


GridPP
UK Computing for Particle Physics

Disk Pool Manager

Greig A Cowan

University of Edinburgh






GridPP
UK Computing for Particle Physics


Outline

1. Site requirements
2. DPM Architecture
3. Installation
4. Usage
5. Publishing in the information system
6. Classic SE → DPM migration
7. Upgrading to a new release
8. Summary




Requirements of GridPP2

- Storage Resource Manager (SRM)
 - SRM interface presents a combined view of a mass storage systems secondary and tertiary storage to grid clients.
 - *Permanent, durable* and *volatile* file types allows for files to be safely stored while they are needed and removed when not.
- Disk pool management
 - A pool is a group of file systems located on one or more disk servers.




Requirements of GridPP2

- Storage Resource Manager (SRM)
 - SRM interface presents a combined view of a mass storage systems secondary and tertiary storage to grid clients.
 - *Permanent, durable* and *volatile* file types allows for files to be safely stored while they are needed and removed when not.
- Disk pool management
 - A pool is a group of file systems located on one or more disk servers.
- Classic SE **not** sufficient.



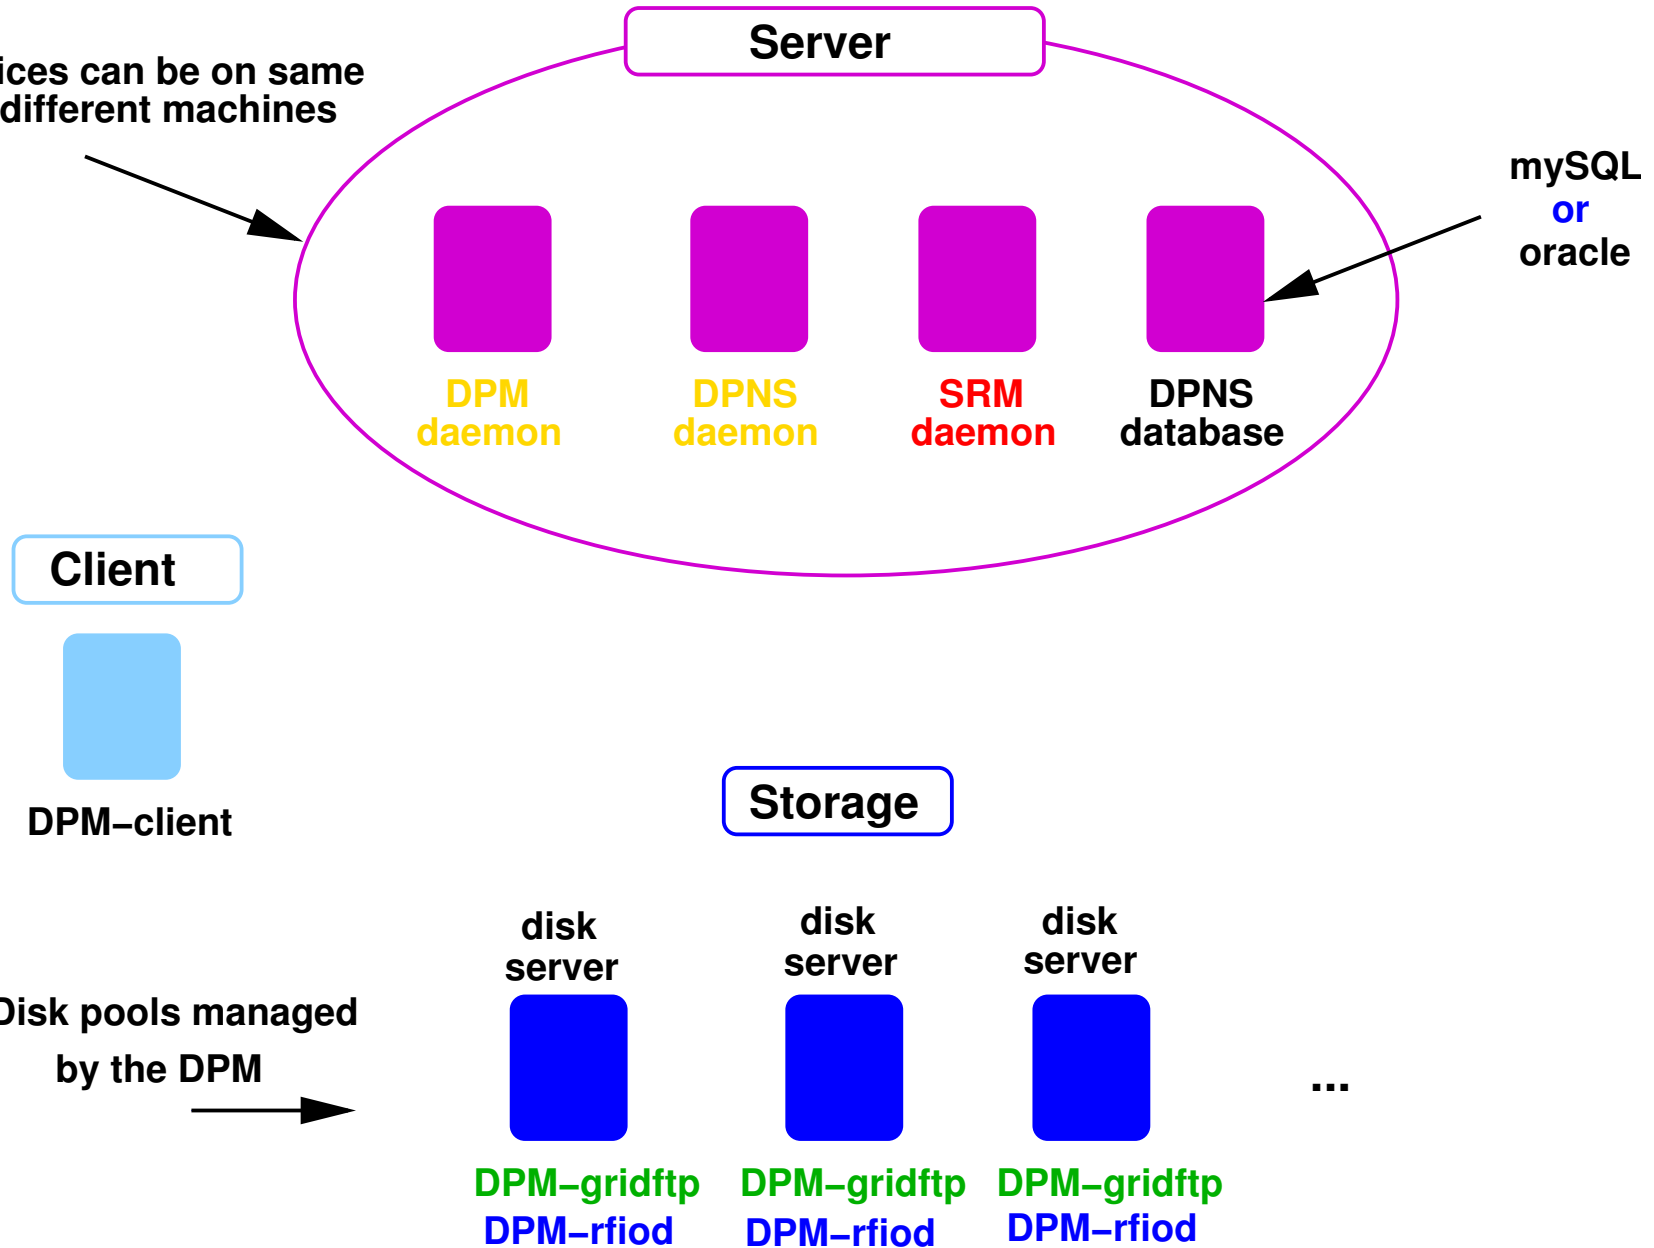
Requirements of GridPP2


- Storage Resource Manager (SRM)
 - SRM interface presents a combined view of a mass storage systems secondary and tertiary storage to grid clients.
 - *Permanent, durable* and *volatile* file types allows for files to be safely stored while they are needed and removed when not.
- Disk pool management
 - A pool is a group of file systems located on one or more disk servers.
- Classic SE **not** sufficient.
- What other solutions exist that provide this functionality?



