



QCDgrid Administration

James Perry
Applications Consultant, EPCC
j.perry@epcc.ed.ac.uk
+44 131 650 5173

- QCDgrid components and node types
- Installation
- General administrative tasks
- Certificates and security
- Dealing with common problems

Three fundamental components:

- Data grid software
 - Globus-based
 - client tools
 - replica catalogue database server

- Metadata catalogue
 - Java-based client
 - XML database server

- Job submission software
 - Globus-based
 - simple client tools

- Client node
 - users access the grid from here but no files hosted
 - needs Globus client installation and QCDgrid software installation
- Storage node
 - stores files for the data grid (may also be a client)
 - needs Globus server installation and QCDgrid software installation, and some special configuration
- Central node
 - special storage node which also hosts the replica catalogue and control thread
 - needs Globus server installation with RLS database and QCDgrid software installation
 - may also host metadata catalogue database and/or security database
 - control thread needs to be running all the time

First step is to install dependencies

- Globus and EDG components can be installed via the Virtual Data Toolkit, available from <http://www.cs.wisc.edu/vdt/>
 - simple installation for common cases
 - central/backup nodes need RLS configured, this extra step is described in the VDT installation guide
 - if QCDgrid software is to be built from source, additional package containing Globus header files is required (also covered in VDT installation guide)
- Host certificate must be obtained (not needed for client machines)
 - available from <http://ca.grid-support.ac.uk/>

- Metadata catalogue requires Java 1.4+ (<http://java.sun.com/>), and eXist XML database (<http://exist.sourceforge.net/>)
 - client install requires libraries from eXist
 - server install requires functioning eXist database, running within Tomcat web application server (<http://jakarta.apache.org/tomcat/>)
 - installation usually straightforward and well documented
- Pooled accounts system may be installed to ease user administration
 - requires Globus patches, these are included with VDT
 - requires some EDG software. Installation documented in QCDgrid security guide

- QCDgrid software may now be installed on top
 - source code available from anonymous CVS at <http://forge.nesc.ac.uk/projects/qcdgrid>
 - official, versioned packages soon to be available
- Client configuration
 - two simple config files must be created:
 - `nodes.conf` – points to the central and backup nodes
 - `nodeprefs.conf` – lists data storage nodes in order of preference
 - Globus installation must trust UK e-Science certificates (works out-of-the-box with VDT)
 - QCDgrid/Globus commands must be in users' `PATH`

In addition to client node setup:

- **Globus/security setup**
 - host certificate must be installed
 - user certs and central node cert must be in gridmap file (either via pooled account system or manually)
 - Globus gatekeeper and GridFTP server must be running
- **Data storage directory setup**
 - symlink the storage directory on RAID to `data/` in the `qcdgrid` installation directory
 - on multidisk machines, symlink the directories to `data/`, `data1/`, `data2/`, etc., and inform a grid administrator that the node has multiple disks
 - create a globally writable `NEW/` directory (the “inbox”) in the storage directory

In addition to storage node setup:

- RLS database must be running
- Control thread must be running
- eXist database for metadata catalogue may need to be running
- User directory for pooled account system may need to be running
- Additional configuration files `mainnodelist.conf` and `qcdgrid.conf` must be created

- Normally just a case of compiling/installing new version of tools over old version
- May be prudent to disable node while upgrade is in progress
- Upgrades normally don't break backwards compatibility
- Proper versioning system in place for future releases

- Adding/removing users
- Adding new nodes
- Retiring old nodes
- Temporary disabling nodes
- Replica catalogue maintenance

- Users are handled at the Globus/UNIX level, not by QCDgrid directly
 - each user's certificate subject must be mapped to a suitable account on each grid storage node
 - may be done manually by sysadmin at each site (requires root permissions to edit grid-mapfile)

