

Table of Contents

The LCG Troubleshooting Guide.....	1
YAIM solutions.....	1
Log messages appear twice.....	1
No configuration target has been found.....	1
Authentication solutions.....	1
7 authentication failed.....	1
530 530 No local mapping for Globus ID.....	2
Proxy expired.....	4
501 501-FTPD GSSAPI error: GSS Major Status: General failure.....	4
Invalid CRL: The available CRL has expired.....	5
Certificate proxy not yet valid.....	5
'Bad certificate' returned instead of 'Unknown CA'.....	6
DPM and LFC solutions.....	6
Cannot map principal to local user.....	6
Problem with Mysql 4.1.....	7
service lfcdaemon stop : No valid credential found.....	7
sendrep: NS003 - illegal function 12.....	8
No user mapping.....	8
How to make srmcopy work.....	9
No space left on device.....	12
globus-url-copy : Connection closed by remote end.....	12
gLite I/O and DPM.....	13
How to restrict a pool to a VO.....	13
globus-url-copy : Permission denied (error 13 on XXX).....	14
rfdirdir : Permission denied (error 13 on XXX).....	15
426 426 Data connection. tmp file_open failed.....	15
Going from a Classic SE to the DPM.....	16
lcg-cr: Permission denied.....	19
CGSI-gSOAPError reading token data: Success.....	19
Error response 550:550 - not a plain file.....	20
LFC daemon crashes with old oracle database 10gR2.....	21
File exists.....	21
VOMS signature error.....	22
grid-proxy-init OK, but voms-proxy-init NOT OK.....	22
lcg_utils : "Invalid Argument" error.....	23
"Could not establish security context: Connection dropped by remote end !".....	23
What to do if the DN of a user changes ?.....	23
What to do if the host certificate expired or going to be changed.....	23
How do ACLs work ?.....	24
How to know all the file residing on a given SE ?.....	24
How to restrict a pool to a given VO ?.....	25
R-GMA solutions.....	25
General, very simple R-GMA test.....	25
Which logs should I back up for accounting purposes?.....	25
Failed to get list of tables from the Schema.....	26
Problems with rgma-client-check.....	26
Unable to source /opt/edg/etc/profile.d/edg-rgma-env.sh.....	26
RGMA_HOME is not set.....	27
No C++ compiler found.....	27
Cannot declareTable: table description not defined in the Schema.....	27
libgcj-java-placeholder.sh.....	28
Connection refused.....	29
HTML returned instead of XML.....	29
No tuples returned.....	30

Table of Contents

The LCG Troubleshooting Guide

Object has been closed: 1949004681.....	30
Unable to locate an available Registry Service.....	31
cannot remove `tmp/cmds.sql': Operation not permitted.....	31
Information System (and BDII) solutions.....	32
General considerations.....	32
Check that the information is in the top level BDII.....	32
Check that the information is in the site level BDII.....	32
Check that the information is is the GRIS.....	33
Check that the information is returned by the information provider.....	33
No information found in BDII.....	33
Entry's missing in the BDII.....	33
Problems updating the BDII configuration file from the web.....	33
Can not connect to the GRIS.....	33
The GRIS fails to start.....	34
No information returned by the GRIS.....	34
There is a problem with the GRIS configuration.....	34
No information was produced by the information provider.....	34
Default values show instead of dynamic values.....	34
New values not shown in GRIS.....	34
How to set up a dns load balanced BDII service.....	35
No such object (32): error message.....	35
How to close the site so it won't receive anymore jobs from the RBs.....	35
Job submission solutions.....	36
10 data transfer to the server failed.....	36
SAM solutions.....	36
VOMS solutions.....	36
Wrong host certificate subject in the vomses file.....	36
Database initialization error with.....	37
WARNING: Unable to verify signature!.....	37
APT solutions.....	38
apt-get update : W: Release file did not contain checksum information for :.....	38
FTS Solutions.....	38
I tried to submit a job and it said: submit: You are not authorised to submit jobs to this service.....	38
I submitted a job from site X to Y but it didn't work. The channel Y-X exists and has a share for my VO!.....	39
Which format should I use for the SURLs?.....	39
I've tried to submit a job but I get back an error saying: SOAP-ENV:Server.userException - org.xml.sax.SAXException.....	40
I've tried to submit a job but I get back an error saying: No match.....	40
I was able to list the channels but I cannot get the channel details.....	40
How do I setup a non-dedicated Channel?.....	41
After upgrading to FTS 1.5 I got "No Channel found or VO not authorized ..." error.....	41
FTA Solutions.....	42
Job always in Submitted state.....	42
Job always in Peding state.....	43
All my transfers fail with a SECURITY_ERROR.....	44
Which MyProxy Server is used?.....	45
I've noticed a warning "Cannot Get Agent DN" in the agent log files.....	45
My srmcopy transfers fail with a dCache MalformedURLException.....	46
I've upgraded to 1.4.1 but srmcopy doesn't seem to work.....	47
I've upgraded to 1.4.1 but the transfer failed with Error in srm_ping: NULL.....	48
The transfer failed with the error: No site found for host	48

Table of Contents

The LCG Troubleshooting Guide

Which Service Types are used?.....	49
I've tried everything, and it still doesn't seem to work.....	49
FTS Channel Administration solutions.....	50
How do I set the number of files transferred per VO instead of per channel?.....	50
General problems.....	51
How to replace host certificates on service nodes.....	51
Where I can find the log files.....	52

The LCG Troubleshooting Guide

WARNING: many of these entries contain **OBSOLETE** information. Please consult the [EGI Wiki](#) instead.

YAIM solutions

Log messages appear twice

Error

Sometimes when running the yaim command log messages appear twice in the screen.

Solution

This is because yaim prints the output messages through a 'tail' command. (This is a workaround for some improperly daemonized soft.). Look for 'tail' processes in your process tree and kill the old ones. This will solve the problem.

No configuration target has been found.

Error

```
ERROR: The node-info for service myservice not found in /opt/glite/yaim/bin/../../node-info.d nor in
```

Solution

You can use `yaim -a` to show you the available configuration targets. Probably you don't have the corresponding yaim module installed for your configuration target.

Authentication solutions

7 authentication failed

Error

This error message can be seen from the job logging information using `edg-job-get-logging-info`:
Something like the following:

```
- reason      =      7 authentication failed: GSS Major Status: Authentication Failed GSS Minor Status:
Chain:init.c:497:
globus_gss_assist_init_sec_context_async: Error during context initialization init_sec_context
```

Solution

- Please refer to [530 530 No local mapping for Globus ID entry in Troubleshooting Guide](#)
- To get more informations, try to list the server files using `gridftp` if possible :

```
edg-gridftp-ls gsiftp://<hostname>/tmp
```

- Please check that your CRLs are up to date (file date must be very recent - less than 6 hours)
- Please check that your host certificate is still valid :

```
openssl x509 -in /etc/grid-security/hostcert.pem -noout -enddate
```

- Please check that your grid-mapfile is up-to-date
- If you get this error when submitting a `globus-job-run <ce-name> /bin/hostname` to the affected:

```
GRAM Job submission failed because authentication failed:  
GSS Major Status: Unexpected Gatekeeper or Service Name  
GSS Minor Status Error Chain:
```

```
init.c:499: globus_gss_assist_init_sec_context_async: Error during context initialization  
init_sec_context.c:251: gss_init_sec_context: Mutual authentication failed: The target name
```

So the reverse resolution of the host IP address(server01.domain.net) is not equivalent to what is found in the host certificate(server02.domain.net)

- Check for the reverse lookup problem in `"/etc/hosts"` on the client side or dns configuration.

530 530 No local mapping for Globus ID

Error

Possible errors could be the following:

- If occurred during job submission, could be credential problem
- Problem in `/etc/grid-security/grid-mapfile`
- Problem with `/opt/edg/etc/edg-mkgridmap.conf`
- Problem with pool accounts
- Problem with `/etc/grid-security/gridmapdir`
- No files about pool accounts in `/etc/grid-security/gridmapdir`
- Variable GRIDMAPDIR is not set correctly

Gatekeeper and gridFTP daemon needs this in order to be able to use pool accounts. No error messages, when starting up the gatekeeper, what's more it even works fine with local accounts (like dteamsgm)!

- All pool accounts were taken
- If the error occurred during job submission, might be related with

```
/opt/edg/etc/lcas/lcas.db or /opt/edg/etc/lcmaps/lcmaps.db files
```

Solution

- Check if

```
globus-url-copy -dbg <from_file> <to_file>
```

complains about CRLs in its long output. If it does, see the topic: Invalid CRL: The available CRL has expired

- Check that it
 - ◆ exists and is updated via cron job

```
30 1,7,13,19 * * * /opt/edg/sbin/edg-mkgridmap --output=/etc/grid-security/g
```

- ◆ it contains right values (entries like: `"/C=CH/O=CERN/OU=GRID/CN=Piotr Nyczyk 9654".dteam`) You should copy a gridmap-file from a service node on the GRID, that you can trust to be configured properly, and compare your node's file with that one.

- Check that it contains correct URLs for the VOs

(like

```
ldap://lcg-vo.cern.ch/ou=lcg1,o=dteam,dc=lcg,dc=org .dteam)
```

- Check that they are existing for each supported VO (like: dteam001, ... , dteam050)
- Check if the directory is on the CE/SE has permissions

```
drwxrwxr-x    2 root    root          8192 Nov 29 15:08 gridmapdir
```

and on the Resource Broker

```
drwxrwxr-T    2 root    edguser      8192 Nov 29 15:08 gridmapdir
```

(instead of 'T' it can be 't' or 'x')

- Touch a file in /etc/grid-security/gridmapdir/ for each pool account like:

```
touch /etc/grid-security/gridmapdir/dteam001
...
touch /etc/grid-security/gridmapdir/dteam050
```

- Set the variable in etc/sysconfig/edg to the following

```
GRIDMAPDIR=/etc/grid-security/gridmapdir/
```

- In /etc/grid-security/gridmapdir/ there are hard links (with strange names like %2fc%3dch%2fo%3dcern%2fou%3dgrid%2fcn%3dpiotr%20nyczyk%209654) to each pool account that is taken. They have the same inode number (`ls -li FILENAME`) as the pool account file they point to. If there's no pool account file left free, run

```
/opt/edg/sbin/lcg-expiregridmapdir.pl
```

- and check if the following crontab entry on the CE exists

```
0 5 * * * /opt/edg/sbin/lcg-expiregridmapdir.pl -v 1>>/var/log/lcg-expiregridmapdir.l
```

- Example files

◆ /opt/edg/etc/lcas/lcas.db

```
# LCAS database/plugin list
#
# Format of each line:
# pluginname="<name/path of plugin>", pluginargs="<arguments>"
#
#
pluginname=lcas_userallow.mod,pluginargs=allowed_users.db
pluginname=lcas_userban.mod,pluginargs=ban_users.db
pluginname=lcas_timeslots.mod,pluginargs=timeslots.db
pluginname=lcas_plugin_example.mod,pluginargs=arguments
```

◆ /opt/edg/etc/lcmaps/lcmaps.db

```
# LCMAPS policyfile generated by LCFG::lcmaps - DO NOT EDIT
# @(#) /opt/edg/etc/lcmaps/lcmaps.db
#
# where to look for modules
path = /opt/edg/lib/lcmaps/modules

# module definitions
localaccount = "lcmaps_localaccount.mod -gridmapfile
/etc/grid-security/grid-mapfile"
```

The LCG Troubleshooting Guide < LCG < TWiki

```
poolaccount = "lcmads_poolaccount.mod -override_inconsistency -gridmapfile
/etc/grid-security/grid-mapfile -gridmapdir /etc/grid-security/gridmapdir/"
posixenf = "lcmads_posix_enf.mod -maxuid 1 -maxpgid 1 -maxsgid 32 "

# policies
standard:
localaccount -> posixenf | poolaccount
poolaccount -> posixenf
```

Proxy expired

Error

(Remaining) lifetime for proxy is less than 30 minutes. After extending with `myproxy-init edg-job-status` returns error for previously submitted jobs, while new job submission results in

```
**** Error: UI_PROXY_EXPIRED ****
Proxy certificate validity expired
```

In the Resource Broker log file (`/var/log/messages`)

```
Apr  6 13:14:45 <rb name> edg-wl-renewd[2567]: Proxy lifetime exceeded value of the Condor limit!
```

Solution

- Check if both proxies are expired

```
grid-proxy-info -text
myproxy-info
```

- How much time was left before issuing `myproxy-init`?
- If there is less than 30 minutes left for your proxy when executing `myproxy-init`, the Work Management System (WMS) will NOT renew your proxy.

501 501-FTPD GSSAPI error: GSS Major Status: General failure

Error

One get the following when using `edg-gridftp-ls`:

```
Error the server sent an error response: 501 501-FTPD GSSAPI error: GSS
Major Status: General failure
501-FTPD GSSAPI error: GSS Minor Status Error Chain:
501-FTPD GSSAPI error:
501-FTPD GSSAPI error: acquire_cred.c:125: gss_acquire_cred: Error with GSI
credential ...
```

```
501-FTPD GSSAPI error: The host key could not be found in:
501-FTPD GSSAPI error: 1) env. var.
X509_USER_KEY=/etc/grid-security/hostkey.pem
501-FTPD GSSAPI error: 2) /etc/grid-security/hostkey.pem
501-FTPD GSSAPI error: 3) /opt/globus/etc/hostkey.pem
501-FTPD GSSAPI error: 4) /root/.globus/hostkey.pem
```

Solution

- Verify validity of host certificate.

