



GLite Middleware

Overview of the LFC Service

Catalin Condurache
RAL Tier1 Grid Services
30 Jun 2010

Outline

- Why the File Catalogs?
- Introduction to LFC
- LFC usage
- LFC deployment at RAL Tier1



LFC ?

~~LFC catalog~~ (better LFC service)

LCG File Catalog

LHC Computing Grid File Catalog

Large Hadron Collider Computing Grid File Catalog



File Catalog Motivations

- Users and programs produce and require data
- Data may be stored in Grid datasets (files)
 - located in Storage Elements (SEs)
 - several replicas of one file in different sites
 - accessible by Grid users and applications from “anywhere”
 - locatable by the WMS (data requirements in Job Description Language)
 - data may be copied from/to local filesystems (WNs, UIs) to/from the Grid
- Why File Catalogs are used?
 - they keep track of the location of copies (replicas) of Grid files
 - they provide Data Management tools and APIs through which users and Grid services can deal with files




Data management example



“User interface”



“Myfile.dat”



LCG File Catalog



Storage Element1



Storage Element2



Science & Technology
Facilities Council

Name conventions

- Logical File Name (LFN)

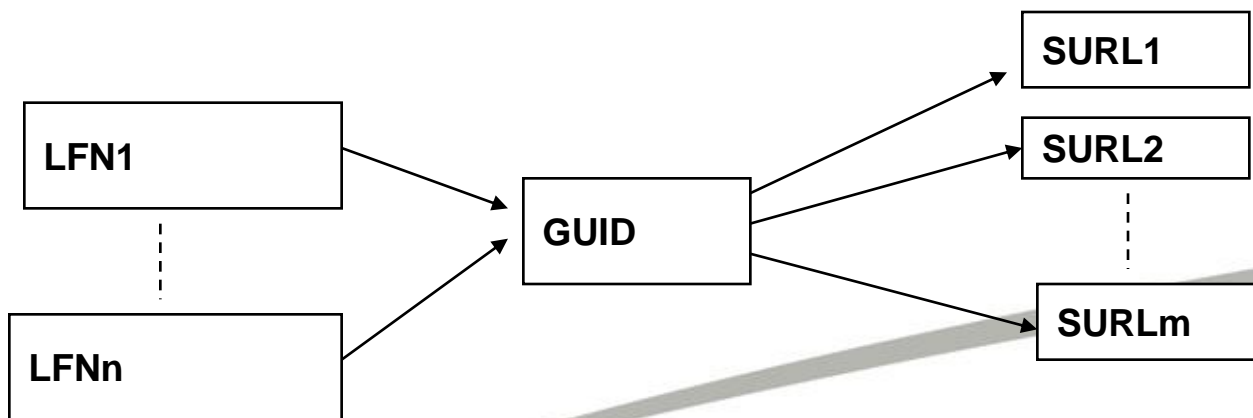
- an alias created by a user to refer to some item of data, e.g.
lfn:/grid/dteam/user/j/jdoe/text_file.txt
- primary LFN + aliases created as symlinks

- Global Unique Identifier (GUID)

- a non-human-readable unique identifier for an item of data, e.g.
guid:b2f3a73a-cd04-4d39-8c4e-3b8d8e84f7e3

- Storage URL (SURL) (or Physical File Name – PFN)

- the location of an actual piece of data on a storage system, e.g.
srm://srm-t2k.gridpp.rl.ac.uk/castor/ads.rl.ac.uk/prod/t2k.org/nd280/raw/nd280_00517.daq.mid.gz




EGEE LCG File Catalog

- Designed by Enabling Grids for E-science project
- Based on lessons learned in Data Challenges 2004
- Implemented entirely in C; based on the CASTOR Name Server code
- It fixes performance and scalability problems seen in EDG Catalogs
 - cursors for large queries
 - timeouts and retries from the client
- Deployed as a centralised service; its endpoint is published on the IS in order to be found by the LCG data management tools and/or other Grid services



LFC features

- All names are in a hierarchical namespace
- LFC has a directory tree structure



- GUID attached to every directory and file
- Integrated GSI Authentication and Authorization
 - mapping done from Client DN to uid/gid pair
 - authorization done in terms of uid/gid
 - VOMS are integrated (VOMS roles appear as a list of gids)
 - ownership of files is stored in catalog
 - permissions implemented
 - Unix (user, group, all) permissions
 - POSIX ACLs (group and users)

