

SRM – dCache - HPSS at IN2P3-CC

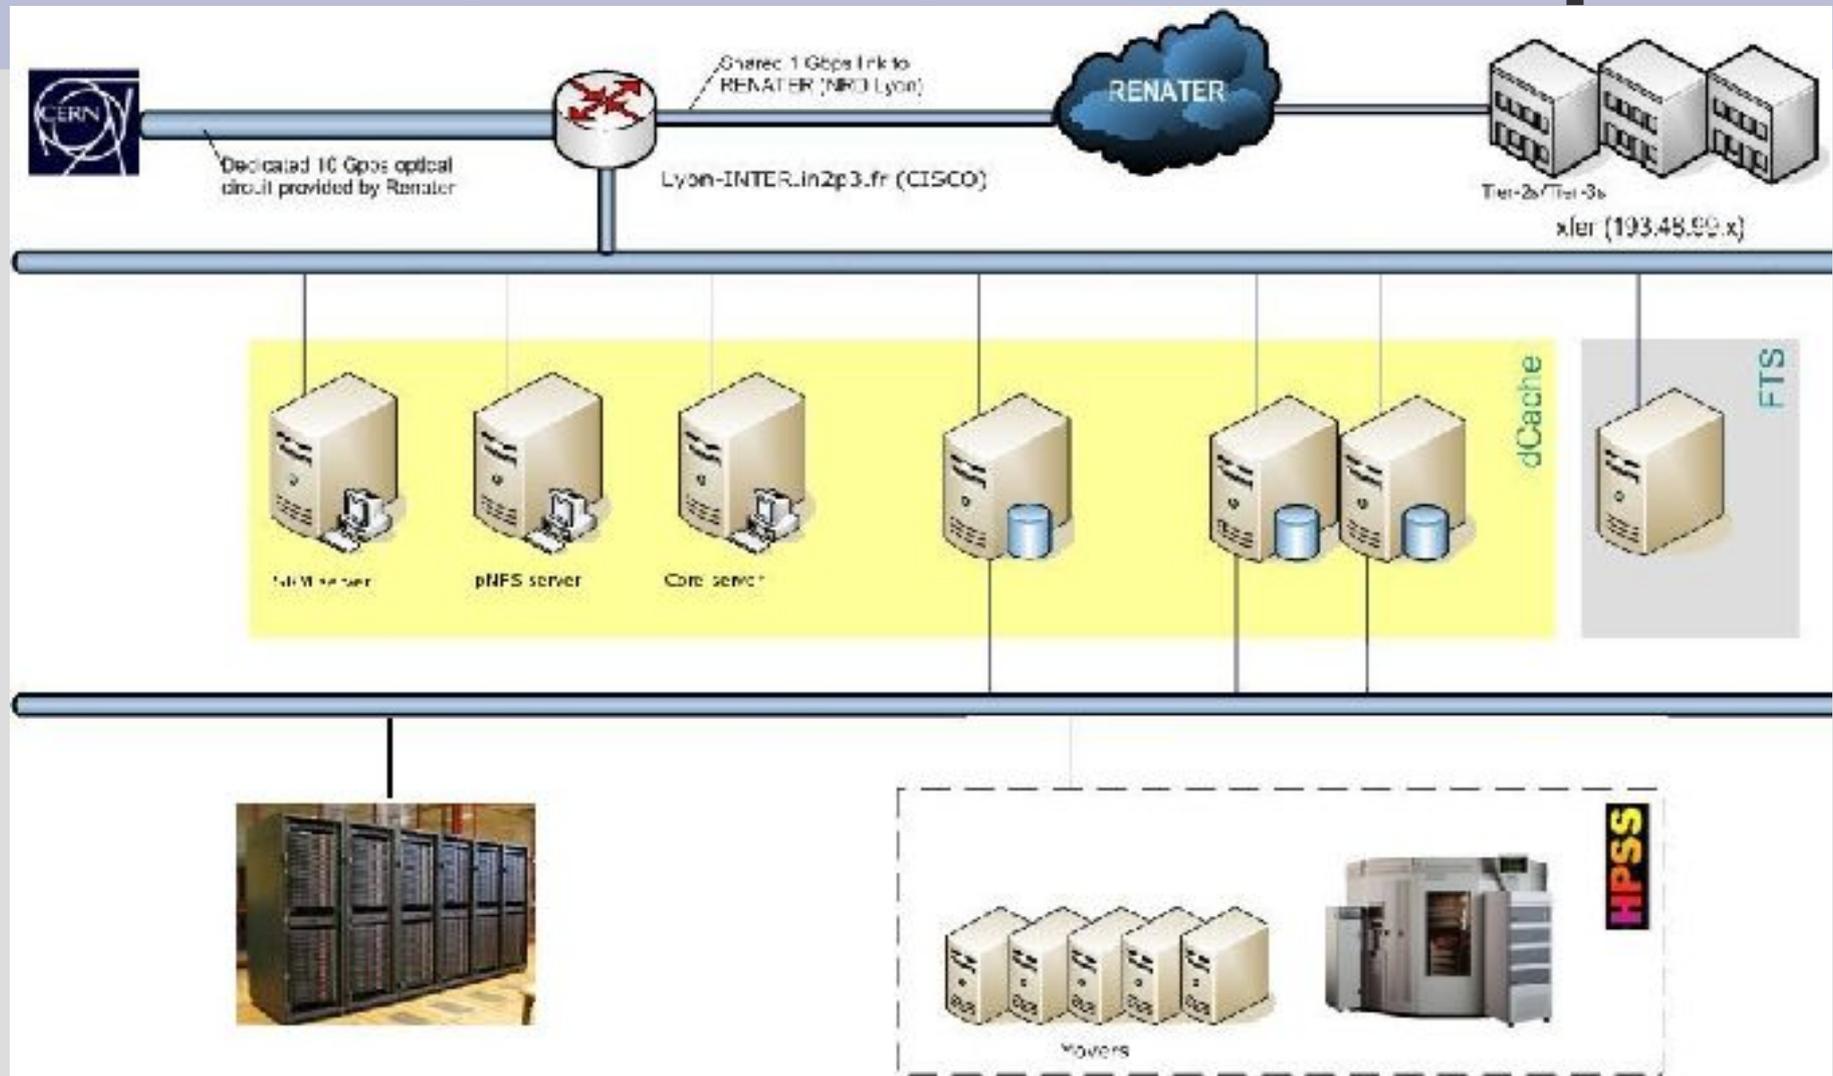
Lionel Schwarz <schwarz@cc.in2p3.fr>

CC-IN2P3

Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules



SRM/dCache/HPSS setup



Hardware setup

- 3 core nodes (PNFS, SRM, dcache)
 - V20Z, AMD bi-opteron 2GHz, 2GB RAM
 - 70GB local disk in RAID 1
- 20 disk servers (pools+GFTP door)
 - Various hardware, latest are V40Z and X4500
 - Typical hardware: quad-pro with 8GB RAM
