



Enabling Grids for E-scienceE

Running reliable services: the LFC at CERN

Sophie Lemaître (Sophie.Lemaître@cern.ch)

*WLCG service Reliability Workshop
CERN – Geneva – 26th November 2007*

www.eu-egee.org



Information Society



- **What is reliability?**
- **LFC usage at CERN**
- **Turning the LFC into a reliable service at CERN**

- **What is reliability?**
 - Ability of a system/component to perform its required functions under stated conditions

- **What does it mean for you?**
 - Ability to meet Tier1 SLA
 - 99% availability with maximum 12 hours downtime

- **Increasing reliability = Increasing time when service is available
= Reducing downtime**

- **Good practice**
 - Don't lose the Experiments data!
 - Automate everything
 - Limit the server downtime
 - Front-end dynamic load balancing
 - Limit the database downtime
 - Oracle RAC
 - Limit time before noticing problem
 - Monitoring
 - Make sure to be able to react as fast as possible
 - Create procedures for operators (first level support)
 - Plan all interventions

- **What is reliability?**
- **LFC usage at CERN**
- **Turning the LFC into a reliable service at CERN**

- **Different experiments decided on different usage patterns**
 - LHCb
 - Central catalog
 - Read-only catalog (for scalability and redundancy)
 - Replica read-only catalogs at Tier1s
 - *Replication done via Oracle Streams*
 - ATLAS
 - Central and local catalog
 - CMS
 - Central catalog
- **Shared catalog**
 - Shared “catch-all” catalog for dteam, unosat, sixt, ...

- **Backend storage is in Oracle**
 - On the LCG Oracle RAC instance
 - Separate database accounts per experiment & catalog type
 - Better VO isolation
- **All front-end nodes are in 2-node load-balanced clusters**
 - Mid-range servers
 - Using automatic load-balancing to isolate “broken” nodes
 - Separate DNS alias for each VO & catalog type
- **Full Service Description:**

