Archive Storage Experiences and Outlook

Tim Bell
CERN

Data Preservation and Long Term Analysis
7th December 2009

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Agenda

- Scale of the problem
- Detecting and resolving issues
- Areas for investigation

Castor Physics Data Archive

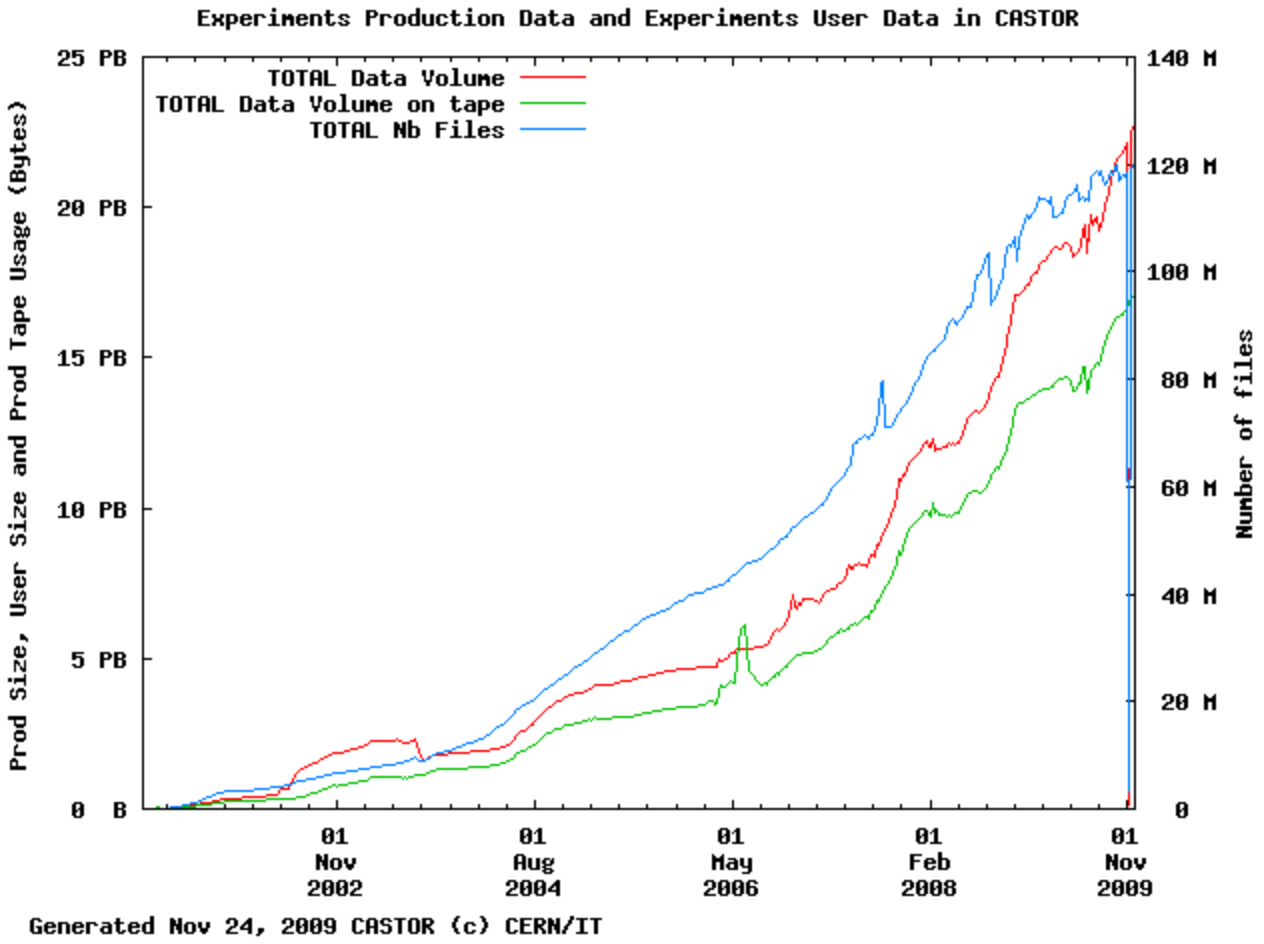