- 150 TB allocated
 - 40 TB on IBM DS8300
 - 100 TB on X4500 (20TB per host)

dCache software setup

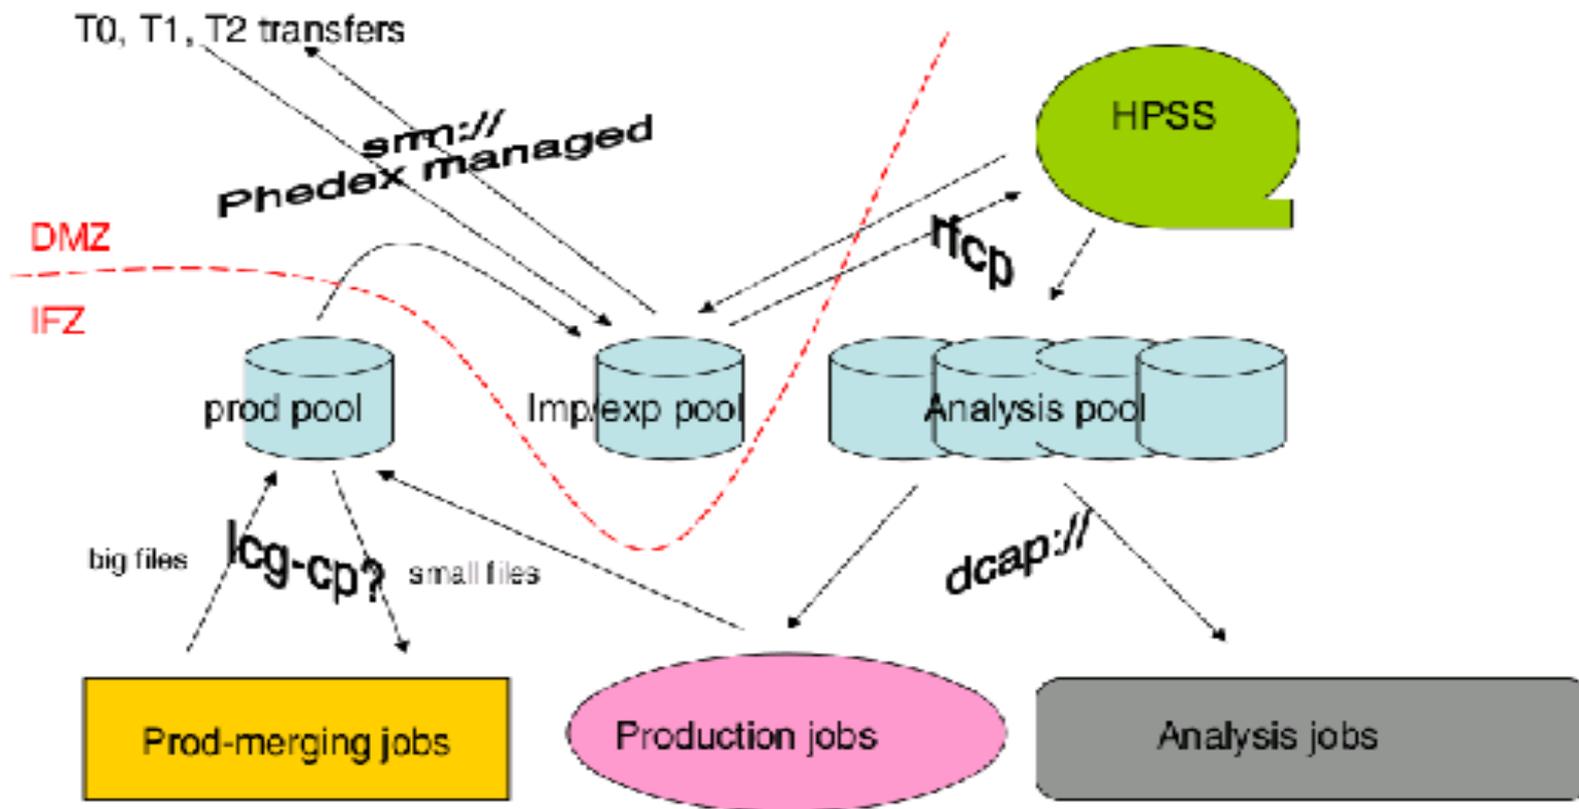
- SL3.0.5 or Solaris10 on X4500
 - GridFTP not yet on Solaris10
 - SL4 not tested yet
- XFS on Linux, ZFS on Solaris
- dCache-1.7 installed in November with latest SRM patch (1.7.0-24)
- PostgreSQL 8.1.0 for pnfs, companion, billing, SRM
- Backup: dump+archive in local system (TSM)

