



Dealing with a decommissioned SE

*Jan Just Keijser
Nikhef / BiG Grid
10 April 2013*

J.J. Keijser
Nikhef
Amsterdam
Grid Group



BiG Grid
the dutch e-science grid

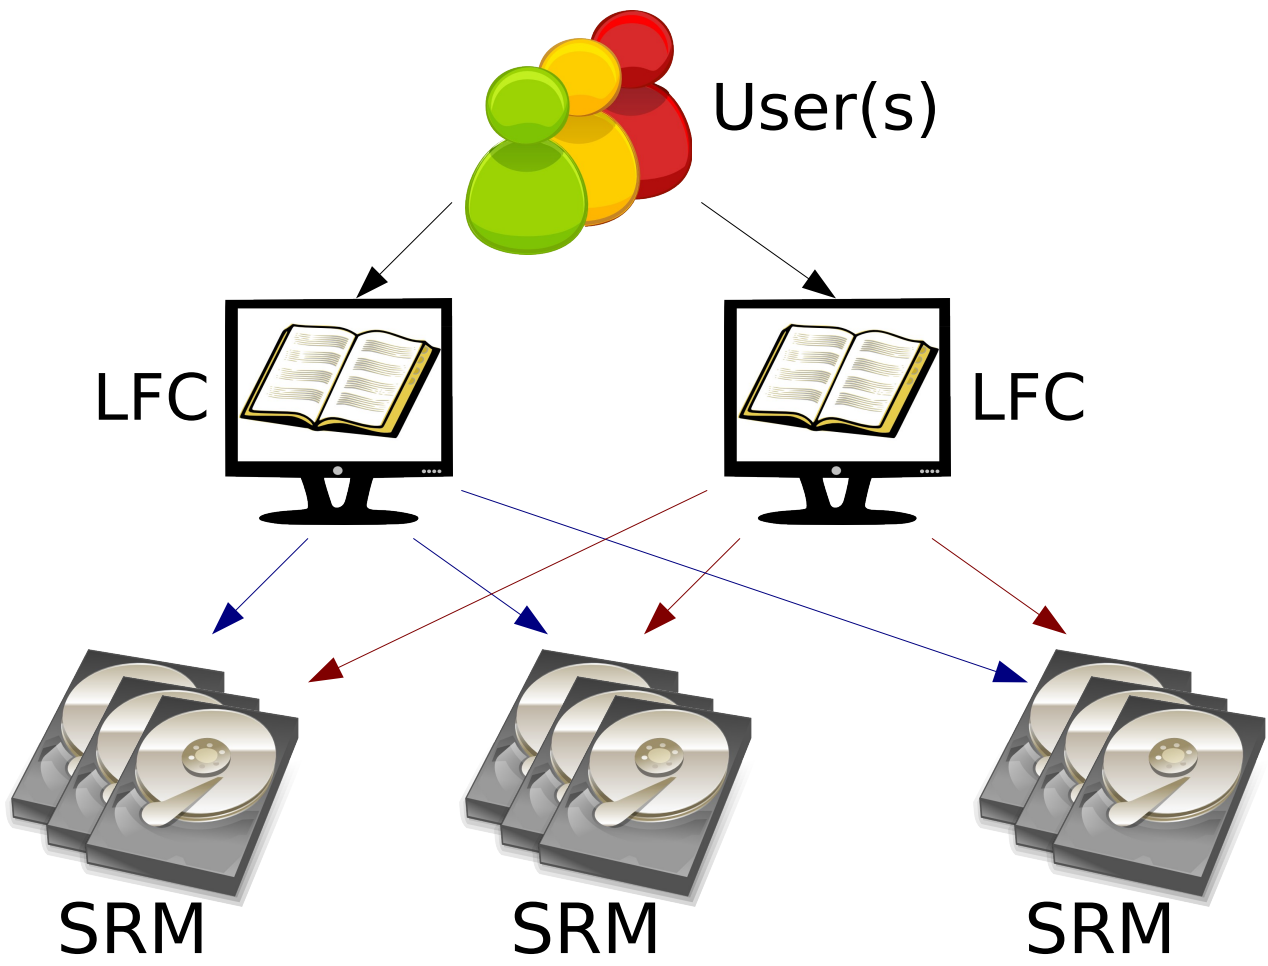
What & Why

- DPM Storage Element (SE, Sun "Thumper") had run out of hardware support
- Admins for all VOs that had data on this SE were contacted to either migrate or remove their files
- As VO admin for the VO 'vlemmed' I was asked to migrate or remove data for this VO