- Expected data rate of 15PB/year
- Keep data for at least 20 years

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Backup in TSM

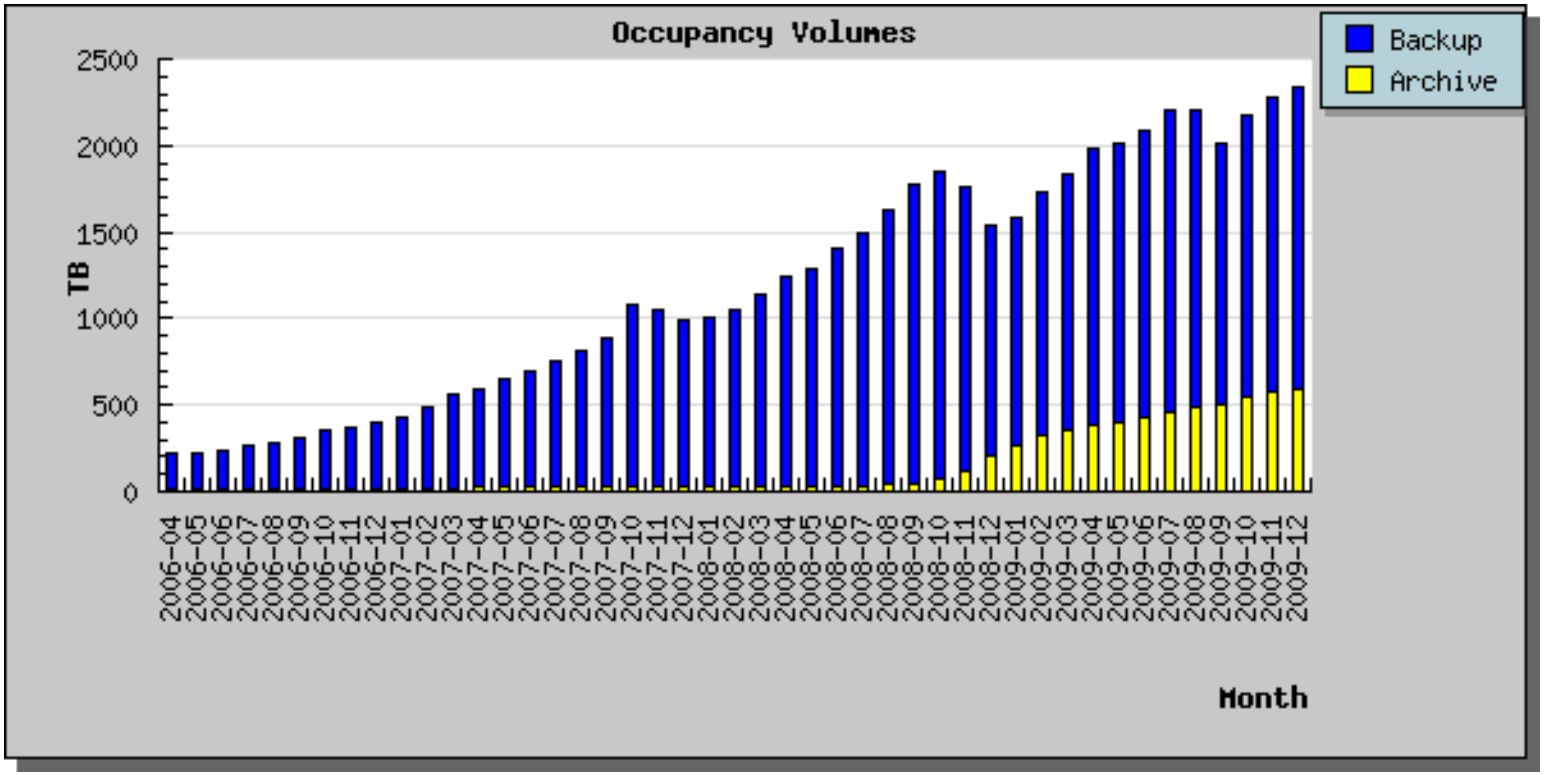

- Backup for Organisations vital data
  - Administration databases, Physics Metadata, User files
- Data volumes are equivalent to 2 LHC experiments
  - 20TB/day sent

