

GridPP

UK Computing for Particle Physics

Maximising the Terabytes in the UK


Philip J Clark, Greig A Cowan, Jamie K Ferguson

University of Edinburgh

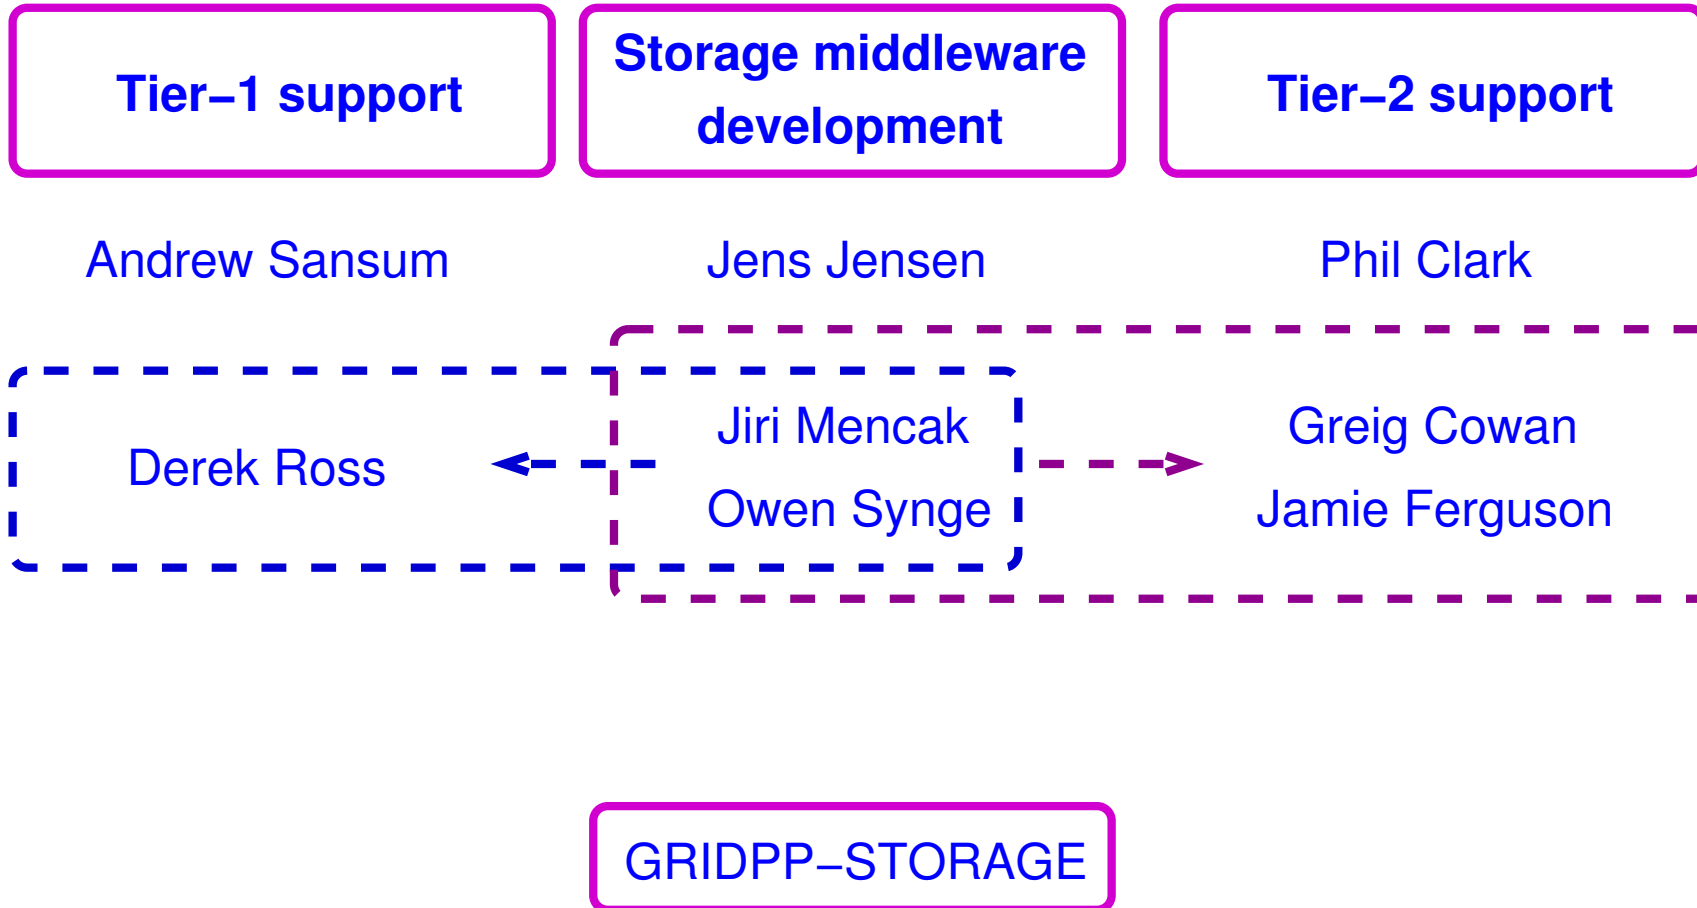



Jens Jensen, Jiri Mencak, Derek Ross, Owen Synge

RAL



Storage personnel






GridPP
UK Computing for Particle Physics

Goals of tier-2 storage

What are we aiming to do?

1. Facilitate large scale tier-2 MC production.
2. Facilitate distributed UK physics analysis.

This requires the following storage goals:




Goals of tier-2 storage

What are we aiming to do?

1. Facilitate large scale tier-2 MC production.
2. Facilitate distributed UK physics analysis.


This requires the following storage goals:

- Make the maximum amount of storage available that the UK can.
- The ability to remove unnecessary files to make sure the available space is always high.
- The ability to dynamically allocate/reserve storage space.




Requirements of GridPP2

- Storage Resource Manager (SRM) v1.1
 - 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) v1.1
 - 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) v1.1
 - 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?




GridPP
UK Computing for Particle Physics

dCache



dCache is a disk pool management system, jointly developed by DESY and Fermilab. Currently being used as the SRM (v1.1) solution for sites wanting a tried and tested platform, particularly if they have tape eg. RAL.

Features



GridPP
UK Computing for Particle Physics


dCache



dCache is a disk pool management system, jointly developed by DESY and Fermilab. Currently being used as the SRM (v1.1) solution for sites wanting a tried and tested platform, particularly if they have tape eg. RAL.

Features

- Single namespace used to combine several hundred pool nodes (disk servers) \Rightarrow able to transparently manage 100s of TBs of data.



GridPP
UK Computing for Particle Physics


dCache



dCache is a disk pool management system, jointly developed by DESY and Fermilab. Currently being used as the SRM (v1.1) solution for sites wanting a tried and tested platform, particularly if they have tape eg. RAL.

Features

- Single namespace used to combine several hundred pool nodes (disk servers) \Rightarrow able to transparently manage 100s of TBs of data.