- Check that the host certificate permissions are set correctly (644)
- Contact CA if certificate has expired.
- Set permissions to 644

Invalid CRL: The available CRL has expired

Error

Invalid CRL: The available CRL has expired

One of the possible error messages (returned by edg-replica-manager command) looks like:

```
GridFTP: exist operation failed. the server sent an error response: 535 535-FTPD GSSAPI error: GS
535-FTPD GSSAPI error: GSS Minor Status Error Chain:
535-FTPD GSSAPI error:
535-FTPD GSSAPI error: accept_sec_context.c:170: gss_accept_sec_context: SSLv3 handshake problems
535-FTPD GSSAPI error: globus_i_gsi_gss_utils.c:881: globus_i_gsi_gss_handshake: Unable to verify
535-FTPD GSSAPI error: globus_i_gsi_gss_utils.c:854: globus_i_gsi_gss_handshake: SSLv3 handshake
535-FTPD GSSAPI error: OpenSSL Error: s3_srvr.c:1816: in library: SSL routines, function SSL3_GET
535-FTPD GSSAPI error: globus_gsi_callback.c:351: globus_i_gsi_callback_handshake_callback: Could
535-FTPD GSSAPI error: globus_gsi_callback.c:477: globus_i_gsi_callback_cred_verify: Could not ve
535-FTPD GSSAPI error: globus_gsi_callback.c:769: globus_i_gsi_callback_check_revoked: Invalid CR
535 FTPD GSSAPI error: accepting context
```

Solution

- Certificates in `/etc/grid-security/certificates/` are outdated Make sure that CA RPMs (called `ca_`, like `ca_CERN`) are installed, and updated to the last CA release.
<http://grid-deployment.web.cern.ch/grid-deployment/lcg2CAList.html>
- Periodic update failed A way to check this is to compare the sizes of the files in `/etc/grid-security/certificates/` with `edg-gridftp-ls` between the node, and a server that surely has the right credentials. Run `edg-fetch-crl` command manually, and see if it produced any error message. Make sure that the following crontab entry exists

```
30 1,7,13,19 * * * /opt/edg/etc/cron/edg-fetch-crl-cron
```

Certificate proxy not yet valid

Error

Following error occurred when using `globus-url-copy` command:

```
error: the server sent an error response: 535 535
Authentication failed: GSSException: Defective credential detected
[Root error message: Certificate C=CH,O=CERN,OU=GRID,CN=Judit Novak 0973,CN=proxy
not yet valid.]
[Root exception is org.globus.gsi.proxy.ProxyPathValidatorException:
Certificate C=CH,O=CERN,OU=GRID,CN=Judit Novak 0973,CN=proxy not yet valid.]
```

Solution

Source and destination nodes weren't synchronized in time. Synchronize the nodes !

'Bad certificate' returned instead of 'Unknown CA'

Error

Couldn't verify the remote certificate !

In SSL, the 'unknown CA' error obtained by the SSL server during the handshake gets translated (by the `ssl3_alert_code` call) into a generic 'bad certificate' error:

```
case SSL_AD_UNKNOWN_CA:          return(SSL3_AD_BAD_CERTIFICATE);
```

This is sent as an alert to the SSL client during the SSL handshake. The Globus GSI handshake callback (`globus_i_gsi_gss_handshake`) always casts a 'bad certificate' error, no matter how it was obtained, into a `GLOBUS_GSI_GSSAPI_ERROR_REMOTE_CERT_VERIFY_FAILED`:

```
839          /* checks for ssl alert 42 */
840          if (ERR_peek_error() ==
841              ERR_PACK(ERR_LIB_SSL, SSL_F_SSL3_READ_BYTES,
842                      SSL_R_SSLV3_ALERT_BAD_CERTIFICATE))
843          {
844              GLOBUS_GSI_GSSAPI_OPENSSL_ERROR_RESULT(
845                  minor_status,
846                  GLOBUS_GSI_GSSAPI_ERROR_REMOTE_CERT_VERIFY_FAILED,
847                  ("Couldn't verify the remote certificate"));
848          }
```

So, the error "Couldn't verify the remote certificate" can also mean (among other things, including its literal meaning) "the SSL client certificate was found by the remote SSL server to be issued by an unknown CA". This is quite misleading.

Solution

The Certification Authority files for the unknown CA are missing in `/etc/grid-security/certificates` or in the directory pointed to by the environmental variable `X509_CERT_DIR`. Instructions on how to upload the CA files for the Certification Authorities accepted by LCG/EGEE can be found here:

<http://grid-deployment.web.cern.ch/grid-deployment/lcg2CAlist.html>

DPM and LFC solutions

Cannot map principal to local user

Error

You get this error: `cannot map principal to local user`

Solution

`/etc/grid-security/gridmapdir` directory should be writable by `lfcmgr` or `dpmmgr`.

If you are using another directory, it also has to be writable, and should be specified in the `/etc/sysconfig/SERVICE_NAME` files.

Problem with Mysql 4.1

Error

When using Mysql 4.1 with either the LFC or the DPM, you get the following error (here in /var/log/dpns/log):

```
09/23 12:19:41 26938 Cns_opendb: CONNECT error: Client does not support =authentication protocol requested by server; consider upgrading Mysql client.
```

Solution

According to the Mysql documentation, paragraph A.2.3, there is a very simple solution to this problem: use the OLD_PASSWORD() function instead of the PASSWORD() function when creating the DB account.

service lfcd daemon stop : No valid credential found

Error

You get this :

- service lfcd daemon start is OK
- but service lfcd daemon stop doesn't work :

```
$ service lfcd daemon stop
Stopping lfcd daemon: send2nsd: NS002 - send error : No valid credential found
nssshutdown: Could not establish context
```

And trying to create /grid as root doesn't work either :

```
$ lfc-mkdir /grid
send2nsd: NS002 - send error : No valid credential found
cannot create /grid: Could not establish context
```

Solution

Check that :

- you have a valid host certificate and key
- you have copied and renamed them to /etc/grid-security/lfcmgr :

```
$ ll /etc/grid-security/ | grep host
-rw-r--r--  1 root    root          5423 May 27 12:35 hostcert.pem
-r-----  1 root    root          1675 May 27 12:35 hostkey.pem
```

- **IMPORTANT : the host certificate and key have to be kept at their original place !!!**

```
$ ll /etc/grid-security/lfcmgr | grep lfc
-rw-r--r--  1 lfcmgr  lfcmgr          5423 May 30 13:58 lfccert.pem
-r-----  1 lfcmgr  lfcmgr          1675 May 30 13:58 lfckey.pem
```

Check that the CA certificates are present :

```
ls /etc/grid-security/certificates/
01621954.0
01621954.crl_url
```

```
01621954.info
01621954.r0
01621954.signing_policy
03aa0ecb.0
03aa0ecb.crl_url
03aa0ecb.info
03aa0ecb.r0
03aa0ecb.signing_policy
...
```

Get more information, with **export CSEC_TRACE=1** :

```
$ export CSEC_TRACE=1
$ lfc-mkdir /grid
```

Further help

If it still doesn't help, send the `/var/log/lfc/log` file to support@ggusNOSPAMPLEASE.org (remove the NONSPAM !).

And send us the output of : `$ cat /proc/lfc_master_pid/environ`

sendrep: NS003 - illegal function 12

Error

You get this :

```
$ tail -f /var/log/lfc/log
...
11/23 09:37:13 12001,0 sendrep: NS003 - illegal function 12
...
```

Solution

It means you are calling a method that is not allowed after another call has failed.

For instance, if an `lfc_opending` fails, you cannot call `lfc_closedirg` afterwards. (In LFC/DPM 1.4.1, this is fixed, and the `lfc_closedirg` is automatically ignored).

The solution is : check the possible failures in your code, so that `lfc_closedirg` isn't called if `lfc_opending` has failed !

No user mapping

Error

You get this error :

```
Could not get virtual id: No user mapping !
```

Solution

Check this :

- permissions/ownership on `/etc/grid-security/gridmapdir` ?
- does the user appear in `/etc/grid-security/grid-mapfile` ?

service lfcdaemon stop : No valid credential found

- aren't all the pool accounts in use ?
- do all the pool accounts exist in /etc/passwd ?
- does /opt/lcg/etc/lcgdm-mapfile exist ?
- if yes, does it contain the user that seems to be missing ?

Further help

If the problem still appears, contact support@ggusNOSPAMPLEASE.org (remove the NONSPAM !) specifying/giving :

- the answers to the previous questions,
- the version of the LFC/DPM server,
- the version of the LFC/DPM client,
- the appropriate logs.

How to make srmcopy work

Here is a recipe from James Casey (James.Casey@cernNOSPAMPLEASE.ch) on how to make `srmcopy` work with the DPM :

- Using `srmcp` to download from `castor2`
- upload that file from local storage to a `dpm`
- copy from `castor2` to `dpm`, in 'pushmode'
- download the file from the `dpm` to local storage.

```
$/opt/d-cache/srm/bin/srmcp srm://castorgridsc:8443/castor/cern.ch/grid/dteam/storage/transfer-t  
file:///tmp/foo
```

```
$ls -l /tmp/foo  
-rw-r--r--  1 jamesc  zg   2364 Sep 27 16:56 /tmp/foo
```

```
$/opt/d-cache/srm/bin/srmcp file:///tmp/foo srm://lxfsrm528:8443/dpm/cern.ch/home/dteam/jamesc-f
```

```
$dpns-ls -l /dpm/cern.ch/home/dteam/jamesc-foo  
-rw-rw-r--  1 dteam002 cg   2364 Sep 27 17:01 /dpm/cern.ch/home/dteam/jamesc-foo
```

```
$/opt/d-cache/srm/bin/srmcp --debug --pushmode=true srm://castorgridsc.cern.ch:8443/castor/cern.
```

```
Storage Resource Manager (SRM) CP Client version 1.16  
Copyright (c) 2002-2005 Fermi National Accelerator Laborarory
```

```
SRM Configuration:  
  debug=true  
  gsissl=true  
  help=false  
  pushmode=true  
  userproxy=true  
  buffer_size=2048  
  tcp_buffer_size=0  
  stream_num=10  
  config_file=/afs/cern.ch/user/j/jamesc/.srmconfig/config.xml  
  glue_mapfile=/opt/d-cache/srm/conf/SRMServerV1.map  
  webservice_path=srm/managerv1.wsdl  
  webservice_protocol=https  
  gsiftpclinet=globus-url-copy  
  protocols_list=gsiftp
```

The LCG Troubleshooting Guide < LCG < TWiki

```
save_config_file=null
srmcphome=/opt/d-cache/srm
urlcopy=bin/urlcopy.sh
x509_user_cert=/afs/cern.ch/user/j/jamesc/.globus/usercert.pem
x509_user_key=/afs/cern.ch/user/j/jamesc/.globus/userkey.pem
x509_user_proxy=/tmp/x509up_u4290
x509_user_trusted_certificates=/afs/cern.ch/user/j/jamesc/.globus/certificates
retry_num=20
retry_timeout=10000
wsdl_url=null
use_urlcopy_script=false
connect_to_wsdl=false
delegate=true
full_delegation=true
from[0]=srm://castorgridsc.cern.ch:8443/castor/cern.ch/grid/dteam/storage/transfer-test/c
to=srm://lxfstrm528:8443/dpm/cern.ch/home/dteam/jamesc-foo-srmcp
```

=

```
Tue Sep 27 17:04:35 CEST 2005: starting SRMCopyPushClient Tue Sep 27 17:04:35 CEST 2005:
SRMClient(https,srm/managerv1.wsdl,true) Tue Sep 27 17:04:35 CEST 2005: connecting to server Tue Sep
27 17:04:35 CEST 2005: connected to server, obtaining proxy SRMClientV1 : connecting to srm at
http://oplapro58.cern.ch:8443/srm/managerv1 Tue Sep 27 17:04:37 CEST 2005: got proxy of type class
org.dcache.srm.client.SRMClientV1 Tue Sep 27 17:04:37 CEST 2005: copying
srm://castorgridsc.cern.ch:8443/castor/cern.ch/grid/dteam/storage/transfer-test/castor2/s00/file-test.dat into
srm://lxfstrm528:8443/dpm/cern.ch/home/dteam/jamesc-foo-srmcp
```

SRMClientV1 : copy,

```
srcSURLS[0]="srm://castorgridsc.cern.ch:8443/castor/cern.ch/grid/dteam/storage/transfer-test/castor2/s00/file-test.dat"
```

```
SRMClientV1 : copy, destSURLS[0]="srm://lxfstrm528:8443/dpm/cern.ch/home/dteam/jamesc-foo-srmcp"
SRMClientV1 : copy, contacting service http://oplapro58.cern.ch:8443/srm/managerv1 Tue Sep 27 17:04:40
CEST 2005: srm returned requestId = 618988755 Tue Sep 27 17:04:40 CEST 2005: sleeping 1 seconds ... Tue
Sep 27 17:04:42 CEST 2005: sleeping 1 seconds ... Tue Sep 27 17:04:44 CEST 2005: sleeping 1 seconds ...
Tue Sep 27 17:04:45 CEST 2005: sleeping 1 seconds ... Tue Sep 27 17:04:46 CEST 2005: FileRequestStatus
fileID = 0 is Done => copying of
srm://castorgridsc.cern.ch:8443/castor/cern.ch/grid/dteam/storage/transfer-test/castor2/s00/file-test.dat is
complete
```

```
$/opt/d-cache/srm/bin/srmcp --debug srm://lxfstrm528:8443/dpm/cern.ch/home/dteam/jamesc-foo-srmcp
```

```
Storage Resource Manager (SRM) CP Client version 1.16
Copyright (c) 2002-2005 Fermi National Accelerator Laboratory
```

```
SRM Configuration:
  debug=true
  gsissl=true
  help=false
  pushmode=false
  userproxy=true
  buffer_size=2048
  tcp_buffer_size=0
  stream_num=10
  config_file=/afs/cern.ch/user/j/jamesc/.srmconfig/config.xml
  glue_mapfile=/opt/d-cache/srm/conf/SRMServerV1.map
  webservice_path=srm/managerv1.wsdl
  webservice_protocol=https
  gsiftpclinet=globus-url-copy
  protocols_list=gsiftp
  save_config_file=null
```