FIO Technologies

CERN IT Department

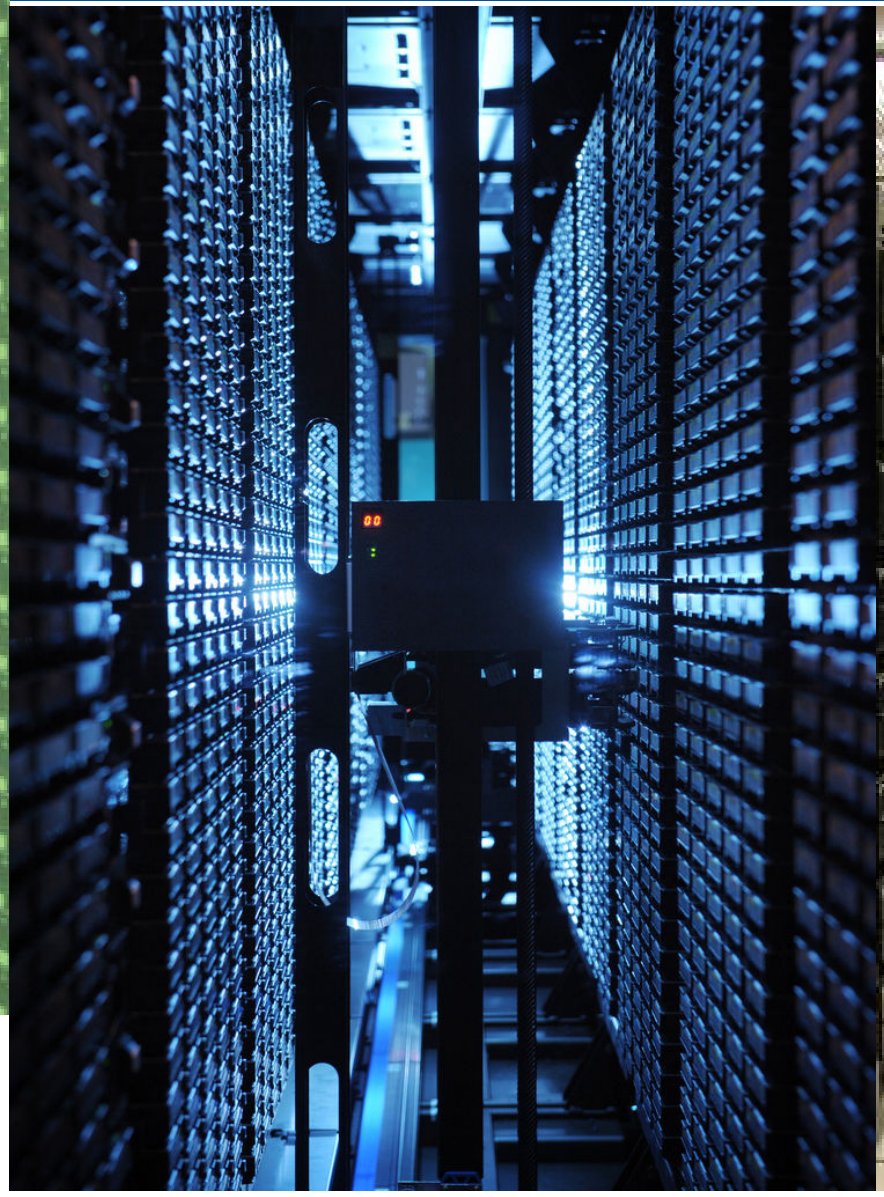

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Disk Data Preservation

- Need to ensure that the data is recorded correctly
  - Checksumming is now end-to-end from experiment pits to tape
  - Sequential I/O only
- RAID array verification at regular intervals to identify bad disks
  - Tuesdays to Thursdays only
- Fsprobe provides a background low I/O check, write and then read back
  - Should never happen.....