- Supports multiple internal and external copies of a single file \Rightarrow redundancy \Rightarrow commodity disk storage components.



GridPP
UK Computing for Particle Physics


dCache



dCache is a disk pool management system, jointly developed by DESY and Fermilab. Currently being used as the SRM (v1.1) solution for sites wanting a tried and tested platform, particularly if they have tape eg. RAL.

Features

- Single namespace used to combine several hundred pool nodes (disk servers) \Rightarrow able to transparently manage 100s of TBs of data.
- Supports multiple internal and external copies of a single file \Rightarrow redundancy \Rightarrow commodity disk storage components.
- Pool to pool transfer upon detection of pool hotspots.



GridPP
UK Computing for Particle Physics


dCache



dCache is a disk pool management system, jointly developed by DESY and Fermilab. Currently being used as the SRM (v1.1) solution for sites wanting a tried and tested platform, particularly if they have tape eg. RAL.

Features

- Single namespace used to combine several hundred pool nodes (disk servers) \Rightarrow able to transparently manage 100s of TBs of data.
- Supports multiple internal and external copies of a single file \Rightarrow redundancy \Rightarrow commodity disk storage components.
- Pool to pool transfer upon detection of pool hotspots.
- Supports multiple access points (doors) on the pool nodes \Rightarrow scalable.



GridPP
UK Computing for Particle Physics


dCache



dCache is a disk pool management system, jointly developed by DESY and Fermilab. Currently being used as the SRM (v1.1) solution for sites wanting a tried and tested platform, particularly if they have tape eg. RAL.

Features


- Single namespace used to combine several hundred pool nodes (disk servers) \Rightarrow able to transparently manage 100s of TBs of data.
- Supports multiple internal and external copies of a single file \Rightarrow redundancy \Rightarrow commodity disk storage components.
- Pool to pool transfer upon detection of pool hotspots.
- Supports multiple access points (doors) on the pool nodes \Rightarrow scalable.
- Administration via ssh and GUI. Monitoring via http.



GridPP
UK Computing for Particle Physics

Doors


dCache has 3 supported access methods, all doors are GSI enabled.



dCache has 3 supported access methods, all doors are GSI enabled.

