



Data Management in EGEE

Ákos Frohner, Jean-Philippe Baud,
Rosa Maria Garcia Rioja, Rémi Mollon,
David Smith, Paolo Tedesco (CERN)
Gilbert Grosdidier (LAL/IN2P3/CNRS)



DM

Outline

- EGEE project
- Data Management Software Stack
- Storage Element: DPM
- File Catalog: LFC
- Transfer: FTS
- Clients: GFAL, lcg_util
- Encrypted Storage: Hydra
- Future directions



DM EGEE Project

- To provide reliable and robust grid facilities for scientists
- Infrastructure
 - 268 sites over 52 countries
 - ~72,000 CPU
 - >20 PB storage in 291 SE instances
- Customers
 - High Energy Physics (LCG)
 - Biomedicine, Earth Sciences, ...
- gLite: lightweight opensource middleware

Client Tools

transfer

encryption

complex
operations

Higher Level Services

file
catalogreliable
transfer

keystore

POSIX
I/O API

Information System

Storage Elements

- Management: SRMv1.1, SRMv2.2
- Data Access: gridftp, http(s), NFSv4(.1)
- Data Access: rfiio, dcap, xrootd

DM

Our part in the Software Stack

Client Tools

FTS client

eds cli

lcg_util

Higher Level Services

LFC

FTS

Hydra

GFAL

Storage Elements

DPM

Castor

dCache

StoRM

BeStMan

DM

Our part in the EGEE Project

- GFAL, lcg_util clients on all worker nodes ~72,000
- 46 instances of LFC is used by two major VOs
- 15 FTS instances manage several 100s of MB/sec (up to 1.5GB/sec) traffic each, via 636 channels
- 190 DPM instances manage up to 360TB data on a single SE



DM Common Features

- Security
 - GSI authentication
 - VOMS authorization, including FQANs
- Information system
 - BDII info providers and queries
- Storage protocols
 - SRM v1.1 and SRMv2.2 support client or server respectively
 - gridftp

DM

Outline

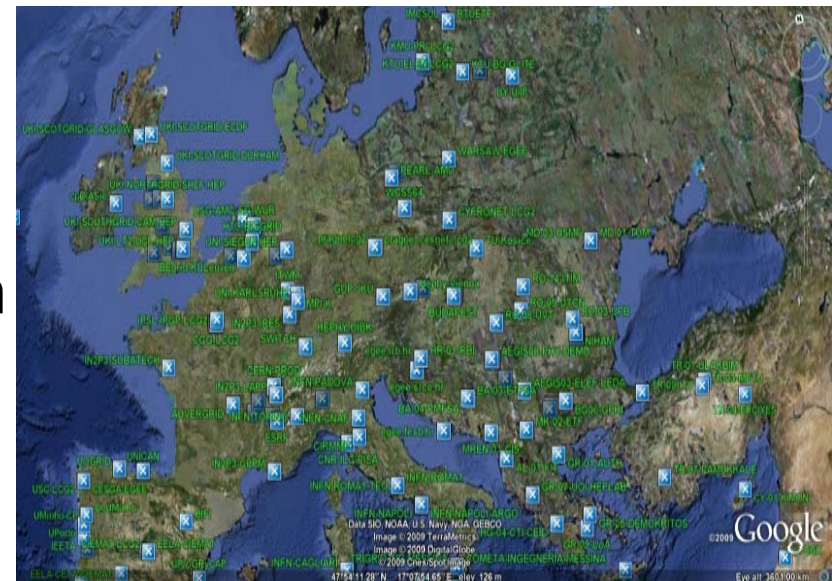
- EGEE project
- Data Management Software Stack
- **Storage Element: DPM**
- File Catalog: LFC
- Transfer: FTS
- Clients: GFAL, lcg_util
- Encrypted Storage: Hydra
- Future directions





Storage Element: DPM

- Manages disk-only storage
 - Deployments up to few 100TB data
 - Easy to install and manage
 - Interoperability with other storage systems
- SRM v1.1 and v2.2 support
- Data access
 - Gridftp
 - Secure rfiio
 - https with re-direction
 - Xroot (special)



DM DPM: features

- POSIX authorization semantics
 - user = certificate DN
 - group = VOMS FQAN
 - Independent of the underlying OS
 - Full ACL and secondary group support
- Disk pool management
 - Pool/space protection
 - Garbage collector
 - Replication of hot files