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Fsprobe results 2007 - 2009

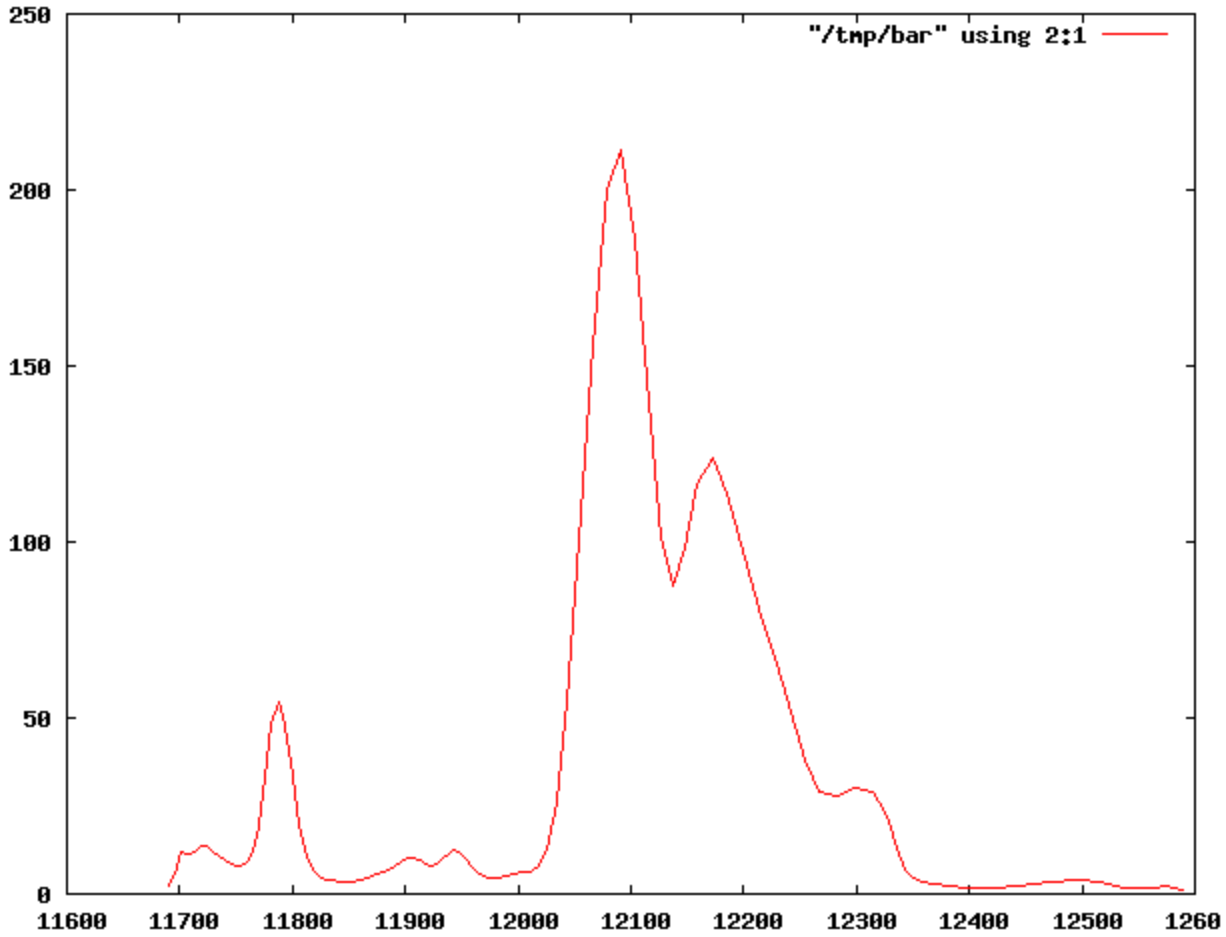

- Incidents/day during past 1000 days
- Problem not often seen for real data

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Fsprobe Root Causes

- Disk firmware
  - 5,000 drives manually upgraded in two campaigns
- Controller firmware
  - Bad disks not identified fast enough
  - RAID-5 cannot fix errors in these cases
- Memory errors
  - ECC memory is supposed to be able to correct/report these
- Operating system bugs
  - Linux and XFS bugs found