dCache configuration

- 58 pools in 17 pool groups
 - No shared pools except Multi-VO transfers
 - Small transfers pools distributed on several machines to achieve nominal throughput
 - Large pools for local access
- Mover queues defined :gridftp, dcap
 - Impossible to define a queue for Xrootd
- HSM migration/staging script on each pool node
- Dcap disabled from SRM
- /pnfs not mounted on client nodes except VO-box

CMS dCache pools setup

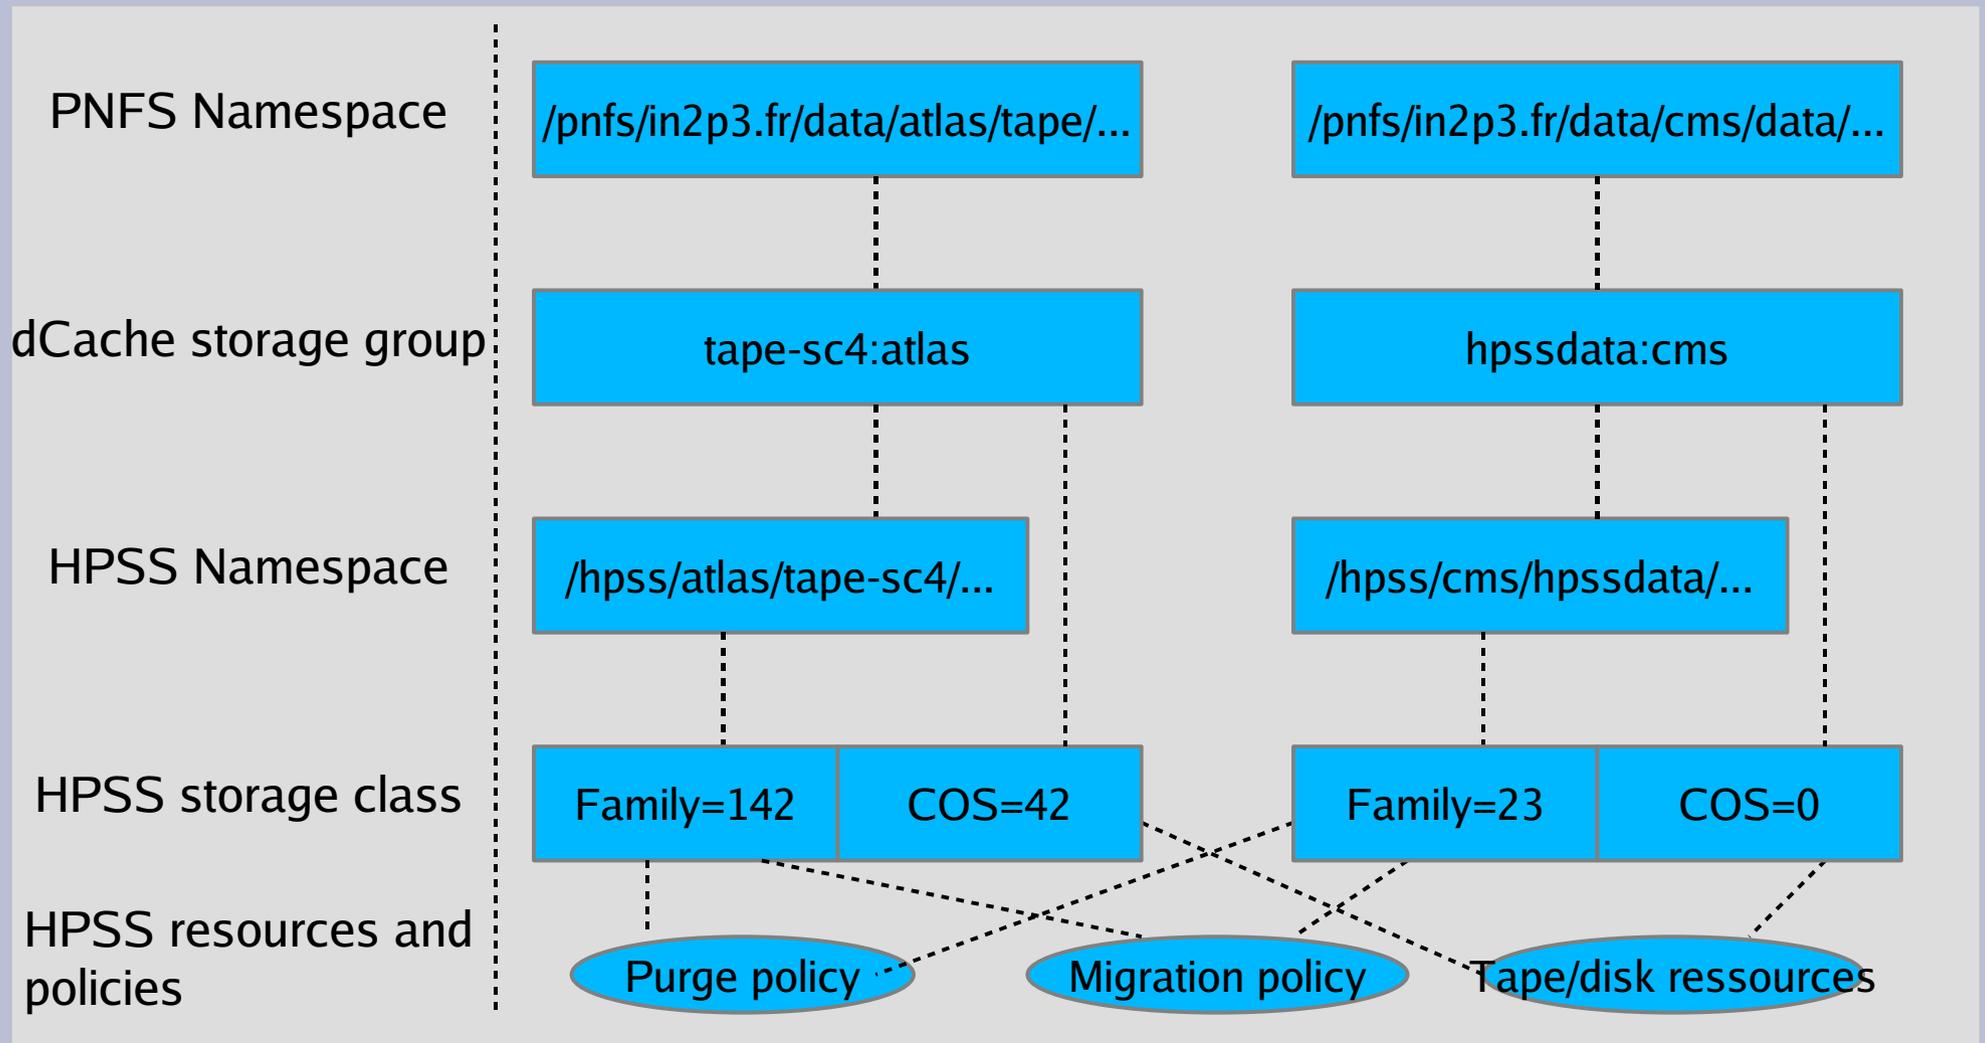
Credit: Artem Trunov



Storage classes in dCache/HPSS

- Storage Group in dCache
 - Defined by **tags** (hidden files) in the PNFS namespace
 - Linked with **pool group(s)**
 - That way, dCache knows where to physically store new files created in the namespace
- Storage Class in HPSS
 - COS (Class Of Service) is not defined in HPSS namespace
 - Associated with resources (type of disk, type of tape)
 - Associated with policies (filesize, migration, purge...)
 - Families are defined in the HPSS namespace
 - Allow to group files from a family on the same tapes

