

Troubleshooting QCDgrid software

Some problems and solutions

Daragh Byrne
Applications Consultant, EPCC
d.byrne@epcc.ed.ac.uk
+44 131 651 3395

- Present some common problems encountered
- Suggest some diagnostic techniques
- Suggest some solutions

This is just our point of view, please feel free to contribute
any ideas to the group!

- QCDgrid is a fairly sophisticated distributed system
 - Multiple components
 - Core
 - Third party
 - Dependence on networks
 - Multiple administrative domains
- It can be hard to diagnose problems
- Hope to provide a checklist that will help
- QCDgrid software problem categories:
 - Configuration problems
 - QCDgrid
 - Environment
 - Globus problems
 - Firewall problems

- How do you know when there is a problem with QCDgrid?
 - Error messages from client tools
 - Command line tools
 - Browser
 - Emails from the control node
- UKQCD installation:
 - Typically bug reports are submitted to NeSCForge
 - Picked up and dealt with by the QCDgrid project team
 - Effort part of the QCDgrid project
 - Can be reasonably time consuming
 - Hopefully sharing this knowledge can help lighten the load!
 - Obviously support will be different if you have your own setup

- `qcdgrid-ping` fails:
 - `%> qcdgrid-ping`
"No response from control thread"
(or hangs)
 - Control thread actually down
 - Other configuration problem stops client from contacting
- `qcdgrid-list` fails
 - Usually just hangs without an error message
- `put-file-on-qcdgrid` fails
 - Variety of errors encountered
 - Could not contact main node
 - Could not connect to metadata catalogue
 - Could not connect to RLS server
 - Could not place file on grid
 - Provides information about the destination node

- **get-file-from-qcdgrid fails**

Error in copy to local: operation timed out

Globus error in data transfer: the operation was aborted

- **Browser fails**

- A variety of ways possible

- “Could not connect to database” most common

- Reasonably confident in the state of the code base
 - Running for three years
 - Thorough development process
 - Time enough to iron out quite a few bugs
- The experience of the team is that most problems are to do with configuration at one of three levels:
 - QCDgrid software configuration
 - Globus configuration
 - Firewall/network configuration
- Typically the problem exists between two particular nodes of the grid
 - Machines can usually be inferred from the error messages
- Although that doesn't rule out software bugs totally!
 - The idea is to rule out common errors before raising a bug

- Check all server components are actually running
 - Control thread process
 - Replica Location Service
 - XML database server
- Sometimes the error messages give a clue
 - Commands not running correctly indicates environment problems
 - Typically Globus problems are flagged in output
 - Hanging/timed out code or connection errors usually indicate firewall/network issues
- Start to diagnose each category of problem
 - Globus
 - Environment
 - Firewall/network

Troubleshooting Globus

- QCDgrid depends on several Globus services
 - GRAM (Job submission)
 - GridFTP (file transfer)
 - GSI (security)
- All of these need to be working correctly for the software to operate successfully
- Several steps that can be taken to diagnose problems with each of these services

- The simplest way is to use `globusrun`
 - Exercises the underlying infrastructure
 - Provides reasonably clear diagnostic information
- Since the problem may be on client or server
 - Try to run a job from client to server:
 - `globusrun -o -r remote.server.name '&(executable=/bin/hostname)'`
 - Try to run a job from client to client
 - `globusrun -o -r localhost '&(executable=/bin/hostname)'`
 - If you have access to the remote machine, try running a job from remote machine to the client machine
- Globusrun should run smoothly
 - Output should simply be hostname of remote/local machine
 - If not, it could be a problem in either the GRAM or GSI Globus components
 - Follow the diagnostics for each of these presented later

- Globus logs can also be useful for diagnosing problems
- `$GLOBUS_LOCATION/var/globus-gatekeeper.log`
 - Main log file for most Globus activities
- Possible to turn on finer grained logging for GRAM
 - Set the `save_logfile` configuration option to `on_error` in `$GLOBUS_LOCATION/etc/globus-job-manager.conf`
 - A detailed log file will be placed in the home directory of the account that executes the job
 - It's possible to figure out which pooled account was used and what the name of the file is by looking at `globus-gatekeeper.log`

- Can usually use `globus-url-copy` command
 - Harnesses the underlying GridFTP infrastructure
 - Any problems with this command usually indicate an underlying issue
- Again, in order to ascertain location of error, must try permutations of machines
 - Copy a file from local to remote
 - Copy a file from remote to local
 - Copy a file from local to local

- A good first step is to try to copy mainodelist.conf from the control node
 - globus-url-copy
gsiftp://main.node.location/path/to/qcdgrid/software/mainodelist.conf
file:///tmp/nodelistcopy.txt
- A second step would be to copy a file from the client machine to the data directory of the remote node:
 - globus-url-copy file:///tmp/test.txt
gsiftp://remote.node/path/to/qcdgrid/software/data/NEW/TEST.txt
- Typical errors include:
 - Security issues, in which case see section on security troubleshooting
 - Unable to write to remote node, in which case there may be a file permissions problem with the remote node