DM DPM v1.7.x

- Next release of DPM includes
 - Compliant with the WLCG SRM MoU
 - Checksum support: Adler32, MD5, Crc32
 - srmCopy support
 - Multi-platform ready
- Upcoming in v1.7.2:
 - Update of the xrootd plug-in for Alice
 - gLite release on Debian

DM

DPM future releases

- v1.8.0:
 - On-the-fly checksum calculation
 - Faster draining of disk servers
 - DB maintenance tools
 - User/VO banning
 - multi-VO xroot support
- v1.9.x
 - Better filesystem selection
 - Admin role for VOs
 - Quota support



DM

Outline

- EGEE project
- Data Management Software Stack
- Storage Element: DPM
- **File Catalog: LFC**
- Transfer: FTS
- Clients: GFAL, lcg_util
- Encrypted Storage: Hydra
- Future directions



- Logical to many physical file (SURL) mapping
- POSIX authorization semantics (see DPM)
- C and Python API (improved)
- Deployment models
 - Central with DB replication
 - Local catalogs
- OS & DB:
 - As with DPM
- IPv6 ready



DM LFC v1.7.1 and after

- Next release of LFC
 - Multi-platform support
- Focus is on performance by bulk and compound methods:
 - **v1.5.x:**
 - lfc_delreplica(replica)
 - lfc_unlink(name)
 - lfc_readdir(directory)
 - **v1.6.x:** up to 10 times speedup
 - lfc_delfilesbyguid(array of guides)
 - lfc_delfilesbyname(array of names)
 - lfc_delfilesbypattern(pattern)
 - lfc_delreplicas(array of replicas)
 - lfc_getreplicas(array of guides)
 - lfc_readdirxp(directory, pattern, SE) reading subset of a directory

Data Management in EGEE, CHEP 2009, Prague - 15



DM

Outline

- EGEE project
- Data Management Software Stack
- Storage Element: DPM
- File Catalog: LFC
- **Transfer: FTS**
- Clients: GFAL, lcg_util
- Encrypted Storage: Hydra
- Future directions



DM Transfer: FTS

- Reliable File Transfer Service
 - Bulk data transfers between SRM compliant storage elements
 - Multi-VO service to balance network/SE utilization
 - Prevent overloading network/SE resources
 - Service monitoring and statistics



DM FTS: Channel

- Single direction queue for transfer jobs
- Between
 - CERN-RAL single sites
 - siteA-[T2region] new management tools
 - CERN-* catch-all channel
- Parameters
 - Detailed timeouts scales with size
 - VO shares
 - Priorities
 - Limit of parallel transfers srm/gridftp split!