The LCG Troubleshooting Guide < LCG < TWiki

```
srmcphome=/opt/d-cache/srm
urlcopy=bin/urlcopy.sh
x509_user_cert=/afs/cern.ch/user/j/jamesc/.globus/usercert.pem
x509_user_key=/afs/cern.ch/user/j/jamesc/.globus/userkey.pem
x509_user_proxy=/tmp/x509up_u4290
x509_user_trusted_certificates=/afs/cern.ch/user/j/jamesc/.globus/certificates
retry_num=20
retry_timeout=10000
wsdl_url=null
use_urlcopy_script=false
connect_to_wsdl=false
delegate=true
full_delegation=true
from[0]=srm://lxfsrm528:8443/dpm/cern.ch/home/dteam/jamesc-foo-srmcp
to=file:///tmp/foo2

Tue Sep 27 18:02:00 CEST 2005: starting SRMGetClient
Tue Sep 27 18:02:00 CEST 2005: SRMClient(https,srm/manager.v1.wsdl,true)
Tue Sep 27 18:02:00 CEST 2005: connecting to server
Tue Sep 27 18:02:00 CEST 2005: connected to server, obtaining proxy
SRMClientV1 : connecting to srm at http://lxfsrm528.cern.ch:8443/srm/manager.v1
Tue Sep 27 18:02:01 CEST 2005: got proxy of type class org.dcache.srm.client.SRMClientV1
SRMClientV1 : get: surls[0]="srm://lxfsrm528:8443/dpm/cern.ch/home/dteam/jamesc-foo-srmcp"
SRMClientV1 : get: protocols[0]="http"
SRMClientV1 : get: protocols[1]="dcap"
SRMClientV1 : get: protocols[2]="gsiftp"
SRMClientV1 : get, contacting service http://lxfsrm528.cern.ch:8443/srm/manager.v1
doneAddingJobs is false
copy_jobs is empty
Tue Sep 27 18:02:09 CEST 2005: srm returned requestId = 27373
Tue Sep 27 18:02:09 CEST 2005: sleeping 1 seconds ...
Tue Sep 27 18:02:11 CEST 2005: FileRequestStatus with SURL=srm://lxfsrm528:8443/dpm/cern.ch/home/

Tue Sep 27 18:02:11 CEST 2005: received TURL=gsiftp://lxfsrm528.cern.ch/lxfsrm528:/shift/1

doneAddingJobs is false
copy_jobs is not empty
Tue Sep 27 18:02:11 CEST 2005: fileIDs is empty, breaking the loop
copying CopyJob, source = gsiftp://lxfsrm528.cern.ch/lxfsrm528:/shift/lxfsrm528/data01/cg/2005-09

GridftpClient: memory buffer size is set to 2048
GridftpClient: connecting to lxfsrm528.cern.ch on port 2811
GridftpClient: gridFTPClient tcp buffer size is set to 0
GridftpClient: gridFTPRead started
GridftpClient: parallelism: 10
GridftpClient: waiting for completion of transfer
GridftpClient: gridFTPWrite: starting the transfer in emode from lxfsrm528:/shift/lxfsrm528/data0

GridftpClient: DiskDataSink.close() called
GridftpClient: gridFTPWrite() wrote 2364bytes
GridftpClient: closing client : org.dcache.srm.util.GridftpClient$FnalGridFTPClient@4be2cc
GridftpClient: closed client
execution of CopyJob, source = gsiftp://lxfsrm528.cern.ch/lxfsrm528:/shift/lxfsrm528/data01/cg/20

setting file request 0 status to Done
doneAddingJobs is true
copy_jobs is empty
stopping copier

$ls -l /tmp/foo2
-rw-r--r-- 1 jamesc zg 2364 Sep 27 18:02 /tmp/foo2
```

No space left on device

Error

You get this with **srmcp**:

```
$ srmcp -debug=true file://localhost//tmp/hello srm://dpm01.pic.es:8443/dpm/pic.es/home/dteam/tes

Exception in thread "main" java.io.IOException: rs.state = Failed rs.error = No space left on dev
    at gov.fnal.srm.util.SRMPutClient.start(SRMPutClient.java:331)
    at gov.fnal.srm.util.SRMCopy.work(SRMCopy.java:409)
    at gov.fnal.srm.util.SRMCopy.main(SRMCopy.java:242)
Tue Oct 18 15:59:17 CEST 2005: setting all remaining file statuses to "Done"
Tue Oct 18 15:59:17 CEST 2005: setting file request 0 status to Done
SRMClientV1 : getRequestStatus: try #0 failed with error
SRMClientV1 : Invalid state
java.lang.RuntimeException: Invalid state
    at org.dcache.srm.client.SRMClientV1.setFileStatus(SRMClientV1.java:1097)
    at gov.fnal.srm.util.SRMPutClient.run(SRMPutClient.java:362)
    at java.lang.Thread.run(Thread.java:534)
```

Or a similar error with **globus-url-copy**, or another utility.

Solution

The problem is that some utilities use Permanent as their default and some others Volatile.

For instance :

- srmcp doesn't work if your pool is of **volatile** type.
- globus-url-copy

You have two possibilities :

- Modify the type of the pool to "-" (this type allows both Volatile and Permanent files):

```
dpm-modifypool --poolname <my_pool> --s_type "-"
```

- Create two pools, one Volatile and one Permanent

Further help

If it still doesn't help, send the relevant DPM log files to support@ggusNOSPAMPLEASE.org (remove the NOSPAM !).

globus-url-copy : Connection closed by remote end

Error

```
globus-url-copy file:/etc/group ↗
gsiftp://DPM_POOL_NODE/dpm/cern.ch/home/dteam/tests.sophie.shift.conf2

error: the server sent an error response: 553 553
/dpm/cern.ch/home/dteam/tests.sophie.shift.conf2: Connection closed by remote end.
```

Is this really what you want to be doing ?

The same command with the `DPM_SERVER` instead of the `DPM_POOL_NODE` will work...

So, this error only occurs if you try to contact a pool node directly. This is not necessarily what you want to be doing, as it can involve an unnecessary copy, if the file finally ends up on another pool node than the one contacted.

So, doing this adds load on the DPM setup.

solution

If you still want to do this, on the DPM server, add this line to `/etc/shift.conf` :

```
RFIOD TRUST DPM_server_short_name DPM_server_long_name disk_server1_short_name
disk_server1_long_name...
```

gLite I/O and DPM

Here is Jean-Philippe's explanation :

All physical files on disk belong to a special user "dpmmgr" and are only accessible by this user.

RFIOD and gsiFTP which are launched as root have been modified to check with the DPNS (DPM Name Server) if the client is authorized to open (or delete or ...). Then RFIOD or gsiFTP does the open on behalf of the user and returns an handle that can be used in `rfio_read/rfio_write` ...

The disk server must be trusted by the DPNS using entries in `shift.conf` of the form :

```
DPNS TRUST disk_server1 disk_server2 ...
```

The users are mapped using the standard `grid-mapfile`.

If the gliteIO daemon runs with a host/service certificate and is modified to be DPM-aware i.e. to contact the DPNS, everything is ok.

If you do not want to modify gliteIO daemon, and gliteIO runs as the client, you may still access data on other disk servers using RFIO, but you cannot access the data residing on the same machine as the gliteIO daemon because in this case the file is seen as local and RFIO does not use RFIOD.

One solution which was explained to Gavin and his successors was: it is possible to modify RFIO to use RFIOD even if the file is local. The cost is an extra copy operation between RFIOD et gliteIO servers. The modification is not very difficult but is not very high on our list of priorities either.

Please note that you will encounter the same problem with CASTOR as soon as the secure version of CASTOR is released.

How to restrict a pool to a VO

How to create a pool dedicated to a VO ?

It is possible to have one pool dedicated to a given VO, with all the authorization behind, using the `dpm-addpool` or `dpm-modifypool` commands.

For instance :

```
dpm-addpool --poolname VOpool --def_filesize 200M --gid the_VO_gid
```

globus-url-copy : Connection closed by remote end

```
dpm-addpool --poolname VOpool --def_filesize 200M --group the_VO_group_name
```

Comment

If you define :

- one pool dedicated to group1 / VO1
- one pool open to all groups / VOs

then, the **dedicated pool will be used until it is full.**

When the dedicated pool is full, the open pool is then be used.

globus-url-copy : Permission denied (error 13 on XXX)

Error

You get this :

```
$globus-url-copy file:///tmp/hello
gsift://<dpm_server>/dpm/<domain.name>/home/dteam/testdir2/test
error: the server sent an error response: 553 553
/dpm/<domain.name>/home/dteam/testdir2/test: Permission denied (error 13 on <disk_server>).
```

Solution

You might want to check that :

- the DPM server and the disk server are not on different subnets. If they are, you should create the `/etc/shift.localhosts` file on the DPM server, containing the disk server subnet (as an IP address).
For instance :

```
$cat /etc/shift.localhosts
212.189.153
```

- the `dpmmgr` user has the same `uid/gid` on each machine (DPM server and disk server). **Important:** if you change the `dpmmgr` `uid/gid`, restart all the daemons afterwards.
- check the permissions on the `/dpm/domain.name/home/dteam/testdir` hierarchy
- `/etc/shift.conf` on the DPM server :

```
DPM TRUST <disk_server1_short_name> <disk_server1_long_name>
<disk_server2_short_name> <disk_server2_long_name>
DPNS TRUST <disk_server1_short_name> <disk_server1_long_name>
<disk_server2_short_name> <disk_server2_long_name>
RFIOD TRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD WTRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD RTRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD XTRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD FTRUST <dpm_server_short_name> <dpm_server_long_name>
```

- `/etc/shift.conf` on the disk server :

```
RFIOD TRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD WTRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD RTRUST <dpm_server_short_name> <dpm_server_long_name>
RFIOD XTRUST <dpm_server_short_name> <dpm_server_long_name>
```

How to restrict a pool to a VO

```
RFIOD FTRUST <dpm_server_short_name> <dpm_server_long_name>
```

- the permissions of the file system on the disk server : the directory and its subdirectories should have

```
ls -lad /data01
drwxrwx--- 365 dpmmgr dpmmgr 8192 Sep 29 09:58 /data01
```

Further help

If it still doesn't help, send the `/var/log/rfiiod/log` file to support@ggusNOSPAMPLEASE.org (remove the NOSPAM !).

rfdir : Permission denied (error 13 on XXX)

Error

You get this :

```
$ rfdir <my_dpm_host>:/storage
opendir(): <my_dpm_host>:/storage: Permission denied (error 13 on <my_dpm_host>)
```

Solution

To use `rfdir` with the DPM, the recipe is :

```
$ export DPNS_HOST=<my_dpns_host>
$ rfdir /dpm/cern.ch/home/dteam/
```

Comment

To use `rfrm`, you need to set `DPM_HOST` and `DPNS_HOST` :

```
$ export DPNS_HOST=<my_dpns_host>
$ export DPM_HOST=<my_dpm_host>

$ rfrm -r /dpm/cern.ch/home/dteam/tests_sophie
```

Furher help

If it still doesn't help, send the `/var/log/rfiiod/log` file to support@ggusNOSPAMPLEASE.org (remove the NONSPAM !).

426 426 Data connection. tmp file_open failed

Error

You get this :

```
$ lcg-cp -v --vo dteam lfn:essai_node08_3 file:/home/cleroy/node08_node02

Source URL:lfn:essai_node08_3
File size: 202
VO name: dteam
Source URL for copy:
gsiftp://MY_DISK_SERVER.cern.ch/MY_DISK_SERVER:/storage/dteam/2005-11-10/file11e39190-5c5a-4a64-b
Destination URL: file:/home/cleroy/node08_node02

# streams: 1
```

globus-url-copy : Permission denied (error 13 on XXX)

The LCG Troubleshooting Guide < LCG < TWiki

```
# set timeout to 0 (seconds)
      0 bytes      0.00 KB/sec avg      0.00 KB/sec instthe server
sent an error response: 426 426 Data connection. tmp file_open failed

lcg_cp: Transport endpoint is not connected
```

Or this :

```
$ globus-url-copy gsiftp://MY_DPM.cern.ch/MY_DPM:/storage/cg/2005-11-14/file356ff811-f30b-412e-bd
error: the server sent an error response: 426 426 Data connection. tmp file_open failed
```

Solution

It seems that the permissions on /tmp are wrong.

They should look like :

```
$ ll -ld /tmp
drwxrwxrwt  14 root      root          4096 Nov 14 17:21 /tmp
```

Further help

If it still doesn't help, send the /var/log/messages file to support@ggusNOSPAMPLEASE.org (remove the NONSPAM !).

Going from a Classic SE to the DPM

Turning your Classic SE into a DPM is easy : it doesn't require to move the data in any way. You only need to make the DPM server aware of the files that are present on your Storage Element. In other words, this is only a metadata operation, and no actual file movement is required at all.

How long will it take ?

To give a time estimate, the tests we have performed at CERN took :

- 4 hours 23 minutes 17 seconds
- for 236546 files

This gives an average of 14.97 files migrated per second.

Possible scenarios

There are two possibilities :

- install the DPM servers on the Classic SE, and consider the Classic SE as a pool node as well,
- install the DPM servers on a different machine, and turn the Classic SE into a DPM pool node.

Preliminary steps

You have to install the DPM servers on a given machine (it can be the Classic SE itself) See the DPM Admin Guide.

If installed on a different machine, the Classic SE will act as a pool node (=disk server) of the DPM.

Important :

426 426 Data connection. tmp file_open failed

Make sure that the VO groups and pool accounts have the same uids/gids on the Classic SE and on the DPM server. Otherwise, the migrated permissions will no be the correct ones.

Permissions

Make sure that the VO groups ids and pool accounts uid/gids correspond on the DPM server and on the Classic SE. Otherwise, the ownership will not be correctly migrated to the DPM Name Server

Get the script

To perform the migration, the IT-GD group provides a migration script. You can find it in the CERN central CVS service (repository [lcgware/migration-classicSE-DPM](#)).

You can also download the following tarball: [migration-classicSE-DPM.tar.gz](#) (*last update: 2005-10-11*).

Note that a new version of this script is currently rewritten in order to manage problems encountered during the migration (for example when migrating the entries to an already existing DPM server (already having entries)).