<https://twiki.cern.ch/twiki/bin/view/LCG/LfcWlcg>

LFC Central Catalogs

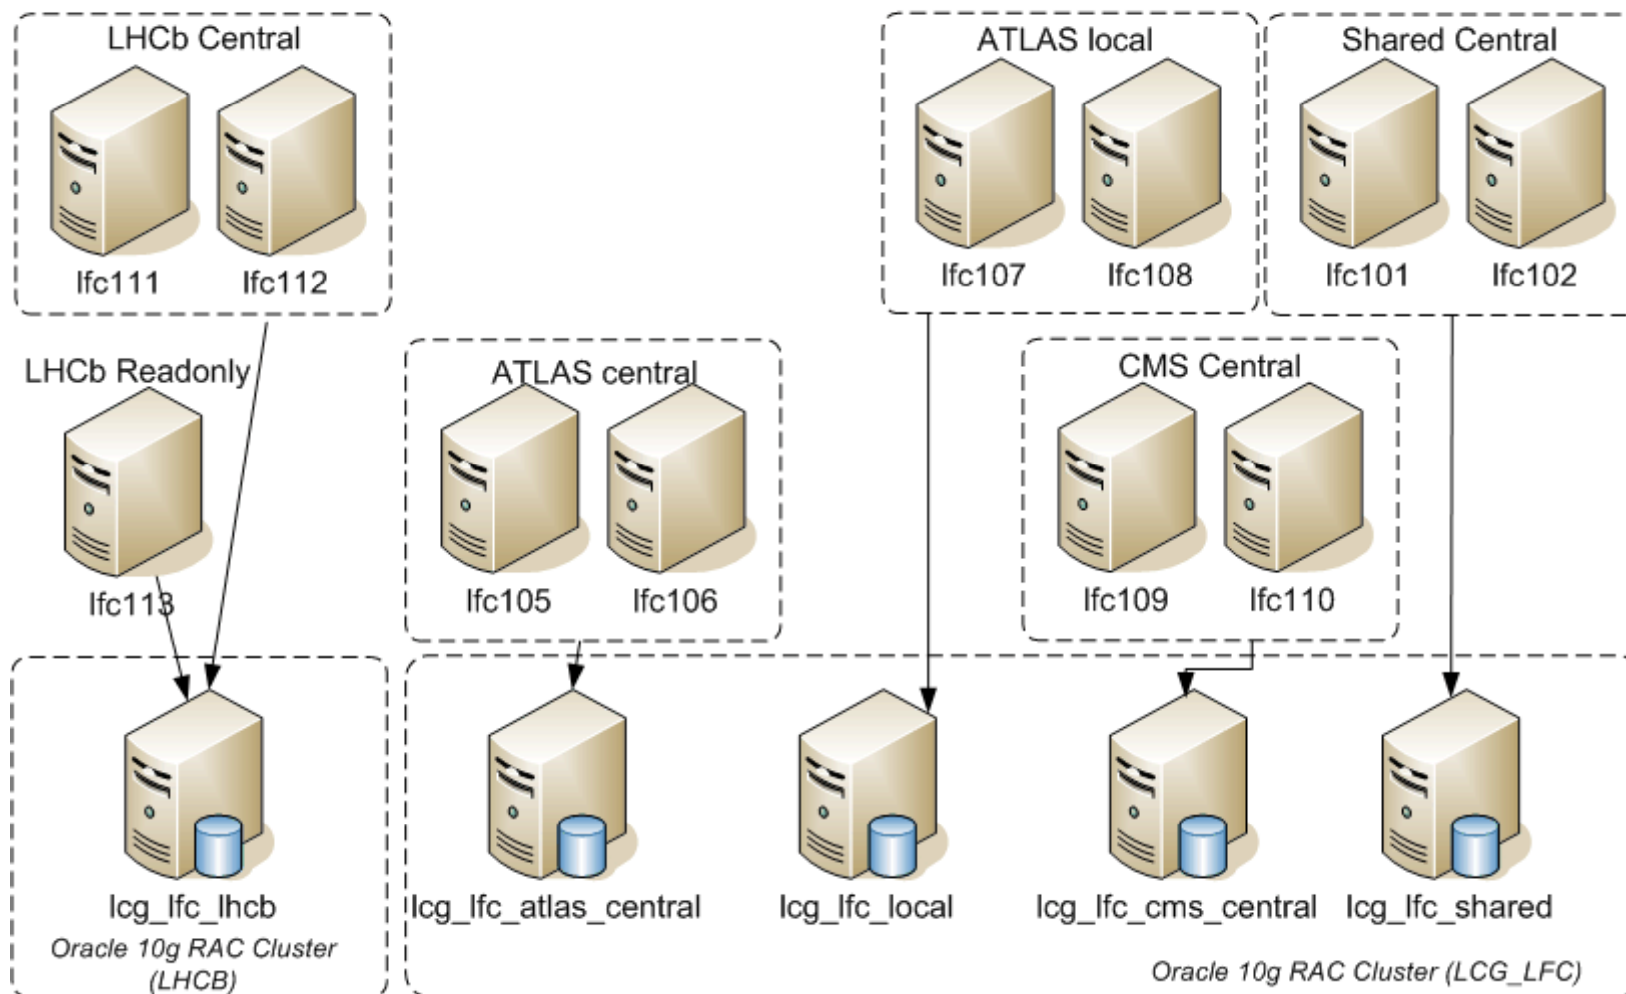
<u>Alias</u>	<u>Supported VOs</u>	<u>Comment</u>
prod-lfc-atlas-central	ATLAS	
prod-lfc-cms-central	CMS	
prod-lfc-shared-central	DTEAM, UNOSAT, GEANT4, GEAR, SIXT, OPS	
prod-lfc-lhcb-central	LHCb	read-write instance
prod-lfc-lhcb-ro	LHCb	read-only instance

LFC Local Catalogs

<u>Alias</u>	<u>Supported VOs</u>	<u>Comment</u>
prod-lfc-atlas-local	ATLAS	

26th November 2007

LFC Service Deployment Layout




- **What is reliability?**
- **LFC usage at CERN**
- **Turning the LFC into a reliable service at CERN**
 - Checklist
 - Database backup
 - Monitoring
 - Dynamic DNS load balancing
 - Firewall
 - Quattor vs. YAIM
 - Operators procedures