1. Native dCache access protocol (dcap) for use within the dCache system.

Uses port [22128](#).




GridPP
UK Computing for Particle Physics

Doors

dCache has 3 supported access methods, all doors are GSI enabled.


1. Native dCache access protocol (dcap) for use within the dCache system.
Uses port [22128](#).
2. GridFTP server on port [2811](#) (interoperable with Globus).



Doors


dCache has 3 supported access methods, all doors are GSI enabled.

1. Native dCache access protocol (dcap) for use within the dCache system.
Uses port [22128](#).
2. GridFTP server on port [2811](#) (interoperable with Globus).
3. SRM v1.1 on port [8443](#).




dCache has 3 supported access methods, all doors are GSI enabled.

1. Native dCache access protocol (dcap) for use within the dCache system.
Uses port [22128](#).
2. GridFTP server on port [2811](#) (interoperable with Globus).
3. SRM v1.1 on port [8443](#).
 - supports queueing for get, put and copy requests.




dCache has 3 supported access methods, all doors are GSI enabled.

1. Native dCache access protocol (dcap) for use within the dCache system.
Uses port [22128](#).
2. GridFTP server on port [2811](#) (interoperable with Globus).
3. SRM v1.1 on port [8443](#).
 - supports queueing for get, put and copy requests.
 - allows you to perform multiple transfers in a single command



dCache has 3 supported access methods, all doors are GSI enabled.

1. Native dCache access protocol (dcap) for use within the dCache system.
Uses port **22128**.
2. GridFTP server on port **2811** (interoperable with Globus).
3. SRM v1.1 on port **8443**.
 - supports queueing for get, put and copy requests.
 - allows you to perform multiple transfers in a single command
 - optimises the parameters of the transfer and allows FTP to scale with user load.



dCache has 3 supported access methods, all doors are GSI enabled.

1. Native dCache access protocol (dcap) for use within the dCache system.

Uses port [22128](#).

2. GridFTP server on port [2811](#) (interoperable with Globus).


3. SRM v1.1 on port [8443](#).

- supports queueing for get, put and copy requests.

- allows you to perform multiple transfers in a single command


- optimises the parameters of the transfer and allows FTP to scale with user load.

- third party transfers allowed.



GridPP
UK Computing for Particle Physics

Documentation/Support



GridPP
UK Computing for Particle Physics


Documentation/Support

- GridPP admin support pages: ← installation, FAQ, usage ...

<http://www.gridpp.ac.uk/deployment/admin/dcache/index.html>

- GRIDPP-STORAGE archive: ← HIGHLY RECOMMENDED!!

<http://www.jiscmail.ac.uk/archives/gridpp-storage.html>




- GridPP admin support pages: ← installation, FAQ, usage ...

<http://www.gridpp.ac.uk/deployment/admin/dcache/index.html>

- GRIDPP-STORAGE archive: ← HIGHLY RECOMMENDED!!

<http://www.jiscmail.ac.uk/archives/gridpp-storage.html>


- RAL storage pages: <http://storage.esc.rl.ac.uk/>
- Weekly storage phone conference (contact Jens Jensen).
 - format of this may soon change to cater for site needs.
- Latest dCache releases: <http://www.dcache.org/>
- SRM specifications: <http://sdm.lbl.gov/srm-wg/index.html>



GridPP
UK Computing for Particle Physics

Disk Pool Manager (DPM)

DPM is an lightweight solution for disk storage management that has been developed at CERN. It currently offers an implementation of SRM v1.1 and v2. Requires the GSI security stack.



Disk Pool Manager (DPM)

DPM is an lightweight solution for disk storage management that has been developed at CERN. It currently offers an implementation of SRM v1.1 and v2. Requires the GSI security stack.


- Installation manual (GOC Wiki):

[http://goc.grid.sinica.edu.tw/gocwiki/ \](http://goc.grid.sinica.edu.tw/gocwiki/How_to_install_the_Disk_Pool_Manager_%28DPM%29)
[How_to_install_the_Disk_Pool_Manager_%28DPM%29](http://goc.grid.sinica.edu.tw/gocwiki/How_to_install_the_Disk_Pool_Manager_%28DPM%29)