Tape Data Preservation

- Data on tape is still at risk
  - 60 tapes last year where some data was lost

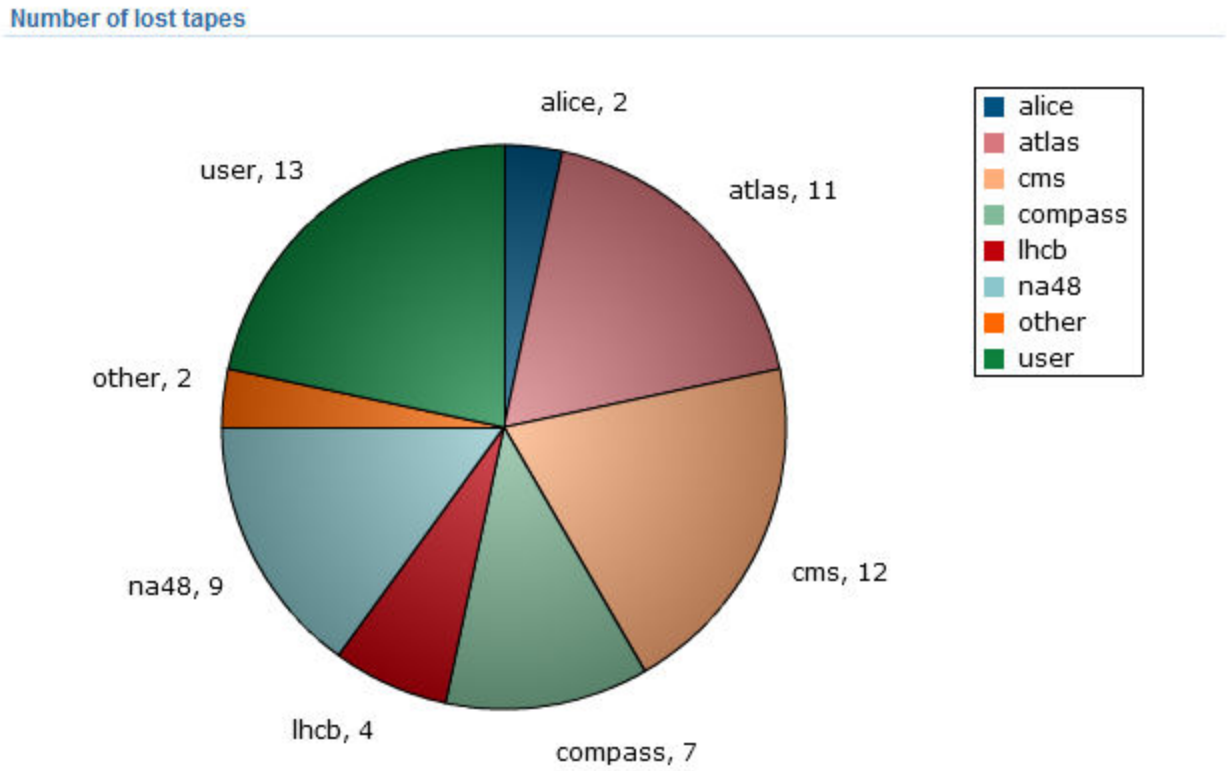

Tape Data Preservation

- However, over 90% of the data on those tapes could be recovered

