Practicals on LFC and gLite DMS

Tony Calanducci

INFN

First gLite tutorial on GILDA, Catania, 13-15.06.2005

Set up your environment

- **Set the following environment variables according to which catalog you want to use:**
  - LFC Catalog
    - export LCG_CATALOG_TYPE=lfc
    - export LFC_HOST=lfc-gilda.cern.ch
  - RLS Catalog
    - export LCG_CATALOG_TYPE=edg
  - FireMan Catalog
    - No particular settings are required if you UI is installed properly
- **Initialize your proxy with grid-proxy-init or voms-proxy-init**

LFC Catalog commands

Listing the entries of a LFC directory

lfc-ls [-cdiLlRTu] [--class] [--comment] [--deleted] [--display_side] [--ds] path…

where *path* specifies the LFC pathname (mandatory)

- Remember that **LFC has a directory tree structure**
- /grid/<VO_name>/<you create it>

- All members of a given VO have read-write permissions under their directory
- Examples:

> *lfc-ls /grid/gilda*
> *lfc-ls -l /grid/gilda*
> *lfc-ls -l -R /grid/gilda*

-l long listing

-R list the contents of directories recursively

LFC Catalog commands

Creating a symbolic link

lfc-ln -s file linkname

lfc-ln -s directory linkname

Create a link to the specified *file* or *directory* with *linkname*

– *Examples:*

> *lfc-ln -s /grid/gilda/iome/prova.txt /grid/gilda/prova/test.txt*

Let's check the link using lfc-ls with long listing (-l)

> *lfc-ls -l /grid/gilda/prova*

LFC Catalog commands

Creating directories in the LFC

lfc-mkdir [-m mode] [-p] path...

- Where *path* specifies the LFC pathname
- Remember that while registering a new file (using lcg-cr, for example) the corresponding destination directory must be created in the catalog before

- Examples:

  > lfc-mkdir /grid/gilda/Examples

  You can just check the directory with:

  > lfc-ls -l /grid/gilda

LFC Catalog commands

Adding/deleting metadata information

lfc-setcomment path comment

lfc-delcomment path

lfc-setcomment adds/replaces a *comment* associated with a file/directory in the LFC Catalog

lfc-delcomment deletes a comment previously added

- Examples:

  > lfc-setcomment /grid/gilda/prova/prova.txt "weather conditions in Catania"

- Let's see what happened:

  > lfc-ls –comment /grid/gilda/prova/prova.txt

Hands-on Session

Exercise No. 1:

- Log onto an UI and initialize your proxy credentials
- set up properly the environment variables to use lfc-gilda.cern.ch catalog
- have a look inside the catalog
- create a directory with your surname
- put inside the just created dir a link to an existing file
- add a comment to that file and verify it

LFC Catalog commands

Summary of the LFC Catalog commands

lfc-chmod	Change access mode of the LFC file/directory
lfc-chown	Change owner and group of the LFC file-directory
lfc-delcomment	Delete the comment associated with the file/directory
lfc-getacl	Get file/directory access control lists
lfc-ln	Make a symbolic link to a file/directory
lfc-ls	List file/directory entries in a directory
lfc-mkdir	Create a directory
lfc-rename	Rename a file/directory
lfc-rm	Remove a file/directory
lfc-setacl	Set file/directory access control lists
lfc-setcomment	Add/replace a comment

lcg-utils and LFC/RLS Catalogs

- The LCG Data Management tools (usually called ***lcg-utils***) allow users to copy files between UI, CE, WN and a SE, to register entries in the File Catalogs and replicate files between SEs.

- Set up LCG_GFAL_INFOSYS environment variable to point to the GILDA Information Index (BDII)
  - **export LCG_GFAL_INFOSYS=grid004.ct.infn.it:2170**
- Again, choose which catalog you want to work with (see slide No 2). The two catalogs are **not** simultaneously accessible !

lcg-utils: lcg-cr

Upload a file to a SE and register it into the catalog

- lcg-cr -d dest_file | dest_host -l lfn [-g guid] [-l lfn]
  [-v | --verbose] --vo vo src_file

where

- ***dest_host*** is the fully qualified hostname of the destination SE
- ***dest_file*** is a valid SURL (both sfn:// or srm:// format are valid)
- ***guid*** specifies the Grid Unique IDentifier. If this option is not present, a GUID is generated internally
- ***lfn*** specifies the Logical File Name associated with the file
- ***vo*** specifies the Virtual Organization the user belongs to
- ***src_file*** specifies the source file name: the protocol can be *file:///* or *gsiftp:///*

lcg-utils: lcg-cr

- **To discover which SEs the user is allowed to use, a useful command is lcg-infosites.**