- GridPP admin support pages: ← early stage of development

<http://www.gridpp.ac.uk/deployment/admin/dpm/index.html>

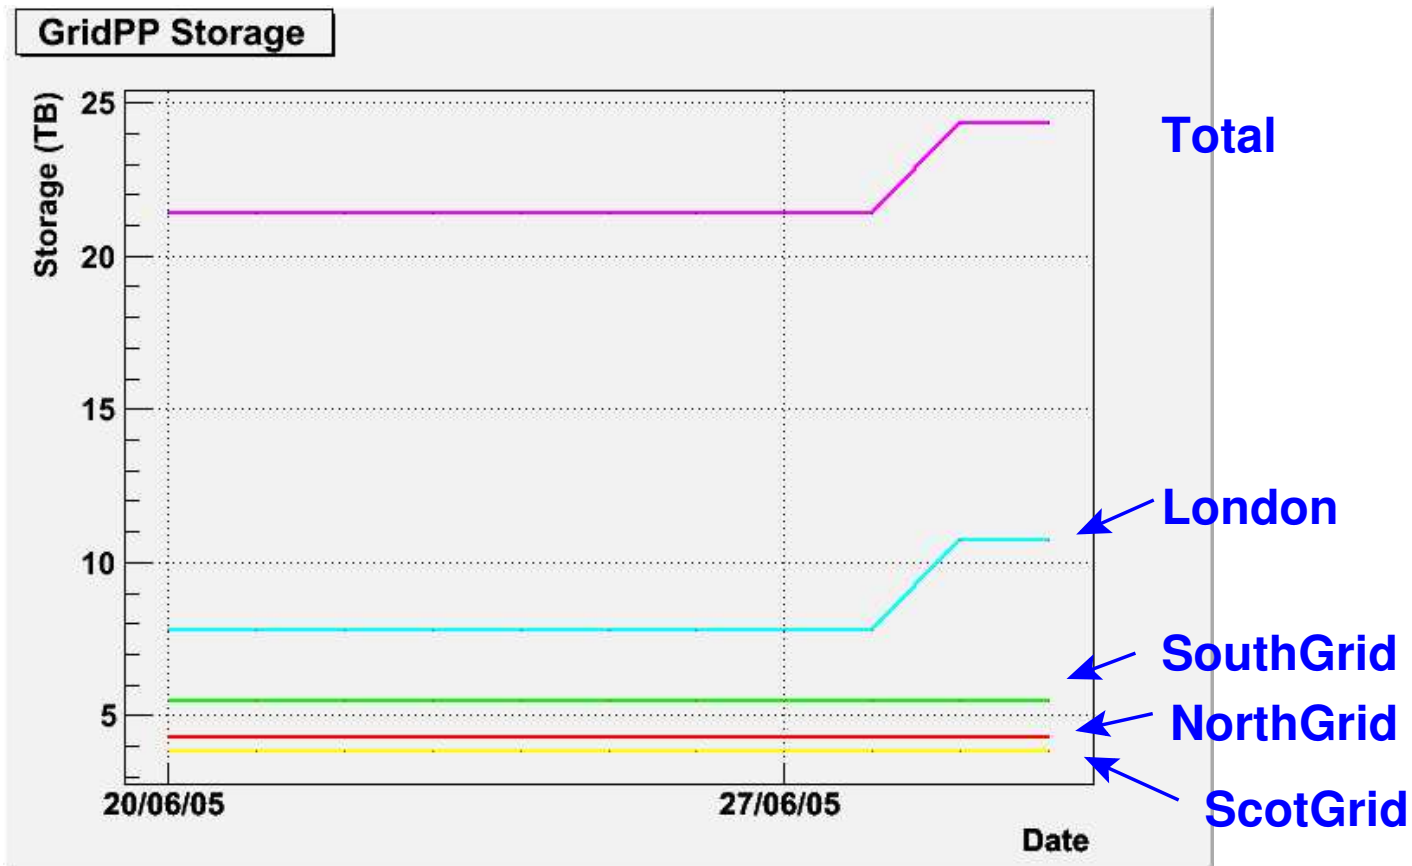
- DPM is currently being evaluated by RAL, Glasgow and Edinburgh.




GridPP storage

- Deployment status:

<http://www.ph.ed.ac.uk/~jfergus7/gridppDiscStatus.html>




- Targets need to be met.