The "grid model"

- The main challenge was to do this without involving the site administrators
- The end-users in the VO typically use GUI-based tools such as 'vbrowser'

LFCs and SRMs



J.J. Keijser
Nikhef
Amsterdam
Grid Group

Step 1: Process all LFC entries

- Recursively list all LFC entries:
 - `$ lfc-ls -lR /grid/$v0 > files.lfc`
- Convert this list into a list of full paths
 - `$./dir2lfns files.lfc > files.lfns`
- Get SURLS for each of these full paths
 - `$ for lfn in `cat files.lfns`
do
 lcg-lr $lfn
done`

Step 2: Process all DPM entries

- Recursively list all DPM entries:
 - `$ dpns-ls -lR ../../$v0 > files.dpm`
- Convert this list into a list of full paths
 - `$./dpns2lfns files.dpm > files.surls`
- Get LFC entries for each of these full paths
 - `$ for surl in `cat files.surls`
do
 lcg-la $surl
done`
- Also tried using 'srmls' and 'srm-get-permissions'; nearly blew up dCache SRM

Sidestep: some statistics

- All Dutch SEs were queried
- Dark storage: files found in DPM but not on the LFC:
 - DPM SEs: on average 4.7% “dark storage”
 - dCache SEs: 9.6 % “dark storage”
 - This includes old data that was stored when DPM & dCache were less stable (2006/2007)
- Missing files: files found in the LFC but not in DPM:
 - Some SEs had been renamed
 - Some SEs had been removed

Step 3: Match LFC&DPM entries

- Match all SURLS found on the LFC against all SURLS found on the DPM SRM
- For each of the missing entries, find the owner using 'lfc-getacl' and 'dpns-getacl'
- For each of the entries that were inaccessible on either the LFC or the DPM (due to permissions), find the corresponding directory owner
- Inform all owners and ask them what to do with the entries found

Step 4: Replicate needed files

- For each of the files (SURLs) that are still needed, create a replica elsewhere:
 - `$ for surl in `cat needed.surls`
do
 newsurl=$newsrm${surl##$old...}.
 lcg-rep -d $newsurl $surl
done`
- Verify that replication was successful:
 - `$ for surl in `cat needed.surls`
do
 lcg-lr `lcg-la $surl`
done`

Step 5: Delete files from old SE

- Delete all files from the to-be-decommissioned DPM SRM using 'lcg-del'. This removes the corresponding LFC entries as well, if no replicas remain
- Delete all empty directories on the SE using 'dpns-rm -r'. This does **not** recursively remove files, but it does remove empty directories.
- This is the easiest way to verify that you are done

Problems encountered

- Listing all files on the LFC for a production VO takes very long. Once done you can start over!
- There is no way to block write access to an SE as VO admin
- Some files were inaccessible due to directory permissions (mostly on LFC). Some files belonged to users that were no longer member of the VO: to overcome this, SRM administrator involvement is required
- Most files that were protected on the LFC were not protected on the SE
- Luckily this was a DPM SE and we could use the 'dpns-*' commands: using standard 'srm-*' commands is much slower/too slow

Recommendations

- If you have multiple LFCs always make sure they're in sync → **Dark Storage**
- Following tools would be useful
 - lfc-find, lfc-chown
 - dpns-find, dpns-chown
 - srm-* equivalents
- For the end-users and VO admins:
 - Run periodic checks on your files! do ***you*** know what you have stored and where?

Conclusions

- Migrating data away from a Storage Element that is to be decommissioned **IS** possible without admin rights on the SE (with 1 caveat)
- Such a migration does require knowledge of both the LFC/GFAL command-line tools and a scripting language (i.c. 'bash')
- A few extra commands would help enormously to make the migration smoother

Final remark

- Data management is too easy to do wrong (at any scale)



J.J. Keijser
Nikhef
Amsterdam
Grid Group