Feedback from Pilot LFC/FTS Usage by Experiments

Patricia Méndez Lorenzo on behalf of the EIS Group
CERN (IT-GD) / CNAF

Pre-GDB Meeting
CERN, 6th-September, 2005

Scope of this talk

- To give a summary of the experiments' experiences during the pilot phase of SC3

- I remind you who we are:

Alice: Patricia Méndez Lorenzo

Atlas: Simone Campana

Biomed: Antonio Delgado Peris

CMS: Andrea Sciabá

LHCb: Roberto Santinelli

Alice: Summary of GDB (July 05)

This is what we said:

Requirements:

- VO-managed node at each site
  - Alice takes the responsibility to deploy and support the software at each site
    - VO Box rpm deployed in LCG2.6.0
    - VO Box prototype already deployed at CERN
- Central Catalog (Metadata Catalog) provided by Alice
  - Local LFC Catalog required at each site
    - Installation of a LFC server in Torino
    - Support of the LFC testing
- FTS server
    - Support for FTS testing
    - Maintenance of the FTS UI

ALICE Feedback

2 months later…

- **Alice is coordinating with SC3 to run their DC05 in the SC3 framework**
- **Already beginning in September**

Principal Goals

- Use of the deployed LCG SC3 infrastructure for the Alice DC05
- Test of the data transfer and storage services
- Test of the distributed reconstruction and calibration model
- Integration of the LCG resources with other resources available
- Analysis of reconstructed data

ALICE Feedback

Current Status: 1st Phase → Production

- Waiting for VO-boxes to be operational, jobs are run by ALICE services (AliEn)
  - Expecting to run 600-800 jobs in parallel in a matter of a day:
  - Reaching 1000 jobs per day after the first week
    - All output data stored in CASTOR2 – registered in AliEnFC and in the LFC of the site
- Expected duration: 3 weeks
- Expecting to continue the production using the SC3 resources