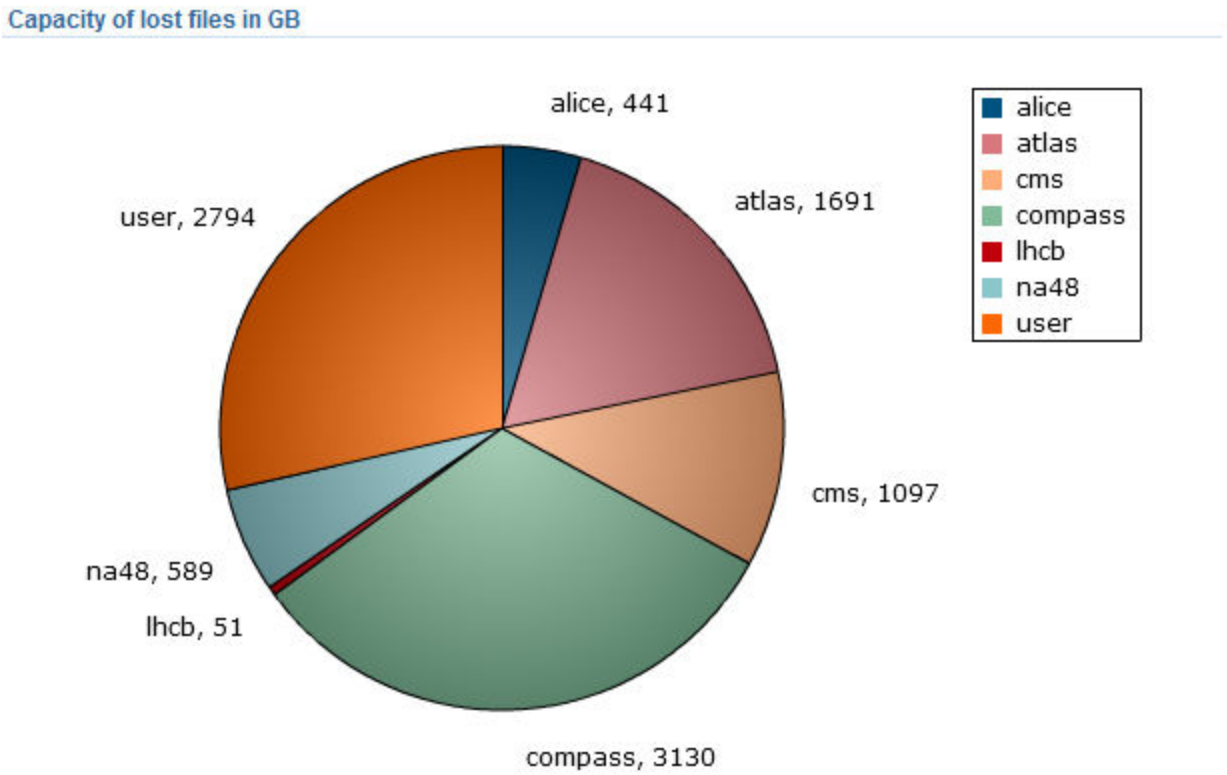

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Improving Tape Reliability

- The most high risk operation is mounting/unmounting the tape
  - Try to avoid using tape in random access
  - Bigger disk caches protect the tapes

