




LCG2 Administrator's Course

Oxford University, 19th – 21st July 2004.

Testing a LCG-2 site


Piotr Nyczyk, CERN IT/GD

Developed in conjunction with GridPP and EGEE




GridPP
UK Computing for Particle Physics

EGEE
Enabling Grids for
E-science in Europe



Outline

- Tests goals
- Testing a UI
- Typical job submission problems (UI)
- Typical RM problems (UI/WN)
- Testing a whole site
- Typical problems
- GIIS/GRIS testing
- PBS batch system testing
- Common mistakes



Reference materials

- LCG-2 Install Notes – appendix F


[http://markusw.home.cern.ch/markusw/
LCG2InstallNotes.html#SECTION00013000000000000000](http://markusw.home.cern.ch/markusw/LCG2InstallNotes.html#SECTION00013000000000000000)

- GOC Wiki Trouble Shooting page

<http://goc.grid.sinica.edu.tw/gocwiki/TroubleShootingHistory>


- FAQ page maintained by Steve Traylen

<http://www.gridpp.ac.uk/tb-support/faq/>




Tests goals

- Hierarchical support structure:
Site – Primary Site(ROC) – Deployment Team
- Installation and configuration are complex
 - tests allow to spot problems and correct them quickly
 - since wrong configuration can affect other sites tests are important for stable operation.
- Learn and understand how it really works!
- Be aware of typical problems and solutions
- Maintain the site in a good condition during its operation – repeat tests, detect problems rapidly




Testing steps

- First test your UI using established sites
 - failures of complex tests are confined to the other services on your site (predictable tests)
- Basic tests of local services are possible even without BDII registration (hosts certificates required):
 - Site GIIS and InfoProviders
 - Globus gatekeeper and fork jobmanager
 - Local batch system (OpenPBS, LSF)
 - Globus job manager with batch system
- Testing simple job submission using RB (BDII registration required)
- Testing a whole site (job submission and data management) with site certification test suite
- **Note!** If there are no problems with data management during certification tests, additional SE tests are not needed.



Testing a UI (1/2)

- Install your user's certificate first!
- Download a set of test scripts and JDL (from course agenda page)
- Check if grid-proxy-init works, and create proxy
- Check if globus-job-run works (with external site, in Oxford please use **t2ce01.physics.ox.ac.uk**):
`globus-job-run lxn1181.cern.ch /bin/pwd`
- Check RB access with default configuration:
`edg-job-list-match --vo dteam testJob.jdl`
- Check custom RB access configuration (different RB):
`edg-job-list-match -c `pwd`/lxn1177.conf --config-vo `pwd`/lxn1177.vo.conf testJob.jdl`
- Test job submission, status, logging, output retrieval:
`edg-job-submit --vo dteam testJob.jdl`
`edg-job-get-status <job_ID>`
`edg-job-get-logging-info -v 1 <job_ID>`
`edg-job-get-output <job_ID>`



Testing a UI (2/2)

- Here in Oxford use **t2se01.physics.ox.ac.uk** as external SE
- Test if you can reach external SE (detect firewall problems):
`edg-gridftp-ls gsiftp://castorgrid.cern.ch/castor/cern.ch/grid/dteam`
- Check storage resources visible by the information system:
`edg-replica-manager -v --vo dteam printInfo`
- Test if you can copy a new file to grid and register it:
 - Create simple text file: `testFile.<your_site_name>.txt`
 - Use `copyAndRegisterFile (cr)` function with external SE


```
edg-rm -v --vo dteam cr file://`pwd`/testFile.<sitename>.txt -l lfn:testFile.<site>.<current_time> -s  
castorgrid.cern.ch
```

- Check if file was registered using *listReplicas (lr)* function:
`edg-rm -v --vo dteam lr lfn:<lfn_from_previous_step>`
- Test if you can copy file (`cp`) back from the grid to your UI:
`edg-rm -v --vo dteam cp lfn:<lfn> file://`pwd`/testBack.txt`
- Test if you can remove the file from your SE and RLS:
`edg-rm -v --vo dteam del lfn:<lfn> -s <your_SE>`

- **Note!** Similar tests are performed by site certification scripts



Typical job submission problems (UI)



grid-proxy-init command fails:

Check if the command and other commands are in your PATH.

globus-job-run to external CE fails – grid security problems:

Check if you are using a reliable CE. Check if your certificate/proxy are valid (*grid-proxy-info*) and if you are registered in your VO correctly (dteam for site admins), the following commands give the lists of users for each LCG VO:

```
ldapsearch -LLL -x -H ldap://lcg-vo.cern.ch/ -b "o=dteam,dc=lcg,dc=org"
```

```
ldapsearch -LLL -x -H ldap://grid-vo.nikhef.nl -b "ou=lcg1,o=alice,dc=eu-datagrid,dc=org"
```

```
ldapsearch -LLL -x -H ldap://grid-vo.nikhef.nl -b "ou=lcg1,o=atlas,dc=eu-datagrid,dc=org"
```

```
ldapsearch -LLL -x -H ldap://grid-vo.nikhef.nl -b "ou=lcg1,o=cms,dc=eu-datagrid,dc=org"
```

```
ldapsearch -LLL -x -H ldap://grid-vo.nikhef.nl -b "ou=lcg1,o=lhcb,dc=eu-datagrid,dc=org"
```

edg-job-list-match fails:


Try the command with `--debug` option. Check your default RB configuration. Try also using a different RB with custom RB configuration.

edg-job-submit fails:

Try the command with `--debug` option and analyse the result.

edg-job-get-output fails:

Check if job is really in “*Done (Success)*” state with `edg-job-status` command. Check if `/tmp/jobOutput` directory exists and if it has correct permissions



Typical RM problems (UI/WN)

Note!!! These problems are also very common on worker nodes.

edg-gridftp-ls at remote SE fails:

Usually this means that your firewall setting is wrong.

printInfo command of RM gives empty or very short output:

Usually you should get a long list of CEs and SEs. If not, check mds.url parameter in file:

/opt/edg/var/etc/edg-replica-manager/edg-replica-manager.conf


The parameter defines BDII server you are using.