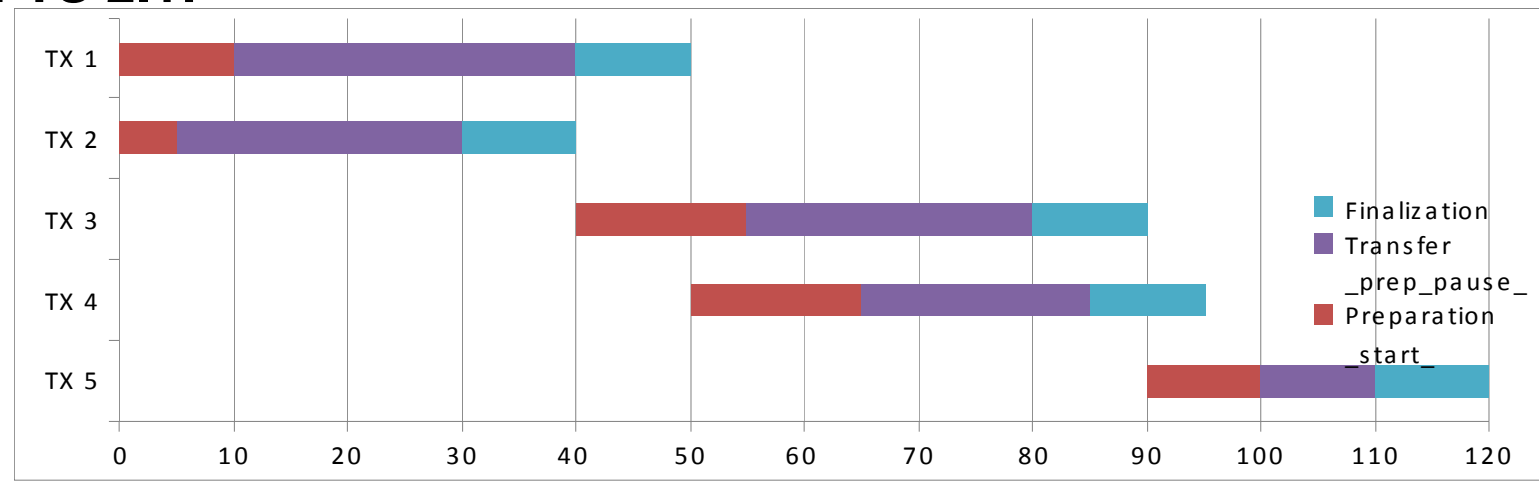




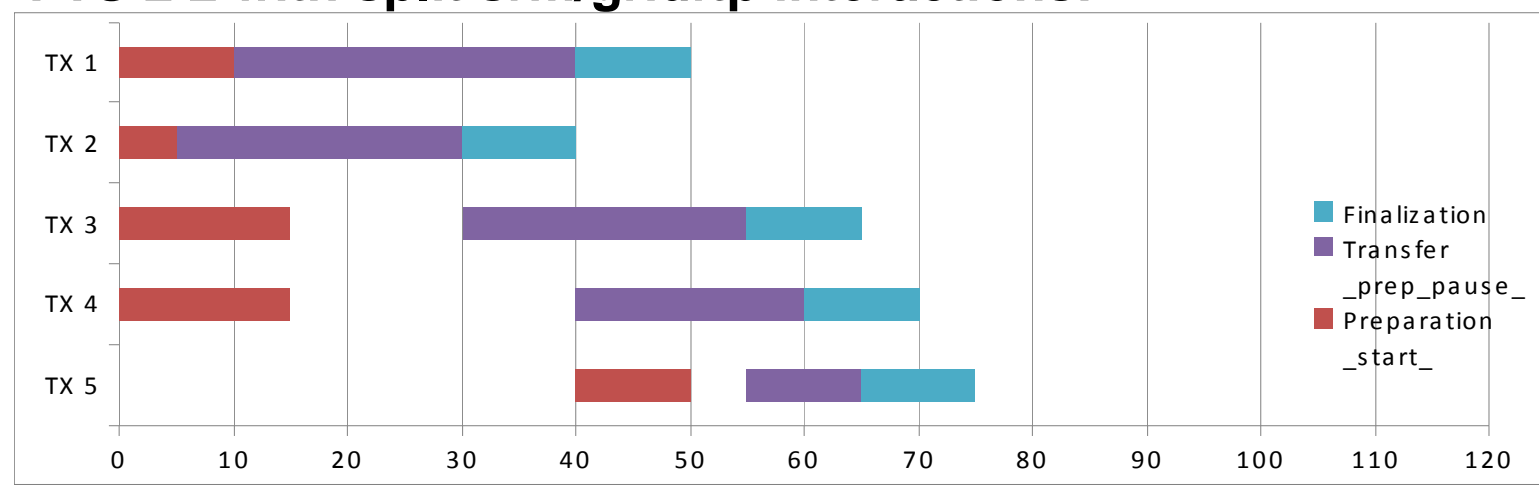
DM

FTS 2.2 speedup

FTS 2.1:



FTS 2.2 with split srm/gridftp interactions:



DM FTS v2.2

- Next release of FTS (48 fixes/features)
 - srm/gridftp split in 'urlcopy' channels
 - CLI tools for channel configuration
 - Implementing the WLCG SRM MoU
 - Smarter timeouts for large files
 - Multi-platform release
- Upcoming releases
 - Checksum support
 - Avoiding SRM overload (aka 'srm busy')
 - New WSDL and Python client library
 - Live transfer information

DM

Outline

- EGEE project
- Data Management Software Stack
- Storage Element: DPM
- File Catalog: LFC
- Transfer: FTS
- **Clients: GFAL, lcg_util**
- Encrypted Storage: Hydra
- Future directions



DM

Client: GFAL, lcg_util

- Hides the complexity for users
- GFAL
 - POSIX-like I/O functions (open(), read(), close())
 - SRM abstraction layer
 - C, Python APIs and CLI
- lcg_util
 - Covering most common use cases:
 - File creation, registration, replication, deletion,...
 - C, Python APIs and CLI



- Next release:
 - Detailed timeout support
 - Fall-back on secondary replica
 - If first is not available
 - Classified replicas
 - Improved Python interface
 - Implementing the WLCG SRM MoU
- Upcoming releases
 - Checksum support
 - Avoiding SRM overload (aka 'srm busy')
- In the long term:
fully synchronize the behavior with FTS

