

Table of Contents

CERN LFC Home Page.....	1
Overview.....	1
LFC Central Catalogs.....	1
Installation and Configuration.....	1
Users and Processes.....	2
LFC daemon configuration.....	2
LFC Sysconfig file creation.....	2
Trusted Hosts.....	3
Oracle RAC Database backend.....	3
Database Connection Configuration File.....	3
Information System.....	3
Management Procedures.....	3
SMS.....	3
Standard Operations Procedures.....	4
Monitoring.....	4
Lemon Alarms.....	4
Load Balancing.....	4
Problem Determination.....	5
Daemons.....	5
DB Configuration Details.....	6
Writer / Reader account.....	6
Oracle accounts used in Production.....	7
Presentations.....	7
CERN LFC Operations guide.....	7
LFC troubleshooting.....	7

CERN LFC Home Page

THIS PAGE NEEDS TO BE REVIEWED

This page defines the installation, configuration and the procedures related to the CERN LFC service.

This page documents the current situation. It does not cover requirements or issues. These are covered in the LfcNotes.

Service Level Status

Overview

The CERN LFC service is defined as a critical service in the services catalog.

The LFC is a core grid component which provides resolution from logical names to physical locations for replicas of files on the Grid. It can be used in two modes:

- Central : Here a single central catalog is used to store pointers to either the site or the actual physical location of a file for all VO files in the grid
- Local : Here there is a catalog per site which stores mappings from logical to physical names for all VO files at that particular site.

We provide a highly available, fault tolerant configuration for both central and local catalog for LHC VOs which require them. We also provide some catch-all central catalogs for other CERN & HEP VOs.

LFC Central Catalogs

Alias	Supported VOs	Database instance	Comment
prod-lfc-atlas	ATLAS, OPS	LCGR	
prod-lfc-shared-central	DTEAM, UNOSAT, GEANT4, GEAR, SIXT, OPS	LCGR	
prod-lfc-lhcb-central	LHCb, OPS	LHCBR	read-write instance
prod-lfc-lhcb-ro	LHCb, OPS	LHCBR	read-only instance

Installation and Configuration

The main cluster in CDB is `gridlfc`.

The base CDB template for this cluster is `prod/cluster/gridlfc/config.tpl` [↗](#)

Configuration specific to subclusters are

- For subcluster atlas corresponding to alias `prod-lfc-atlas` `prod/cluster/gridlfc/subcluster-atlas.tpl` [↗](#)
- For subcluster lhcb-ro corresponding to alias `prod-lfc-lhcb-ro`

prod/cluster/gridlfc/subcluster-lhcb-ro.tpl

- For subcluster lhcb-rw corresponding to alias prod-lfc-lhcb-central
prod/cluster/gridlfc/subcluster-lhcb-rw.tpl

The production servers are currently all running SLC5.

Users and Processes

The LFC processes run under the lfcmgr account and group. The reserved accounts and uid/gid values for grid server processes are here. These are delivered to the node via SINDES.

LFC daemon configuration

There is no specific NCM Component for the LFC, but we instead use some other generic components, like **exportconf** and **SINDES**. There is a CDB component description `"/software/components/lfc/` in which you put the LFC configuration. this must be set **before** the `pro_system_gridlfc` template is included. The following values are currently supported:

Name	Values	Description
alias		This is one of the aliases listed above. It is used to extract a suitable DB connect string from the SINDES LFCnsconfig component
readonly	true, false	Is this catalog readonly ? This will updated the <code>/etc/sysconfig/lfcdaemon</code> file appropriately

```
# LFC Sysconfig configuration
include pro_declaration_component_lfc;
"/software/components/lfc/active" = true;
```

LFC Sysconfig file creation

We use the NCM `exportconf` component to re-write the `lfcdaemon` and `lfc-dli` sysconfig files. An example is :

```
"/software/components/exportconf/active" = true;
"/software/components/exportconf/dispatch" = default(true);

"/software/components/exportconf/lfc-dli/rules" = push(nlist(
    "file",      "/etc/sysconfig/lfc-dli",
    "template",  "/etc/sysconfig/lfc-dli.templ",
    "rules",     nlist("LFC_HOST",hostname)));

"/software/components/exportconf/lfcdaemon/rules" = push(nlist(
    "file",      "/etc/sysconfig/lfcdaemon",
    "template",  "/etc/sysconfig/lfcdaemon.templ",
    "rules",     nlist("NB_THREADS", "40",
                      "RUN_LFCDAEMON", "yes",
                      "ORACLE_HOME", "/usr/lib/oracle/10.2.0.1/client",
                      "TNS_ADMIN", "/etc")));

"/software/components/exportconf/lfcdaemon/rules/0/rules" =
if( exists ("software/components/lfc/readonly") && (value("/software/components/lfc/readonly")
    merge(value("/software/components/exportconf/lfcdaemon/rules/0/rules"),
          nlist("RUN_READONLY", "yes")));
} else {
    value("/software/components/exportconf/lfcdaemon/rules/0/rules");
};
```