Check if the BDII server is working, and if it is really the TestZone BDII.

Try using *lxn1189.cern.ch* for Test Zone if you do not know the right one.

copyAndRegisterFile, replicateFile or copyFile commands fail:


- Check if you are using correct BDII server in RM configuration
- Use “-v” parameter to get detailed output
- Keep the output for further analysis (contact support, rollout-list)



Testing a whole site


- All tests (site GUIS, job submission, data replication) are included in site certification test suite (download from course agenda page)
- Just login to your UI, download the test suite and start it:

```
$ ./site-test <your_CE> <your_VO>
```
- The script will test your site starting from the UI, then it will submit test jobs to all queues of your CE and will test CE, WNs and SE
- The test sequence is similar as in previous slides, but it contains additional replication tests with different combinations of SE's.
- It takes several minutes but you are constantly informed about the status
- If it fails with basic tests (eg. GUIS) it will display the reason and information what to check and how.
- After the script finishes you will get the detailed report in several files. **Look carefully through them to spot problems.**



What to do in case of problems?

- Look carefully through the report and test jobs' output files.
- Try to reproduce the problem by finding the command which failed in the report, invoking it manually on your UI and analysing the output.
- Remember that Replica Manager configuration must be correct not only on the UI but also on all WNs.
- Look at GOC Wiki Trouble Shooting page:
<http://goc.grid.sinica.edu.tw/gocwiki/TroubleShootingHistory>
- Get help from next level (Primary Site, ROC, Deployment Team)



Typical site related problems

ldapsearch command fails, site is not present in BDII:

Check if globus-mds service is running both on CE and SE. Simple “*service globus-mds status*” command is not sufficient in many cases. Check the running processes to see if “*slapd*” process is there. If not, restart the service.

site GLIS works, but SE information is missing:

Check if globus-mds service is running on SE.

site GLIS works, but is not publishing correct information about queues:

Login to your CE and run *qstat* command. If it does not return there are problems with PBS or one of the worker nodes.


globus-job-run command fails:

Check if *globus-gatekeeper* service is running on your CE.

In case of authentication problems check if your *grid-mapfile*, *gridmapdir* and pooled accounts are configured correctly. Check *globus-gatekeeper.log* on CE

jobmanager-lcgpbs is unable to stage out output files:

Check if password-less authentication SSH from your WNs to your CE is configured properly (ssh keys, known hosts etc.). If not, check ssh configuration according to “LCG-2 Install Notes”



Typical SE related problems

Replica Manager is unable to find your SE:

Check if *globus-mds* service is running both on your CE and SE. Simple “*service globus-mds status*” command is not sufficient in many cases. Check the running processes to see if “*slapd*” process is there. If not, restart the service.


Access problems for members of one or more VOs:

Check if information provided by your information system reflects the reality. Look at the section that describes your SE in the report generated by `printInfo`. Are all the supported VOs properly configured on your SE? (grid-mapfile, pooled accounts, storage directories, permissions...)

Replica Manager is unable to copy and register or replicate file to your SE (default SE):


Check if *globus-gridftp* service is running on your SE.
Check if the storage directory (eg. */storage/dteam*) exists and has correct permissions.

NOTE! Before you start to suspect your SE configuration, check if Replica Manager is configured correctly on your UI and all your worker nodes – please refer to UI testing and common problems.




GIIS/GRIS testing

- First, check if slapd demon is running
(*service globus-mds status* command is not sufficient)
- You can check infoproviders manually:
 - (CE) `/opt/edg/libexec/ceinfo-wrapper.sh`
 - (SE) `/opt/lcg/libexec/lcg-info-wrapper`
- Use `ldapsearch` command line interface:
 - `ldapsearch -x -h <CE_or_SE> -p 2135 -b mds-vo-name=local,o=grid`
 - `ldapsearch -x -h <CE> -p 2135 -b mds-vo-name=<site-giis-name>,o=grid`
 - `ldapsearch -x -h <BDII> -p 2170 -b o=grid`
- Use graphical LDAP browser (GQ...)
- Interesting entries:
 - GlueCEUniqueID, GlueSEUniqueID, GlueSEName, GlueCESEBindSEUniqueID, ...



PBS batch system testing

- Check if *pbs_server* and *pbs_sched* are running on your CE
- Check if *qstat* returns without delays
- Check with “*pbsnodes -a*” if all WNs are alive
- Try submitting a simple “Hello world” script using *qsub* command from regular user account on CE
- In case of problems with getting job output:
 - check ssh keys on CE and WNs (refer LCG Install Notes)
 - try to login from regular user account on WN to CE – ssh should not ask for password!



Common mistakes

- Shared directories for Experiments' software installation
 - If they are defined (other than “.”) they have to exist
 - for the HEP VOs they are required
 - Correct ownership and permissions: readable by all VO members and writable by VO software managers
- Certification tests currently test only dteam VO environment. But are other VOs also configured correctly? (UI, SE, grid-mapfile, etc.)
- Test your site not only once after installation, but keep track of Test Zone reports and react in case new problems are detected
- Be sure that you keep your local settings while updating your site