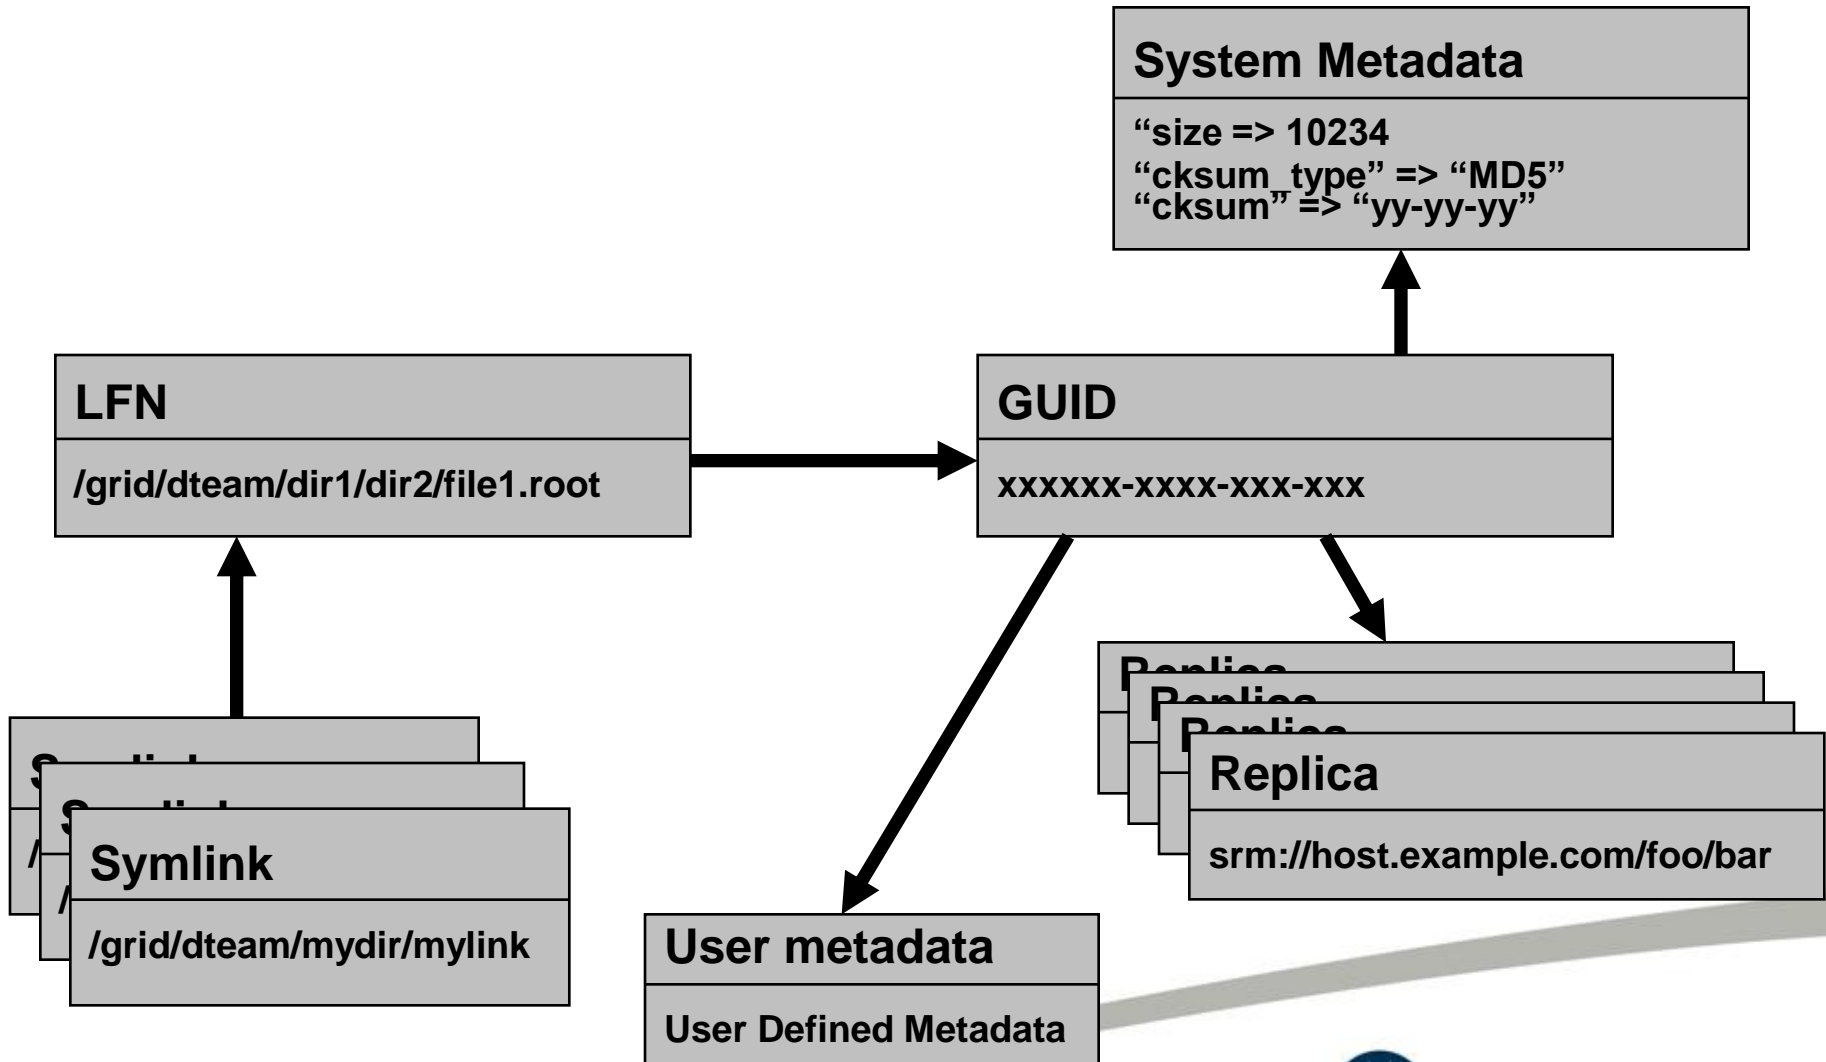


LFC features (2)