Configuration

- on the classic SE

- Stop the GridFTP server :

```
service globus-gridftp stop
chkconfig globus-gridftp off
```

- Install the DPM-client package.
- Set the environment variable DPNS_HOST with the DPNS hostname :

```
export DPNS_HOST=DPNS_HOSTNAME
```

- Put in the `/etc/shift.conf` the following lines:

```
RFIOD RTRUST SHORT_DPNS_HOSTNAME LONG_DPNS_HOSTNAME
RFIOD WTRUST SHORT_DPNS_HOSTNAME LONG_DPNS_HOSTNAME
RFIOD XTRUST SHORT_DPNS_HOSTNAME LONG_DPNS_HOSTNAME
RFIOD FTRUST SHORT_DPNS_HOSTNAME LONG_DPNS_HOSTNAME
```

- Compile the migration.c file using the Makefile :

```
make all
```

- on the DPNS server

- Put in the `/etc/shift.conf` the following line:

```
DPNS TRUST SHORT_CLASSIC_SE_HOSTNAME LONG_CLASSIC_SE_HOSTNAME
```

Migration

Run the following command on the classic SE host:

```
./migration classicSE_hostname classicSE_directory dpm_hostname dpm_directory dpm_poolname
```

where:

- `classicSE_hostname` is the (short) name of the classic SE (i.e. without the domain name).
- `classicSE_directory` is the name of the directory where are stored all the files (for example `/storage`).
- `dpm_hostname` is the (short) name of the DPM (i.e. without the domain name).
- `dpm_directory` is the name of the directory where will be stored all the files (for example `/dpm/DOMAIN_NAME/home`).
- `dpm_poolname` is the name of the pool (obtained by using `dpm_qryconf`) on the DPM.

Important : Note that you have to put short hostname (i.e. do not add the domain name) on the command line.

Post migration steps

If the Classic SE is a separate machine, make sure you turn it into a DPM pool node :

- on the Classic SE :

Attention : before doing that, make sure that the entries appear in DPM Name Server as expected

Configure the Classic SE to be a pool node :

- remove the `CASTOR-client` RPM
- install the `DPM-client`, `DPM-rfiio-server` and `DPM-gsiftp-server` RPMs
- configure security (globus, grid-mapfile, gridmapdir, pool accounts)
- create the `dpmmgr` user/group (with the same uid/gid as on the DPM server)
- `chown root:dpmmgr /etc/grid-security/gridmapdir`
- create `/etc/grid-security/dpmmgr`
- `chown dpmmgr:dpmmgr /etc/grid-security/dpmmgr`
- `cp -p /etc/grid-security/hostcert.pem /etc/grid-security/dpmmgr/dpmcert.pem`
- `cp -p /etc/grid-security/hostkey.pem /etc/grid-security/dpmmgr/dpmkey.pem`
- `service rfiiod start`
- `service dpm-gsiftp start`

VERY IMPORTANT : Change the ownership of all the Classic SE files/directories :

WARNING : before changing the permissions, make sure that all the files have been properly migrated in the DPNS. Once the permissions changed, you cannot get the old permissions back...

- `chown -R dpmmgr:dpmmgr /YOUR_PARTITION`
- `chmod -R 660 /YOUR_PARTITION`
- `find /storage -type d -exec chmod 770 {} \;` to have the correct permissions on directories

- on the DPM server :

Create the pool and add the Classic SE file system to it :

- `export DPM_HOST=YOUR_DPM_SERVER`
- `dpm-addpool --poolname POOL_NAME --def_filesize 200M (if the pool doesn't exist yet !)`
- `dpm-addfs --poolname POOL_NAME --server CLASSIC_SE_SHORT_NAME --fs CLASSIC_SE_FILE_SYSTEM`

For more details, refer to the DPM Admin Guide.

Catalog

The entries that exist already in a catalog (RLS or LFC) won't be migrated.

The corresponding entries can still be accessed in the same way as before the migration. For instance :

```
lcg-cp --vo dteam
sfn://ClassiSE_hostname/storage/dteam/generated/2005-03-29/filef70996ba-ba4e-42dc-9bae-03a3d7e7ac
file:/tmp/test.classic.se.migration.1
```

Information System

You have to publish the DPM as an SRM in the Information System.

There is no need to publish the Classic SE as such in the Information System.

Further help

Please send all your questions/comments to hep-service-dpm@cernNOSPAMPLEASE.ch (remove the NONSPAM !) or to yvan.calas@cernNOSPAMPLEASE.ch.

lcg-cr: Permission denied

Error

You get this **when targetting a DPM Storage Element** :

```
$ lcg-cr -v --vo dteam -d se.polgrid.pl -l lfn:/grid/dteam/apadee/test-file-polgrid.pl.1 file:///
Using grid catalog type: lfc
Using grid catalog : lfc-dteam.cern.ch
lcg_cr: Permission denied
```

Solution

It can be that one of the partitions on one Disk Server is not properly configured.

The permissions on **all partitions** should be :

```
$ ll -ld /storage
drwxrwx---    3  dpmmgr      dpmmgr      4096 Nov 14 17:21 /storage
```

CGSI-gSOAPError reading token data: Success

Error

You get this error:

```
CGSI-gSOAP: Error reading token data: Success
```

This means that the SRM server has dropped the connection.

Solution

Try to restart the srm server:

```
service srmv1 restart
```

If it doesn't help, the reasons can be :

lcg-cr: Permission denied

- a security handshake problem
- a `grid-mapfile` or `gridmapdir` problem
- one of the server thread crashed (but, it has never been seen in production...)

Check :

- the `/var/log/srmv1/log` and `/var/log/srmv2/log` log files
- the permissions/contents of `grid-mapfile` and `gridmapdir`
- that all the DPM ports are open

Set the following environment variables :

```
$CGSI_TRACE=1
$CGSI_TRACEFILE=/tmp/tracefile
```

and see if the error messages contained in `/tmp/tracefile` help.

Error response 550:550 - not a plain file

Error

For instance, you get this :

```
$ lcg_cp srm://grid05.lal.in2p3.fr:8443/dpm/lal.in2p3.fr/home/atlas/dq2/file.11 /tmp/test --vo dt
the server sent an error response: 550 550 grid07.lal.in2p3.fr:/dpm/part/part1/atlas/2006-04-29/fi
```

```
lcg_cp: Invalid argument
```

But the file exists in the DPM Name Server :

```
$ dpns-ls -l /dpm/lal.in2p3.fr/home/atlas/dq2/csc11.root.11
-rw-rw-r--  1 19478   20008                28472534 Apr 29 23:23 /dpm/lal.in2p3.fr/home/atlas/dq2
```

Solution 1

Although it appears in the DPM namespace, the file doesn't **physically** exist on disk anymore.

You should un-register the file from the namespace, to avoid this inconsistency.

Solution 2

Check that, **on all disk servers** you are actually running :

- the DPM RFIO server, and not the CASTOR one,
- the DPM GRIDFTP server, and not the Classic SE GRIDFTP one :

```
$ ps -ef|grep rfio
root 20313 1 0 Sep19 ? 00:00:10 /opt/lcg/bin/rfiod -sl -f /var/log/rfio/log
```

```
$ ps -ef|grep ftp
root 20291 1 0 Sep19 ? 00:00:03 /opt/lcg/sbin/dpm.ftpd -i -X -L -l -S -p 2811 -u 002 -o -a -Z /va
```

Also check that :

- the `dpmmgr` user has been created before `rfiod` and `dpm-gsift` were started,
- the `dpmmgr` user has the same uid and gid **on all disk servers**.

LFC daemon crashes with old oracle database 10gR2

Error

The LFC daemon crashes regularly with Oracle 10gR2 database backend.

What can I do ?

Solution

You have to use the 10gR2 Oracle Instant Client, instead of the 10gR1 one.

Remember to change `$ORACLE_HOME` in `/etc/sysconfig/lfcdaemon` to point to the right directory.

And restart the service :

```
$ service lfcdaemon restart
```

For further help: Get a core dump, by uncommenting the following line in `/etc/sysconfig/lfcdaemon` :

```
#ALLOW_COREDUMP="yes"
```

And restarting the service :

```
$ service lfcdaemon restart
```

The core dump will appear under `/home/lfcmgr/lfc`.

Put the core dump in a public location, and send this location to helpdesk@ggusNOSPAMPLEASE.org (remove the NOSPAM!) : your ROC will help you, and contact the appropriate experts if needed.

File exists

Error

You get this error :

```
lfc-rm /grid/atlas/tests/file1
/grid/atlas/tests/file1: File exists
```

or this

```
dpns-rm /dpm/in2p3.fr/home/auvergrid/tests/file1
/dpm/in2p3.fr/home/auvergrid/tests/file1: File exists
```

Solution

`lfc-rm` and `dpns-rm` remove the entry in the Name Server only, but not the physical file itself.

The `File exists` error means that there are still physical replicas attached to the Name Server entry.

To remove both physical and logical files, you can :

- use `lcg_util`
- use `rfrm` (in the DPM case)

VOMS signature error

Error

You get this error in `/var/log/lfc/log` or `/var/log/dpns/log` :

```
05/19 12:05:13 16051,0 Cns_serv: Could not establish security context: _Csec_get_voms_creds: VOMS
```

Solution

On the LFC/DPNS machine, the host certificate of your VO VOMS server is missing in `/etc/grid-security/vomsdir`.

For instance :

```
$ ls /etc/grid-security/vomsdir | sort
cclcgvomsli01.in2p3.fr.43
lcg-voms.cern.ch.1265
voms.cern.ch.1877
voms.cern.ch.963
```

grid-proxy-init OK, but voms-proxy-init NOT OK

Problem

For a given user, usage of LFC/DPM with:

- `grid-proxy-init` or simple `voms-proxy-init` works fine,
- `voms-proxy-init -voms` doesn't work fine

Solutions

Wrong VOMS setup

Check the VOMS setup on:

- the UI
- the LFC / DPM server

On LFC & UI, `/etc/grid-security/vomsdir` contains VO VOMS server

```
$ ls -ld /etc/grid-security/vomsdir/
drwxr-xr-x  2 root  root  4096 Jun  8 15:07 /etc/grid-security/vomsdir/

$ ls /etc/grid-security/vomsdir
cclcgvomsli01.in2p3.fr.43
lcg-voms.cern.ch.1265
```

On the UI (client), `/opt/glite/etc/vomses` should contain :

```
$ ls /opt/glite/etc/vomses
alice-lcg-voms.cern.ch
alice-voms.cern.ch
```

User uses several different VOMS roles

For details, see LFC and DPM internal virtual ids

The same user with two different VOMS roles will be mapped to two different internal virtual gids. To grant privileges to other VOMS roles on given directories/files, use `lfc-setacl` (see `man lfc-setacl`).

lcg_utils : "Invalid Argument" error

Error

An `lcg_util` command returns the `Invalid Argument` error.

Solution

It usually means that there is a problem with the information published by the Information System. Either :

- for the LFC, or
- for the Storage Element

"Could not establish security context: Connection dropped by remote end !"

Error

This error appears in the LFC/DPM log file.

```
07/28 10:08:22 18550,0 Cns_serv: Could not establish security context: _Csec_rcv_token: Connecti
```

Explanation

This is not a problem.

This warning only means that the LFC/DPM client dropped the connection itself.

For instance, it appears in the server log file, if a user doesn't have a valid proxy :

```
$ lfc-ls /
send2nsd: NS002 - send error : No valid credential found
/: Bad credentials
```

What to do if the DN of a user changes ?

Problem

The DN of a user changes. What does the LFC/DPM admin have to do, so that the user can still access her files ? **Problem**

The name of a group/VO changes. What does the LFC/DPM admin have to do, so that the permissions remain correct ?

Solution Use the `lfc-modifyusrmap` or `lfc-modifygrpmap` commands. See `man lfc-modifyusrmap` and `man lfc-modifygrpmap`.

What to do if the host certificate expired or going to be changed

Problem

grid-proxy-init OK, but voms-proxy-init NOT OK

The LFC or DPM server host certificate will expire soon.

Solution

Replace the old host certificate and key :

```
$ ll /etc/grid-security/ | grep host
-rw-r--r--    1 root    root          5423 May 27 12:35 hostcert.pem
-r-----    1 root    root          1675 May 27 12:35 hostkey.pem
```

At the same time, a renamed copy of them has to be put under :

```
$ ll /etc/grid-security/lfcmgr | grep lfc
-rw-r--r--    1 lfcmgr  lfcmgr          5423 May 30 13:58 lfccert.pem
-r-----    1 lfcmgr  lfcmgr          1675 May 30 13:58 lfckey.pem
```

You don't need to restart any of the services then.

Note : replace `lfcmgr` with `dpnmgr` for the DPM.

How do ACLs work ?

Question

How do ACLs work in the LFC or DPM Name Server ?

Answer

ACLs are standard POSIX ACLs.

For details, see `man lfc-setacl` or `man dpns-setacl`.

If a same file has several Logical File Names (LFNs), this file has :

- a primary LFN,
- secondary LFNs : they are implemented as symlinks, and have dummy `777` permissions.

When an LFN (primary or secondary) is accessed, the permissions/ACLs on the primary LFN are checked.

How to know all the file residing on a given SE ?

Question

How can I know all the replicas stored on a given Storage Element ?

Answer

The `"lfc_listreplicax"` method allows to do this : it lists all the replica entries stored in the LFC for a given server.

It is available in :

- the LFC C API,
- the LFC Python interface,
- the LFC Perl interface

See `man lfc_listreplicax`.

Warning

This method is based on the `host` field in the `Cns_file_replica` table.

But be aware that **some VOs don't store the actual server machine name in the `host` field !**

For instance, in its LFC central server, LHCb stores `CERN_Castor` instead of `castorsrm.cern.ch...`

In the future, `srmLs` can be used too. But it has to be implemented for all Storage Element types first.

How to restrict a pool to a given VO ?

It is possible to have one pool dedicated to a given VO, with all the authorization behind, using the `dpm-addpool` or `dpm-modifypool` commands.