- See <https://twiki.cern.ch/twiki/bin/view/LCG/WlcgScDashLfc>

Requirements


Assess	Question to set assessment color if not available
Green	Service class defined for calendar periods: class C in all periods
Green	Disk space requirements defined: minimal local disk
Green	Database requirements defined: Oracle 10g on RAC
Green	Middleware components defined: LFC Server, BDII
Green	Connectivity requirements defined : incoming on 5010 (lfcdaemon), 8085 (lfc-dli)
Green	Backup objects defined (directories, databases): Oracle Database. All other state in CDB
Green	High level diagram of services: See LfcWlcg for diagram

 Edit

Development

Assess	Question to set assessment color if not available
Green	Code delivered and installation images available: part of gLite-3.0.0 release
Green	Code certified
Green	Administration Guide available: LfcAdminGuide
Green	Problem Determination Guide available: LfcTroubleshooting
Green	Middleware pre-requisites (web server, LDAP, ...) defined: VDT

Hardware

Assess	Question to set assessment color if not available
Green	CPU power required for application: mid-range server machine OK
Green	Memory requirement for application: low-memory requirement - 2GB more than enough
Green	Inventory of hardware components defined v: Specified in LfcWlcg
Yellow	Machines on order or available: Machines on order, batch-nodes in place until available as replacements
Green	Machines installed in appropriate location within data centre (UPS,Network) : Installed in LCG network
Green	Database server and disk space allocated : Database backend on LCG Oracle RAC Cluster
 Edit	

Assess	Question to set assessment color if not available
Green	2nd level support organisation defined (who to call when there is a problem with the application or middleware) : Support teams defined in GGUS
Green	Mechanism to contact 2nd level organisation : GGUS
Yellow	Response time for 2nd level organisation : Best Effort
Green	List of machines where service is running defined: See LfcWlcg#Overview
Green	List of configuration parameters and their values for the software components : Expressed in CDB template pro_system_gridlfc.tpl
Green	List of processes to monitor : Defined (lfcdaemon, lfc-dli, bdii)
Green	List of file systems and their emergency thresholds for alarms : Defined (standard)
Green	Application status check script requirements defined : Integrated into LEMON
Green	Definition of scheduled processes (e.g. cron) : Managed by QUATTOR
Green	Test environment defined and available : Certification nodes allocated
Green	Problem determination procedures including how to determine application vs middleware vs database issues : Operator and Sysadmin procedures in OPM
Green	Procedures for start/stop/drain/check status defined : Integrated with SMS
Green	Automatic monitoring of the application in place : Integrated into LEMON
Green	Backup procedures defined and tested: database backups and recoveries are handled by IT-PSS-Physics Databases section : BackupPolicyManagers

- **Database backups**
 - Don't lose the Experiments data!
 - Ask your friendly database administrator to backup the LFC database

- **Plan all interventions**
 - Even the transparent ones
 - To prevent coordination / communication problems as much as possible

- **What to monitor?**

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	<code>ORA-number</code> 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>

- **Note: avoid monitoring based on creating an LFC file/directory**
 - To avoid using file ids in the LFC database tables
 - Update file utime instead (LFC_NOWRITE)

- **Existing tools**

- Generic LFC probe

- See Monitoring Working Group web page

- <https://twiki.cern.ch/twiki/bin/view/LCG/GridServiceMonitoringInfo>

- Direct link to LFC probe

- <https://www.sysadmin.hep.ac.uk/svn/grid-monitoring/trunk/probe/ch.cern/src/LFC-probe>


- At CERN, a Lemon sensor is used

- “lemon-sensor-grid-lfc” package
 - See corresponding source in CVS

- <http://isscvcs.cern.ch:8180/cgi-bin/cvsweb.cgi/elfms/lemon/sensors/sensor-grid-lfc/?cvsroot=elfms>