  > lcg-infosites --vo gilda se

  The output is a list of SEs and related information on available/used space

- **lcg-cr usage example:**

  > lcg-cr -v --vo gilda –d gildase.ct.astro.it

  -l lfn:/grid/gilda/Example/release.txt file:`pwd`/release

  – P.S.: Here it's supposed you are using an LFC catalog as you can notice from the format of the used lfn. Remember, RLS catalog entries are not organized in a hierarchical structure.

lcg-utils: lcg-aa and lcg-la

Adding an alias for a given GUID

lcg-aa --vo vo guid lfn

where

- ***vo*** specifies the Virtual Organization the user belongs to
- ***guid*** specifies the Grid Unique Identifier of the file you want to add the alias to
- ***lfn*** specifies the new alias

- *Example:*

  > lcg-aa --vo gilda guid:6f67888c-4628-48c4-b91c-1e6b0564bfd8 lfn:/grid/gilda/prova/release.txt

- To check if the previous command was successful, you can use **lcg-la** command to **list the aliases for a given LFN, GUID or SURL**

  > lcg-la --vo gilda lfn:/grid/gilda/prova/release.txt

Hands-on session

Exercise No.2:

- verify that your **LCG_GFAL_INFOSYS** is correctly set up
- create a local dummy file
- check the available storage elements
- copy into one SE and register the previous created file into your previous created dir
- add an alias to the just uploaded file
- check if the alias was assigned correctly

lcg-utils commands for replicas (I)

Copying a file from one SE to another one and register it in the Catalog

lcg-rep -d dest_file | dest_host [-v | --verbose] --vo vo src_file

where

- ***dest_host*** is the fully qualified hostname of the destination SE
- ***dest_file*** is a valid SURL (both sfn:// or srm:// are valid)
- ***vo*** specifies the Virtual Organization the user belongs to
- ***src_file*** specifies the source file name: the protocol can be LFN, GUID or SURL. An SURL scheme can be sfn: for a classical SE or srm:

- *Example:*

> lcg-rep -d gildase.ct.astro.it –vo gilda lfn:/grid/gilda/prova/myprova6.txt

lcg-utils commands for replicas (II)

Listing of replicas for a given LFN, GUID or SURL

lcg-lr --vo vo file

where

- ***vo*** specifies the Virtual Organization the user belongs to
- **file** specifies the Logical File Name, the Grid Unique IDentifier or the Site URL. An SURL scheme can be sfn: for a classical SE or srm:

Example:

> lcg-lr --vo gilda lfn:/grid/gilda/prova/myprova6.txt

> lcg-lr --vo gilda guid:00ec1459-6154-4d89-b016-503a23c445c7

lcg-utils commands for replicas (III)

Deleting replicas

- **lcg-del [-a] [-s se] [-v | --verbose] --vo vo file**

where

- ***a*** is used to delete all replicas of the given file
- ***se*** specifies the SE from which you want to remove the replica
- ***vo*** specifies the Virtual Organization the user belongs to
- ***file*** specifies the Logical File Name, the Grid Unique IDentifier or the Site URL. An SURL scheme can be sfn: for a classical SE or srm:.

Example:

- delete one replica

  >lcg-del --vo gilda -s gildase.ct.astro.it lfn:/grid/gilda/prova/myprova6.txt

- delete all the replicas

  > lcg-del -a --vo gilda lfn:/grid/gilda/prova/myprova6.txt

lcg-utils: lcg-cp

Downloading a Grid file in a SE to a local destination

lcg-cp [-v | --verbose] --vo vo src_file dest_file

where

- ***vo*** specifies the Virtual Organization the user belongs to
- ***src_file*** specifies the source file name: the protocol can be LFN, GUID, SURL or local file. An SURL scheme can be sfn: for a classical SE or srm:
- ***dest_file*** specifies the destination. The protocol can be file:/// or gsiftp:///

Example:

> lcg-cp --vo gilda lfn:/grid/gilda/prova/tony file:/home/tcaland/mpi.jdl

Hands-on session

Exercise No. 3:

- Create two replicas of the file you previously uploaded (you could also use the alias to indicate the input file)
- Check if the operation was successful
- Download the file back in your UI
- Delete just one replica and verify that
- Delete all the replicas and verify that
- Verify if the entry is still into the catalog