DM

Outline

- EGEE project
- Data Management Software Stack
- Storage Element: DPM
- File Catalog: LFC
- Transfer: FTS
- Clients: GFAL, lcg_util
- **Encrypted Storage: Hydra**
- Future directions



DM

Encrypted Storage: Hydra

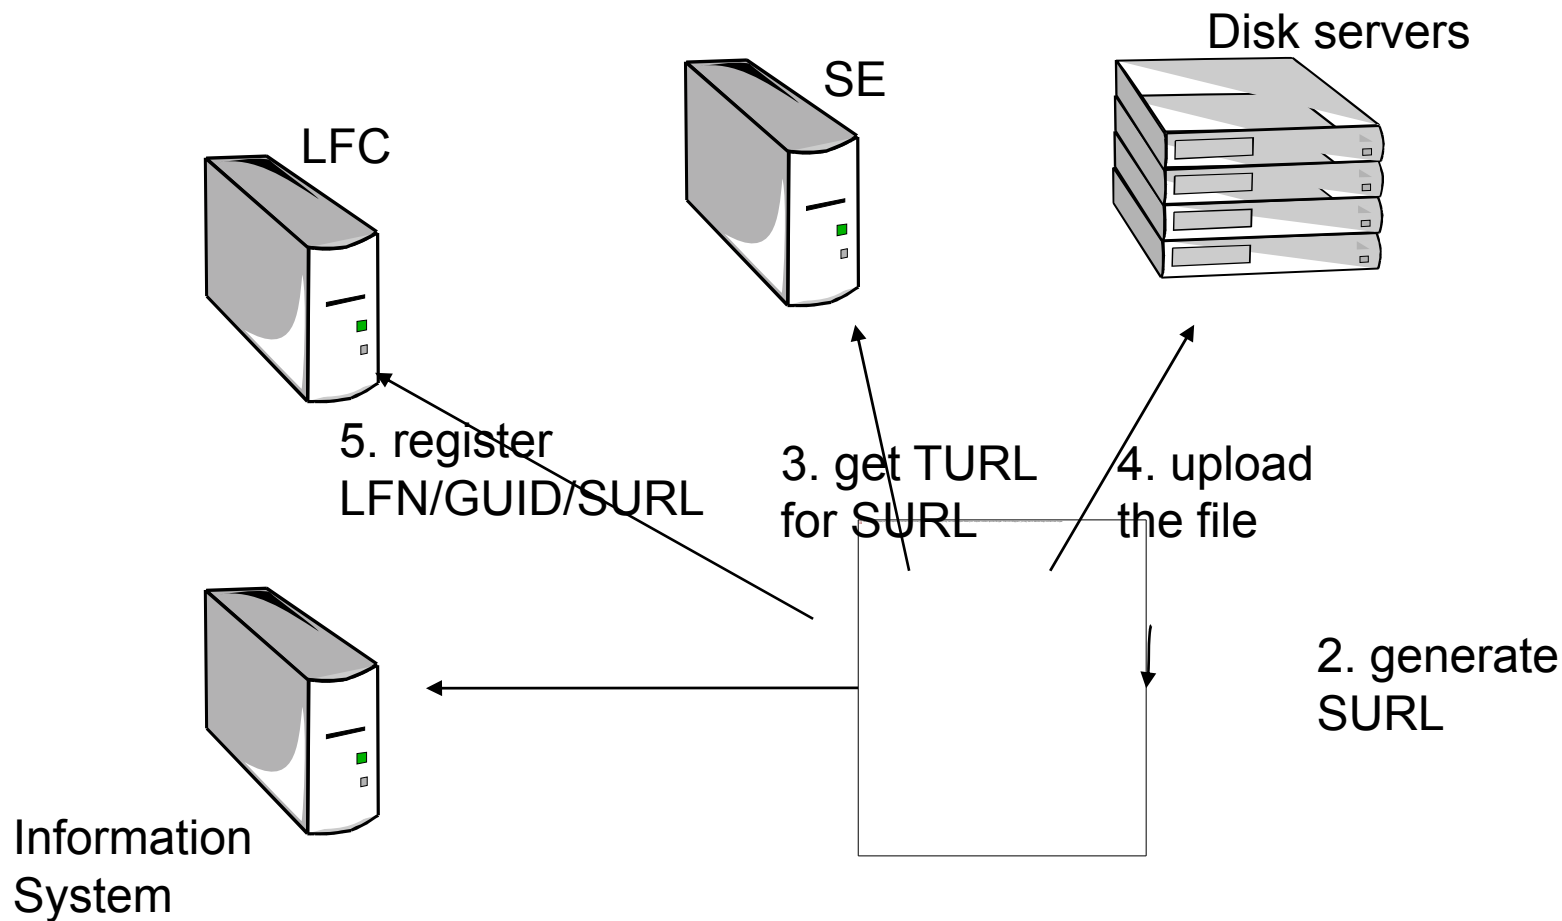
- Encrypted file storage solution
 - Data integrity and confidentiality
 - Encrypted files are in existing components (LFC, DPM, GFAL)
 - Quick symmetric encryption
- Hydra Keystores
 - M-of-N pieces of the key is stored (Shamir's Secret Sharing Scheme)
 - Web service interface