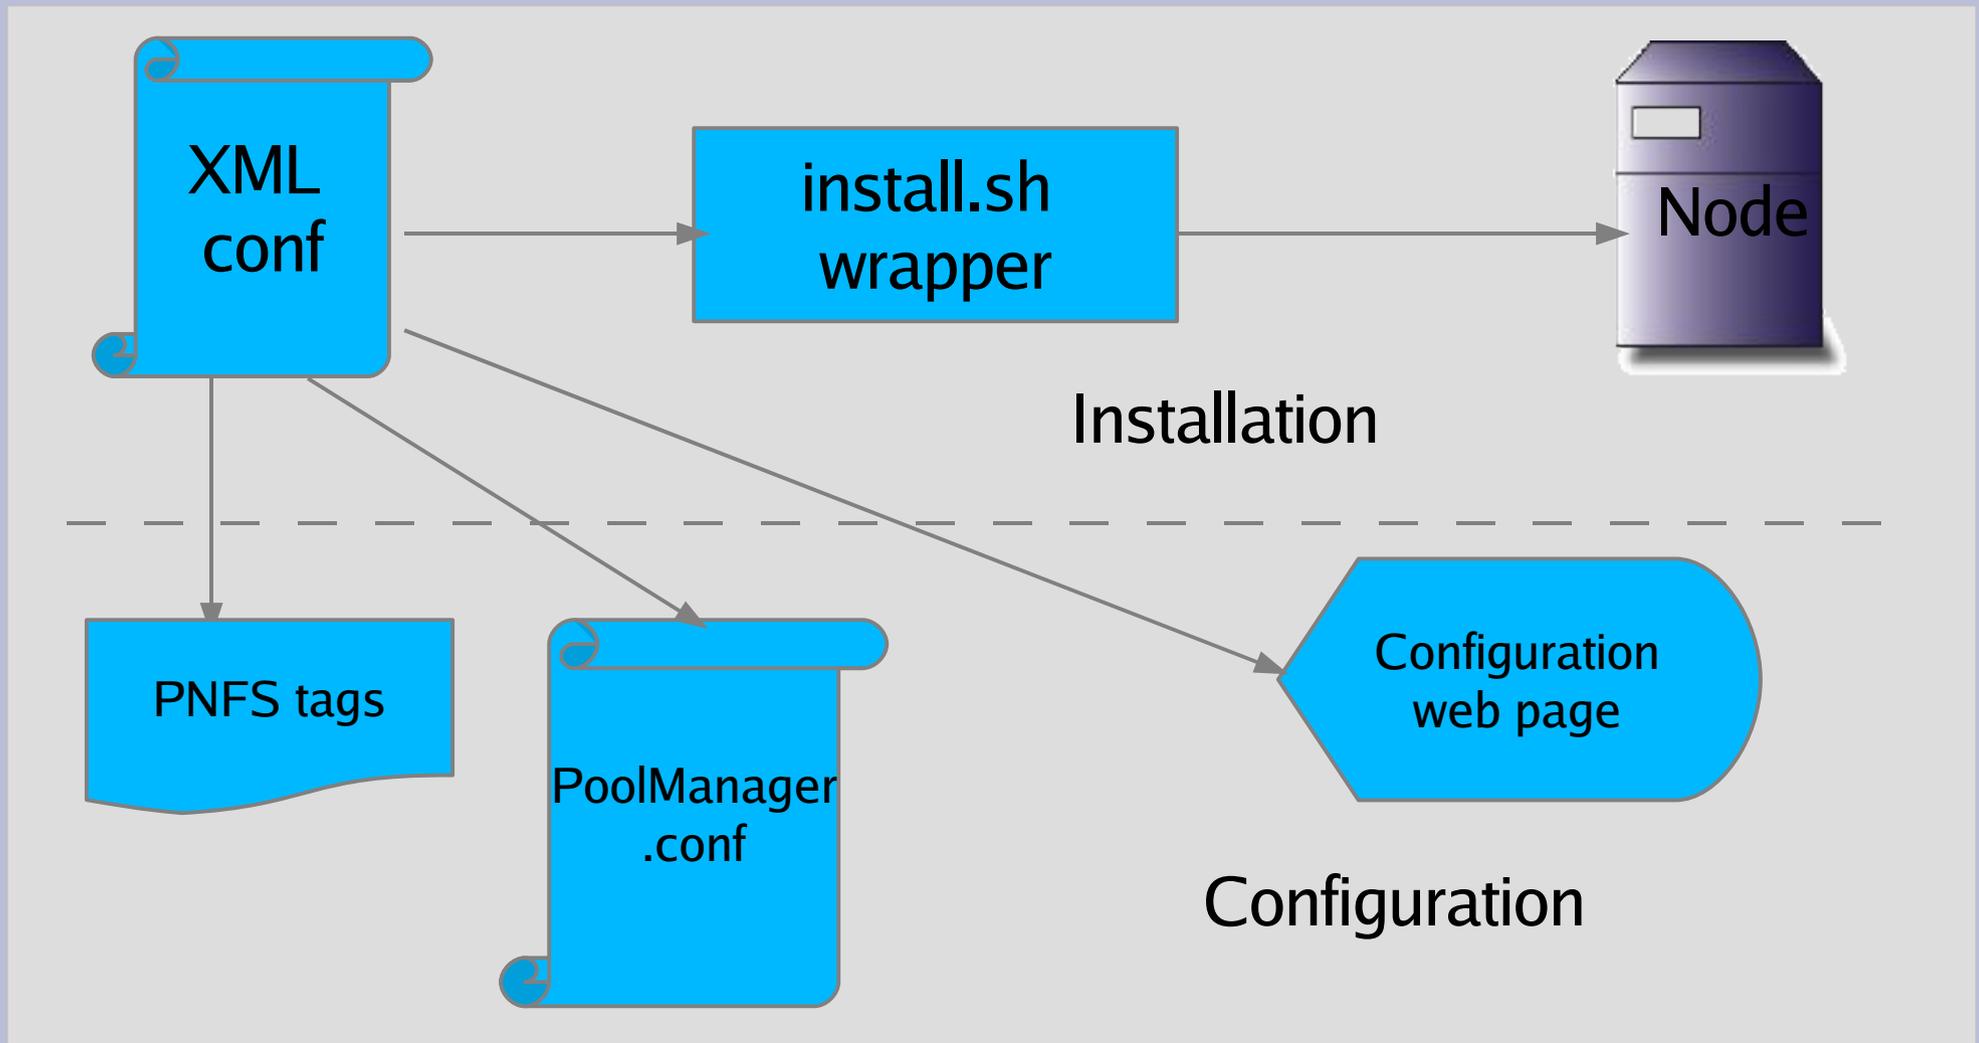
Storage Classes levels



Plans for SRM-2.2

- Not clear!!! Need collaboration with local vo-support
- Disk0Tape1
 - Separate input buffers from output buffers
 - Small size (depends on throughput required)
- Disk1Tape0
 - Large volumes distributed on several servers regarding to the aggregated throughput required
 - Separate production pools from “user” pools
- Disk1Tape1
 - could be implemented by automatically moving files from transfer pools to disk pools (“file hopping”) ?

Installation and configuration



Monitoring and operation

- dCache team: 2 people (part-time)
- Web pages (pool usage, current SRM requests, current errors, srmwatch...)
- dCache portal
- GUI for daily operation
- Integrated with local monitoring and logging systems
- Database query web pages integrated with local FTS database information
- Operation scripts for on-duty people
- Need a way to follow current dcap requests