Trusted Hosts

The LFC uses the `shift.conf` file to specify external hosts on which the `root` account should be considered as the `root` user within the LFC. This is used for admin tasks, and also by LHCb to allow their DIRAC nodes to have access directly to the catalog. This is controlled by the `castorconf` NCM component:

```
# Enable the trusted hosts for the LFC
# LHCb hosts have extra on their central R/W and R/O catalogs
define variable lhcb_trusted_hosts = "lxgate03 lxgate03.cern.ch lxgate05 lxgate05.cern.ch lxgate07 lxgate07.cern.ch"
define variable admin_trusted_hosts = "lxadm01 lxadm01.cern.ch lxadm02 lxadm02.cern.ch lxadm03 lxadm03.cern.ch"

"/software/components/castorconf/LFC/TRUST" =
  if ( exists("/system/vo/lhcb/services/LFC") && value("/system/vo/lhcb/services/LFC") == "central" ) {
    admin_trusted_hosts + " " + lhcb_trusted_hosts;
  } else {
    admin_trusted_hosts;
  };
```

To Add another host for either `admin` or LHCb purposes, simply update the appropriate variable, and re-run the `castorconf` NCM component.

Oracle RAC Database backend

The database backend for the Production LFCs (central and local) at CERN is **Oracle 10g on RAC**. The database / service name is `lcg_lfc` at CERN.

Database Connection Configuration File

The only LFC configuration file is `/opt/lcg/etc/NSCONFIG` contains the database connection parameters :

```
cat /opt/lcg/etc/NSCONFIG
my_account_w/XXXXXX@lcg_lfc
```

This file is delivered by SINDES, along with the host certificates, configured in `pro_system_gridlfc.tpl`.

```
# SINDES config - used to deliver the LFC DB connect string
"/software/components/sindes/items/lfcNSCONFIG" = nlist("method", "file", "scope", "cluster");
"/software/components/sindes/items/grid-host-certificates" = nlist("method", "file", "scope", "node");
"/software/components/sindes/all" = "passwd-header,group-header,lfcNSCONFIG,grid-host-certificates";
```

Information System

Currently we use a BDII instead of `globus-mds` to run the GRIS. We also publish into the info sys the LFC alias, rather than the hostname. The BDII is currently hand-configured by using the `run_function yaim` script on `config_bdii`, but will be in `yaim` after `glite 3.0` is released.

Management Procedures

SMS

For the LFC, we need to remove the nodes from the load-balanced alias when in standby or maintenance. Currently we use `/usr/libexec/SetToDesiredState.gridbdii` is used to put the nodes into production/maintenance. On maintenance, there is NO `/etc/nologin` file, otherwise the `bdii` daemon cannot be

started.

NOTE : We should, either rename this script to something more general, or create a LFC specific one.

Standard Operations Procedures

How to split a database backend

Monitoring

Lemon Alarms

In addition to the OS standard alarms, specific Lemon Alarms have been defined for the LFC:

Alarm name	Description	Comment
LFCDAEMON_WRONG	No <code>lfcdaemon</code> process running	
LFC_DLI_WRONG	No <code>lfc-dli</code> process running	
LFC_DB_ERROR	ORA-number string detected in <code>/var/log/lfc/log</code>	
LFC_NOREAD	can't stat given directory	trying to read <code>/grid/ops/</code>
LFC_NOWRITE	can't utime on file	
LFC_SLOWREADDIR	excessive time taken to read directory	time > 10 s
LFC_ACTIVE_CONN	number of active connections to LFC	use <code>netstat</code>

To configure this for a machine, there are two CDB profiles

The `pro_monitoring_cos_gridlfc` profile defines the templates for the monitors.

Within the profile `pro_system_gridlfc`, the `pro_monitoring_cos_grid_lfc` template is included and the metrics are set to active.

The data will be stored in the Lemon database and visible through the lemon interface. An example is Number of LFC Processes.