Example manually created grid-mapfile:

```
"/C=UK/O=eScience/OU=Liverpool/L=CSD/CN=steve downing" steve
"/C=UK/O=eScience/OU=Edinburgh/L=NeSC/CN=james perry" jamesp
"/C=UK/O=eScience/OU=Liverpool/L=CSD/CN=craig mcneile" mcneile
"/C=UK/O=eScience/OU=Liverpool/L=CSD/CN=alan irving" aci
"/C=UK/O=eScience/OU=Liverpool/L=Physics/CN=paul rakow" rakow
"/C=UK/O=eScience/OU=Edinburgh/L=NeSC/CN=balint joo" bj
"/C=UK/O=eScience/OU=Edinburgh/L=NeSC/CN=trumpton.ph.ed.ac.uk/Email=cmaynard@ph.ed.ac.uk" qcdgrid
```

- Pooled account system generates each gridmap file automatically from a central database
- Users only need to be added once to the central database
- A script for administering the database is provided

Example script usage:

```
./setup-security.sh adduser James Perry \  
  "/C=UK/O=eScience/OU=Edinburgh/L=NeSC/CN=james perry" \  
  jamesp@epcc.ed.ac.uk
```

```
./setup-security.sh rmuser James Perry
```

- Commands exist for adding, removing, enabling, disabling nodes, e.g.:

```
add-qcdgrid-node pygrid1.swan.ac.uk Swansea /home/qcdgrid  
disable-qcdgrid-node dylan.amtp.liv.ac.uk  
enable-qcdgrid-node ukgrid0.phys.columbia.edu
```

- Nodes should be retired for a while before removal, to give the system time to duplicate their data elsewhere

```
retire-qcdgrid-node trumpton.ph.ed.ac.uk
```

- Hardware failures can cause catalogue to become inconsistent
- `verify-qcdgrid-rc` command scans for inconsistencies and prompts the user for whether to correct them
- `rebuild-qcdgrid-rc` rebuilds catalogue from scratch based on what's actually on the nodes (can recover catalogue if it's lost or severely corrupted)
- `create-qcdgrid-rc` and `delete-qcdgrid-rc` also available

- Certificate issues
 - expiry
 - users missing from grid map files
 - trust
- Firewall issues
- Miscellaneous problems
 - nodes not responding
 - checksum failures

- Many problems are caused by certificates
- Expiry
 - certificates expire every year and must be renewed
 - if a host certificate expires, that machine becomes inaccessible to the grid until it is renewed
 - UK e-Science CA send out reminders when certificate is due to expire
 - expiry time can also be checked using Globus:

```
grid-cert-info -file /etc/grid-security/hostcert.pem
```

- Users missing from grid map files
 - each user must be in every grid map file
 - this must be done manually by the sysadmin at sites without pooled accounts
 - failure to do so will result in error 530 “No local mapping for Globus ID” when user attempts to access grid
- Trust
 - all machines (including clients) must trust all certificate authorities used on grid
 - currently just the UK e-Science CA
 - will happen automatically if installed using VDT, otherwise must obtain CA certificate and signing policy files and install them manually
 - ILDG may necessitate trusting additional CAs

- Second major cause of grid problems
- Usually manifests as connection failures or time-outs
- Storage nodes must have Globus TCP ports open to all grid machines, including clients
 - gatekeeper (usually port 2119)
 - GridFTP server (usually port 2811)
 - dynamic port range (usually 65000-65255 or 50000-52000)
 - additionally RLS server port, control thread port (and possibly MDC and user directory ports) on central node only
- Client nodes must have dynamic port range open
- `GLOBUS_TCP_PORT_RANGE` environment variable must be set, both for client tools and Globus services

- Node not responding
 - firewall issue? (see previous slide)
 - certificate issue? (see earlier slide)
 - Globus services not running?
 - clock skew?
- Checksum failure in file transfer
 - corrupted RAID disk? check and replace
 - false alarm? (software error reporting to be improved to eliminate this)

Grid administration tasks include:

- Installing software and dependencies
- Upgrading software
- Configuring new nodes
- Adding/removing users
- Adding/removing nodes
- Dealing with problems (e.g. certificate and firewall issues)

- <http://forge.nesc.ac.uk/projects/qcdgrid>
- <http://www.globus.org/>
- <http://exist.sourceforge.net/>
- <http://www.cs.wisc.edu/vdt/>