3)

- Software (cont.)
  - Tape staging system
    - Still used by a few experiments (~7)
    - Also used for copying data on cartridges for import/export
  - Tivoli Storage Management for backup
    - 4 TB/month
    - ~800 cartridges used by TSM (~300 in the robot)
    - User's and group's data on AFS, HPSS metadata, desktops,…

Plans (1/4)

- As HPSS-managed data volume doubles every year, the robot installed capacity will be fully used by 2007
  - Or 2008 if we replace all 9840 cartridges
- Constraint: not too much floor space left
  - Increase the density: looking for cartridges with higher storage capacity
  - IBM and STK both have proprietary hardware designed for random access, while keeping (more or less) the access time to data constant

Plans (2/4)

- The selected tape drive has direct impact on the cartridge library
  - IBM 3592 tape drives cannot be integrated in the latest STK library (STK L8500)
  - STK T10000 tape drives imply (either or both):
    - Modification of the current library, which is not longer manufactured (but maintenance guaranteed up to 2010)
    - Progressive replacement by a new one
- Is LTO-3 an option?

Plans (3/4)

- Throughput of new tape drives (~100 MB/sec) suppose modifications in the storage infrastructure
  - More powerful tape servers
    - Currently, HPSS tape movers are configured with 2 or 3 drives for getting ~60 MB/sec per server
  - Modifications in the networking infrastructure
    - Infiniband or 10 GigE?

Plans (4/4)

- During first half of 2006
  - Purchase at least 10 more STK 9940 tape drives
    - 16+ tape drives will be needed for reaching nominal throughput goals of LCG SC4 (200 MB/sec)
  - Study our options for the evolution of the tape storage infrastructure (tape drives, cartridge library, tape servers and storage network infrastructure)
    - Testing of STK T10000 tape drives planned for late January 2006
  - Goal: launch the procurement procedure of the selected hardware during 3Q2006
- We would be happy to share the results of our tests, and we are very interested in knowing what other sites with similar needs are planning/testing

Questions