



dCache at RAL




Derek Ross
E-Science Department

Setup

- 8 GridFtp/GSIDcap door systems
- 10 Disk servers (more soon)
 - Some small partitions on another
 - 4*2 pools to act as tape cache to ADS
- “Head node”
 - PNFS
 - PoolManager etc, Main SRM door

Setup 2

- Separate Postgres system
 - Mainly for CPU load reasons

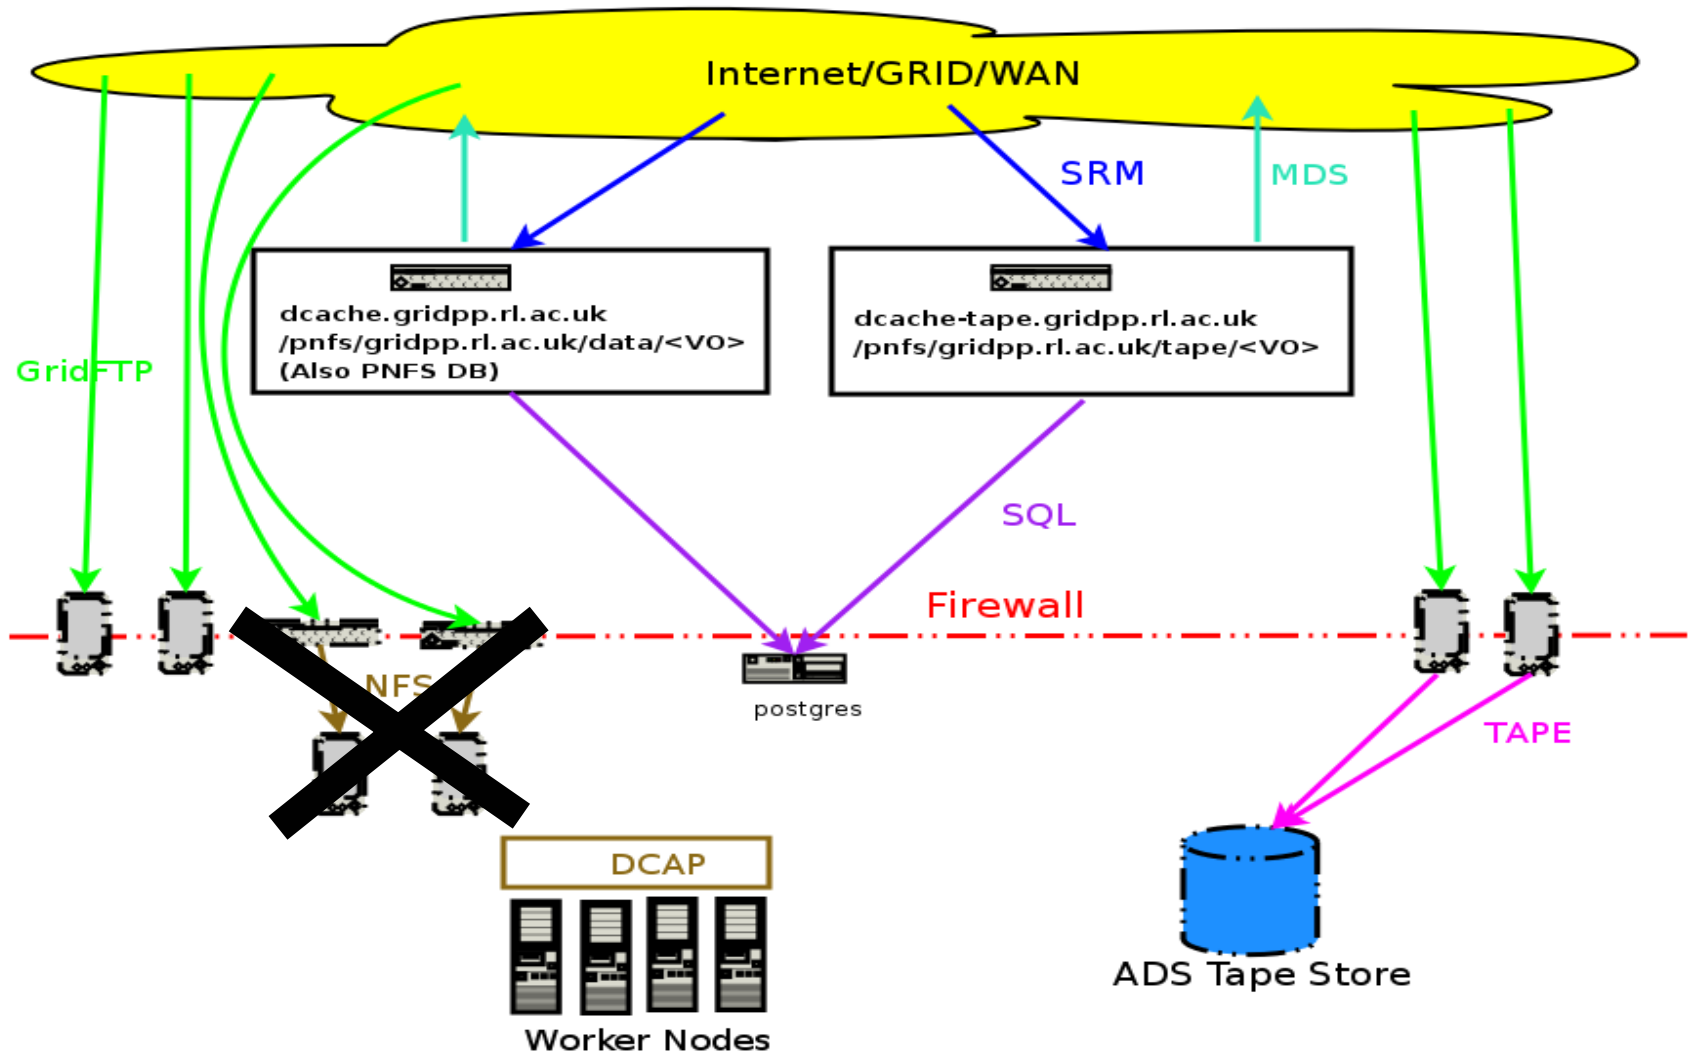


Setup 3

- Second SRM door on separate system for tape access
 - Allows us to have different lifetimes for tape access versus disk access
 - LCG info system only allows an SE to publish one StorageRoot per VO

