Storage issues for end-user analysis

Bernd Panzer-Steindel, CERN/IT
08 July 2008

High level Data Flow

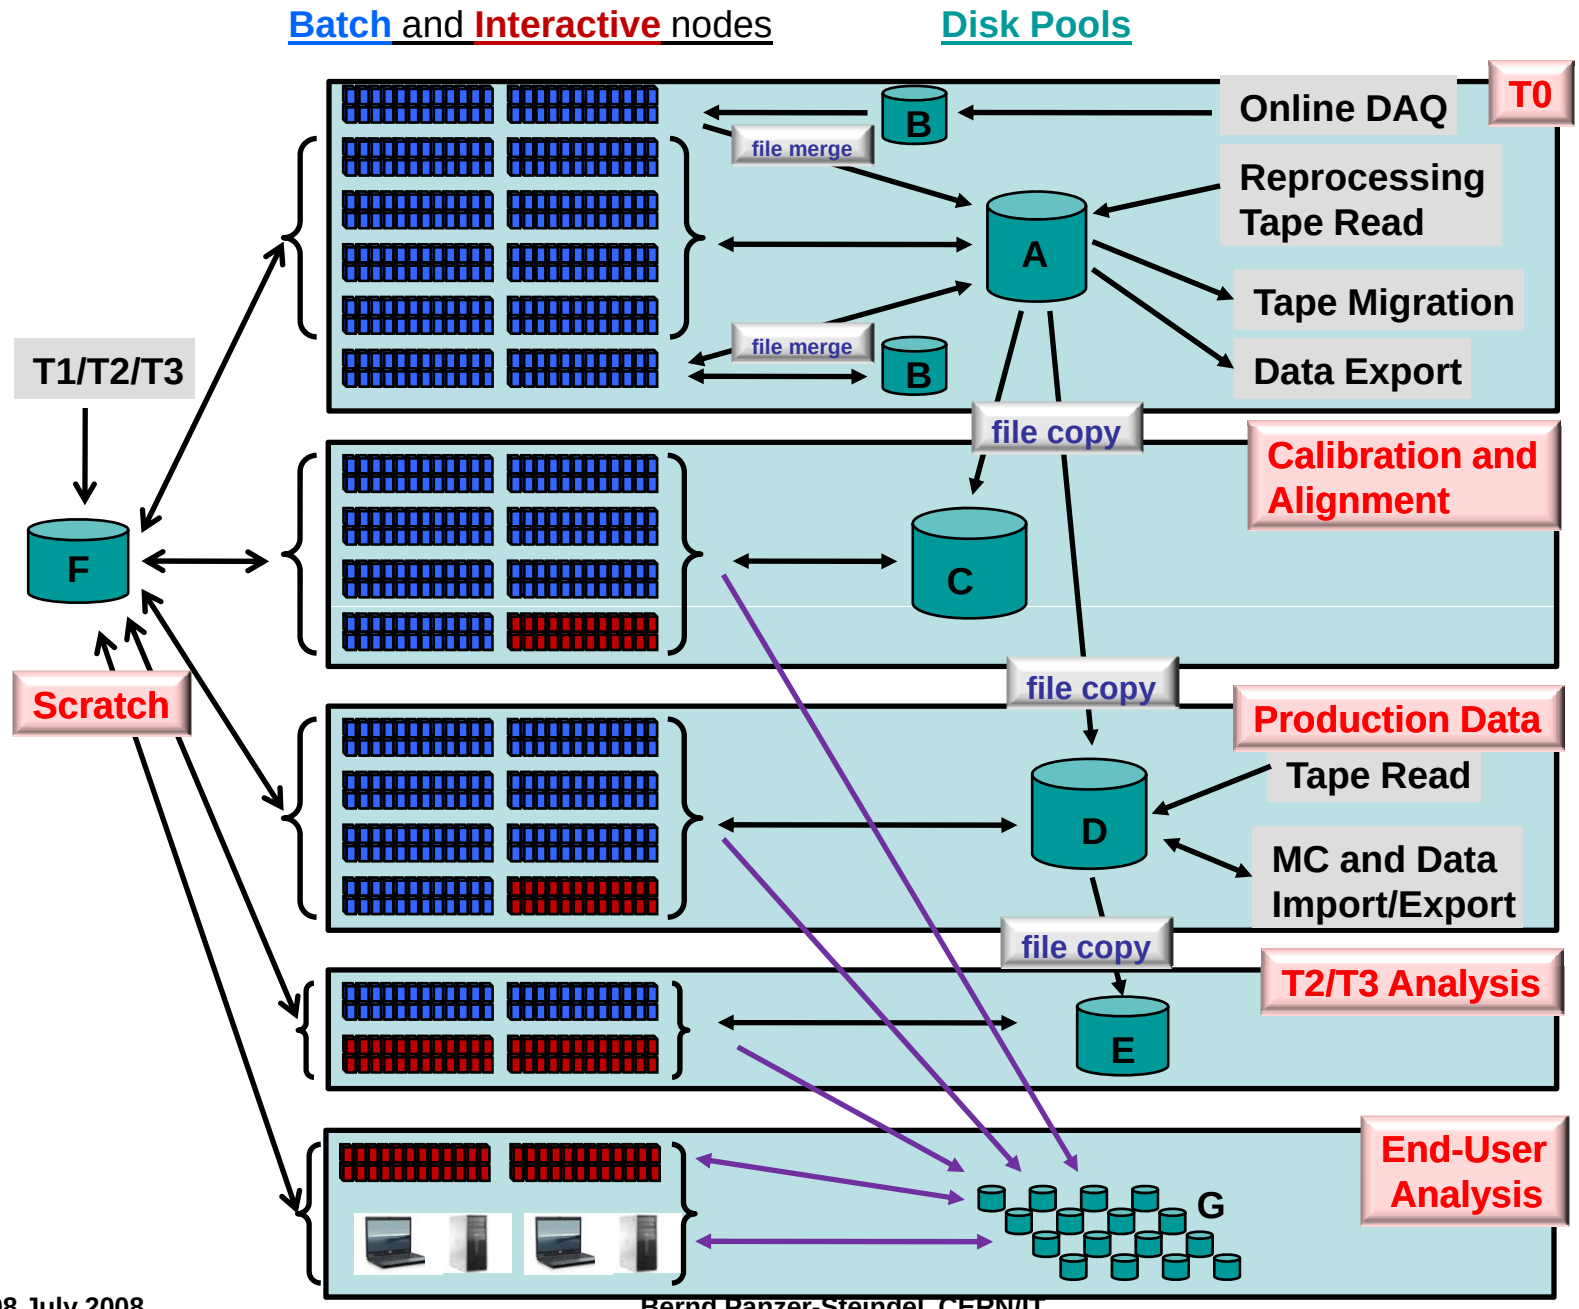

Tentative 'requirements' for end-user analysis storage

- Storage capacity of 1-2 TB per user (→ assume 1.5 TB)
  (Ntuple, data samples, logfiles, ….)
- Reliable storage, 'server-mirroring'
  99.9% availability → 4 times per year unavailable for 4 h each
- No tape access → too many small files
- Some backup possibility (archive ?)
  Backup → 5% changes per day (75 GB/d) + 4 month retention time
  = 9 TB backup space per user
- Quota system
- Easy accessibility from batch and interactive worker nodes
  and the notebook
- POSIX access type → distributed file system
- World-wide access
- High file access read/write performance
- User identity and security
- ……

Cost and Technology Scenarios

Estimated 500 users at CERN
→750 TB of analysis disk storage
+ backup and archive?

Distributed File System implementations

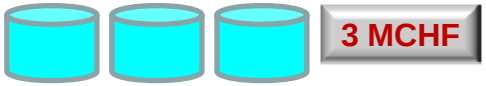

3 MCHF

Isilon, BlueArc, Exanet, DataDirect

4.5 MCHF

AFS

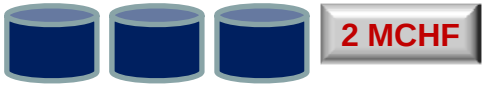

2 MCHF

NFS4, Lustre, xrootd,

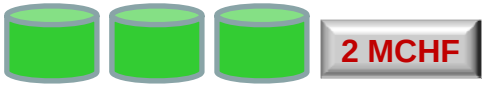

2 MCHF

HSM , disk-only

1.5 TB USB disk

0.4 MCHF

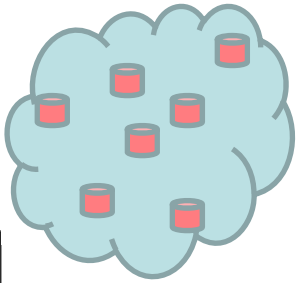

7 MCHF

