

Backup system migration

From a local infrastructure
to IN2P3 centralized system


Outline

- LAPP – IN2P3
- The local backup system
- The centralized infrastructure
- Migration steps
- Pros and cons
- Plans
- Conclusion

LAPP

Laboratoire d'Annecy le vieux de Physique des particules

- Located in Annecy-le-Vieux (France)
- Part of IN2P3 (*) / CNRS
- Around 150 people on premises



(*) IN2P3 = 20 laboratories in France
CC-IN2P3 in Lyon is the main IN2P3
Computing Center

LAPP

- Involved in LHC (Atlas and LHCb) experiments and others (Opera, Virgo, Hess ...)
- IT : 20 people (developeppers + 8 in support team)
- WLCG T2 for Atlas and LHCb



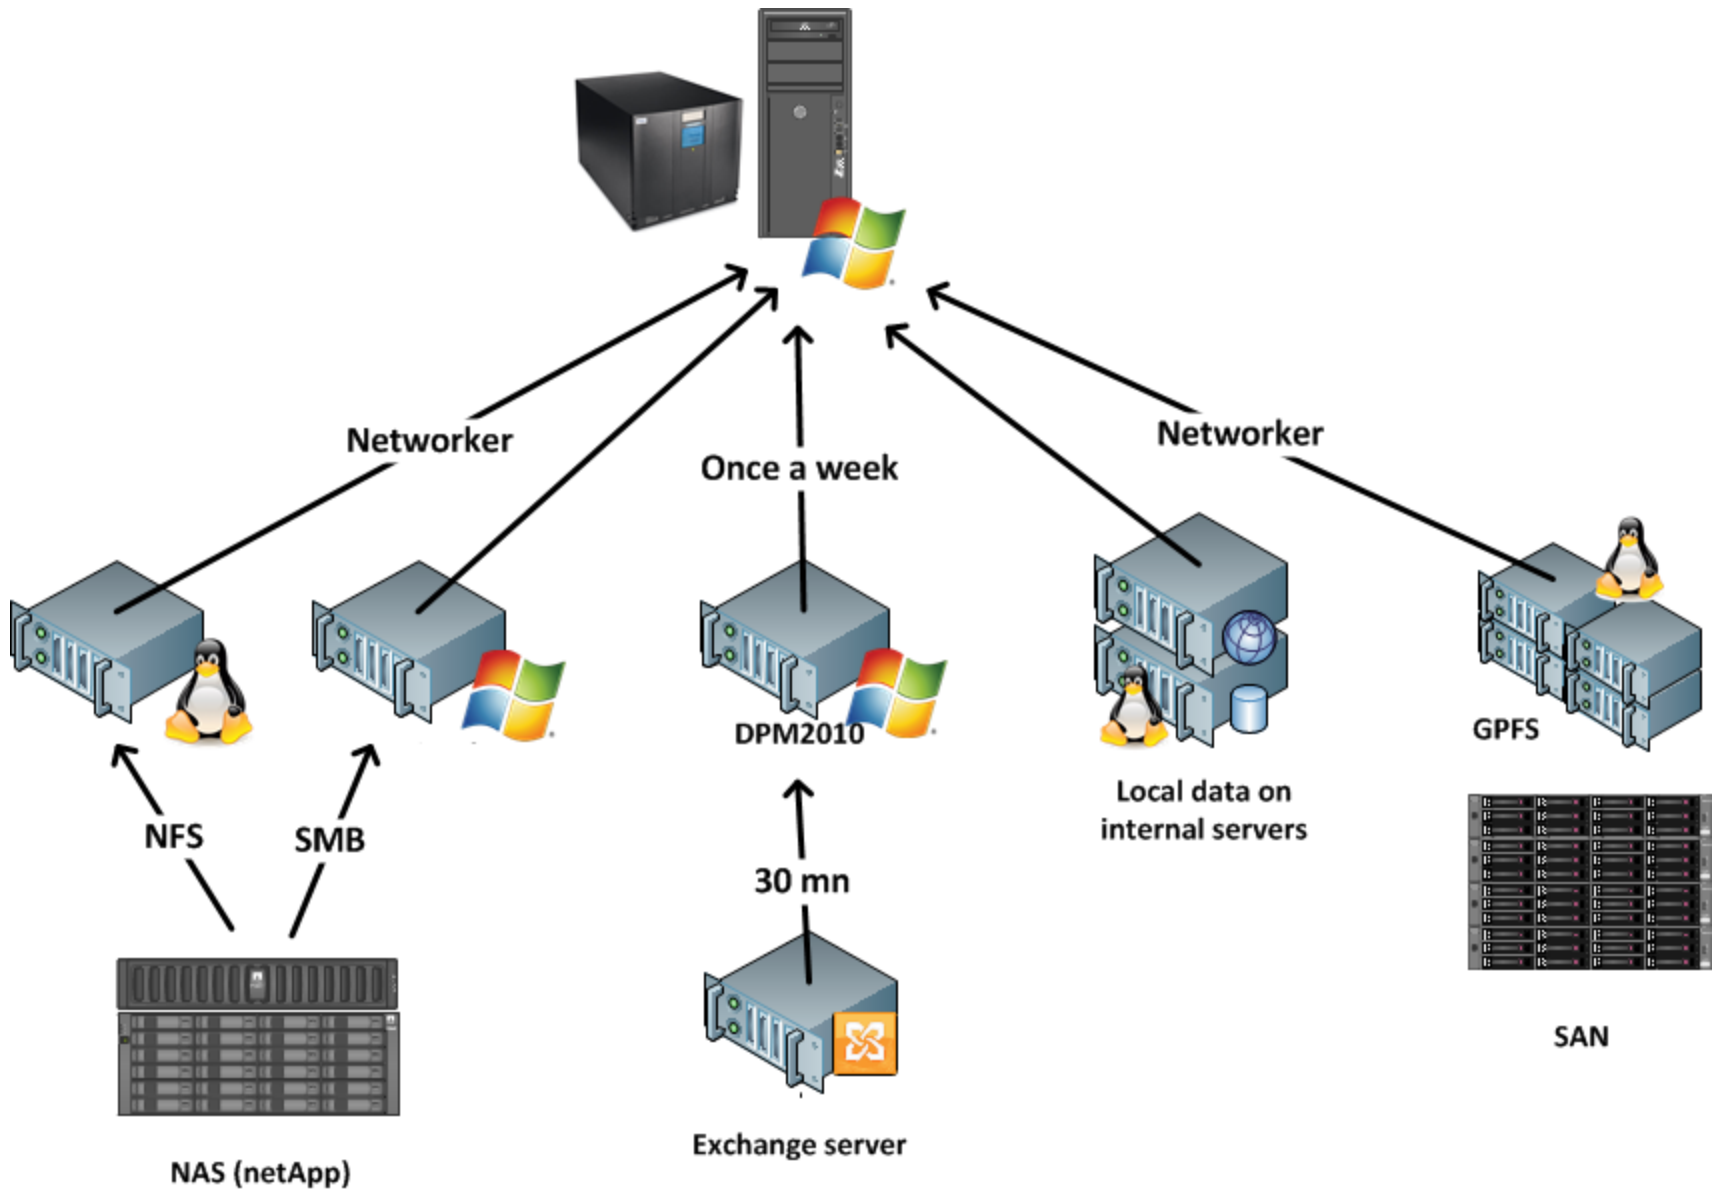
Hardware

- Internal servers (intensive use of vmware Esxi) :
 - Web and DB servers
 - Network services
 - Exchange ...
- Netapp filer : home directories for users and shared files for groups (16 TO) , Windows and/or Linux access
- Grid node :
 - 1298 jobs slots for local batch and grid
 - Servers for grid services
 - GPFS storage on HP SAN : 90 TO
 - Grid storage (DPM) : 650 TO