dCache portal

Current activity

ALL ▾

| <u>SRM INFO</u> | <u>Movers</u> <u>Active/Queued</u> | <u>Store/Restore</u> <u>Active/Queued</u> | <u>Restore</u> <u>Handlers</u> | <u>GFTP transfers</u> <u>hanged</u> | <u>Transfers errors (last 3h)</u> | | | | | | | | |
|---|---|--|-----------------------------------|--|--|--------|--------------|--------------------------|---------------|-------------------------|-------------------------|-----------------------|-----------|
| COPY: Done:1190 RunningWithoutThread:9 RetryWait:1 Failed:1235 Canceled:64 PUT: Ready:8 AsyncWait:28 Failed:6335 Done:4147 Transferring:1246 GET: Ready:246 AsyncWait:4 Transferring:33 Failed:681 Canceled:18 Done:8672 | atlas-dq2:34/10 atlas-mc12:7/0 dteam-monitoring:1/0 lhcb-dst:5/0 | Restore: Store: | []:5 Suspended:3 | GFTP-ccxfer03:1 GFTP-ccxfer05:2 GFTP-ccxfer07:3 GFTP-ccxfer08:2 GFTP-ccxfer09:1 GFTP-ccxfer10:5 GFTP-ccxfer11:1 GFTP-ccxfer15:2 GFTP-ccxfer17:1 GFTP-ccxfer18:3 | <table border="1"> <thead> <tr> <th>Action</th> <th>Errors/total</th> </tr> </thead> <tbody> <tr> <td>transfer</td> <td>25/2350 (1 %)</td> </tr> <tr> <td>request</td> <td>1515/3808 (39 %)</td> </tr> <tr> <td>store</td> <td>0/1 (0 %)</td> </tr> </tbody> </table> | Action | Errors/total | transfer | 25/2350 (1 %) | request | 1515/3808 (39 %) | store | 0/1 (0 %) |
| Action | Errors/total | | | | | | | | | | | | |
| transfer | 25/2350 (1 %) | | | | | | | | | | | | |
| request | 1515/3808 (39 %) | | | | | | | | | | | | |
| store | 0/1 (0 %) | | | | | | | | | | | | |

Monitoring errors

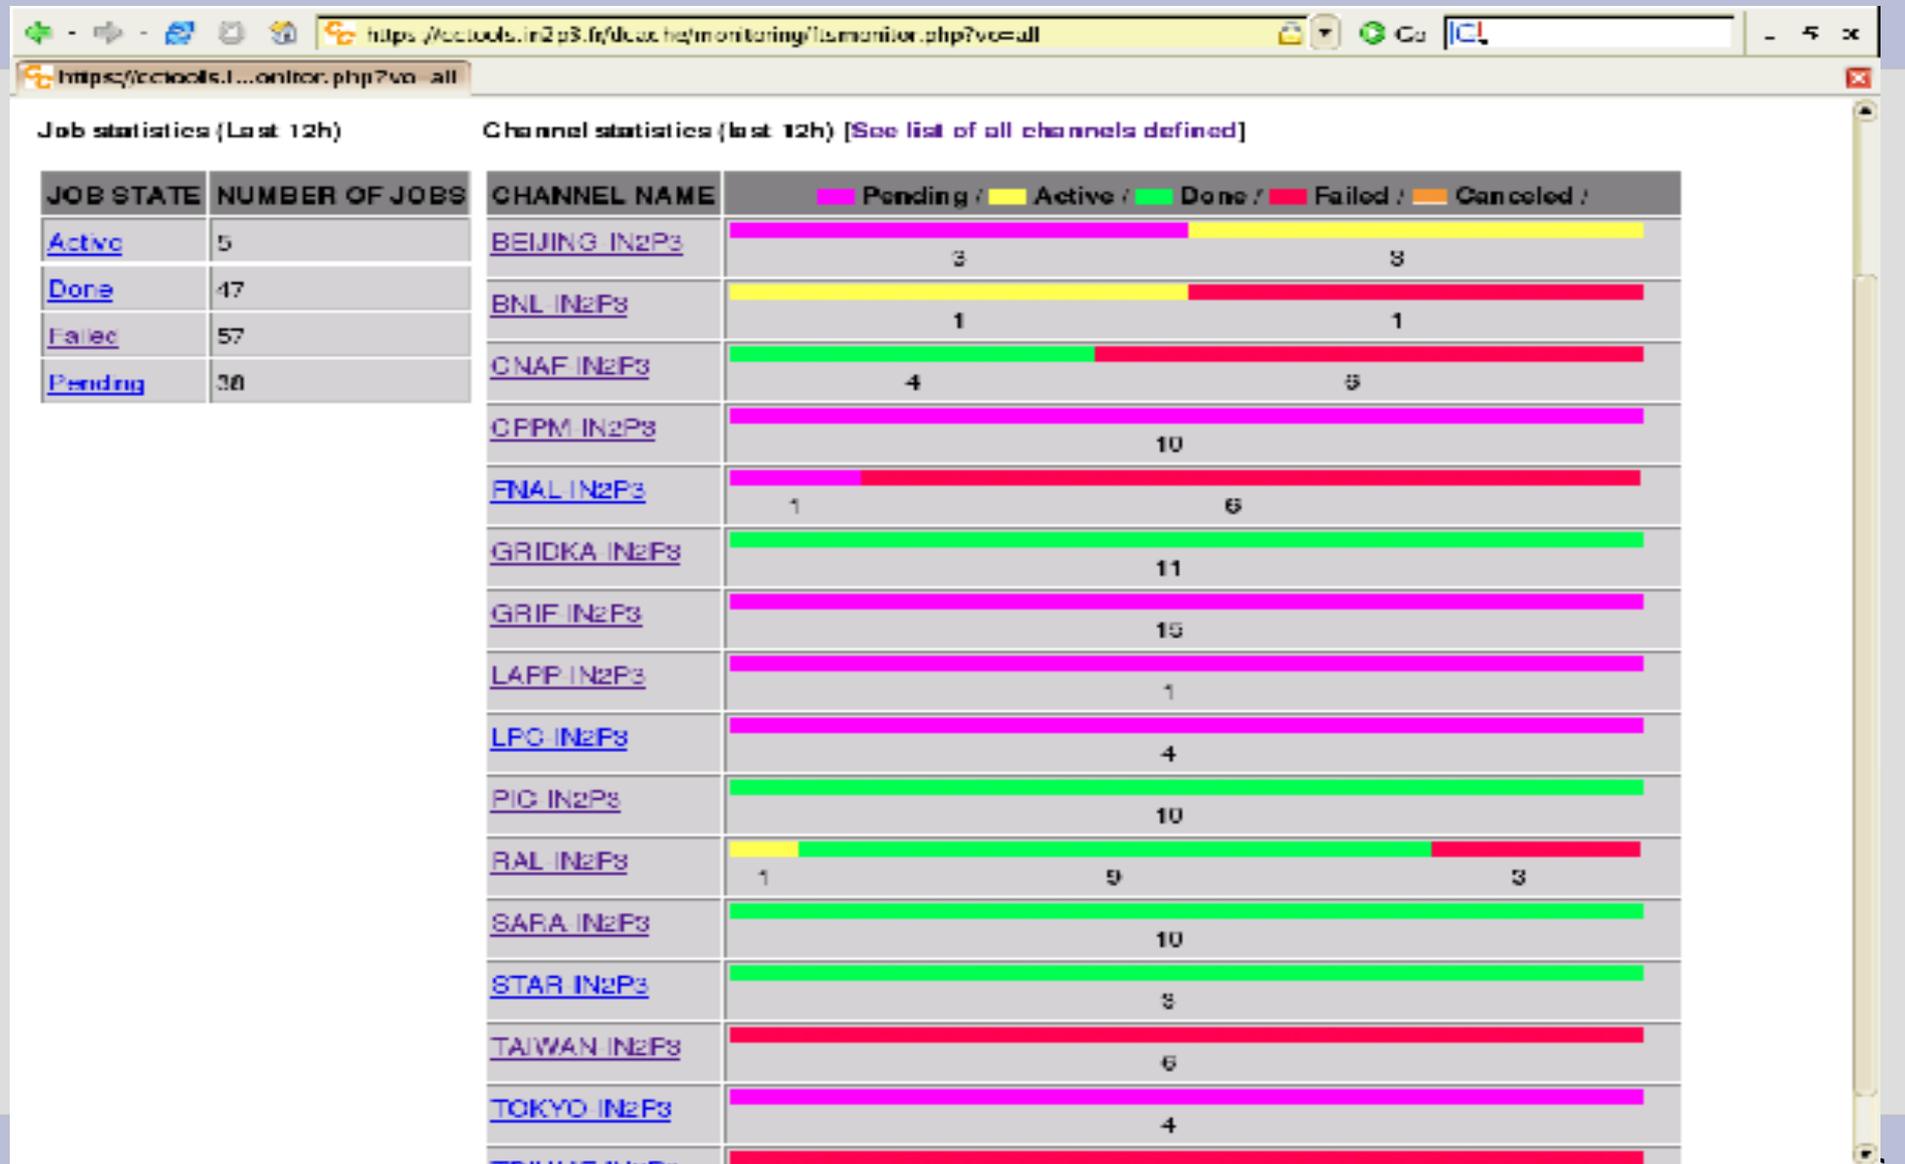
<https://cc-tools.in2p3.fr/portal.php?vo=all>
<https://cc-tools.in2p3.fr/vo=all&type=all>

