



# RAL Storage

J Jensen <j.jensen@rl.ac.uk>

Contributions from Tier 1 (D Ross, A Sansum, S Traylen)

Storage Workshop, CERN

5-7 April 2005

Jens G Jensen  
RAL

RAL – <http://www.clrc.ac.uk/Activity/RAL;>

GridPP – <http://www.gridpp.ac.uk/>

- GridPP Tier 1
  - <http://www.gridpp.ac.uk/tier1a/>
- GridPP Storage middleware group
  - Supporting deployment to T2s throughout UK
  - <http://storage.esc.rl.ac.uk/>
- RAL Atlas Datastore
  - <http://www.e-science.clrc.ac.uk/web/services/datastore>

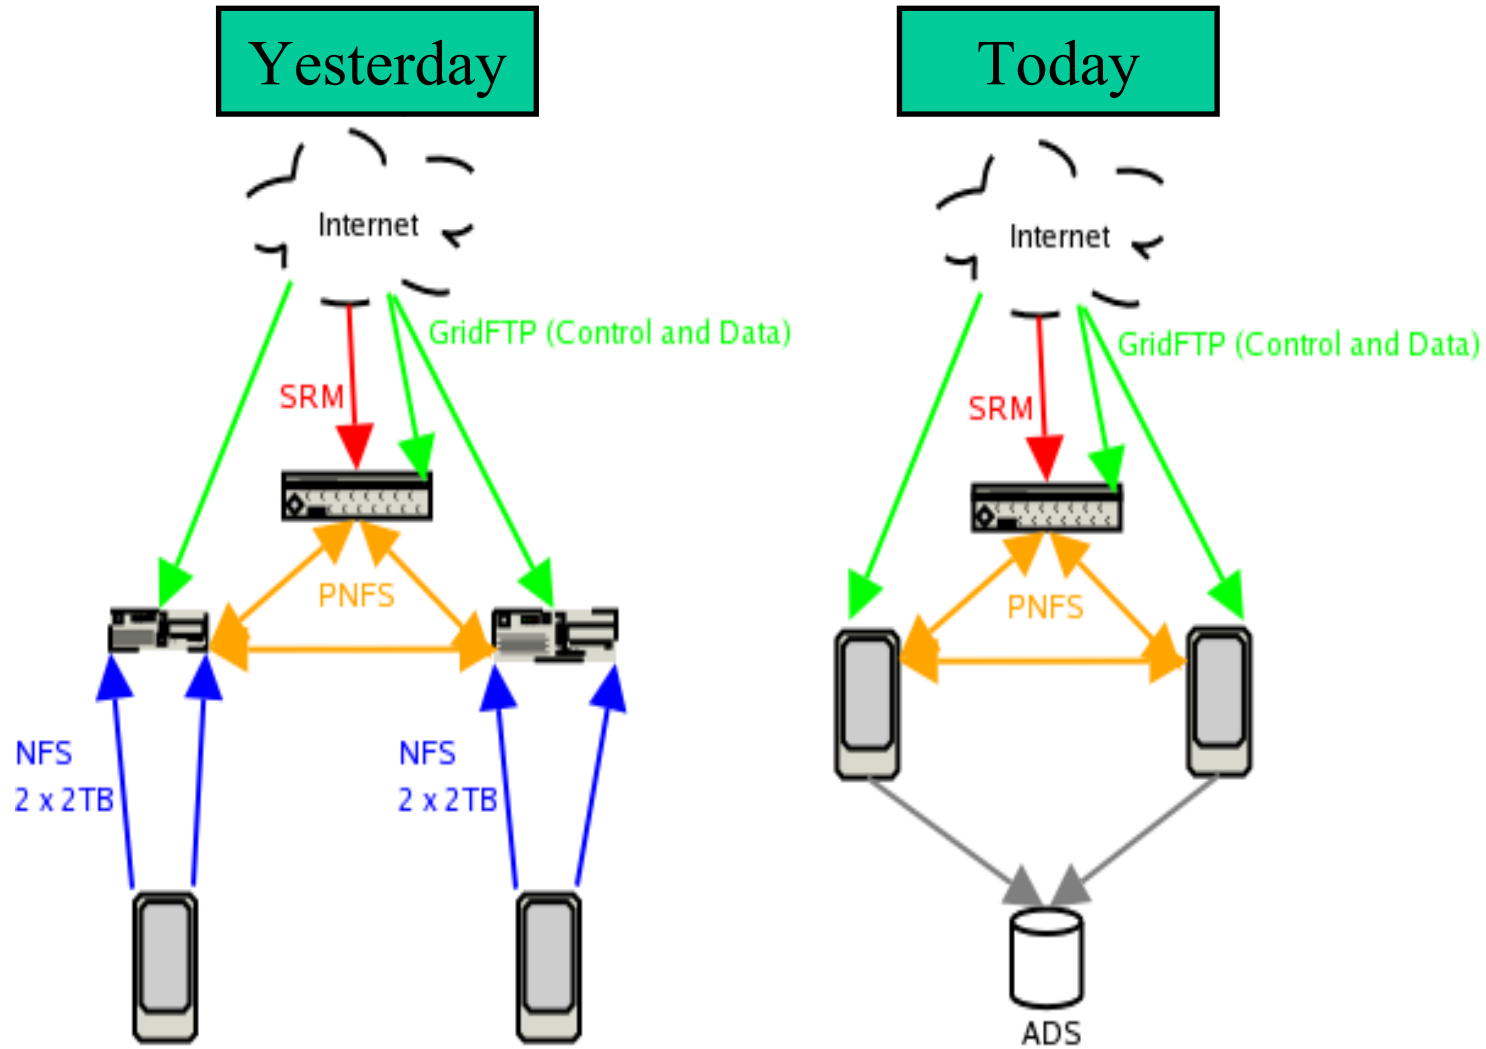
- SRM 1 to tape and disk at Tier 1
- SRM 1 to disk at Tier 2s
- Disk pool management