The local backup system

- Dedicated hardware :
 - NetWorker (Legato) software (used at LAPP for more than 10 years)
 - A Windows 2008 server
 - Overland Neo library : 60 slots LTO2 (2 tape devices) , end of support = 01/2013
- Networker client on :
 - « Bridge » servers for NAS (Netapp) backup (Windows shares via SMB and Linux file systems via NFS)
 - Linux servers (GPFS HeadNodes) for SAN areas
 - Windows and Linux servers for local disk backup (data, logs, conf...)
- File retention :
 - daily incremental = 6 months
 - montly full = 2 years
 - 3 years archives on demand
- PLUS : dedicated Microsoft DPM2010 server for Exchange server backup (live backup on disk every 30 minutes + weekly .edb files backup via Networker for extra security)



Observations

- High cost of our solution (tapes, software maintenance...)
- End of life reached : no IN2P3 budget for library renewal
- Fewer requests for restoration (use of snapshots on the NAS)
- Storage hardware more reliable
→ **EVOLUTION to centralized solution powered by CC-IN2P3 in Lyon**

The new infrastructure

- At CC-IN2P3:

- IBM TIVOLI Storage Manager software (used at CC since 1994)
- In the datacenter :
 - 4 SUN SL8500 libraries (shared with HPSS system)
 - 4 IBM servers (10Gbps Eth, 16 Go Ram) for backup: Aix 6.1, TSM 6.2
 - one "library manager"
 - one server for monitoring
- From the support team :
 - Tools (wrapper)
 - Help and documentation
- Process :
 - Backup is initiated by client (file list defined on TSM server)
 - Data written on disk
 - Data migration to primary tape
 - Copy on secondary tape
 - Transport secondary tapes to a secure room (manually)



The new infrastructure

- At LAPP:
 - For NAS and SAN files :
 - Install TSM client on « bridge » servers (linux)
 - NFS mounts for NAS areas
 - GPFS mounts for SAN areas
 - For local data on linux servers :
 - Duplicity software
 - For Exchange server:
 - Microsoft DPM2010 live backup on disk
 - Add a second server with DAG (databases replication)