SRM pin for A-1 show A-1 errors

12 pull files requests to ERRORS

| creation date | client host | url | url | final state | error | request error |
|---------------------------|--------------------|---|--------------------------|-------------|--|--|
| 2006-12-19 14:06:01+01 | b-102.com.fr | dcache-backup-c07/multi_vo_jestore-c07/ser02/207211 | | Success | Success: GetStorageInfoFile.ec file.ec: cannot write | Success: RequestFileStatList-21 info.ec: cannot write 14:08:01 CET 2006 stc GetStorageInfoFile.ec ec: cannot write! |
| 2006-12-19 14:00:21+01 | b-102.com.fr | dcache-backup-c07/multi_vo_jestore-c07/ser02/207211 | | Failure | Failure: GetStorageInfoFile.ec file.ec: cannot write | Failure: RequestFileStatList-21 info.ec: cannot write 14:00:01 CET 2006 stc GetStorageInfoFile.ec ec: cannot write! |
| 2006-12-19 14:56:12+01 | nono006.com.fr | dcache-backup-c07/2006-12-19/14:56:12/063e4ac3 | gsftp://box1018.in2p3.fr | Failure | Failure: changing file state because request state changed | Failure: Bad request |
| 2006-12-19 14:26:47+01 | 5-cc-01.m.in2p3.fr | dcache-backup-c07/disklog2/disklog2_msa_1_msa_1_msa_2 | gsftp://box1009.in2p3.fr | Failure | Failure: changing file state because request state changed | Failure: Bad request |
| 2006-12-19 14:25:16+01 | 5-cc-01.m.in2p3.fr | dcache-backup-c07/disklog2/disklog2_msa_1_msa_1_msa_2 | gsftp://box1007.in2p3.fr | Failure | Failure: changing file state because request state changed | Failure: Bad request |
| 2006-12-19 14:22:00+01 | b-102.com.fr | dcache-backup-c07/multi_vo_jestore-c07/ser02/207211 | | Success | Success: GetStorageInfoFile.ec file.ec: cannot write | Success: RequestFileStatList-21 info.ec: cannot write 14:22:00 CET 2006 stc GetStorageInfoFile.ec ec: cannot write! |
| 2006-12-19 | 5-cc-01.m.in2p3.fr | dcache-backup-c07/disklog2/disklog2_msa_1_msa_1_msa_2 | gsftp://box1007.in2p3.fr | Success | Success: changing file | Success: Bad request |