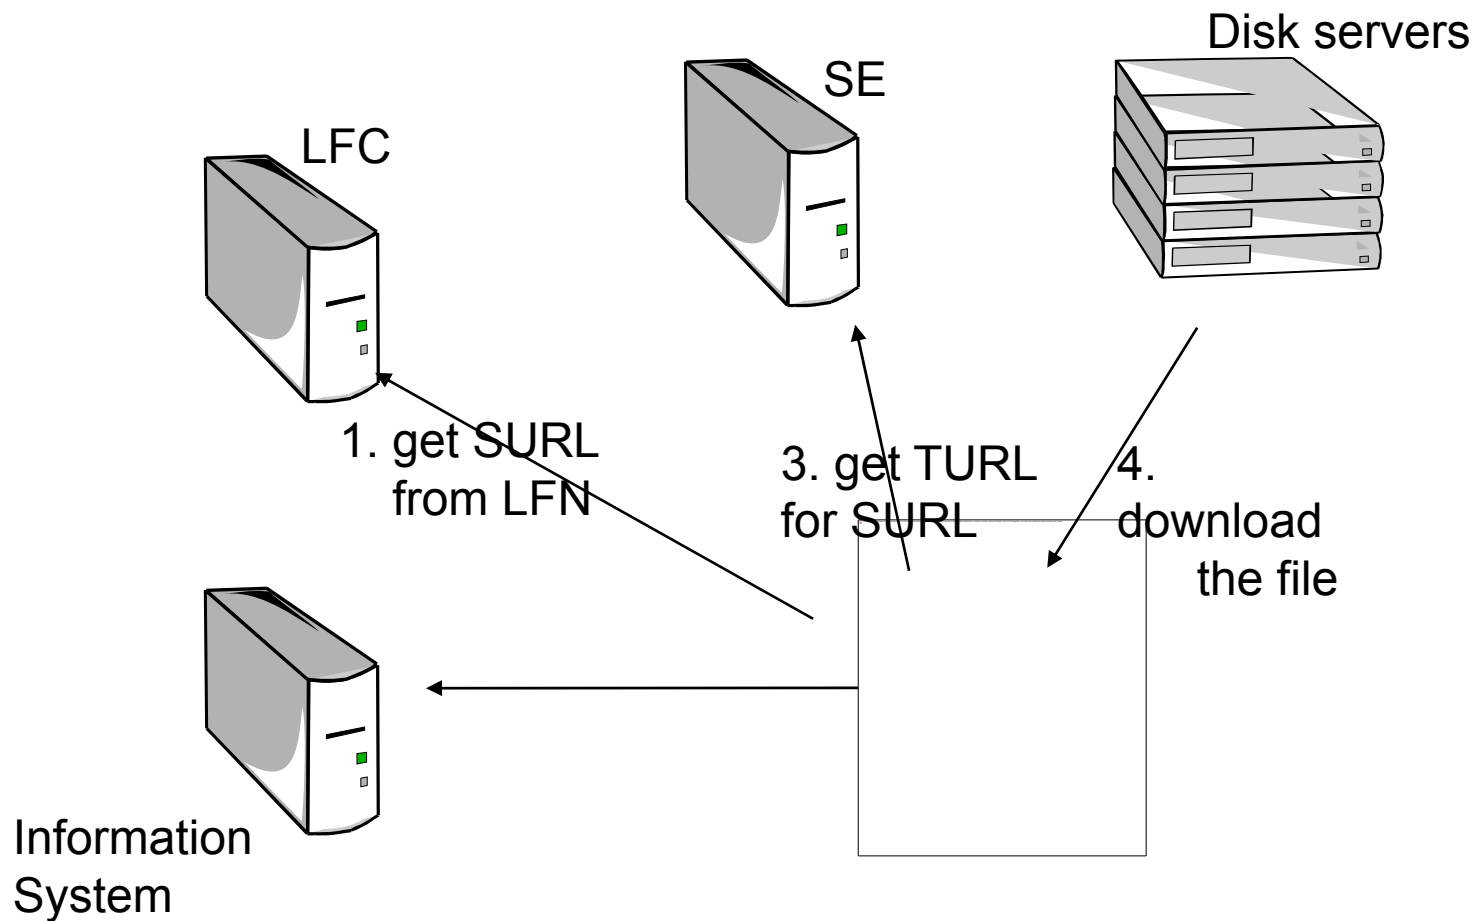
DM Future directions

- Priorities
 - Stability
 - Reliability
 - Maintainability
- Improved administrative tools
 - Real time monitoring of services
 - Automating regular procedures (e.g. cleanup)
 - Friendly (web) interface for configuration