Steps (1) : NAS

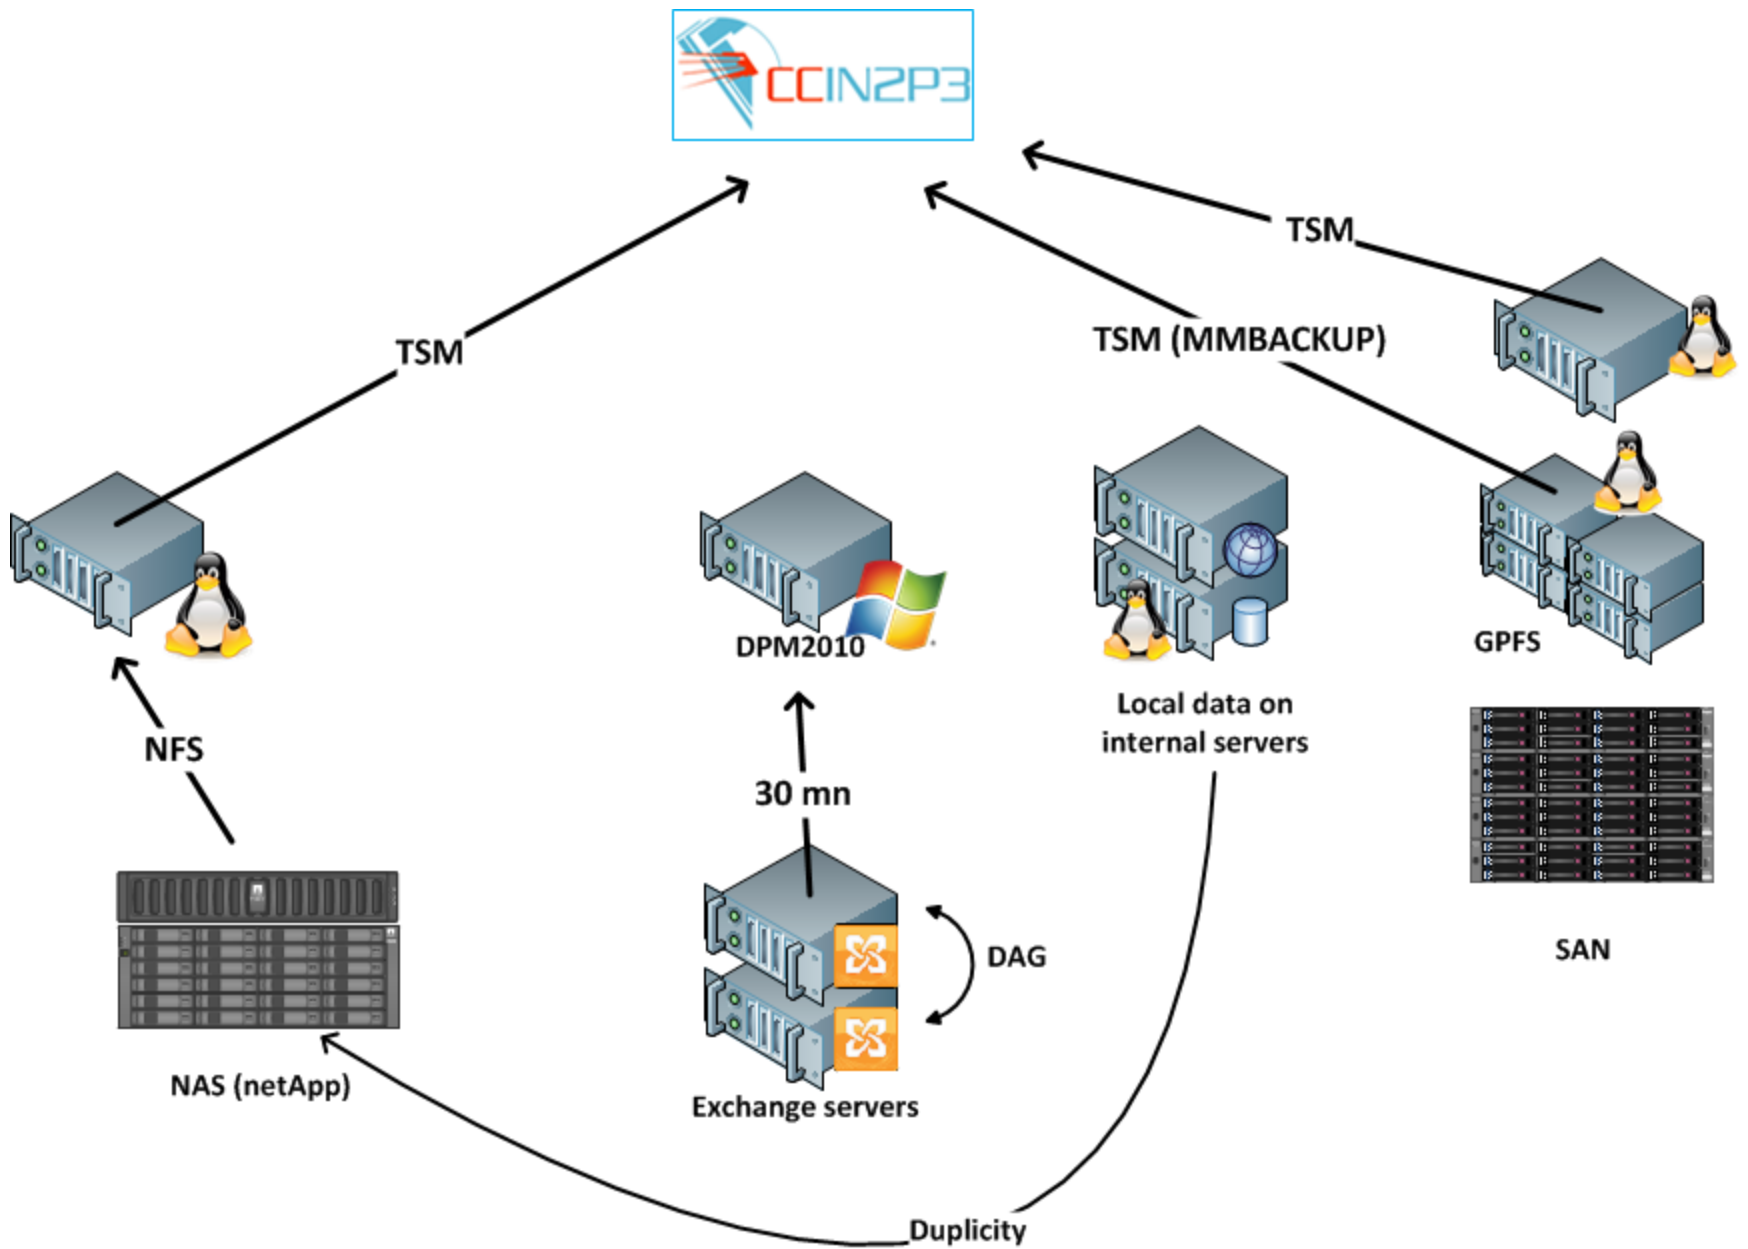
- January 2012 : TSM client installation and configuration, tests on one server
- Full backup of the NAS areas (add 2 or 3 areas per week, average 5 hours for 200 GO) : total 3 TO
- Tuning : increase number of cores on bridge server (virtual machine)
- April : daily incremental NAS backup OK : 10 000 000 files scanned , < 100 GO, 5 hours
- Old system still used in parallel