- Host certificates (server and client)

- Check that it's in date
 - Use the `grid-cert-info -f` command
- Check the file permissions

```
-rw-r--r--      1 root      root          2109 Feb 13 11:08 hostcert.pem
-rw-----      1 root      root          1124 Feb 13 11:09 hostkey.pem
```

- User certificate

- Date again
- File permissions

```
-r--r--r--      1 daragh    users          1947 Jul 12 2005
  usercert.pem
-r-----      1 daragh    users          1190 Jul 12 2005
  userkey.pem
```

- UK eScience CA files
 - Need to ensure that CA certificate and signing policy files are present on all machines
 - UK eScience files are prefixed by: 01621954
 - i.e. 01621954.signing_policy, 01621954.0
- grid-mapfile entries
 - Certificate subject for user needs to be mapped to an account on the server
 - Typically an entry of the form

```
"/C=UK/O=eScience/OU=Edinburgh/L=NeSC/CN=daragh byrne" qcdgrid
"/C=UK/O=eScience/OU=Edinburgh/L=NeSC/CN=daragh byrne" .
```
 - Usually taken care of by pooled accounts system, although worth checking

- Check to ensure that the Globus services are set up to run correctly via xinetd
 - Check that `/etc/xinetd.d/gsiftp` and `/etc/xinetd.d/globus-gatekeeper` both exist and nothing odd about them
 - Check that both GridFTP and `globus-gatekeeper` are configured to run automatically via `/etc/services`

- RLS setup
 - Ensure that RLS is running

```
[daragh@qcdgrid1 daragh]$ ps -ax|grep rls
14887 ?          S      483:41 globus-rls-server
```

 - It is possible to set this up to run via xinetd

Diagnosing environment problems

- Necessary environment settings
 - Globus related
 - QCDgrid software related

- Globus related:
 - GLOBUS_HOSTNAME
 - LD_ASSUME_KERNEL
 - GLOBUS_TCP_PORT_RANGE

- The `globus-hostname` command should return the fully qualified domain name of the machine it's run on

```
[daragh@edqcdgrid daragh]$ globus-hostname
edqcdgrid.epcc.ed.ac.uk
```
- If not, this can be set via the `GLOBUS_HOSTNAME` environment variable
 - `export GLOBUS_HOSTNAME=fqdn.of.host`
- Can be done server side too
 - Edit `/etc/xinetd.d/gsiftp` to include the line

```
env+=GLOBUS_HOSTNAME=fqdn.of.host
```

- LD_ASSUME_KERNEL
 - Used to correct a problem in the threading libraries that causes bugs
 - Usually needs to be set to 2.2.4
 - There are dependencies on the version of `glibc` and the `kernel` version
- QCDgrid CLASSPATH
 - In the binary distribution needs to contain all JAR files in `~qcdgrid/exist_libraries`
 - Also should contain `~qcdgrid/java`
- LD_LIBRARY_PATH
 - Needs to contain `libqcdgrid.so` and `libQCDgridJavaClient.so` (located in `~qcdgrid/java/jni`)
- GLOBUS_TCP_PORT_RANGE
 - A range of ports that are dynamically used by Globus
 - Usually set to `50000,52000`

- We provide a setup script (customisable)
 - `example_setenv.sh`
- Should be possible to set up environment by using this script

Diagnosing network/firewall problems

- QCDgrid needs a number of ports open to operate correctly
 - 2811 – GridFTP – between all machines
 - 2119 – GRAM – between all machines
 - QCDgrid port as configured in qcdgrid.conf (typically 51001) – between clients and the control node
 - RLS port – typically 39281 – between clients and the node running the RLS
 - GLOBUS_TCP_PORT_RANGE – typically 50001,52000
- Can typically use telnet to diagnose firewall problems

```
[daragh@edqcdgrid daragh]$ telnet qcdgrid1.epcc.ed.ac.uk 2811
Trying 129.215.63.161...
Connected to qcdgrid1.epcc.ed.ac.uk (129.215.63.161).
Escape character is '^]'.
220 qcdgrid1.epcc.ed.ac.uk GridFTP Server 1.12 GSSAPI type Globus/GSI wu-
2.6.2 (gcc32dbg, 1069715860-42) ready.
```
- Can check dynamic port range (GLOBUS_TCP_PORT_RANGE) by starting a Globus GASS server on a port within this range trying to connect via telnet

Still have problems?

- Contact the QCDgrid team

<http://forge.nesc.ac.uk/projects/qcdgrid>

- Include all output/debug messages
- We'll try to get back to you!

- QCDgrid is a sophisticated distributed application
 - Can be difficult to debug
- The root cause of many problems is usually one of
 - Globus
 - Environment
 - Network/firewall
- Use the checklist we have provided to try to solve the problem
 - Otherwise get in touch.