 - Resource protection (quota, limits)
- Better integration
 - Status feedback from services (i.e. “srm busy”)
 - Framework integration (notifications, libraries)

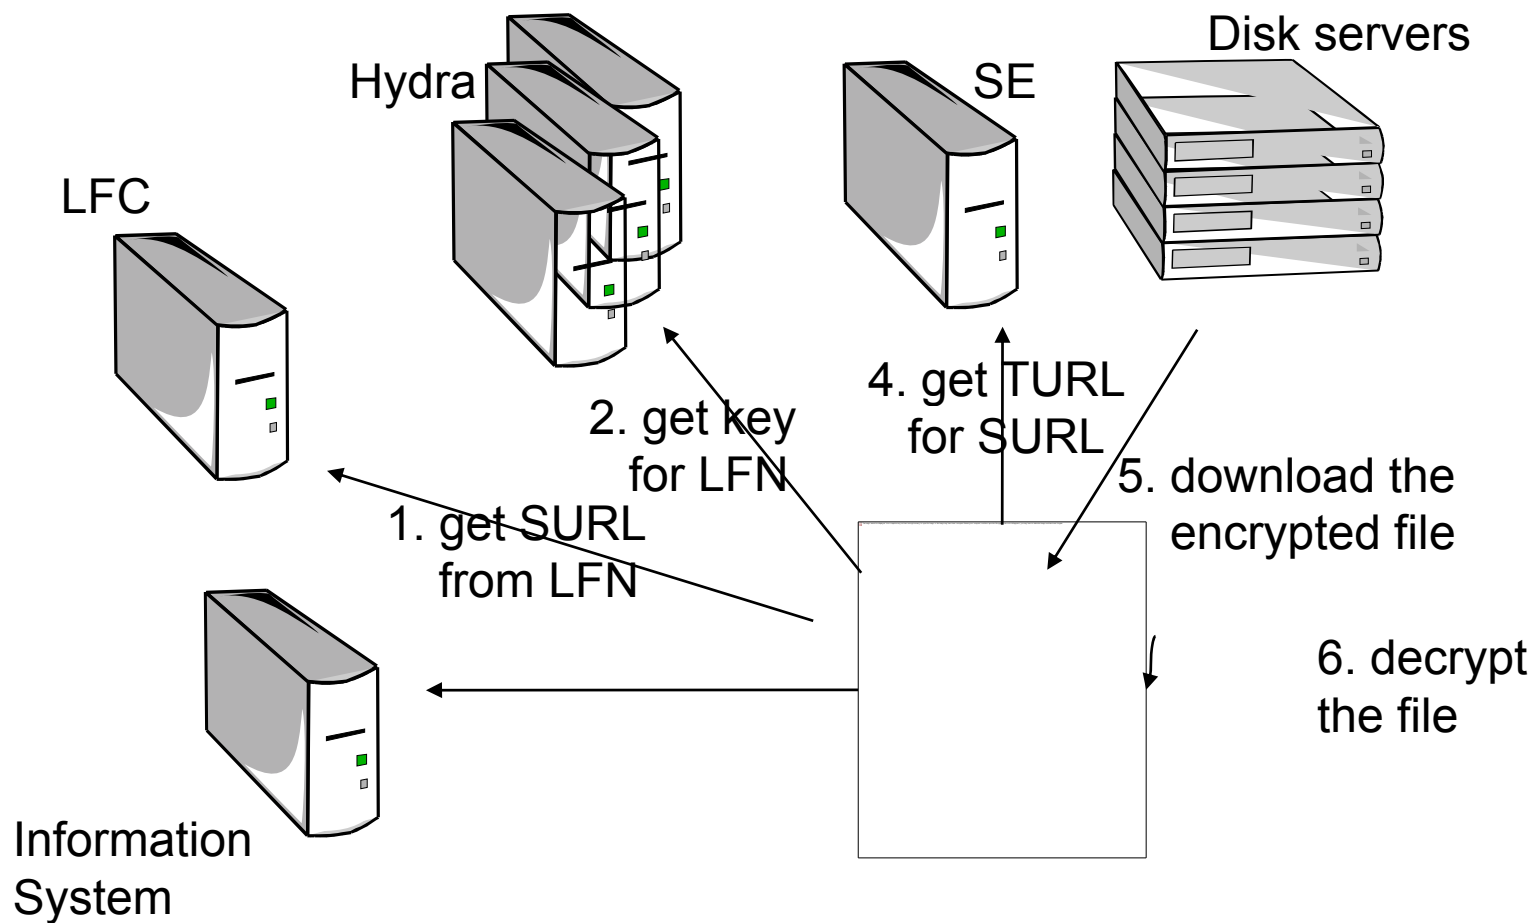
```
lcg-cr -d myse.cern.ch -l lfn:/grid/dtem/myfile /path/to/my/local/file
```