Steps (2) : SAN

- Add another bridge server for SAN areas via GPFS :
2,5 TO full, 750 000 files, < 50 GO incr per day, 1 hour
- September : add the last SAN area
18 000 000 files, full 700GO
- Incremental per day : < 100 files , <1 GO
→ more than 24 hours !!!
- Tuning : integration with GPFS :
 - Install TSM client on one of the GPFS headNodes (direct FC acces to the SAN)
 - Adapt wrapper in order to use GPFS « mmbackup » command
→ only 35 minutes !!!

Steps (3)

- Find a solution for local areas on linux servers
 - List of files to backup centrally defined in TSM server → too static, more flexibility needed
 - TSM client to be installed on each server
 - Many different OS versions
 - Limited number of clients (Misunderstanding)
- duplicity software



Advantages

- Externalisation de facto
- Low cost : only 2 virtual servers versus very expensive software (support) + dedicated server + library + tapes
- No need to manage tapes
- No server to manage
- Free choice of backup and restore window

Disadvantages

- Total time for a full restoration could be very high (node colocation on tapes)
 - Solution : perform a new full backup every 6 months? → impact on volumetry
- Less flexibility : parameters and policies statically defined :
 - List of files to backup centrally defined in TSM server
 - Retention policy : one year retention for deleted files, 1 to 5 versions for active files, no long life archives
 - Solution : increase number of snapshots on NAS server
- NFS access for NAS → Windows ACL loss
- New user interface for restaurations → document the internal procedures and train team members to support this new tool
- Need another sotware for linux local areas → complicates the recovery procedure
- No solution for Windows servers
- Impact on network load
 - Solution : allready planned bandwidth increase (from 1Gb/s to 10) within a few months

Plans at LAPP

- Use mmbackup for all gpfs file systems
- Test Tivoli Client for Windows
- Get rid of the old system (one redondant powersupply out of order on 15th of january ...) : one year with both systems reached

Plans on the CC side

- Replace LTO4 with LTO6 (less tapes)
- Migrate to TSM 6.3.3
- Usage statistics for labs (web access)
- Thinking of using SSL for transfers : impact on CPU ?
- 2014 : Add a (small) LTO6 library in the secured room for second copy
- Decrease file retention (choice of policies based on file “class”) → impact on labs
- Fixed backup windows
- Archive solution

Conclusion

- From CC point of view:
 - Free service for about 15 labs, for 8 years
 - Support on a best effort basis (no SLA)
 - In 2012 : 15 days off (availability = 98%)
 - 800 TO for labs = CC internal use
- From labs :
 - Some constraints
 - Very good disaster recovery solution
 - Mutualization is a good thing