AFS Backup at CERN

Jakub T. Moscicki
("Kuba")

05/10/2011 EAKC, DESY, HAMBURG

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

Numbers

- Numbers (~55 file servers):
  - Total backup size on tape (TSM): ~ 10 x 55TB
  - Daily transfer to tape: 1-2TB
  - Retention policy: 6 months (was 12 months)
  - Granularity: daily (yesterday on disk as “BK”)
- Servers are backed-up in parallel
  - partitions are backup-up in sequence
  - sometimes 24h is not enough, will be more of a problem in the future
- Backup schedule:
  - interleaving full and partial backups between servers (to level the TSM load)
- All volumes backed up
  - Currently being introduced

Implementation

- Using “native” AFS backup system
  - backup volsets: one volset per partition
  - dump hierarchy:
    - simply defined as a chain: full-i1-...-i52 (no differentials)
    - differential backups based on additional state information stored locally via perl Catalog module
  - butc linked against XBSA/TSM library
  - butc wrapped by expect script and fired on demand
  - TSM server (one TSM account at the moment, archive mode)
  - afs_admin recover in “synchronous” mode
    - underlying arc call keeping connection alive until user leaves a temporary shell which is position where recovered volumes are temporarily mounted

Issues/questions

- Any alternatives to “native” AFS backup system?
- 24h time window for large volsets:
  - parallelization of dumps and TSM transfer
  - optimization of diffs?
- Butcx is complex and error-prone: butc in non-interactive mode?
- Synchronous restore: can we do in a simpler way?