- Mapping between LFN, GUID and SURL
- Transactions, Sessions, Bulk methods
 - bulk delete (files, replicas)
 - bulk queries
- It runs as a multi-threaded daemon, with a relational (Oracle and MySQL) database backend.
- LFN acts as main key in the database. It has:
 - Symbolic links to it
 - Unique Identifier (GUID)
 - System metadata
 - Information on physical replicas
 - Single string user metadata



Relationships in the Catalog



The LFC – sets of commands

lfc commands

- used to interact with the catalog only
 - to create catalog directory
 - to list files
- used by user and by lcg-utils commands

lcg-utils

- couple catalog operation with the file management
 - keep SEs and catalog in step
- copy files to/from/between Ses

\$LFC_HOST – the LFC front-end



LFC Command Line Interface

<code>lfc-chmod</code>	change access mode of a LFC directory/file in the name server
<code>lfc-chgrp</code>	change group ownership of a LFC directory/file in the name server
<code>lfc-chown</code>	change owner and group of a LFC directory/file in the name
<code>lfc-ls</code>	list LFC name server directory/file entries
<code>lfc-ln</code>	make a symbolic link to a file or a directory in the LFC Name Server
<code>lfc-setacl</code>	set LFC directory/file access control lists
<code>lfc-getacl</code>	get LFC directory/file access control lists
<code>lfc-rm</code>	remove LFC files or directories in the name server
<code>lfc-mkdir</code>	make LFC directory in the name server
<code>lfc-rename</code>	rename a LFC file or directory in the name server
<code>lfc-setcomment</code>	add/replace a comment associated with a file/directory
<code>lfc-delcomment</code>	delete the comment associated with a file/directory

+ few others – for more details, please refer to the man pages.



lcg-utils commands

lcg-cp	copy a Grid file to a local destination
lcg-cr	copy and register a file
lcg-del	delete one file (either one replica or all replicas)
lcg-rep	copy a file from one SE to another SE and registers it in the catalog
lcg-gt	get the TURL for a given SURL and transfer protocol
lcg-aa	add an alias in the catalog for a given GUID
lcg-la	lists the aliases for a given LFN, GUID or SURL
lcg-lg	get the GUID for a given LFN or SURL
lcg-lr	lists the replicas for a given LFN, GUID or SURL
lcg-ra	remove an alias from the catalog for a given GUID
lcg-rf	register in the LFC a file residing on an SE
lcg-uf	unregister in the catalog a file residing on an SE

for more details, please refer to the man pages



LFC deployment at RAL-Tier1

- Remote Oracle backend
- Two LFC front-ends deployed for ATLAS
 - load-balancing (alias *lfc-atlas.gridpp.rl.ac.uk*)
 - serves as local file catalog
- One LFC front-end deployed for LHCb
 - read-only replicas of the CERN LHCb LFC
 - the LFC database content is replicated
(from CERN to the LHCb Tier1s, via Oracle Streams)
- One LFC front-end deployed for non-LHC VOs
 - *lfc.gridpp.rl.ac.uk*
 - serves as local file catalog for most VOs
 - exceptions: *minos.vo.gridpp.ac.uk*, *ralpp*, *mice*, *ngs.ac.uk*



LFC deployment at RAL-Tier1 (2)

- ORACLE implementation - remote Oracle server
 - tables in the FILECATALOG schema

```
CNS_FILE_METADATA  
CNS_USER_METADATA  
CNS_SYMLINKS  
CNS_FILE_REPLICA  
CNS_CLASS_METADATA  
CNS_GROUPINFO  
CNS_USERINFO  
CNS_UNIQUE_GID  
CNS_UNIQUE_UID  
SCHEMA_VERSION
```



LFC deployment at RAL-Tier1 (3)

```
SQL> desc FILECATALOG.CNS_FILE_METADATA
```

Name	Null?	Type
-----	-----	-----
FILEID	NOT NULL	NUMBER
PARENT_FILEID		NUMBER
GUID		CHAR(36)
NAME		VARCHAR2(255)
FILEMODE		NUMBER(6)
NLINK		NUMBER(6)
OWNER_UID		NUMBER(6)
GID		NUMBER(6)
FILESIZE		NUMBER
ATIME		NUMBER(10)
MTIME		NUMBER(10)
CTIME		NUMBER(10)
FILECLASS		NUMBER(5)
STATUS		CHAR(1)
CSUMTYPE		VARCHAR2(2)
CSUMVALUE		VARCHAR2(32)
ACL		VARCHAR2(3900)





Questions?

Thank you!