Architecture of DPM


Services can be on same or different machines






Architecture of DPM

- DPM, DPNS and SRM servers can either be installed on the same machine or on different ones.
- If using mySQL database, it is recommended to install on the same machine as the DPM/DPNS server. Oracle database likely to be provided by institute, so different machine. Make *backups* of DPNS.
- DPM client has to be installed on the SRM servers machine, and should be installed on WNs and UIs. It should be installed on the DPM and DPNS machine(s) as well.
- DPM Gridftp and RFIO servers have to be installed on each disk server managed by the DPM.



- Ports that must be opened in firewall:

DPM server	tcp	5015	
DPNS server	tcp	5010	← same as LFC
SRM server v1	tcp	8443	
SRM server v2	tcp	8444	
RFIO server	tcp	5001	
GridFTP server	tcp	2811	← control port
GridFTP server	tcp	GLOBAL_TCP_PORT_RANGE	← transfer port range



Installation of DPM

- Well integrated with YAIM.

- `site-info.def`

`$ MY_DOMAIN` → your domain

`$ SE_HOST` → the DPM hostname


`$ SE_TYPE` → must be `SE_TYPE=srm_v1`

`$ DPMDATA` → the file system on the separate partition (ex: /storage)

`$ DPMUSER_PWD` → the password for the Mysql DPM user account

`$ DPMPOOL` → the name of the pool including the DPMDATA file system (ex: Permanent)

`$ DPMFSIZE` → the space to be reserved by default for a file stored in the DPM (ex: 200M)



Installation of DPM

- Install node

```
/opt/lcg/yaim/scripts/install_node /path/to/site-info.def lcg-SE_dpm_mysql
```

- Configure node

```
/opt/lcg/yaim/scripts/configure_node /path/to/site-info.def SE_dpm_mysql
```

- Should give a simple setup that will be suitable to most Tier-2s.
- If DPM server and disk pools on different subnets, need to create/modify `/etc/shift.localhosts`.

http://wiki.gridpp.ac.uk/wiki/DPM_Yaim_Install

- Setup environment

```
$ export DPNS_HOST=dpm.epcc.ed.ac.uk
```

```
$ export DPM_HOST=dpm.epcc.ed.ac.uk
```

- DPNS commands

```
$ dpns-ls -l /dpm/epcc.ed.ac.uk/home
```

```
$ dpns-mkdir /dpm/epcc.ed.ac.uk/home/dteam/test_dir
```

- Modify DPM setup


```
$ dpm-qryconf
```

```
$ dpm-addpool --poolname Vol --def_filesize 20M --s_type V
```

```
$ dpm-modifyfs --server sehost --fs /dpm-data --st RDONLY
```

- RFIO

```
$ rfdir dpm.epcc.ed.ac.uk:/tmp
```




Publishing DPM storage

- GIP plugin provided (and being developed) by Graeme Stewart:

<http://www.physics.gla.ac.uk/~graeme/scripts/>


- Easy to install, just place the script in the correct location.
- Now publishes correct storage values per VO (if you have set DPM up to map VOs to particular pools).
- This should get included in the next official LCG release.