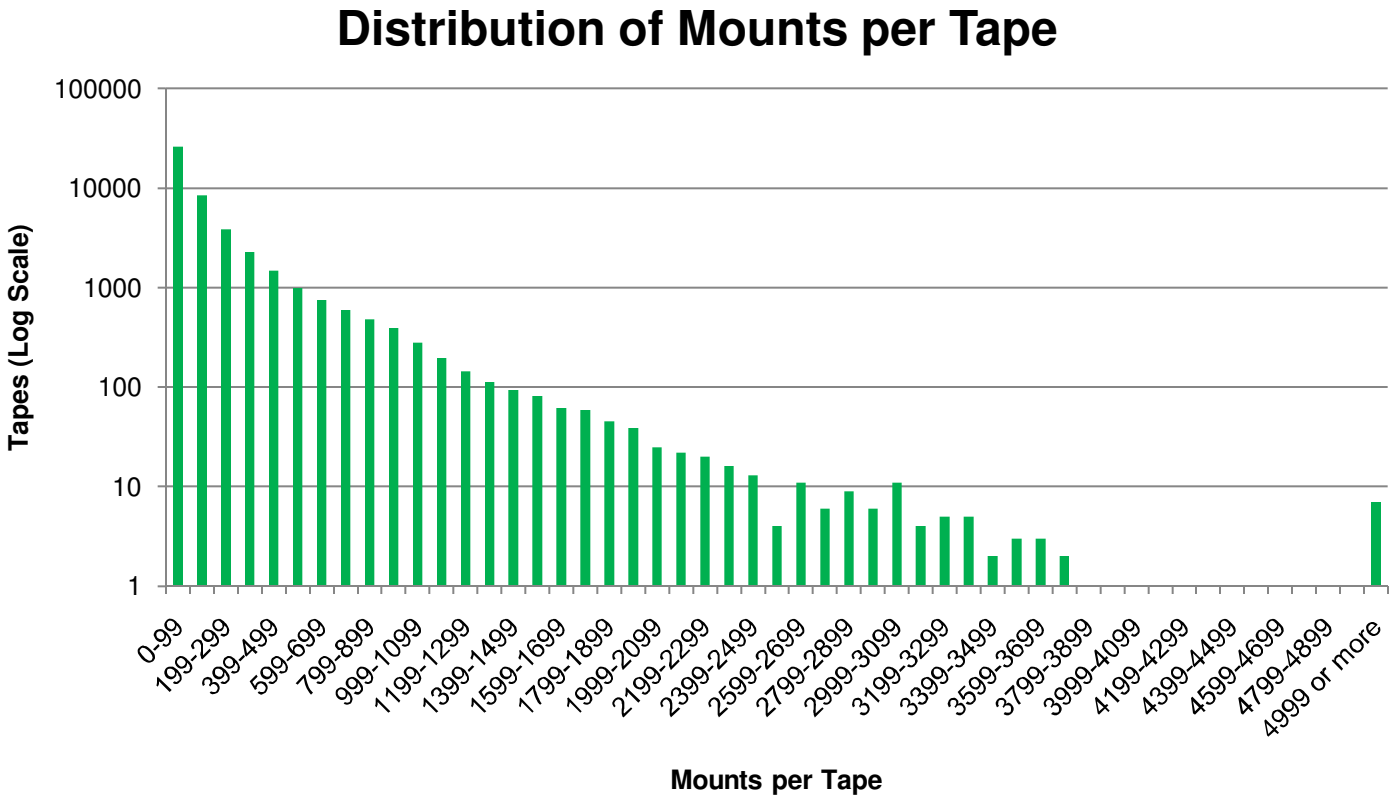

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Improving Tape Reliability

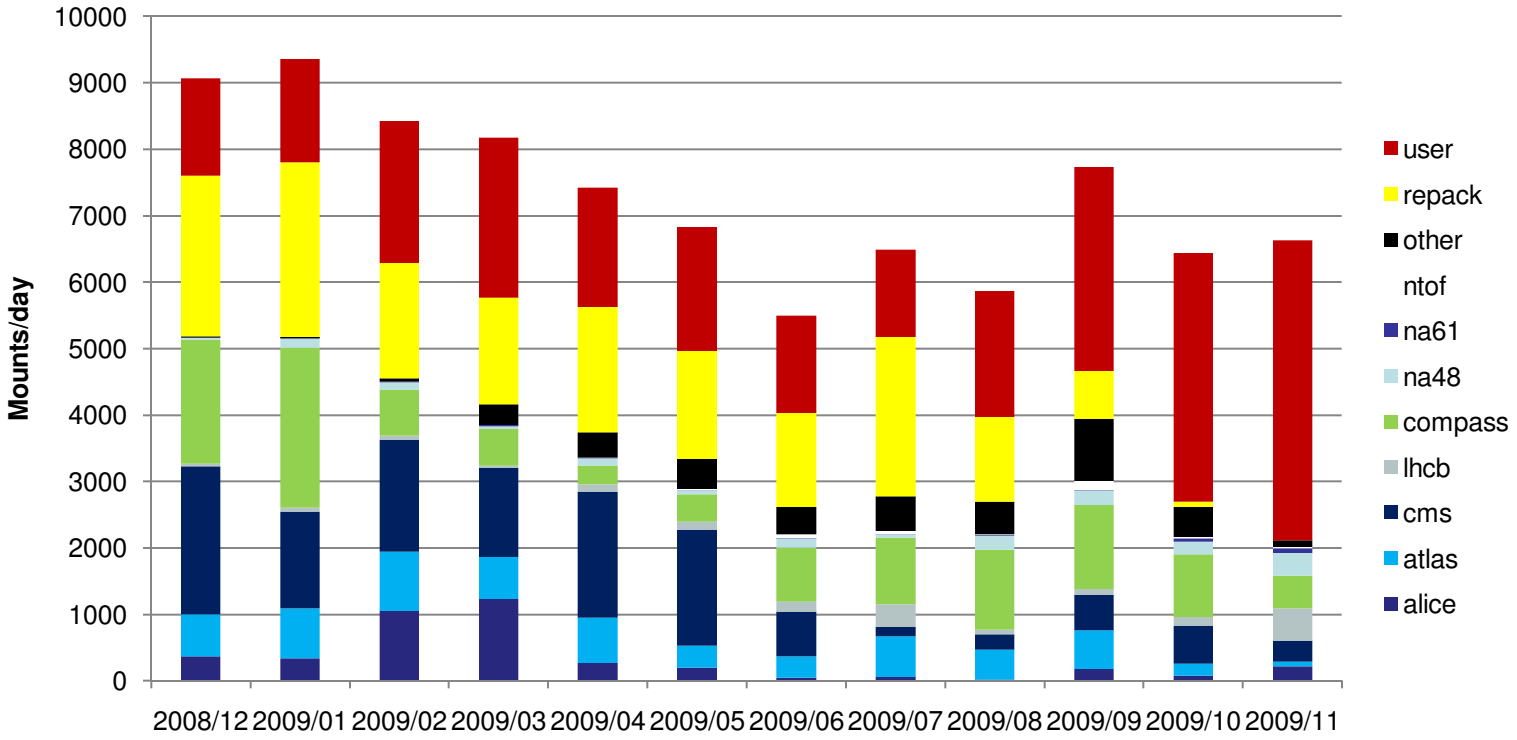

- 9,000 mounts per day during peak activity over 120 drives
- 60 tapes failed out of 2.6 million mounts reasonable reliability

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Impact of Capacity

- Tapes are getting larger
  - Currently 1TB doubling every 2 years
- All of LEP experiment data on 134 tapes
- Cast, AMS, HARP are around 20 tapes
  - Single tape break on load could lose 5% of experiment data
- Second copy of this data is being made to different tapes to reduce risk
  - Low cost as volumes are low

Tape Copying (Repack)

- Copying all files from one tape to another and updating the mass storage system catalogues
  - Validates checksums at the same time
- Used for tape data recovery such as media errors along with migrating to higher density tapes
  - Recently completed migration from 500GB to 1TB tapes
- Copying 45,000 tapes took around a year using 1.5 FTE and up to 40 tape drives
  - Next round in 2010 will take 18 months
  - New drives appear every two years...
- Continuing in the background would be desireable
  - But depends on resource availability

Future Directions?

- Find more efficient mechanism to handle data loss / unavailability
  - It is inevitable but quick action can save data
  - Copy still on disk even if tape copy is unreadable
  - Copy on the grid