For instance :

```
dpm-addpool --poolname VOpool --def_filesize 200M --gid the_VO_gid
```

```
dpm-addpool --poolname VOpool --def_filesize 200M --group the_VO_group_name
```

Comment:

If you define :

- one pool dedicated to `group1 / VO1`
- one pool open to all groups / VOs

then, the **dedicated pool will be used until it is full**.

When the dedicated pool is full, the open pool is then be used.

R-GMA solutions

General, very simple R-GMA test

Question

How can I test if I've set up RGMA correctly?

Answer

R-GMA developers provide 2 scripts for testing the installation.

```
/opt/edg/bin/rgma-client-check  
/opt/edg/bin/rgma-server-check
```

Which logs should I back up for accounting purposes?

Question

I need to know which logs to back up for accounting purposes.

How to know all the file residing on a given SE ?

Answer 

This question is answered on the Accounting FAQ page at the UK GOC and the list, in short, comprises:

- Gatekeeper logs: /var/log/globus-gatekeeper.log.*
- Job Manager logs: /var/spool/pbs/server_priv/accounting/*
- System logs: /var/log/messages*

Note 

Note that there may be other logs that it is necessary to retain for security audit reasons.

Failed to get list of tables from the Schema

Error

Something like this one:

```
=====
You are connected to the following R-GMA Schema service:
  https://lcgic01.gridpp.rl.ac.uk:8443/R-GMA/SchemaServlet
WARNING: failed to get list of tables from the Schema
=====
```

Solution 

Generally this error message appears when one would like to connect to a secure R-GMA server a.) without a user proxy or b.) having a user proxy but the X509_USER_PROXY enviromental variable is not pointing to the proxy.

Comment 

Note, that the `grid-proxy-init` does not set the value of the `X509_USER_PROXY` variable.

Problems with `rgma-client-check`

Unable to source `/opt/edg/etc/profile.d/edg-rgma-env.sh`

Error

Running R-GMA client checking script

```
/opt/edg/sbin/test/edg-rgma-run-examples
Unable to source /opt/edg/etc/profile.d/edg-rgma-env.sh
```

Solution 

R-GMA has not been configured. Configure R-GMA.

RGMA_HOME is not set**Error**

Running R-GMA client checking script

```
/opt/edg/bin/rgma-client-check
RGMA_HOME is not set
```

Solution 

R-GMA is not configured. Configure R-GMA or set the environment variable RGMA_HOME

No C++ compiler found**Error**

Running `rgma-client-check` gives:

```
/opt/edg/sbin/test/edg-rgma-run-examples

Configuring...
No C++ compiler found
```

Solution 

This testing script requires a C++ compiler to complete successfully. Install both the `gcc-c++` and `openssl-devel` packages for the operating system.

Cannot declareTable: table description not defined in the Schema**Error**

Running `rgma-client-check` gives:

```
/opt/edg/bin/rgma-client-check

*** Running R-GMA client tests on cmsfarmbl12.lnl.infn.it ***

Checking C API: Failed to declare table.

Failure
Checking C++ API: R-GMA application error in PrimaryProducer.
Cannot declareTable: table description not defined in the Schema
Success
Checking Python API: RGMA Error          StreamProducer__declareTable_StringString:Cannot declareT

Failure
Checking Java API: R-GMA application error in PrimaryProducer.
org.glite.rgma.RGMAException: Unknown RGMA Exception: Cannot declareTable: table description not
    at org.glite.rgma.stubs.PrimaryProducerStub.declareTable(Unknown Source)    at Primar

Failure
Checking for safe arrival of tuples, please wait... ERROR: Failed to instantiate Consumer
There should be 4 tuples, there was only:
```

Solution 

The Registry servlet has a hosts allow file and the site R-GMA server machine is not registered in this file.

Running:

```
wget http://lcgic01.gridpp.rl.ac.uk:8080/R-GMA/SchemaServlet
cat SchemaServlet

<?xml version = '1.0' encoding='UTF-8' standalone='no'?>
<edg:XMLResponse xmlns:edg='http://www.edg.org'>
<XMLException type="SchemaException" source="Servlet" isRecoverable="false">
<message>cannot service request, client hostname is currently being blocked</message>
</XMLException>
</edg:XMLResponse>
```

This shows that the host you running this command on is currently blocked. Send a mail to lcg-support@gridppNOSPAMPLEASE.rl.ac.uk for the allow list to included the machine running the R-GMA server. In the email, specify the full machine name as well as the full domain. For instance:

Hi,

Please could you add MY-SITE to the R-GMA Registry.

R-GMA Server : mon.my-site.my-domain

Domain : my-domain

libgcj-java-placeholder.sh

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```
/opt/edg/bin/rgma-client-check
```

```
Checking C API: Done.
```

```
Success
```

```
Checking C++ API: Success
```

```
Checking Python API: Success
```

```
Checking Java API: libgcj-java-placeholder.sh
```

This script is a placeholder for the `/usr/bin/java` and `/usr/bin/javac` master links required by `jpackage.org` conventions. `libgcj's` `rmiregistry`, `rmic` and `jar` tools are now slave symlinks to these masters, and are managed by the `alternatives(8)` system.

This change was necessary because the `rmiregistry`, `rmic` and `jar` tools installed by previous versions of `libgcj` conflicted with symlinks installed by `jpackage.org` JVM packages.

```
Success
```

```
Checking for safe arrival of tuples, please wait... There should be 4 tuples, there was only:
```

```
| C producer      |
| C++ producer   |
| Python producer|
```

Solution

The default installation of linux puts a placeholder for the `java` command. This is being pick up instead of the proper `java` command.

Make sure that Java has been installed and that the `java` command is found in the path before the placeholder.

Connection refused

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```

*** Running R-GMA client tests on alifarm19.ct.infn.it ***

Checking C API: Failed to create producer.

Failure
Checking C++ API: R-GMA application error in PrimaryProducer.
Cannot open connection to servlet: Connection refused
Success
Checking Python API: RGMA Error Failed to instantiate StreamProducer
Failure
Checking Java API: Failed to contact PrimaryProducer service.
org.glite.rgma.RemoteException
    at org.glite.rgma.stubs.ProducerFactoryStub.createPrimaryProducer(Unknown Source)
    at PrimaryProducerExample.main(Unknown Source)
Failure

Checking for safe arrival of tuples, please wait... ERROR: Failed to instantiate Consumer
There should be 4 tuples, there was only:

```

Solution

The tomcat and the servlets are not up and running. Restart Tomcat and check the Tomcat logs for errors. As root do the following:

```

/etc/rc.d/init.d/tomcat5 stop (use Ctrl-C if this hangs.)
su - tomcat4 -c 'killall -9 java'
rm -f /var/log/tomcat5/catalina.out
/etc/rc.d/init.d/tomcat5 start
tail -f /var/log/tomcat5/catalina.out

```

Note

Note: tomcat5 runs as user tomcat4 !!!

HTML returned instead of XML

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```

/opt/edg/bin/rgma-client-check

*** Running R-GMA client tests on node064.lancs.pygrid ***

Checking C API: Failed to create producer.

Failure
Checking C++ API: R-GMA application error in PrimaryProducer.
HTML returned instead of XML. This usually means either there is a problem with the proxy cache,
Success
Checking Python API: RGMA Error Failed to instantiate StreamProducer
Failure
Checking Java API: Failed to contact PrimaryProducer service.
org.glite.rgma.RemoteException

```

The LCG Troubleshooting Guide < LCG < TWiki

```
at org.glite.rgma.stubs.ProducerFactoryStub.createPrimaryProducer(Unknown Source)
at PrimaryProducerExample.main(Unknown Source)
```

Failure

Checking for safe arrival of tuples, please wait... ERROR: Failed to instantiate Consumer
There should be 4 tuples, there was only:

Solution

A previous configuration script for R-GMA removed some jar files that were in deployed in the Tomcat rpm.
Checking the rpm shows the error:

```
rpm -V tomcat4
.....GT c /etc/tomcat4/server.xml
SM5..U.T c /etc/tomcat4/tomcat-users.xml
S.5....T c /etc/tomcat4/tomcat4.conf
missing  /var/tomcat4/common/endorsed/jaxp_parser_impl.jar
missing  /var/tomcat4/common/endorsed/xml-commons-apis.jar
```

Re-install tomcat4 !

No tuples returned

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```
/opt/edg/bin/rgma-client-check
```

```
*** Running R-GMA client tests on bf35.tier2.hep.man.ac.uk ***
```

```
Checking C API: Done.
Success
Checking C++ API: Success
Checking Python API: Success
Checking Java API: Success
```

Checking for safe arrival of tuples, please wait... There should be 4 tuples, there was only:

Solution

- The clocks could be out and the producers are probably being cleaned up as soon as they have been created. Check that the time is correct. NTP needs to be running on all nodes.
- Port 8088 could be blocked by a firewall. Run the `rgma-server-check` on the R-GMA server and open port 8088 in the firewall if it reports that it is blocked.

Object has been closed: 1949004681

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```
+ /opt/edg/bin/rgma-client-check
```

```
*** Running R-GMA client tests on egeewn14.ifca.org.es ***
```

```
Checking C API: Done.
Success
Checking C++ API: Success
Checking Python API: Success
```

HTML returned instead of XML

Checking Java API: Success

```
Checking for safe arrival of tuples, please wait... ERROR: Consumer__isExecuting:Servlet not
  Caused by:
  Object has been closed: 1949004681
```

There should be 4 tuples, there was only:

Solution

The clocks could be out and the producers are probably being cleaned up as soon as they have been created. Check that the time is correct. NTP needs to be running on all nodes including the R-GMA servlet box.

Unable to locate an available Registry Service

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```
/opt/edg/bin/rgma-client-check
*** Running R-GMA client tests on PAKWN1.pakgrid.org.pk ***
Checking C API: Failed to create producer.
Failure
Checking C++ API: R-GMA application error in PrimaryProducer.
Unable to locate an available Registry Service
Success
Checking Python API: RGMA Error Failed to instantiate StreamProducer
Failure
Checking Java API: R-GMA application error in PrimaryProducer.
org.glite.rgma.RGMAException: Unable to locate an available Registry Service
    at org.glite.rgma.stubs.ProducerFactoryStub.createPrimaryProducer(Unknown Source)
    at PrimaryProducerExample.main(Unknown Source)
Failure
Checking for safe arrival of tuples, please wait... ERROR: Failed to instantiate Consumer
There should be 4 tuples, there was only:
*** R-GMA client test failed ***
```

Solution

The configuration on the R-GMA server is incorrect. Using the R-GMA browser on the R-GMA server and looking at "Table Sets" should show and error message.

```
Cannot connect to servlet:
```

Correctly configure the R-GMA server to point to the correct Registry and Schema.

cannot remove `/tmp/cmds.sql': Operation not permitted

Error

Running `/opt/edg/bin/rgma-client-check` gives:

```
Checking for safe arrival of tuples, please wait... /opt/edg/bin/rgma-client-check: line 99: /tmp
There should be 4 tuples, there was only:
rm: cannot remove `/tmp/cmds.sql': Operation not permitted
```

Object has been closed: 1949004681

Solution 

The file has probably been created when the client check script command was run as root or as a pool account. A new pool account is now unable to delete the file. Delete the file. A fix is in the latest version of R-GMA which will be deployed with the next R-GMA version to be deployed.

Information System (and BDII) solutions

General considerations

LCG uses an LDAP based information system. Click [here](#) for a quick introduction to LDAP.

The LCG information system consists of four distinct parts. The Generic Information Provider (GIP), the MDS, GRIS, the site BDII and the top level BDII.

All the information is produced by the information provider, everything else is the transport mechanism. If there are any problems with the information then the information provider will need to be investigated. Each site should produce the following information.

- One `SiteInfo` entry.
- One `GlueCluster` and `GlueSubCluster` entry per cluster.
- One `GlueCE`, `GlueCESEBind` and `GlueCESEBindGroup` entry per queue.
- One `GlueSE` and `GlueSL` entry per Storage Element.
- One `GlueSA` entry per VO.

If the correct information for the site is in the top level BDII then there is usually no problem. For this reason we can take a top down approach for trouble shooting. See the following 4 entries in the topic.

Check that the information is in the top level BDII

The following query can be used to extract the information about the site from the top level BDII. Replace `bdii-host.invalid` with the BDII host and `domain.invalid` with the domain name of the site. An assumption has been made in the query where the mail address for the `sysAdminContact` contains the domain name of the site.

```
ldapsearch -LLL -x -h bdii-host.invalid -p 2170 -b o=grid\
'(|(GlueChunkKey=*domain.invalid)(GlueForeignKey=*domain.invalid)(GlueInformationServiceURL=*domain.invalid)\
(GlueCESEBindSEUniqueID=*domain.invalid)(GlueCESEBindSEUniqueID=*domain.invalid)\
(GlueCESEBindGroupSEUniqueID=*domain.invalid)(sysAdminContact=*domain.invalid))'
```

Adding to the end of the command,

```
dn | grep dn | cut -d "," -f 1
```

will show just the entries.

Check that the information is in the site level BDII

To check that the information for the site is in the site `bdii`, do the following `ldapsearch`, replacing `site-bdii.invalid` with the hostname of the machine running the site BDII.

```
ldapsearch -x -h site-bdii.invalid -p 2170 -b o=grid.
```

Check that the information is is the GRIS

To check that the information for is in a GRIS, do the following ldapsearch, replacing gris-host.invalid with the hostname of the machine running the GRIS.

```
ldapsearch -x -h gris-host.invalid -p 2135 -b mds-vo-name=local,o=grid.
```

Check that the information is returned by the information provider

Run the following command to check the output of the information provider.

```
/opt/lcg/libexec/lcg-info-wrapper.
```

No information found in BDII

If there is no information returned, then there is a problem with either the URL used to obtain the information or the information source itself. The URLs are found in the file `/opt/lcg/var/bdii/lcg-bdii-update.conf`. Find the URL in the file and transform it into and ldapsearch.

```
NAME ldap://host.invalid:port/bind
ldapsearch -x -h host.invalid -p port -b bind
```

Entry's missing in the BDII

If invalid LDIF is produced, then the entry will be rejected when it is being inserted in to the LDAP database. To see if any entries are being rejected run the BDII update script.

```
/opt/lcg/libexec/lcg-bdii-update /opt/lcg/var/bdii/lcg-bdii.conf
```

The dn of any rejected entries will be shown along with the error. This will also show if any problems with the ldap URLs.