Classic SE migration

- What do you do with your Classic SE machine?
 1. Convert it to a DPM head node
 2. Convert it to a DPM pool node (that gets added into an existing DPM setup)
 3. Decommission node completely (announce to VO, give time for them to move data from it, switch off)
- Benefit of 1,2: sfn:// LFNs are still valid when talking to DPM's gridftp server.


http://wiki.gridpp.ac.uk/wiki/Classic_SE_Migration



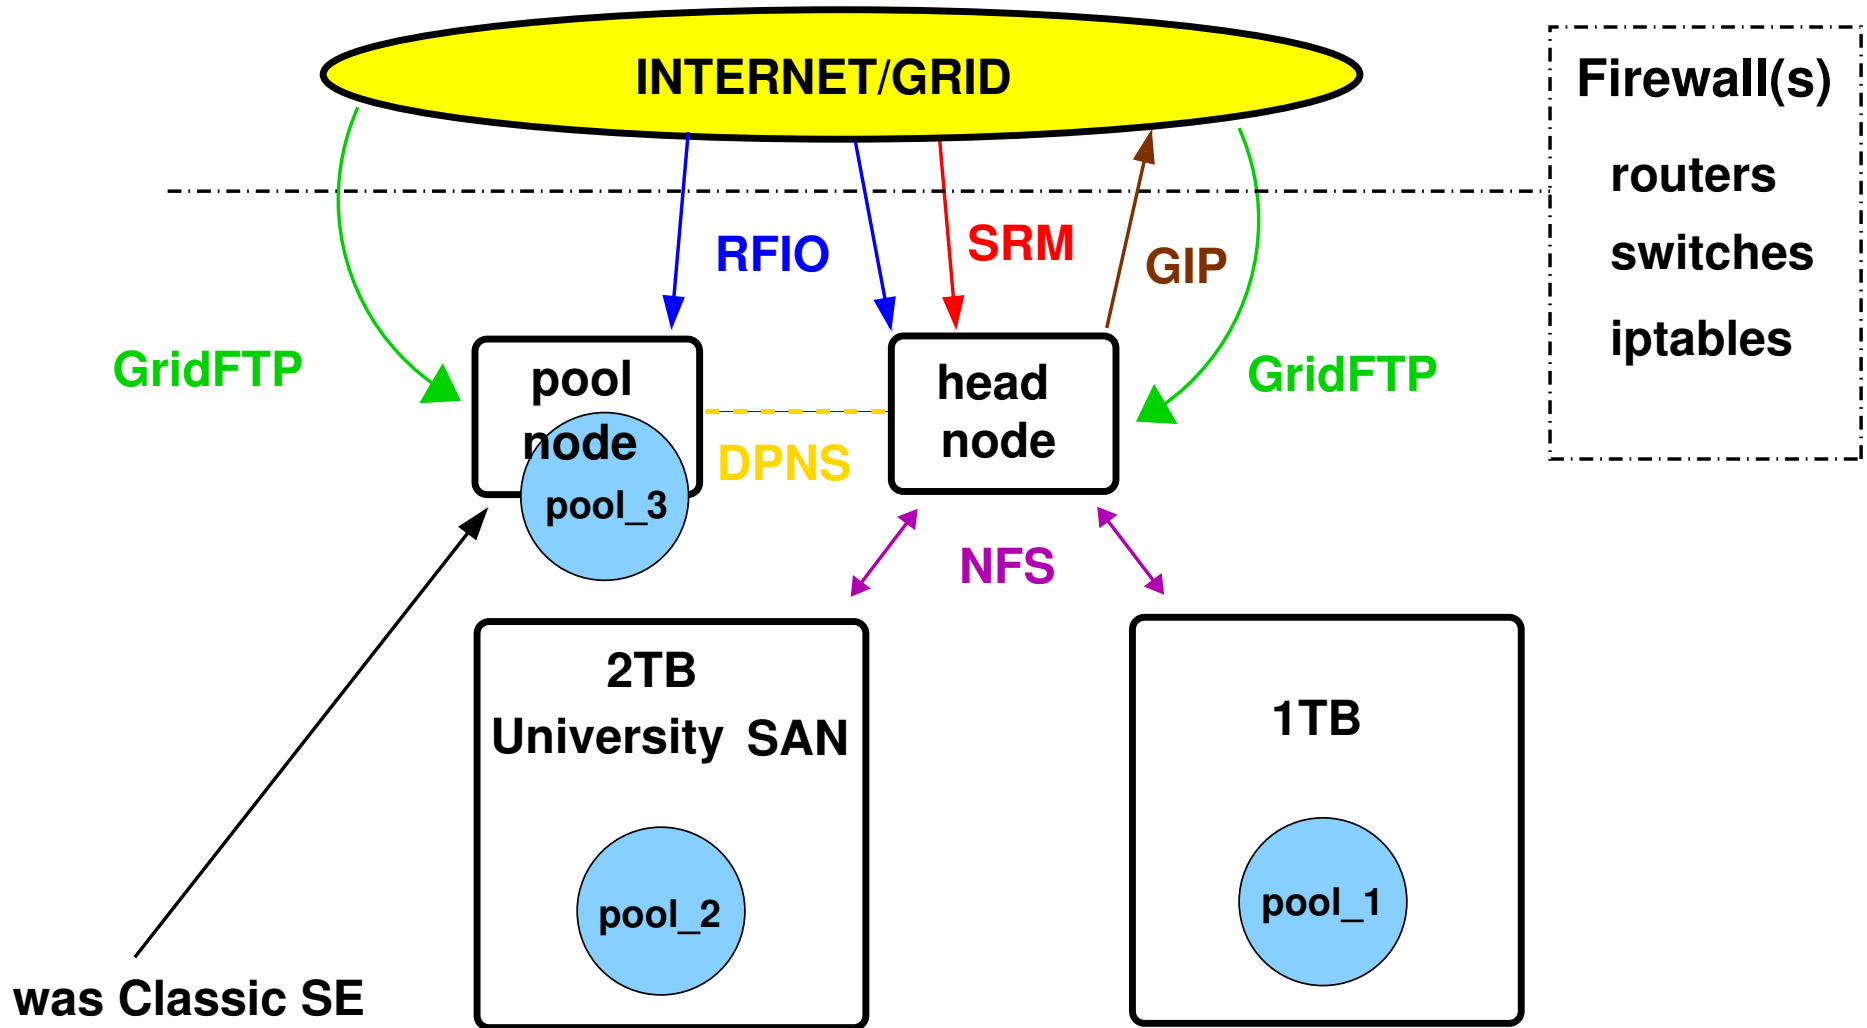
Upgrading DPM


- DPM has (fairly) regular updates.
- v1.4.1 released last week - bug fixes, perl/python API bindings. srmCopy still not working.
- Just upgrade the relevant rpms. In addition to the DPM rpms, you will also need the latest version of `lcg-dm-common`.
- Configuration will not be touched after upgrade.
- Performed last week on Edinburghs test machine. No problems observed with SRM functionality but further testing required (adding pool nodes).

http://wiki.gridpp.ac.uk/wiki/DPM_Upgrades




Edinburgh DPM setup





- SC4 requires all sites to have an SRM interface to storage.
- This needs to be in place by end of January to meet milestones.
- Tier-2's need to decide if they want to use dCache or DPM.
 - Within GridPP, 10 sites using DPM (5 using dCache).
 - ScotGrid, SouthGrid and London.
 - ~ 33TB of disk space (150TB total).
- Shown repeatedly to be relatively easy to install and administer.
- srmCopy still not available, though should be coming soon.



GridPP
UK Computing for Particle Physics

Storage Support

- Provided by the GridPP storage community.
 - Mailing list
 - Weekly (half hour!) phone conferences
 - GridPP wiki

http://wiki.gridpp.ac.uk/wiki/Disk_Pool_Manager

- DPM developers
 - `hep-service-dpm@cern.ch`
 - twiki admin guide