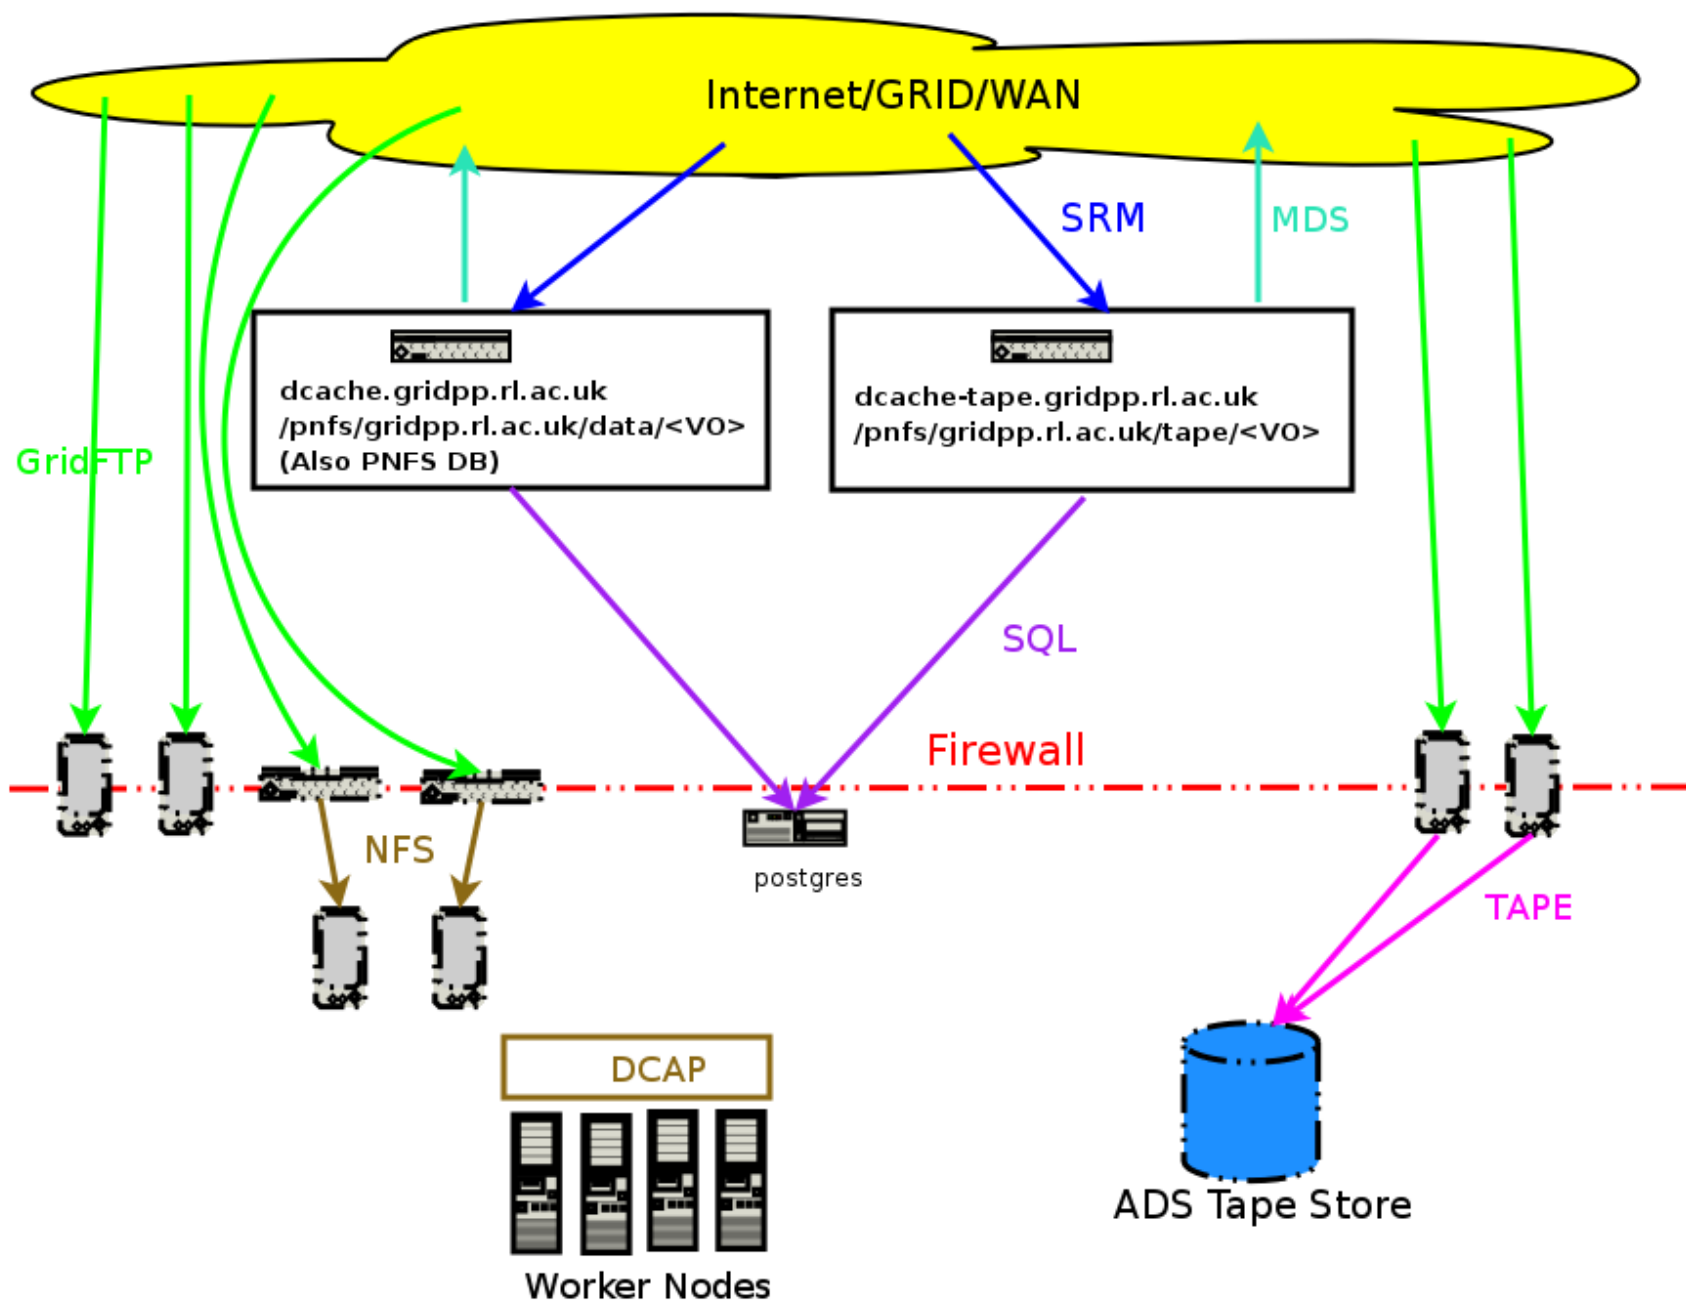


# RAL DCache Status

Derek Ross/Steve Traylen  
April 2005

# Outline

- Proposed setup for production system.
- Multi-homing for SC2.
- Issues, problems and confusions.
- DCache Errors?



# Questions about Production Setup

- Is this sensible, any better ideas.
- CMS and others look to be assuming they can delete files immediately.
  - This includes files on tapes.
  - There is no delete method to interact with the HSM.
- Truncate option does not work for us?
- We don't know how to drain a disk server for maintenance.

# Quotas on Production Setup

- Currently we quota VOs by having pools uniquely attracted to say.
  - /pnfs/gridpp.rl.ac.uk/data/dteam
- This should allow us to drain a disk partition and move it to a different VO to change allocations.
- We can also have a shared pool to introduce soft limits. (The tape backed disk will be like this)
- But we reduce the number of servers that can do I/O though. Anything better?
- Storage space advertised needs to reflect this.

# Multi Homing for SC2

- Aim
  - Have disk servers on production network.
    - These were running the pool software.
  - Expose this to UK light network via multi-homed GridFTP servers.
  - We could not get it to work.
    - FTP protocol returned i.p. for data channel on wrong interface. (Can we bind the gridFTP doors?)
    - Pool nodes tried to contact external interfaces of GridFTP Nodes.
- For SC3 this problem may go away but is very likely relevant for other T2 sites.

# Data Flow

- We have general confusion about what happens with.
  - SRMGet
  - SRMPut
  - SRMCopy (Push and Pull)
- In particular what initiates the various connections at GridFTP and dCap levels.
- Choice of GridFTP server is not influenced by location of data?

# DCache Errors

- Some errors encountered during running of dCache
  - All seen in latest release
  - Full java backtraces omitted for brevity