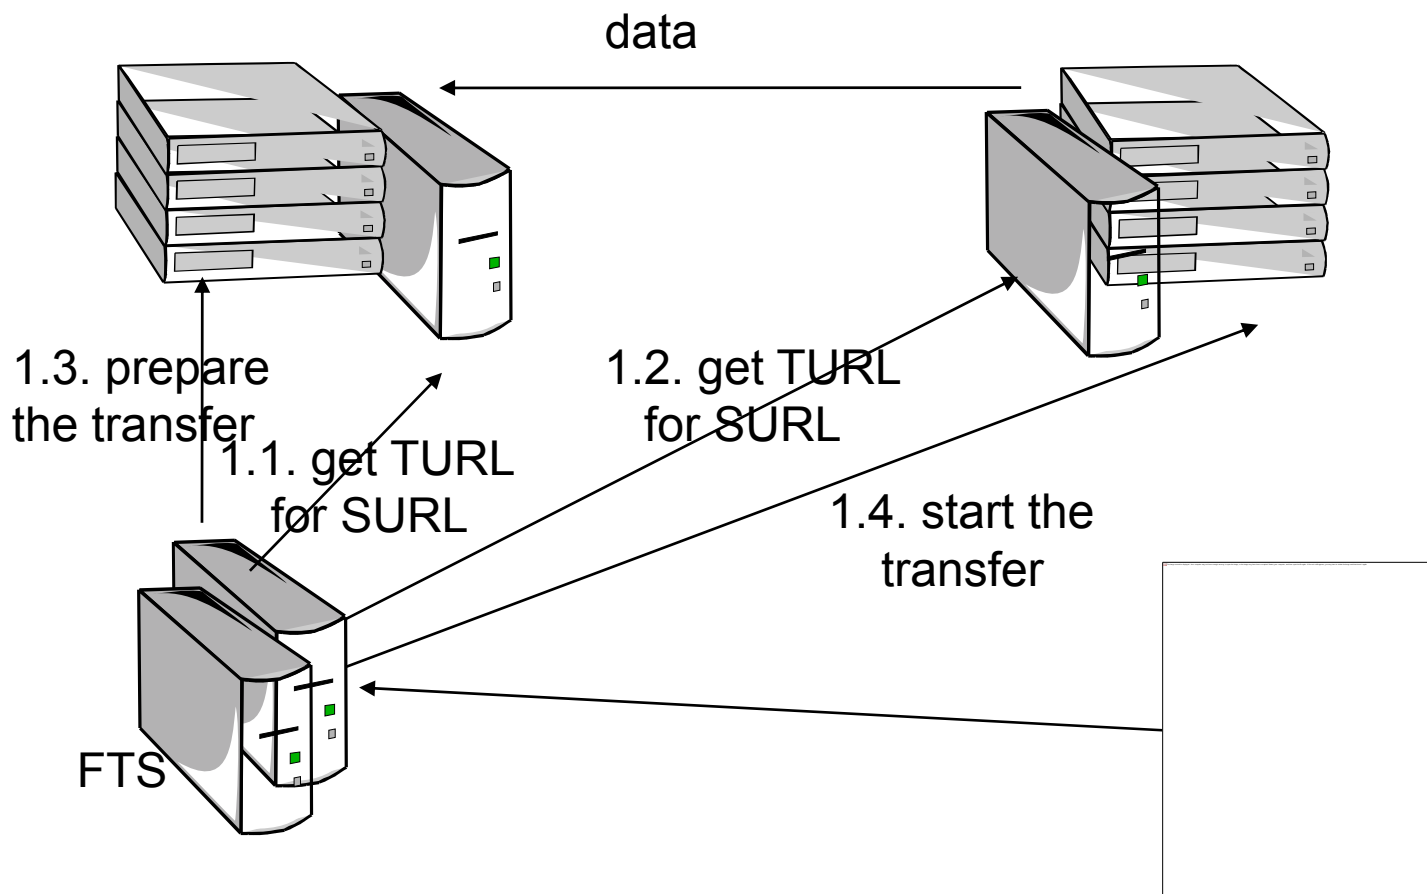
```
lcg-cp lfn:/grid/dtem/myfile /path/to/my/local/file
```