- A Lemon wrapper to the LFC probe will be written in the future

- **Dynamic DNS load balancing**
 - Solution implemented by Vlado Bahyl and Nick Garfield
 - See Vlado's presentation this morning
 - <http://indico.cern.ch/materialDisplay.py?subContId=0&contribId=2&sessionId=0&materialId=slides&confId=20080>
- **For the LFC nodes, checking whether:**
 - Cannot ssh, or
 - Node in maintenance, or
 - /tmp full, or
 - Alarm among
 - LFC_NOREAD, LFC_NOWRITE, LFCDAEMON_WRONG



Machine removed from DNS aliases

- **Advantage:**
 - All LFC software upgrades are transparent for the users
 - Except when database schema changes
- **Ex: two DNS aliased nodes A and B**
 - 1) Put node A in maintenance
 - Wait for node A to be taken out of production by dynamic DNS load balancing
 - 2) Stop + upgrade + start LFC on node A
 - 3) Take node A out of maintenance
 - Wait for node A to be put back into production by dynamic DNS load balancing
 - 4) Start at step 1) with node B



**Transparent
upgrade done!**

- **At CERN, using LANDB sets to control firewall access**
 - Common firewall settings for the whole LFC cluster
 - If a change is made, it is applied to all machines automatically
 - No more “oooops, forgot to configure this one”...

- **The LFC log files are stored under**
 - /var/log/lfc
- **Make sure the /var partition is big enough!**
 - This problem hit us at CERN... ☹
 - WLCG log retention policy: 90 days
 - Now, /var size is 200G on the LFC nodes

- **Yaim is used to configure the grid aspects of the service**
- **Some duplication between QUATTOR and yaim**
 - Approach is to use QUATTOR for system functionality
 - E.g. chkconfig, FS, accounts, access control, package management
 - And yaim for grid functionality

```
# yaim
include pro_software_components_lcg_yaim_3_0;
"/software/components/yaim/active"=true;
"/software/components/yaim/configure" = true;
"/software/components/yaim/nodetype/LFC_oracle" = true;
"/software/components/yaim/conf/LFC_HOST" = hostname;
"/software/components/yaim/conf/LFC_HOST_ALIAS" = aliasname;

# We need to run sindes and exportconf before yaim, so that we
# have host certs, LFC DB password file and sysconfig
"/software/components/yaim/dependencies/pre" = list("sindes", "exportconf");
```

- For pieces of system functionality configured by both yaim and quattor, we remove them from YAIM
- **Simple mechanism**
 - place an empty config_XXX file in the yaim **functions/local** directory
 - Handled by an rpm 'CERN-CC-glite-yaim-local'

- **React to problems as fast as possible**
 - Get operators to respond directly to known issues
 - At CERN, most alarms open standard tickets
 - LFC_ACTIVE_CONN, LFC_SLOW_READDIR don't
 - ... because there's no immediate action you can take ;(
 - *It's a sign of heavy usage of the service by the VO*

- **What is reliability?**
 - Ability of a system/component to perform its required functions under stated conditions

- **What does it mean for you?**
 - Ability to meet Tier1 SLA
 - 99% availability with maximum 12 hours downtime

- **Increasing reliability = Increasing time when service is available
= Reducing downtime**

- **Good practice**
 - Don't lose the Experiments data!
 - Automate everything
 - Limit the server downtime
 - Front-end dynamic load balancing
 - Limit the database downtime
 - Oracle RAC
 - Limit time before noticing problem
 - Monitoring
 - Make sure to be able to react as fast as possible
 - Create procedures for operators (first level support)
 - Plan all interventions



**This will save you
from a lot of trouble**

- **Integration with CERN fabric infrastructure**
 - James Casey
- **LFC administrators at CERN**
 - Jan Van Eldik, Miguel Coelho, Ignacio Reguero, David Collados, Diana Bosio
- **Dynamic DNS load balancing**
 - Vlado Bahyl, Nick Garfield
- **LFC expert**
 - Jean-Philippe Baud

Questions ?