FTS monitor



Tools

- Scripts connect to AdminNode with ssh public key
 - Clean non PNFS entries on pool nodes
 - Enable/disable HPSS access
 - Print GFTP stuck cells
 - Clean PoolManager suspended requests
- Register HPSS files in dCache (migration of classic SE)
- On-duty scripts (enable pools, pnfs register...)
- Remote start/stop of GFTP doors and pools (pexpect)
- pgFouine (pgfouine.projects.postgresql.org/): DB stats

The 1.7 upgrade story 1/2

- Upgrade 1.6_ \rightarrow 1.7 on 11/21/2006
 - Various configuration problems quickly fixed
- SRM problems started few days later
 - PUT requests stay in QUEUED state, the machine was overloaded, gPlazma crashed every 30 min.
 - Disabled gPlazma: did not help
 - Installed SRM on a separate node: did not help
 - Dropping the DB helped but for a few days only
 - Changed dCache, Tomcat, Postgres conf files: did not help
 - Used ThreadManager: did not help
 - Pray God, sacrifice a virgin, immolate myself: did not help

The 1.7 upgrade story 2/2

- The solution
 - Timur finally installed a patch on 12/09 (optimized access to DB for “nextrequestid” + Tomcat 5.5)
 - Thanks for good support by Timur and Tigran
- We still see the following issues
 - Lots of UPDATE requests on “srmrequestcredentials” makes VACUUM ANALYZE impossible, so we have disabled this table from VACUUM
 - *It seems it is not the case with latest patch 1.7.0-24*
 - DPM->dCache authentication problems with Atlas proxy not yet understood
 - SRMcopy incompatibility with NorduGrid SE

Operational issues - SRM

- In 1.7, SRM seem much more resource-consumer than 1.6
 - Even with SRM-only, the machine seem not powerful enough
 - srmwatch is disabled, need to install on another node?
- Still need to restart the service to change parameters
- Difficult to manage transfers (kill, list...)
 - Killall “.*badse.badsite.xx.*” would be useful
- How to blacklist a site, a user DN?

Operational issues - DB

- No archiving of billing DB and SRM DB
 - Plans to export “old” entries to another DB for statistics and debug purpose
 - Not clear the role of SRM tables (XXX and XXX_b)
- Maintenance operations (vacuum and analyze) are important for heavily used (modified) DB
 - Pb with “srmrequestcredentials”: vacuum analyze lasts 10h!
=> was disabled (can likely be enabled with latest patch)
- Default parameters are likely wrong for heavy usage
 - Need to re-run Multi-VO transfers to test performances
- Slow requests (see pgfouine_report)

Operational issues - params

- Lots of different levels of parameters for queues
 - SRM queues
 - RemoteGsiFTPTransferManager (“set max transfers”)
 - Doors (“set max logins”)
 - Pools (“mover set max active”)
- Lots of different levels of timeouts
 - SRM lifetime
 - RemoteGsiFTPTransferManager (“set pool manager timeout”, “set pool timeout”, “set pnfs manager timeout”)
 - Pools (“jtm set timeout”)
 - PoolManager (“set timeout fetch”, “set timeout pool”, “set timeout pnfs”)

Operational issues - errors

- Difficult to extract serious errors when thousands are displayed!
 - Plans to avoid duplicate displays of the same type of errors.
- Difficult to extract local errors from remote errors
 - Timeouts, “Server error”: difficult to know why
- We have developed a local “error repository”
 - Plan to migrate repository in relational DB
 - Would it be interesting to build a world-wide repository for all dCache admins?
- What to do with suspended requests in PoolManager?
 - Would be nice to separate permanent errors (file disappeared) from temporary errors (pool disabled)

Operational issues – misc

- Data checking:
 - How to compare LFC catalog with pnfs catalog?
 - How to compare pnfs catalog with pools catalogs?
- Downtimes
 - How to nicely shutdown all/some services?
 - SRM: no commands to suspend commands
 - GridFTP: “set max logins 0” OK but sometimes new requests are accepted
 - dCap/gsidCap: “set max logins 0”: no effect
 - Remote downtimes
 - OK when FTS is used (channel is closed)
 - Impossible to temporarily “blacklist” remote sites

SRM/dCache service

- Plans to describe different levels of intervention on the SRM/dCache system for example:
 - “administrator”: all commands allowed (operations and configuration + installation)
 - “operator”: operations commands allowed only
 - “vo-manager”: basic operations commands allowed for VO only
 - “on-duty”: minimum operations commands for week-end fixes
- Need to provide better diagnosis to map with other systems (HPSS, batch, FTS...)
- Need to automate operation commands (“mover kill”...)

What next ?

- Integration with local batch system BQS
- Improve monitoring and operation tools
- Modify the HSM staging process to group HPSS requests
- Test and implement file hopping, central flushing, resilient pools...
- Deployment of SRM v2.2
- Setup a dCache service for EGEE (upgrade of old “Classical SE”) + Phenix (transfers from BNL)
- FNAL-CC throughput tests (using direct 2*1Gbs connexion)