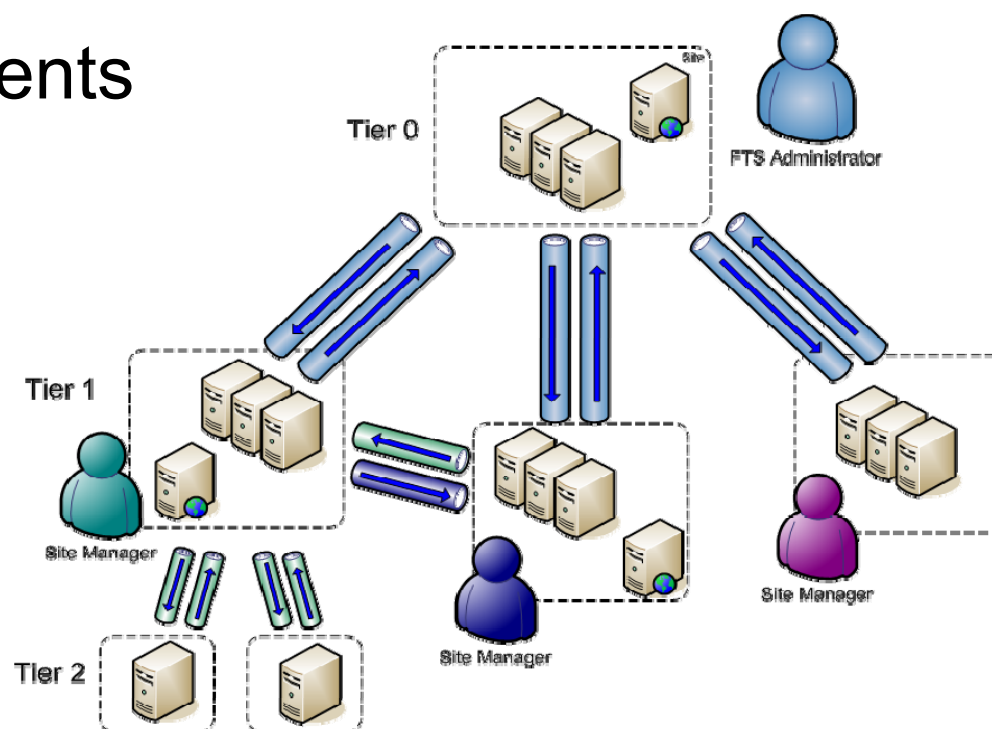
```
glite-eds-get lfn:/grid/dtem/myfile /path/to/my/local/file
```



```
glite-transfer-submit srm://myse.cern.ch:8443/srm/managerv2?SFN=/foo  
srm://myse.example.org:8443/srm/managerv2?SFN=/foo
```



- Decoupled components
 - Web service for client interaction
 - VO and Channel agents
 - Monitoring service
- Service components
 - Stateless
 - Load balanced
 - Highly available



DM

FTS: srm/gridftp split

FTS <= 2.1: channel = sum(Prepare + Transfer + Finalize)
FTS >= 2.2: channel = sum(Transfer)