Problems updating the BDII configuration file from the web

Check that the attribute `BDII_AUTO_UPDATE` in the configuration file `/opt/lcg/var/bdii/lcg-bdii.conf` is set to "yes". If this value is set to "no" the BDII will not attempt to update the configuration file from the web. Next check that the value for the attribute `BDII_HTTP_URL` points to an existing web page and that this web page is the file that contains the URLs that you want to use for the BDII.

Can not connect to the GRIS

Check the status of the GRIS.

```
/etc/rc.d/init.d/globus-mds status
```

If the GRIS failed to start, try to restart it.

```
/etc/rc.d/init.d/globus-mds restart.
```

Repeat this this command a few times. If it fails on stopping the GRIS then it usually means that it failed to start.

The GRIS fails to start

The GRIS sometimes fails to start due to stale slapd processes being left around. Try to removed all these.

```
kill -9 slapd.
```

Note that if the BDII is on the same machine this will now need to be restarted. Try re-starting the GRIS a few times.

```
/etc/rc.d/init.d/globus-mds restart.
```

If it fails on stopping the GRIS then it usually means that it failed to start. Try starting the GRIS by hand with debugging turned on. This should show up any errors.

```
/opt/globus/libexec/slapd -h ldap://localhost:2135 -f /opt/globus/etc/grid-info-slapd.conf -d 255
```

No information returned by the GRIS

If no information is returned, then either the information provider is not working or there is a problem with the GRIS configuration.

There is a problem with the GRIS configuration

Check that the entry for the information provider is in the GRIS configuration file `/opt/globus/etc/grid-info-resource-ldif.conf`. This file is automatically created from the `globus-mds init.d` script. It uses the file `/opt/edg/var/info/edg-globus.ldif` get the entry.

No information was produced by the information provider

Check that the static ldif file has been created. The static ldif file location is defined in the file `/opt/lcg/var/lcg-info-generic.conf` and by default is `/opt/lcg/var/lcg-info-static.ldif`. If this file does not exist try to re-run the configuration to create it.

```
/opt/lcg/sbin/lcg-info-generic-config /opt/lcg/var/lcg-info-generic.conf
```

If this does not create the ldif file check the contents of the file `/opt/lcg/var/lcg-info-generic.conf`. There should be at least one template and one dn specified in this file.

Default values show instead of dynamic values

The dynamic plug has a problem or there is a miss-match with the dn's. The command used to run the dynamic plug-in is in the file `/opt/lcg/var/lcg-info-generic.conf`. Copy and paste the command on to the command line and execute it. This should show up any errors. Check that the dn's produced by the dynamic plug-in are the same as in the static ldif file.

New values not shown in GRIS

This can occur because a stale slapd processes is left around and is still serving the data even after a restart. This error can usually be found be doing `globus-mds stop` . The command will fail and you should still be able to do a query. The solution is to kill all the slapd process and restart the GRIS.

```
kill -9 slapd.
```

Note that if the BDII is on the same machine this will now need to be restarted.

How to set up a dns load balanced BDII service.

Question

How to use several BDII and load sharing ?

Solution

Multiple BDIIs can be used behind a "round robin" dns alias to provide a load balance BDII Service.

No such object (32): error message

Error

Gstat BDIIUpdate Check gives following error:

```
No such object (32)
```

Solution

BDIIUpdate Check tries to update the bdii database by contacting each GIIS listed at:

```
http://grid-deployment.web.cern.ch/grid-deployment/gis/lcg2-bdii/dteam/lcg2-all-sites.conf
```

If your site has this error, you should check try to query the contact string listed in the bdii config above and verify that it is functioning properly. If the contact string is incorrect please email the ROLLOUT list to request a change. A search example:

```
ldapsearch -x -H ldap://<giis host>:2170 -b mds-vo-name=<sitename>,o=grid
```

How to close the site so it won't receive anymore jobs from the RBs

Question

How to close the site so it won't receive anymore jobs from the RBs

If you want to stop the RB from sending you jobs (for example as you want to do some update on your CE), an attribute exists in the Idif Schema which is consulted by the RB to check the availability of your site. This page explains how to publish a closed status on your farm. It's about the information system. The right place

The attributes GlueCEStateStaus can take some values for which the RB will look. These attributes may be :

- **Queueing**: the queue can accept job submission, but can't be served by the scheduler
- **Production**: the queue can accept job submissions and is served by a scheduler
- **Closed**: The queue can't accept job submission and can't be served by a scheduler
- **Draining**: the queue can't accept job submission, but can be served by a scheduler

This attribute is published under the dn : GlueCEUniqueId\=hostname... And such a dn exists for each queue.

Answer

Now we are going to change the value of this attribute. You'll have to edit the /opt/lcg/var/gip/lcg-info-generic.conf Find the line whith the right dn. If it doesn't allready exist, add the line :

GlueCEStateStatus: Closed

for closing your site.

else, you'll only have to change the value of this attribute. Be careful to remove any space at the end of the line. Do this for each queue you have to change. You should find a dn for each of these queues. To activate the changes use the command:

```
/opt/lcg/sbin/lcg-info-generic-config /opt/lcg/var/gip/lcg-info-generic.conf=
```

Don't forget that, if you're using a BDII as GIIS, you have to wait until the BDII refreshes itself or refresh it manually. **Note** 

If you want to remove the closed status of your site, simply remove the line you added or change the value at will.

Job submission solutions

10 data transfer to the server failed

Error

Globus job manager on the CE cannot call back RB (or UI in tests)

Solution

- Check if the account to which the DN is mapped has a writable home directory. A globus-job-run (instead of edg-job-get-logging-info) may report this error:

```
GRAM Job submission failed because cannot access cache files in
~/globus/.gass_cache, check permissions, quota, and disk space
(error code 76)
```

- Check contents of \$GLOBUS_LOCATION/etc/grid-services/jobmanager-* files.
- Check contents of \$GLOBUS_LOCATION/etc/globus-job-manager.conf.
- Ensure /etc/grid-security is world-readable (only hostkey.pem must be protected).
- Ensure outgoing connections are allowed from the CE to the GLOBUS_TCP_PORT_RANGE on RB (or UI).

SAM solutions

VOMS solutions

Wrong host certificate subject in the vomses file

It is possible that after renewing a host certificate, the host certificate subject changes and the vomses file containing the VOMS server information is not updated accordingly.

The client side message is like in the following example:

```
bash-2.05b$ voms-proxy-init -voms mysql_vol -userconf ~/vomses
Your identity: /C=CH/O=CERN/OU=GRID/CN=Maria Alandes Pradillo 5561 Enter GRID p
Creating temporary proxy ..... Done
Contacting lxb0769.cern.ch:15001 [/C=CH/O=CERN/OU=GRID/CN=lxb0769.cern.ch] "my
```

```
Error: Could not establish authenticated connection with the server.
GSS Major Status: Unexpected Gatekeeper or Service Name GSS Minor Status Error

an unknown error occurred

Failed to contact servers for mysql_vol.
```

The server log file contains the following lines:

```
Wed Aug 16 11:04:48 2006:lxb0769.cern.ch:vomsd(4341):ERROR:REQUEST:AcceptGSIAuth
home/glbuid/GLITE_3_0_0_final/org.glite.security.voms/src/socklib/Server.cpp:25
security context (accept):.GSS Major Status: General failure.GSS Minor Status Er
Chain:..accept_sec_context.c:305:gss_accept_sec_context: Error during delegation
violation
```

In this case it's good that you check whether the vomses file contains the correct host certificate subject. To check what's your VOMS host certificate subject, run the following command:

```
[root@lxb0769 root]# openssl x509 -in /etc/grid-security/hostcert.pem -noout -su
subject= /C=CH/O=CERN/OU=GRID/CN=host/lxb0769.cern.ch
```

And check in the vomses file that the certificate subject is correct:

```
bash-2.05b$ more vomses
...
"mysql_vol" "lxb0769.cern.ch" "15001" "/C=CH/O=CERN/OU=GRID/CN=host/lxb0769.cern
...
```

Database initialization error with

When installing VOMS MySQL sometimes the following error appears just after starting the VOMS server: Database initialization error.

This could be caused because before the configuration of the server, the following commands were not executed:

```
/usr/bin/mysqladmin -u root password 'yourPassword'
/usr/bin/mysqladmin -u root -h yourHostname password 'yourPassword'
```

When installing VOMS MySQL it is extremely important to execute the mentioned commands before configuring VOMS. Although this is specified in the Installation guide that can be found here [here](#) many people don't read it.

It is also mentioned when VOMS MySQL rpms are installed using APT. However, since many messages and warnings appear it is easy to miss the message that warns about the need of executing the above mentioned commands.

WARNING: Unable to verify signature!

Error

Running `voms-proxy-info` gives the following error:

```
error = 5025
WARNING: Unable to verify signature!
subject : /O=GermanGrid/OU=LMU/CN=John Kennedy/CN=proxy
...
..
```

Wrong host certificate subject in the vomses file

While `voms-proxy-init` is OK:

```
voms-proxy-init -voms atlas
```

```
Your identity: /O=GermanGrid/OU=LMU/CN=John Kennedy
Enter GRID pass phrase:
Creating temporary proxy .....
Done
Contacting voms.cern.ch:15001 [/C=CH/O=CERN/OU=GRID/CN=host/voms.cern.ch]
"atlas" Error: VERR_NOSOCKET Failed.
Trying next server for atlas.
Creating temporary proxy .....
Done
Contacting lcg-voms.cern.ch:15001
[/C=CH/O=CERN/OU=GRID/CN=host/lcg-voms.cern.ch] "atlas"
Creating proxy ..... Done
Your proxy is valid until Mon Jul 17 13:36:56 2006
```

Solution

It just means that you don't have the VOMS server host certificate (or at least v-p-i can't find it) so the code can't verify that the VO signature is valid. It doesn't matter if you just want to see the info.

APT solutions

`apt-get update` : W: Release file did not contain checksum information for

Error

Running `apt-get update` gives a message similar to this one:

```
W: Release file did not contain checksum information for http://grid-
deployment.web.cern.ch/grid-deployment/gis/apt/LCG-2_7_0/sl3/en/i386/base/pkglist.lcg_sl3
W: Release file did not contain checksum information for http://grid-
deployment.web.cern.ch/grid-deployment/gis/apt/LCG-2_7_0/sl3/en/i386/base/release.lcg_sl3
W: Release file did not contain checksum information for http://grid-
deployment.web.cern.ch/grid-deployment/gis/apt/LCG-2_7_0/sl3/en/i386/base/pkglist.lcg_sl3.securit
W: Release file did not contain checksum information for http://grid-
deployment.web.cern.ch/grid-deployment/gis/apt/LCG-2_7_0/sl3/en/i386/base/release.lcg_sl3.securit
W: You may want to run apt-get update to correct these problems
```

Solution

There is a problem on the server side, thus please send an e-mail to `lcg-rollout@listservNOSPAMPLEASE.cclrc.ac.uk` including the error message.

FTS Solutions

I tried to submit a job and it said: `submit: You are not authorised to submit jobs to this service`

The user is not authorised to submit jobs to the FTS service. In order to authorize him/her, you have to add his/her DN in the `submit-mapfile` on the FTS server. You can have a look at `FtsServerInstall112` in the `Mapfile` section and at `FtsServerSubmitMapfile13`

However, due to bug in the FTS (#10362), if the user has a double or more delegated proxy (i.e. the DN ends with /CN=proxy/CN=proxy), a parsing error will cause a authorization denied. This bug has being solved in FTS version 1.4 and in the latest QuickFix for 1.3

If the user is still not authorized to submit request, check his/her DN is not in the veto-mapfile

I submitted a job from site X to Y but it didn't work. The channel Y-X exists and has a share for my VO!

From version 1.3 onwards the channel definitions are mono-directional. You have to create another channel in the opposite direction (`glite-transfer-channel-add`), set the share for the VO interested in using the channel (`glite-transfer-channel-setvoshare`) and install an Channel Agent that will managed it

Which format should I use for the SURLs?