Amazon Cloud Storage

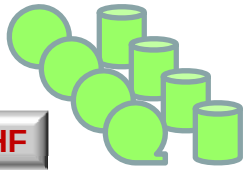

1 MCHF

TSM Backup

Hardware investments over 2 years 'guestimates'

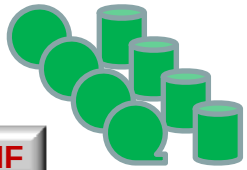

2.5 MCHF

HSM Backup, Archive

+ software operation, support, functionality,..

Questions

- Where are these ‘extra’ resources coming from ?

- Is there only one unique storage per user world-wide ?
  → What about users working on different sites ?
  → Do they have multiple end-user storage instances ?
  →How is data transferred between instances ?

- The difference between the ‘home-directory’ storage and end-user analysis space is small.
  → Analysis tools/programs and the data itself must be accessed at the same time.

- Who decides which user gets how much space where ?
  →Experiment specific policies

- What is the data flow model ?
  → Notebook disk + site local file system + global file system
  → Notebook disk + site local scratch + cloud storage
  → Global file system only
  → ….. More combinations…….

- Notebook issues
  → OS support, virtual analysis infrastructure, network connectivity = data ‘gas station’

- ……many more questions……

Is there some common interest to solve this problem ?

Need/interest for the creation of a working group to investigate in more detail?

Experiments, Sites ?