Summary of lcg-utils commands

Replica Management

lcg-cp	Copies a grid file to a local destination
lcg-cr	Copies a file to a SE and registers the file in the catalog
lcg-del	Delete one file
lcg-rep	Replication between SEs and registration of the replica
lcg-gt	Gets the TURL for a given SURL and transfer protocol
lcg-sd	Sets file status to "Done" for a given SURL in a SRM request

File Catalog Interaction

lcg-aa	Add an alias in LFC for a given GUID
lcg-ra	Remove an alias in LFC for a given GUID
lcg-rf	Registers in LFC a file placed in a SE
lcg-uf	Unregisters in LFC a file placed in a SE
lcg-la	Lists the alias for a given SURL, GUID or LFN
lcg-lg	Get the GUID for a given LFN or SURL
lcg-lr	Lists the replicas for a given GUID, SURL or LFN

RLS Interaction Commands

- The edg-local-replica-catalog (edg-lrc) and edg-replica-metadata-catalog (edg-rmc) commands are low level tools that allow users to browse and directly manipulate the LRC and the RMC catalogs.
- In normal operation, a user should **preferably** use the high level LCG Data Management tools previously described.
- However, lcg-utils do not offer a way to browse and search in RLS Catalog. To do this you still need to use edg-rmc command.

Searching into the RMC Catalog

The Replica Metadata Catalog keeps the mappings between LFNs and GUID, while, Guid – SURLs are stored into the Local Replica Catalog (LRC), as shown below:

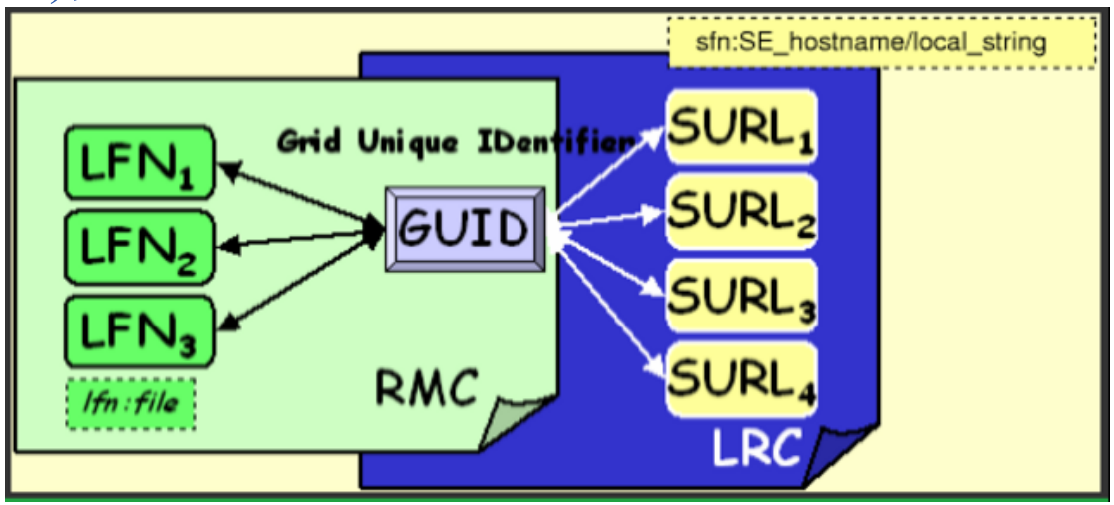

To find out if a particular LFN or LFNs satisfying a pattern are into the Catalog, we need to use edg-rmc command. For example, if we want to find all the entries containing the word “my” we should issue the following command:

> edg-rmc mappingsByAlias *my* --endpoint $RMC_ENDPOINT

where $RMC_ENDPOINT contains the endpoint of the RMC Catalog

Searching into the RMC Catalog

- **Actually there is no particular command to list the whole contents of the RMC Catalog. A way to reach this goal is to make a search on GUID with the wildcard *.**
- **Example:**

```
> edg-rmc -i mappingsByGuid \* -l 100000 --endpoint
      $RMC_ENDPOINT
```

-l flag is used to limit the number of results.

Hands-on session

Exercise No.4:

- Set your env variables to point to the RLS catalog (check also that you have a valid JAVA_HOME variable, otherwise you will not able to use edg-rmc command)
- Copy and register a file into the catalog
- Ask your neighbour the name of the LFN he/she put in
- Search that file (with edg-rmc), download locally and delete it from the catalog

Reference

- LFC : LCG File Catalog
- RLS : Replica Location Services
- RMC: Replica Metadata Catalog
- LRC : Local Replica Catalog
- LFN : Logical File Name
- GUID: Grid Unique Identifier
- SURL: Storage URL
- FiReMan: File Replica Manager