These alarms, along with all standard alarms on the nodes, are handled by the operator and sysadmin teams. the procedures are all stored in OPM

- [Operator Procedure](#)
 - [Sysadmin Information](#)
-

Load Balancing

We use the standard DNS load-balancing mechanism provided at CERN (DnsAliases). The alias to be used for a particular host is specified in the CDB variable `/software/components/lfc/alias`. This is then used to configure the `loadbalancing` component on the node:

```
# DNS Alias name in FQDN
define variable aliasname = if(exists("/software/components/lfc/alias")) {
    value("/software/components/lfc/alias") + "." + value("/system/network/domainname");
} else {
    "";
};
```

...
...

```
"/software/components/loadbalancing/clustername" =
  if(exists("/software/components/lfc/alias")) {
    value("/software/components/lfc/alias");
  };
```

The LEMON exception which takes the node out of the alias is 30075. This is a alarm which merges together the three possible error alarms *LFC_NOREAD*, *LFC_NOWRITE* and *LFCDAEMON_WRONG*

```
#
# JC - This alarm is only an aggregate for the lbclient system, and should
# not be raised to the operator
#
"/system/monitoring/exception/_30075" = nlist(
  "name",          "lfc_noservice",
  "descr",         "LFC Service not available",
  "active",        true,
  "latestonly",   false,
  "importance",   2,
  "correlation",  "39:1 != 1 || 5202:1 != 0 || 5203:1 != 0"
);
```

Problem Determination

Here is what to do in case of a problem with the LFC :

LFC Smoke Tests and Actions

Daemons

There are 2 daemons running on an LFC machine :

- lfcdaemon
- lfc-dli

To start/stop and get the status of a daemon, use :

- service lfcdaemon start|stop|status
- service lfc-dli start|stop|status

The cluster is configured so that `lfcdaemon` and `lfc-dli` are automatically started at boot.

There should be 40 LFC threads running under the `lfcmgr` account:

Note: there can be more `lfcdaemon` threads (see the `-t` number option and `/etc/sysconfig/lfcdaemon`)...

The status check should return OK:

```
service lfcdaemon status
lfcdaemon (pid 2632) is running...           [ OK ]

service lfc-dli status
lfc-dli (pid 2656) is running...           [ OK ]
```

Daemons should start after boot (chkconfig mechanism) with the rolling logs in :

```
/var/log/lfc/log
/var/log/lfc-dli/log
```

If the daemons are not running after reboot look at the log. You could try to start them using :

```
service lfcdaemon start
service lfc-dli start
```

If the load is high, all the 20 threads might be occupied, the `LFC_NOREADDIR` error will occur, and the users might see this :

```
$ lfc-ls /grid
send2nsd: NS002 - connect error : Connection timed out
/grid/atlas: Communication error
```

You can check if all the threads are often in use by checking the `/var/log/lfc/log` file :

```
tail -f /var/log/lfc/log

03/23 13:51:01 2631,0 Cns_srv_mkdir: NS092 - mkdir request by /C=CH/O=CERN/OU=GRID/CN=Sophie Lem
03/23 13:51:01 2631,0 Cns_srv_mkdir: NS098 - mkdir /grid/dteam/testsl 777 22
03/23 13:51:01 2631,0 Cns_srv_mkdir: returns 0
                ^
                |
                |
                here: thread #0 used
```

DB Configuration Details

Writer / Reader account

For security reasons, the Physics Database team at CERN requires the use of **writer / reader accounts** by applications.

The writer / reader accounts have limited privileges on the LFC Oracle tables, sequences and views - compared to the owner account.

The scripts granting the appropriate privileges for the LFC accounts are in :

```
ls /afs/cern.ch/project/gd/SC3/LFC-DB-Accounts/

create-reader-account.sql
create-writer-account.sql
create-synonym.sql
```

Everytime there is a schema change, you have to run them for each account in use :

- set the correct user name in `create-reader-account.sql`, `create-writer-account.sql` and `create-synonym.sql`.
- run the `create-reader-account.sql` script :

```
sqlplus lfc_account/XXXXX@lcg_lfc < create-reader-account.sql
```

- execute the output in the reader account.

```
sqlplus lfc_account_r/XXXXX@lcg_lfc
```

- run the `create-synonym.sql` script :

```
sqlplus lfc_account_r/XXXXX@lcg_lfc < create-synonym.sql
```

- execute the output in the reader account :

```
sqlplus lfc_account_r/XXXXX@lcg_lfc
```

Same steps for the **writer account**.

See Writer / Reader accounts for details.

Oracle accounts used in Production

Several Oracle accounts are used, but some VOs share the same Oracle account.

Check the `/usr/etc/NSCONFIG` file on all LFC servers to know the current configuration :

```
lxplus003# wssh -h "root@lfc[001-011]" cat /opt/lcg/etc/NSCONFIG
```

Presentations

- CERN-LFC-admins-tutorial-27_07_2006.ppt: LFC Overview and Debugging
- 060727-FIO-LFC-Training.ppt: 060727-FIO-LFC-Training.ppt

CERN LFC Operations guide

See LfcOperations.

The OPM guide can be found here [↗](#)

LFC troubleshooting

See the developers DataManagementDocumentation pages.

This topic: LCG > LfcWlcg

Topic revision: r42 - 2014-04-16 - AlbertoRodriguezPeon



Copyright &© 2008-2022 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