  - None appear to affect operation, but disconcerting at first
  - Does make it harder to see actual errors in logs
  - Ignoring errors is not a good habit to get into



# Odd Errors 1

- Extra blank line in <host>.poollist leads to following in <host>Domain.log

```
Cell(${0}@jra1dch01Domain) : CellAdapter :  
prepareRemoval : got java.lang.NullPointerException
```

```
...
```

# Odd Errors 2

- **From pnfs.log**

```
Cell(PnfsManager@pnfsDomain) : Error obtaining 'l' flag  
for getSimulatedFilesize :  
java.io.FileNotFoundException:  
/pnfs/fs/.(puse)(0000000000000000000000001690)(2) (Is a  
directory)
```

# Odd Errors 3

- **From srm.log during startup**

```
Cell(SRM@srmDomain) : java.sql.SQLException: ERROR:  
  Relation 'srmnextrequestid' already exists  
Cell(SRM@srmDomain) :      at  
  org.postgresql.core.QueryExecutor.executeV2 (QueryExec  
  utor.java:287)
```

...

# Odd Errors 4

- From gridftpdoor.log for every transfer

```
Cell (GFTP-Unknown-101@gridftpdoorDomain) :  
  java.io.IOException: Log root directory /tmp/dcache-  
  ftp-tlog not found  
Cell (GFTP-Unknown-101@gridftpdoorDomain) :          at  
  diskCacheV111.doors.FTPTransactionLog.begin(FTPTransa  
  ctionLog.java:98)  
Cell (GFTP-Unknown-101@gridftpdoorDomain) :          at  
  diskCacheV111.doors.GsiFtpDoorV1.startTlog(GsiFtpDoor  
  V1.java:109)
```

...

- Does go away if directory is created

# Odd Errors 5

- **From http.log:**

```
Cell(topo@httpdDomain) : Exception in Loop :  
    java.lang.ClassCastException
```

- **From dcache.log:**

```
Cell(1-100-Unknown-80270@dCacheDomain) :  
    java.io.InvalidClassException:  
    dmg.cells.nucleus.CellMessage; local class  
    incompatible: stream classdesc serialVersionUID =  
    2455460467756494833, local class serialVersionUID = -  
    5559658187264201731
```

```
Cell(1-100-Unknown-80270@dCacheDomain) : at  
    java.io.ObjectStreamClass.initNonProxy(Unknown  
    Source)
```

...