Starting from gLite 1.4.1, the FTA implements the enhancement request #8364, that allows a user to specify any format he prefers: the agent would then convert each SURL before transferring or registering into the catalog to either a fully qualified format

```
srm://<host>:<port>/srm/managerv1?SFN=<file_path>
```

or a compact one

```
srm://<host>/<file_path>
```

depending on the configuration. By default it would use the compact format. In case you want to change this parameter, you have to set the related ChannelAgent configuration parameter `transfer-agent-channel-actions.SurlNormalization` to one of the following values:

- `compact` all the SURLs will be converted to the format:

```
srm://<host>/<file_path>
```

- `compact-with-port` all the SURLs will be converted to the format:

```
srm://<host>:<port>/<file_path>
```

- `fully-qualified` all the SURLs will be converted to the format:

```
srm://<host>:<port>/srm/managerv1?SFN=<file_path>
```

- `disabled` no SURL conversion will be performed

If you're using a previous version, for interoperability reasons we suggest to use fully qualified SURLs, i.e. in the format

```
srm://<srm_host>:<srm_port>/srm/managerv1/?SFN=<file_path>
```

If you know the type of the SRM that would be involved in the transfer, you can also specify one of the supported compact format. For Castor, as example, you can use

```
srm://<castorsrm>:8443/srm/managerv1?SFN=<file_path>  
srm://<castorsrm>:8443//srm/managerv1?SFN=<file_path>  
srm://<castorsrm>:8443/?SFN=<file_path>
```

I tried to submit a job and it said: submit: You are not authorised to submit jobs to this service

```
srm://<castorsrm>:8443/<file_path>  
srm://<castorsrm>/<file_path>
```

In case the transfer is processed by a channel configured to use `srmcopy`, the fully qualified format may not work. Please have a look here for a workaround

I've tried to submit a job but I get back an error saying: SOAP-ENV:Server.userException - org.xml.sax.SAXException

Usually this issue is related to an endpoint pointing to the wrong server (typically `ChannelManagement` instead on `FileTransfer`): when you observe an error similar to

```
submit: SOAP fault: SOAP-ENV:Server.userException -  
org.xml.sax.SAXException: Deserializing parameter 'job': could not find deserializer for type {h
```

please ask the user to look at the command he just submitted and to check that the specified endpoint is correct; all the CLIs commands that start with `glite-transfer-channel-*` require to use a `ChannelManagement` interface, while the ones that start with `glite-transfer-*` require the `FileTransfer` interface. In order to check if the endpoint is correct, the user can also re-run the command with the `-v` option and checks if the line `Using Endpoint` ends with `FileTransfer` or `ChannelManagement`

I've tried to submit a job but I get back an error saying: No match

When the user submit a transfer job, he usually specify some SURLs that may contains a question mark (?). In some shells this character has to be escaped by simply quoting it ('?'): for example, if the SURLs are

```
srm://castorgridsc.cern.ch:8443/srm/managerv1?SFN=/castor/cern.ch/grid/dteam/src_file  
srm://castorgridsc.cern.ch:8443/srm/managerv1?SFN=/castor/cern.ch/grid/dteam/dst_file
```

please make sure you run `glite-transfer-submit` in this way

```
glite-transfer-submit \  
  srm://castorgridsc.cern.ch:8443/srm/managerv1?'SFN=/castor/cern.ch/grid/dteam/src_file \  
  srm://castorgridsc.cern.ch:8443/srm/managerv1?'SFN=/castor/cern.ch/grid/dteam/dst_file
```

I was able to list the channels but I cannot get the channel details

Listing channels is open to any user as long as he/she is not in the veto mapfile - you only get the channel name from this call.

However, getting the details of a channel - source, destination, bandwidth, etc is restricted. For this you need to be:

- an admin
- manager of the channel being queried
- manager of any VO on the given FTS

You can check your roles on a given FTS by running `glite-transfer-getroles`. Information on channel and VO managers can be managed by a service admin or other managers by using the appropriate client tools. Information on service ADMINS is stored inside the `admin-mapfile`.

How do I setup a non-dedicated Channel?

Non-dedicated channels (a.k.a. "catch-all" channels) are a special channel configuration that allows matching any site as source or destination, therefore not coupled with the underlying network. Using "catch-all" channels allows to limit the number of channels you need to manage, but also limits the degree of control you have over what is coming into your site (although it still provides the other advantages like queueing, policy enforcement and error recovery). The usage of these channels is mainly recommended in Tier1 for providing full connectivity to all other sites, where the suggested channels definition is:

- Dedicated channels from any other Tier1 to the T1
- Non-dedicated channels to each of the related Tier2
- A non-dedicated channel to the T1

You can setup a non-dedicated channel that will manage all the transfers from any site to your site by issuing a `glite-transfer-channel-add` and using `*` and source site name, like:

```
glite-transfer-channel-add -f NUM_OF_FILES -S CHANNEL_STATE [...] CHANNEL_NAME "*"
YOUR_SITE
```

Of course, you have then to issue a `glite-transfer-channel-setvoshare` for each VO that should be authorized to use the channel and then configure a ChannelAgent for that channel.

Please note that if a VO is not authorized to use a channel between site A and B but has privileges on a `*-B` channel, transfer requests for that VO from site A to B are denied since the non-dedicated channel is evaluated *after* all the dedicated ones.

In addition, please also note that the default ChannelAgent configuration for that channel requires that all the SRM that would be involved in the managed transfers should be listed in the information system. In case a VO needs to relax this constraint, for example in order to transfer files to/from Classic SEs not included in the information system, the following parameters should be added to the VOAgent configuration:

- `transfer-agent-vo-actions.EnableUnknownSource` should be set to `true` if SEs not known to the InfoSys should be allowed as valid source (these would be matched by the `*-Site` catch-all channels)
- `transfer-agent-vo-actions.EnableUnknownDest` should be set to `true` if SEs not known to the InfoSys should be allowed as valid destination (these would be matched by the `Site-*` catch-all channels)

In case a VO needs these parameters, it would be better to turn off the SURL Normalization, or at least set it to `fully-qualified`, for all the ChannelAgents associated to non-dedicated channels, since it would be impossible to resolve the correct endpoint for the SRM not listed in the InformationSystem. It will also be worth to recommend the users to use fully-qualified SURLs for transfers that should be processed through these channels.

Use of the `*-*` 'catch everything' channel is not recommended for production grids.

After upgrading to FTS 1.5 I got "No Channel found or VO not authorized ..." error

Symptom: After upgrading to FTS 1.5 I got "No Channel found or VO not authorized ..." error

Running the FTS service we encountered many inconsistencies in the way the information was published in BDII, especially related to the case used to publish the site name. This is not a problem when BDII is used directly, since it's case insensitive, but creates some interoperability issues when used via ServiceDiscovery (that is case sensitive). We therefore decided to apply a convention, within the FTS

boundaries, in order to have all the site names uppercase in the channel definitions. Starting from version 1.5, the FTS WebService forces the case when you create a new channel, but when upgrading from previous versions, this convention may conflict with already defined channels. In order to fix this, we have provided an admin pack that allows changing the channel definitions. The instructions on how to use those tools are available [here](#).

Therefore, if you hit this problem, download the `glite-data-transfer-scripts` RPM and follow the instructions reported above in order to replace all the site names that contain lowercase letters in all the channel definitions (you may need the support of your DBA).

Note: If this RPM is not yet available in the repository, please contact `fts-support`

FTA Solutions

Job always in Submitted state

The first action that is executed on a transfer request is the Allocation, performed by the VO agent associated with the VO of the submitter. This action checks the source and destination URLs of the job request, finds the sites of the involved SEs using ServiceDiscovery and then looks up in the registered channels for a matching one. When this operation succeeds, the job is moved to Pending and the `channel_name` property is filled with the name of the found channel.

Due to a bug in FTA 1.3 and 1.4 ([#10076](#)) a job stays in Submitted state instead of going to Failed in one of the following cases

- The channel doesn't exist but the source and destination SE are registered in ServiceDiscovery or the VO is configured to accept unknown source and destination
- The VO of the user who submitted the job has no valid share on the channel
- The channel is in Stopped, Drain or Halted (actually, when the channel status is Halted, a job should go in Pending and not in Failed)

Usually this problem is due to a configuration error. The first thing to do is to retrieve the status of the channel that should be involved in the transfer

```
glite-transfer-channel-list CHANNEL_NAME
```

check the channel state, that the VO has a share and that the names of the source and destination sites match the ones retrieved using ServiceDiscovery: in case the file plugin is used, look at the `site` element of the SRM services reported into the `services.xml` file

```
<service name='CERN-SC-SRM'>
  <parameters>
    <endpoint>http://castor.gridsc.cern.ch:8443/srm/manager/v1</endpoint>
    <type>SRM</type>
    <version>1.1.0</version>
    <site>CERN-SC</site>
    <param name='SEMountPoint'>/castor/cern.ch/grid/dteam/storage</param>
  </parameters>
</service>
```

and compare them with the value returned by `glite-transfer-channel-list`

In case this doesn't fix the problem, check that a VO agent is configured and running for that VO. Do

```
glite-transfer-status --verbose JOB_ID
```

After upgrading to FTS 1.5 I got "No Channel found or VO not authorized..." error

And check that the value of the `VOName` property is correct; in case is not, it's a problem with the FTS `glite-data-transfer-submit-mapfile`: edit that file manually or regenerate it following the procedures reported by `FtsServerSubmitMapfile13`, cancel the job, wait that the files is reloaded by the FTS and ask the user to resubmit the request.

In case the VO is set correctly, check on the agents node that an agent is configured:

- if you're using gLite 1.3, please have a look at `/opt/glite/etc/config/glite-data-transfer-agents-oracle.cfg.xml` and see if there is an instance for the VO:

```
<instance name="YOUR_VO-fts">
  <parameters>
    <transfer-vo-agent.Name value="YOUR_VO"/>
    <!-- Other parameter -->
    <!-- ... -->
  </parameters>
</instance>
```

- if you're using gLite 1.4, open the file `/opt/glite/etc/config/glite-file-transfer-agents-oracle.cfg.xml` and look for an instance:

```
<instance name="YOUR_VO" service="transfer-vo-agent-fts"/>
```

If the instance is missing, or the naming convention is not correct, edit the appropriate file and rerun the configuration script.

If the instance is there, check if it's running, using the command

```
/opt/glite/etc/init.d/glite-data-transfer-agents --instance
glite-transfer-vo-agent-YOUR_VO status
```

or

```
service glite-data-transfer-agents --instance glite-transfer-vo-agent-YOUR_VO status
```

If the job is still Submitted, follow the procedure reported here

Job always in Pending state

After the a transfer request is allocation to a channel, its status is moved to Pending. The ChannelAgent will then process this request based on its internal inter-VO scheduling.

In case the job state remaining Pending forever, you have to check the following things:

- The related ChannelAgent daemon should be running
- The Channel state should be set to Active
- The VO should have a share on the channel that is greater than 0

In order to check if the agent is running, use the command

```
/opt/glite/etc/init.d/glite-data-transfer-agents --instance
glite-transfer-channel-agent-CHANNEL_NAME status
```

or

```
service glite-data-transfer-agents --instance glite-transfer-channel-agent-CHANNEL_NAME
status
```

You can check the Channel state and VO share using the command:

```
glite-transfer-channel-list CHANNEL_NAME
```

If the job is still Pending, follow the procedure reported here

All my transfers fail with a SECURITY_ERROR

This issue is usually due to a problem in the interaction from a FTA and the MyProxy server. This mainly happens in the following cases:

- User is mistyping the MyProxy passphrase when submitting the job
- User has an invalid or expired certificate in MyProxy
- The agent is not an authorized retriever for MyProxy
- There is a authentication problem (expired certificate or crl)

In the first two cases, all the transfers of this user should fail while the ones of other users succeed, while in the others all the transfers would fail, independently of the user.

Usually, you can detect the type of the error by having a look at the agent log file in

```
/opt/log/glite/glite-transfer-channel-agent-CHANNEL_NAME.log or
/opt/log/glite/glite-transfer-vo-agent-VO_NAME.log
```

- If the problem is due to a wrong passphrase, you'll see

```
2005-08-26 07:25:52,281 ERROR transfer-agent-myproxy - Failed to get the proxy from the
Reason is Error in bind()
ERROR from server: invalid pass phrase
```

Ask then the user to resubmit his/her file, possibly using the `-p` option of `glite-transfer-submit`. In case the problem persists, maybe the user forgot the passphrase, so ask him/her to restore the credential in myproxy using

```
myproxy-init -s MYPROXY_SERVER -d
```

- In case the agent is not an authorized retriever, you'll see the a similar entry

```
2005-08-26 07:25:52,281 ERROR transfer-agent-myproxy - Failed to get the proxy from the
ERROR from server: "<anonymous>" not authorized by server's authorized_retriever policy
```

If that is the case, you have to contact the MyProxy server administrator and ask him to add the DN of the certificate of the account used to run the agent. If it still doesn't work, please also check the the agent is running with a valid certificate, following what described here

- in case the entry is similar to

```
2005-08-26 07:25:52,281 ERROR transfer-agent-myproxy - Failed to get the proxy from the
Error authenticating: GSS Major Status: Authentication Failed
GSS Minor Status Error Chain: (null)
```

Job always in Pending state

This problem is usually due to an expired certificate or to an expired certificate revocation list (crl). Please check the validity of the certificates and update the crl in both the agent and MyProxy nodes

- In the other cases, ask the user to store again his/her certificate in MyProxy, running the command
`myproxy-init -s MYPROXY_SERVER -d`

Please note that the `-d` option is required in order to associate the credentials to the DN of the user instead of the account name

If you need to know which MyProxy server is used, have a look here

Which MyProxy Server is used?

When an agent has to perform an operation in behalf of the user, it retrieves the user's delegated credentials from the configured MyProxy server, cache it in the local file system and then impersonate the user by setting the environment variable `X509_USER_PROXY`. The operations where this is required are:

- Retrieve services endpoints and information from ServiceDiscovery
- Perform the transfer (unless the property `transfer.vo-agent.DisableDelegationForTransfers` is set to true)
- Contact the catalog for retrieving the list of replicas and registering the new ones when the transfer is finished (only in case of FPS VO Agent)

The endpoint of the MyProxy server is usually retrieved using ServiceDiscovery, so in case of the file plugin, you need to have an entry in `/opt/glite/etc/services.xml` like

```
<service name='MyProxy'>
  <parameters>
    <endpoint>myproxy://myproxy.cern.ch</endpoint>
    <type>MyProxy</type>
    <version>1.14</version>
  </parameters>
</service>
```

You can query the InfoSys using the command

```
glite-sd-query -t MyProxy
```

In order to resolve which MyProxy server should be used, the FileTransferAgent looks into the associated services of the FileTransferService who received the user's request (available from gLite 1.3 QF23) or, if not found, takes the first MyProxy server returned by the InformationSystem; you can also force the server to use a specific instance by setting the agent configuration property `transfer-agent-myproxy.Server`. In case this property is not set and there is no MyProxy entry registered in the InfoSys, the environment variable `$MYPROXY_SERVER` is used.

Starting from version gLite 1.3 QF23, the user is also allowed to specify the myproxy he want to use by providing the option `-m myproxy_hostname` in the `glite-transfer-submit` command line.

I've noticed a warning "Cannot Get Agent DN" in the agent log files

You can see this entry in case the agent doesn't run with a valid certificate. When an FTA starts, it put an logs the DN of the certificate the agent will use. This certificate is used to perform the following actions:

All my transfers fail with a `SECURITY_ERROR`

- Retrieve the user delegated credentials from MyProxy using the passphrase provided by the user. This happens both on the Channel and the VO Agents
- Perform the transfer if the `transfer.vo-agent.DisableDelegationForTransfers` property is set to `true`. This happens only in the VO Agent and it's the default behavior the FPS configuration

If the agent doesn't have a valid certificate, it's likely that these operations would fail.