Historical Setup

- Pool disk nfs mounted
 - Not successful, nfs hanging under transfers
 - Pool was running but disk inaccessible



Future Setup

- Continue to add disk servers
 - New LHC experiments disk allocation
- Move SRM door off admin node
- Put admin node on more robust system

Hardware

- Disk Servers
 - Dual 2.4GHz P4 Xeon, 1GB memory
 - 2x 1.7TB partitions RAID 5
 - Dual 2.8GHz P4 Xeon, 2GB memory
 - 4 x 1.5TB partitions RAID 5

Hardware 2

- Other Systems
 - Head node, Postgres node
 - Dual 2.66GHz P4 Xeon, 2GB memory
 - Gftp servers
 - Dual 2.8GHz P4 Xeon, 2GB memory

Network

- Everything on Gb network
 - But not in same room, 4Gb between rooms
- Two routes out
 - Default through site router to JANET
 - Secondary via OPN to CERN, Lancaster

Configuration

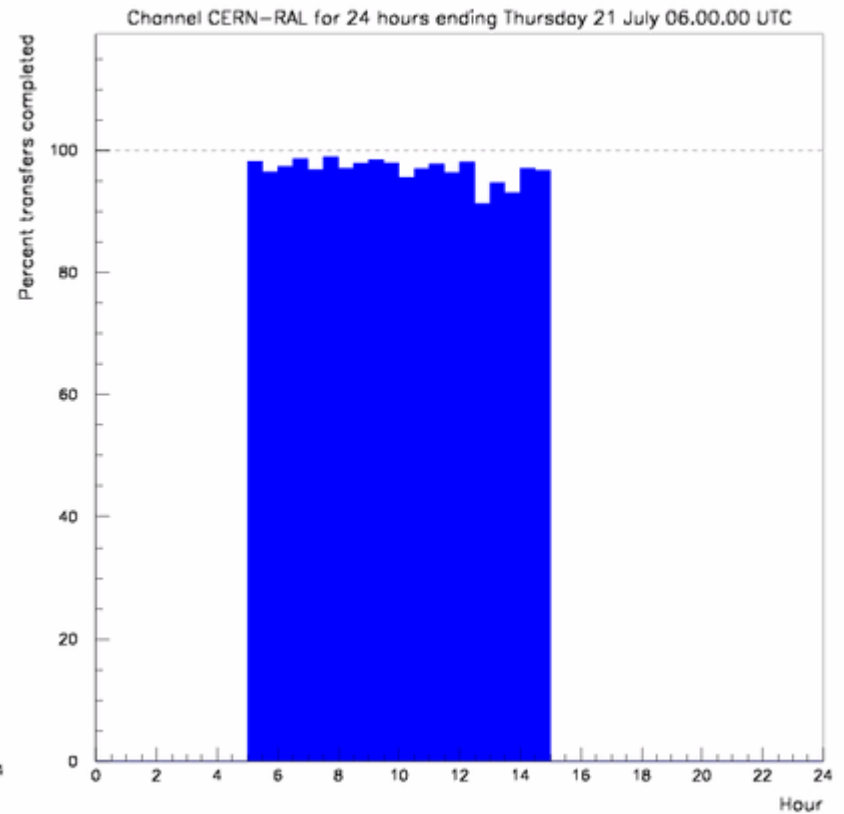

- Experiments/VOs are allocated disk
 - Need to be able to show that they have it
 - Use PNFS tags to send files to VO's pool group

Problems

- lcg-gt as file existence check
 - no corresponding lcg-sd
 - blocks transfer slot (for 24hrs in default config)
- FTS gave up on transfers, but didn't inform SRM
- ***Pool logs filling***
 - *removeFiles : invalid syntax in remove filespec >null<*

Problems 2

- PinManager hangs
- Can't SrmCopy from castorgrid.cern.ch
- 07/21 10:05:40 Cell(RemoteGsiftpTransferManager@utilityDomain) : [id=3419 store src=gsiftp://castorgrid.cern.ch:2811//shift/lxfs6013/data02/zh/stage/ZippedEVD.63300016.63300030.eg_Hit245_2_g133.eg03_jets_1e_pt2550.zip.981194 dest=///pnfs/gridpp.rl.ac.uk/data/cms/TEST FILES GO IN HERE/heartbeat]:sending error reply, reply code=8 errorObject=transfer failed :org.globus.common.ChainedIOException: Authentication failed [Caused by: Operation unauthorized (Mechanism level: Authorization failed. Expected "/CN=host/castorgrid.cern.ch" target but received "/C=CH/O=CERN/OU=GRID/CN=host/castorgrid02.cern.ch")]
- PostGres
 - Seems to slow down over time
 - Have to drop tables



Performance

- SC2
 - 75MB/s to disk
 - 50MB/s to tape
- SC3
 - disk rate varied wildly
 - 60MB/s to disk

Performance 2

- Seen better
 - Achieved 3000Mb/s into dCache locally(Pre-SC2)
 - Need to rerun test

- Comment/Questions?