  - Effort to maintain data catalogues could be reduced by automation
- Need for low impact background checks on archived data
  - Checksums still OK?
  - Media still OK?
  - Early detection before unavailable
  - Identify good candidates
    - Check when tape is full
    - Check on round-robin and on error thresholds
- Investigate large disk caches to reduce load on tape
  - Move to tape for backup only, not HSM
  - Investigate disk reliability/power as a potential archive media

A Sample Case

- Some data in our backup system has been migrated over the past 18 years
- A user came to IT in 2008 asking for the contents of a Wylbur data disk (the mainframe installed in 1976)
- The image of the disk had been copied to CERNVM and compressed (tersed) before sending to TSM
  - Terse was VM's equivalent of zip
- TSM had been migrated from CERNVM to AIX to Linux
- The data had been migrated from IBM 3480, IBM 3490, STK 9940 and IBM 3592
- Terse was not available on Linux so data could not be read
- Luckily, we found an open source 'terse' 16bit OS/2 program and ported it to 64bit Linux to allow us to recover the files

Conclusions

- Regular probes, scanning existing data and checksums allows losses to be identified and metrics established
- Maintaining the stream of bits is possible for an extended period of time with regular media refreshes
  - It does take man power to keep refreshing

References

- Fsprobe and silent data corruptions
  http://fuji.web.cern.ch/fuji/talk/2007/kelemen-2007-C5-Silent_Corruptions.pdf