In order to fix this problem, check first that the user running the agents has a valid certificate: usually this certificate are installed in `$HOME/.globus/usercert.pem` and `$HOME/.globus/userkey.pem` and should be owned by the user. In case the certificate is installed in a different place, the environment variables `X509_USER_CERT` and `X509_USER_KEY` should be set accordingly. You should also check that the certificate is not expired, by running:

```
openssl x509 -text -in ~/.globus/usercert.pem
```

or

```
openssl x509 -text -in $X509_USER_CERT
```

In case the certificate is valid but the agent always reports the warning, check if there is an expired proxy certificate in `/tmp/x509up_uUSER_ID` (where `USER_ID` is the user id of the account used to run the agent) and delete it.

My srmcopy transfers fail with a dCache MalformedURLException exception

You may notice this error when a user is transferring files to a dCache SE using a channel configured to perform `srmcopy` transfers. This is due to a bug in dCache version `<= 1.6.5` in parsing the URL. You have to ask the user to resubmit his/her requests using the following conventions:

- In case the destination SE is dCache, and the source is Castor or DPM
 - ◆ Source SURL can be

```
srm://<castorsrm>:<port>//srm/managerv1?SFN=<path>  
srm://<castorsrm>:<port>/?SFN=<path>  
srm://<castorsrm>/<path>
```

- ◆ Destination SURL should be

```
srm://<dcachesrm>:<port>/srm/managerv1?SFN=<path>  
srm://<dcachesrm>/<path>
```

- In case the source SE is dCache and the destination one is Castor or DPM
 - ◆ Source SURL should be

```
srm://<dcachesrm>:<port>/srm/managerv1?SFN=<path>  
srm://<dcachesrm>/<path>
```

- ◆ Destination SURL can be

```
srm://<castorsrm>:<port>/srm/managerv1?SFN=<path>  
srm://<castorsrm>:<port>//srm/managerv1?SFN=<path>  
srm://<castorsrm>:<port>/?SFN=<path>  
srm://<castorsrm>:<port>/<path>  
srm://<castorsrm>/<path>
```

- In case both the source and destination SE are dCache

◆ Source SURL should be

```
srm://<dcachesrm>:<port>//srm/managerv1?SFN=<path>  
srm://<dcachesrm>/<path>
```

◆ Destination SURL should be

```
srm://<dcachesrm>:<port>/srm/managerv1?SFN=<path>  
srm://<dcachesrm>/<path>
```

This problem is fixed in dCache v 1.6.6, however this new version doesn't seem to accept the compact SURL format

```
srm://<srmhost>/<path>
```

If the destination SE is then dCache and it's version is 1.6.6, we suggest to use for both source and destination SURLs either:

```
srm://<srmhost>:<port>/<path>
```

or the fully qualified one:

```
srm://<srmhost>:<port>/srm/managerv1?SFN=<path>
```

I've upgraded to 1.4.1 but srmcopy doesn't seem to work

Starting from version 1.3QF23, the FileTransferAgent normalize the SURLs before executing all the SRM get, put and copy requests and the default normalization is to convert them into the compact format

```
srm://<srmhost>/<path>
```

As illustrated here, we observed a problem with dCache srmcopy in version 1.6.6 not working with this format: after ~30 minutes the error returned is

```
number of retries exceeded:org.dcache.srm.scheduler.NonFatalJobFailure: java.io.IOException: both
```

In order to workaround this problem, you have to change the configuration of FileTransferAgent normalization to use a different format, by setting the ChannelAgent configuration property `transfer-agent-channel-actions.SurlNormalization` to either `compact-with-port` for converting to the format

```
srm://<srmhost>:<port>/<path>
```

or fully-qualified for the format

```
srm://<srmhost>:<port>/srm/managerv1?SFN=<path>
```

Please note that this is not a bug in FTS, but a problem in dCache; you might have observed after upgrading to 1.4.1 because this version of FTS has been release more or less at the same time as dCache 1.6.6

I've upgraded to 1.4.1 but the transfer failed with Error in srm__ping: NULL

Starting from version 1.4.1, FTS retrieves the srm endpoint from the information system, instead of parsing the SURL and, in case one of the compact formats are used, using the default port (8443) and service path (srm/managerv1). In case your transfers start failing after the upgrade with an error:

```
Cannot Contact SRM Service. Error in srm__ping: NULL
```

probably the entry in the information system is not correct: in fact, a common error that has been observed is that the SRM endpoint is stored as

```
srm://<srmhost>:<port>/srm/managerv1
```

instead of

```
httpg://<srmhost>:<port>/srm/managerv1
```

You can also check by looking into the transfer log files (located in `/var/tmp/glite-transfer-url-copy-UID/CHANNEL_NAMEfailed` in the related ChannelAgent box) and check the endpoint that is used for the SRM calls

The transfer failed with the error: No site found for host ...

During the allocation phase the VOAgent needs to resolve what are the sites that will be involved during the transfer. In order to do that, the agent will look up in the information system the site names of the source and destination SRMs, querying by the hostname retrieved from the provided SURLs.

In case the user gets an error like:

```
Failed to Get Channel Name: No site found for host ...
```

You have to look at the following things:

- The entry concerning the SRM services should be listed in the information system
- The SD library plugins are defined and configured properly (environment variables, files, etc)
- If the file-based plugin is chosen, the `/opt/glite/etc/services.xml` file is properly formatted

In order to do detect errors, it's useful to run the command:

```
su - ACCOUNT_USED_TO_RUN_THE_VOAGENT -c '/opt/glite/bin/glite-sd-query -t SRM --host SRM_HOSTNAME
```

and check the result (this command execute the same query as the agent).

In the problem still persists, it may be worth to have a look at the `/proc tanle` and see if the

```
/proc/VOAGENT_PROCESS_ID/environ
```

contains the correct values for the `GLITE_LOCATION` and `GLITE_SD_*` environment variables.

In case the StorageElement should not be listed in the information system, you may want to have a look here

Which Service Types are used?

The File Transfer Agent needs to interact with external services in order to accomplish its tasks and used the gLite ServiceDiscovery API in order to discover their properties. The involved services are:

- MyProxy: used to retrieve the clients' delegated credentials
- SRM & GridFtp: the site information is used to allocate a transfer job to a channel
- FileCatalog: used by the vo-agent in FPS mode in order to retrieve the source replicas to be used for a transfer and registered the new replicas when the transfer is finished

In order to discover that information the File Transfer Agent used the service types listed in [Glue Service Types](#)

As reported in [bug #12961](#), however, the service type for a GridFtp server is set to `GridFTP` instead of `gsiftp` and a backward compatible fix is foreseen for a future release. As a temporary workaround you could follow the comments reported on the bug.

I've tried everything, and it still doesn't seem to work

In case your problem is listed in this page, but none of proposed solutions doesn't seem to work, you can generate verbose log files and send them to fts-support. In order to generate these files, please follow the procedure:

For each agent involved (the VO one responsible to allocate a transfer to a channel and retry failed transfer; and the Channel one, responsible to transfer the files and monitor the status), please edit the file `glite-transfer-vo-agent-VO_NAME.log-properties` (in case of VO FTA) or and `glite-transfer-channel-agent-CHANNEL_NAME.log-properties` (in case of Channel FTA) and replace the lines

```
log4j.rootCategory=INFO, file
```

with

```
log4j.rootCategory=DEBUG, file
```

and

```
log4j.appender.file.fileName=/var/log/glite/glite-transfer-channel-agent-CHANNEL_NAME.log
```

or

```
log4j.appender.file.fileName=/var/log/glite/glite-transfer-vo-agent-VO_NAME.log
```

with

```
log4j.appender.file.fileName=/var/log/glite/glite-transfer-channel-agent-CHANNEL_NAME.debug.log
```

or

```
log4j.appender.file.fileName=/var/log/glite/glite-transfer-vo-agent-VO_NAME.debug.log
```

Restart the agents and let them running for ~ 1 minute; then stop the agents, restore the original values of the modified files, start the agents again and mail these `/var/log/glite/*.debug.log` files to fts-support

FTS Channel Administration solutions

How do I set the number of files transferred per VO instead of per channel?

In the FTS Channel Agent you have three parameters you can act on in order to tune the inter-vo scheduling: the channel VO share, the numbers of files that the channel can process concurrently and the `transfer-channel-agent.VOShareType` configuration property. The purpose of this configuration parameter is to define a policy how the VO share should be interpreted for a channel and you can add it to the instance that corresponds to the related channel agent in the `glite-file-transfer-agents.cfg.xml` configuration file. The allowed values are:

- **normalized:** the share is the value of the channel `voshare` property for the given VO, normalized to the sum of all the shares for all the VOs in the same channel. This option could be used when channel administrators want to guarantee slots for certain VOs, in order to implement some sort of QoS, accepting to eventually penalize the total throughput (transfer slots would be reserved to a VO even if that VO has no job to process)
- **absolute:** the share is the value on the channel `voshare` property expressed as a percentage. No normalization is performed, that means that the sum of all the shares on the same channel can exceed 100%. This option could be used when channel administrators want to balance the share between the VOs, without allowing that a single VO fully allocate a channel but minimizing the risk to allocate slots to VOs that don't have any job to process. This option implies some tuning on the VO share values based on experience, but it would allow to have a compromise between throughput and QoS.
- **normalized-on-active:** the share is the value of the channel `voshare` property for the given VO, normalized to the sum of all the share for all the VOs in the same channel that has at least one job that can be processed by the Channel Agent (job state should be Active, Pending or Canceling). This option is the default and should be used when the channel administrators want to optimize the throughput of the channel (the channel can be fully allocated even by one VO), but with a lower QoS

As an example, supposing you have a channel that has 30 files and 3 VOs, you could have:

		Normalized	Absolute	Normalized-on-active*
VO	Share	Max Files	Max Files	Max Files
VO_1	50	15	15	0
VO_2	30	9	9	18
VO_3	20	6	6	12

(* supposing VO_1 has no job to submit)

As you can notice, in case the sum of the VO share is 100, there's no difference between the "normalized" and "absolute" setup. But if this constraint is not respected, you can have:

		Normalized	Absolute	Normalized-on-active*
VO	Share	Max Files	Max Files	Max Files
VO_1	70	14	21	0
VO_2	50	10	15	19
VO_3	30	6	9	11

(* supposing VO_1 has no job to submit)

Please note that the value of the column "Max Files" correspond to the maximum number of files a VO is authorized to submit at the same time. In any case the constraint imposed by the "files" channel property is always respected.

If you want to start with two VOs, setting them each to be able to perform up to 15 transfers concurrently: Set the `transfer-channel-agent.VOShareType` to *normalized* (or *absolute*), having the VO share set to 50 and the channel files set to 30: you'll allow then up to 30 parallel transfers on the channel, but each VO would not be able to submit more than 15 at the same time. In case you'll have to support other VOs, you'll need to adjust these percentages.

General problems

How to replace host certificates on service nodes

Problem

The host certificate is expired or going to be changed.

Solution

- On **DPM** and **LFC** machines

See the corresponding section in the 'DPM and LFC' section of this troubleshooting guide: What to do if host certificate expired or going to be changed [↗](#)

- On **dCache** node
 - ◆ copy in the new certs to `/etc/grid-security/`
 - ◆ run the following line

```
/opt/d-cache/bin/dcache-core restart
```

The connections will be interrupted, this is unfortunately unavoidable at present. It could be minimized with the individual domains being restarted eg

```
/opt/d-cache/jobs/gsidcapdoor stop
/opt/d-cache/jobs/gsidcapdoor start
```

for all of the following domains

```
gPlazma
gridftpdoor
srm
xrootdDoor
gsidcapdoor
```

- On **FTS** node

The new host certificate has to be put to the usual place (`/etc/grid-security/`), All FTS daemons need to be reconfigured (with YAIM) to copy the hostcerts to where the (non-root) user running the daemon can see it. You should restart all the daemons using the standard procedure for this (which gives no user-visible downtime).

- On **VOMS** node

Copy the new host certificate to `/etc/grid-security`, and restart the service: `/etc/init.d/gLite restart` Pay attention that on all node that refer to this VOMS server, the server host certificate has to be changed, as well. In the

```
/etc/grid-security/vomses
```

directory. Furthermore the entries under

```
~.glite/vomses/  
/opt/glite/etc/vomses/  
/opt/edg/etc/vomses
```

has to be changed correspondingly.

- On **lcg-CE** node

Put the new certificates under

```
/etc/grid-security/
```

and restart the services.

- On **glite-CE** node

Put the new certificates under

```
/etc/grid-security/
```

and copy also to `/home/glite/.certs` and restart the services.

- On **lcg-RB** node

Put the new certificates under

```
/etc/grid-security/
```

and restart the services.

- On **glite-RB (WMS)** node

Put the new certificates under

```
/etc/grid-security/
```

and copy also to `/home/glite/.certs` and restart the services.

Where I can find the log files

- On **DPM** node

- ◆ `/var/log/dpns/log`
- ◆ `/var/log/dpm/log`
- ◆ `/var/log/dpm-gsiftp/dpm-gsiftp.log`
- ◆ `/var/log/rfio/log`
- ◆ `/var/log/srmv1/log`
- ◆ `/var/log/srmv2/log`
- ◆ `/var/log/srmv2.2/log`
- ◆ `/var/log/lcgdm-mkgridmap.log`

- On **LFC** node
 - ◆ /var/log/dli/log
 - ◆ /var/log/lfc/log
 - ◆ /var/log/lcgdm-mkgridmap.log
- On **BDII** node
 - ◆ /opt/bdii/var/bdii-fwd.log
 - ◆ /opt/bdii/var/bdii.log

Last edit:

Number of topics: 0

Maintainer: Gergely Debreczeni

This topic: LCG > TheLCGTroubleshootingGuide

Topic revision: r23 - 2011-10-17 - MaartenLitmaath



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)