GridPP

Location	Target Capacities (TB)		Current Capacities (TB)				Classic SE/SRM	FileSpace Type
	2004	2007	Capacity	Available	Used	Used (%)		
ScotGrid								
Durham	5	5	1.915	1.912	0.003	0	Classic SE	permanent
Edinburgh	24	70.5	0.034	0.021	0.013	38	Classic SE, SRM	permanent
Glasgow	12.3	14.8	1.873	0.955	0.918	49	Classic SE	permanent
ScotGrid Totals	41.3	90.3	3.822	2.888	0.934	24		
NorthGrid								
Lancaster	4.9	86.7	1.995	1.881	0.114	5	Classic SE	permanent
Liverpool	0.8	80.3	0.009	0.002	0.007	77	Classic SE	permanent
Manchester	12.6	372.6	0	0	0	0	SRM	volatile
Sheffield	0	3	2.315	2.146	0.169	7	Classic SE	permanent
NorthGrid Totals	18.3	542.6	4.319	4.029	0.29	6		
SouthGrid								
Birmingham*	9	9.3	0	0	0	0	Classic SE	permanent
Bristol	1.9	1.9	0.064	0.063	0.001	1	Classic SE	permanent
Cambridge	3.5	4.4	2.006	2.006	0	0	Classic SE	permanent
Oxford*	9.5	24.5	3.21	3.126	0.084	2	Classic SE	permanent
RAL PPD*	0.8	17.4	0.223	0.113	0.11	49	Classic SE	permanent
SouthGrid Totals	24.7	57.5	5.503	5.308	0.195	3		
London								
Brunel University	1	21	0.437	0.43	0.007	1	Classic SE	permanent
Imperial College	13.6	93.3	0.31	0.288	0.022	7	Classic SE, SRM	permanent
Queen Mary University	28.5	58.5	0.108	0.069	0.039	36	Classic SE	permanent
Royal Holloway University	3.2	23.2	8.772	8.241	0.531	6	Classic SE	volatile
University College	0	0.7	1.073	0.855	0.218	20	Classic SE	volatile
London Totals	46.3	196.7	10.733	9.913	0.82	7		
GridPP Totals	130.6	887.1	24.344	22.108	2.236	9		



GridPP
UK Computing for Particle Physics


Conclusions



GridPP
UK Computing for Particle Physics

Conclusions


- dCache is mature and been tested. Yaim install recommended.



GridPP
UK Computing for Particle Physics


Conclusions

- dCache is mature and been tested. Yaim install recommended.
- dCache is the SRM solution of choice for SC3 Tier-2 sites.




Conclusions

- dCache is mature and been tested. Yaim install recommended.
- dCache is the SRM solution of choice for SC3 Tier-2 sites.
- Testing currently underway with those sites involved in SC3 (T1 ↔ T2 transfers).




Conclusions

- dCache is mature and been tested. Yaim install recommended.
- dCache is the SRM solution of choice for SC3 Tier-2 sites.
- Testing currently underway with those sites involved in SC3 (T1 ↔ T2 transfers).
- This will provide us with valuable information when we it comes to rolling dCache out at all sites:
 - firewall configurations
 - hardware requirements and optimal configurations
 - updates




Conclusions

- dCache is mature and been tested. Yaim install recommended.
- dCache is the SRM solution of choice for SC3 Tier-2 sites.
- Testing currently underway with those sites involved in SC3 (T1 ↔ T2 transfers).
- This will provide us with valuable information when we it comes to rolling dCache out at all sites:
 - firewall configurations
 - hardware requirements and optimal configurations
 - updates
- DPM is alternative solution but not as thoroughly tested. Knowledge base not as good. Some sites may migrate to DPM in future?




GridPP
UK Computing for Particle Physics

dCache: HowTo get started




dCache: HowTo get started

1. Basic install of SL3 + apt
2. Place host certificates in `/etc/grid-security`
3. Install java (`j2sdk 1.4.2`)
4. Point system to the LCG/CA rpm repositories
5. Install `yaim`




dCache: HowTo get started

1. Basic install of SL3 + apt
2. Place host certificates in `/etc/grid-security`
3. Install java (`j2sdk 1.4.2`)
4. Point system to the LCG/CA rpm repositories
5. Install `yaim`
6. Configure `site-info.def` file. Most importantly:
`DCACHE_ADMIN=<FQDN of admin node>`
`DCACHE_POOLS=<FQDN of pool node>:/path/to/pool`



GridPP
UK Computing for Particle Physics

dCache: HowTo get started




dCache: HowTo get started

7. Use the yaim install target lcg-SEDCache

```
/opt/lcg/yaim/scripts/install_node \  
/root/yaim/site-info.def lcg-SEDCache
```

8. Run the YAIM configure script for SE_dcach

```
/opt/lcg/yaim/scripts/configure_node \  
/root/yaim/site-info.def SE_dcach
```



dCache: HowTo get started

7. Use the yaim install target lcg-SEDCache

```
/opt/lcg/yaim/scripts/install_node \  
/root/yaim/site-info.def lcg-SEDCache
```

8. Run the YAIM configure script for SE_dcach

```
/opt/lcg/yaim/scripts/configure_node \  
/root/yaim/site-info.def SE_dcach
```

The yaim install scripts must be run on both the admin and pool nodes.

Yaim can figure out which components to install from the values specified in `site-info.def`.