- SRB used for DC04
- SRB also currently used to tapestore
  - Used by many non-HEP experiments
- CMS and ATLAS now using the SRM to Tier 1

- Getting 75 MB/s from CERN
  - But seeing 3-3.5 Gb/s over LAN
  - Investigating...
- Lightpath
  - Got it at last moment
  - Underprovisioned?

- GridPP Tier 1
  - Derek ROSS
  - Deployed three dCache-with-SRM
    - Production
    - Service Challenge
    - JRA1
- GridPP Storage
  - Owen SYNGE and Jiří MENCÁK
  - Deployed three dCache-with-SRM
    - Test
    - Test
    - Test

# Current Deployment at RAL



- Several Tier 2s installing dCache
  - Edinburgh
  - Manchester
  - Lancaster
- Planning to install
  - Imperial (London) (staff shortage)
  - RAL PP (due to start soon)
- Others have expressed interest



- Running production dCache
  - 25 TB disk
  - Interfaces to permanent disk
  - Interfaces to tape (dteam only atm) – no delete yet
- SC dCache
  - 4 disk servers
  - 13 GridFTP servers
  - Aggregated 500 MB/s

- Stability is good 😊
- Performance is good 😊
  - Optimising by adding more hardware
  - Head node, SRM, and SRM db all on different machines for SC2
- Satisfied with support from dCache 😊

- SRM requests left open
  - If interrupted, lcg-gt leaves requests open
  - Users must call lcg-sd to close
- Running dCache GridFTP on multihomed host
  - TURL points to default IP address
- Some minor VO related problems
  - E.g. total available space published, not space per VO

- One size doesn't necessarily fit all...
  - Investigating DPM (result next week)...
  - xrootd... (BaBar using it)
- Collaborating with CASTOR group
- Own SRM implementation
- RAL Tapestore plans
  - Investigating running CASTOR
  - ...or need to upgrade current software

- dCache
  - currently managing access to storage at Tier 1
  - Deployed or deploying to 4-5 major T2 sites
- Interfacing to disks and tape at RAL
  - But tape only opened for dteam right now