ALICE Feedback

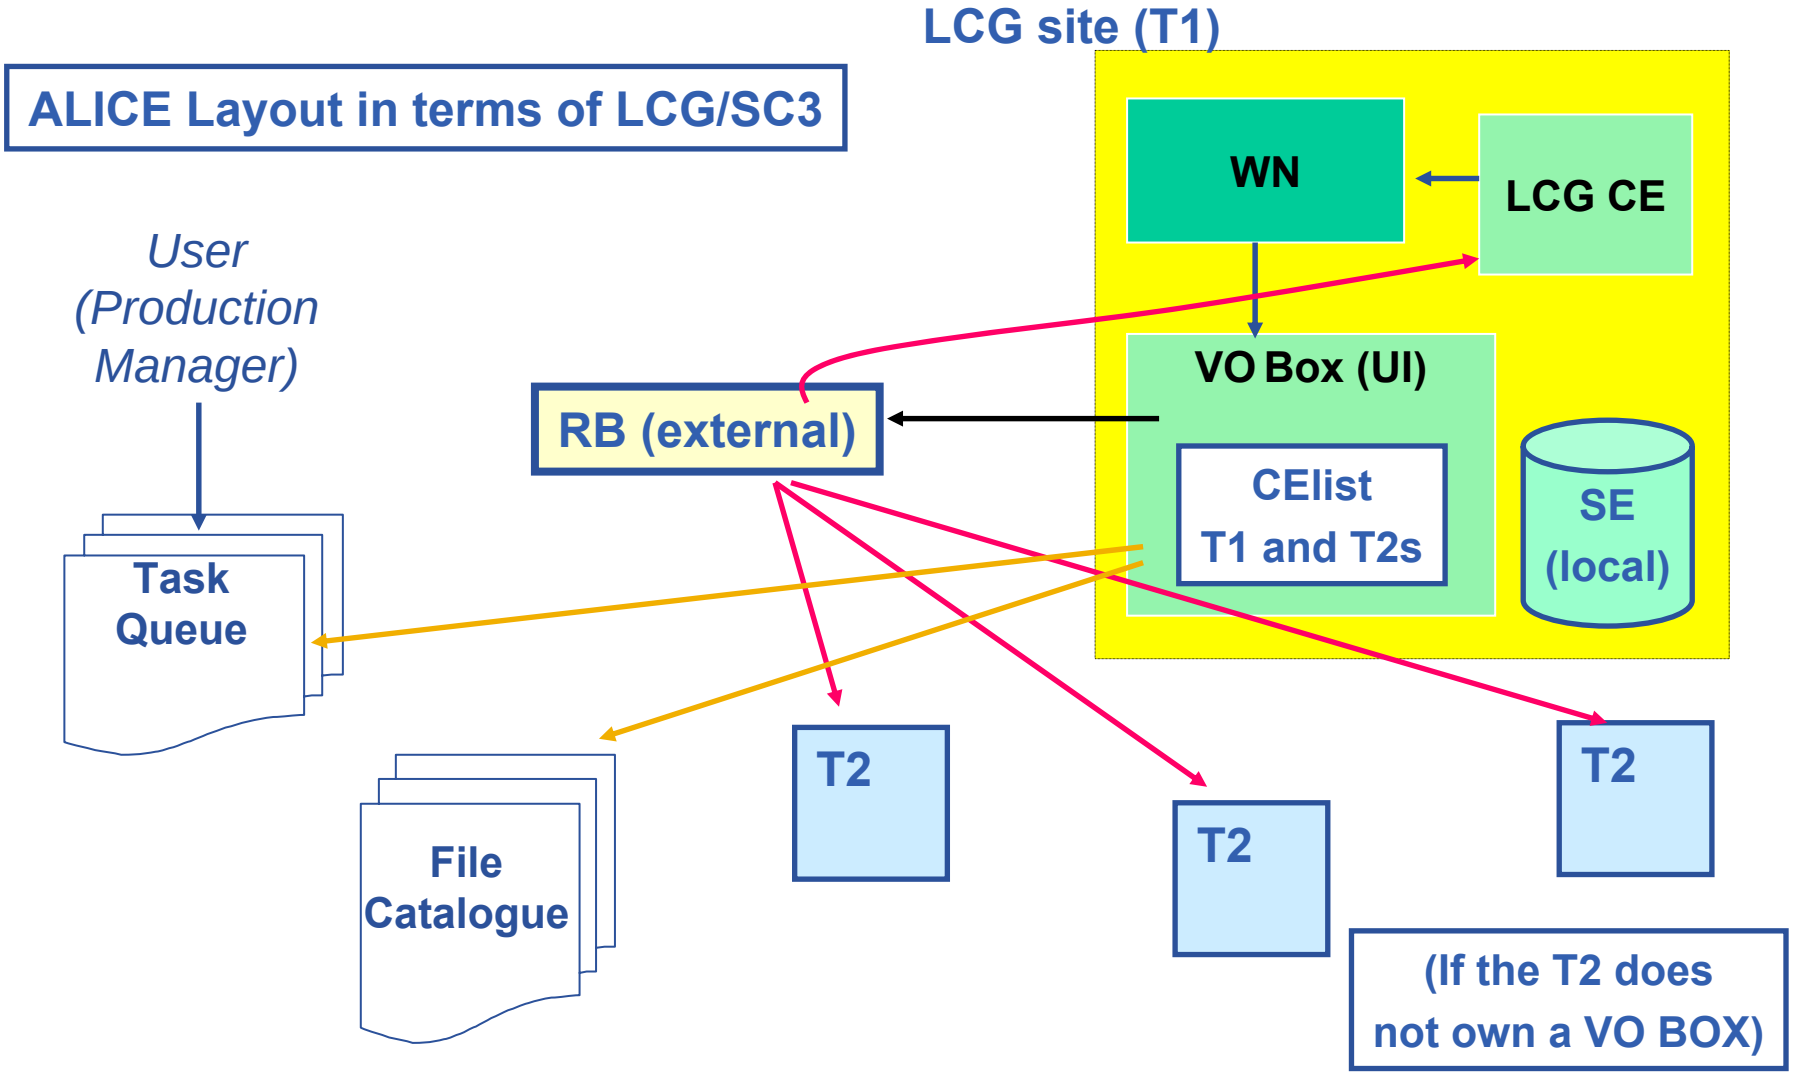

Alice Requirements and Status

- VO Boxes deployed in all sites
  - ➳ CERN, CNAF, NiKHEF/SARA, IN2P3, Catania Torino, Bari, GSI, GridKA.
  - ➳ AliEN specific services and software deployed in all VO Boxes

    Services: CE (job submission to the RB through the UI deployed to the VO BOX), SE (xrootd, LFC), PACKMAN (Software Installation), FTD (transfers)
  - ➳ Submissions through VO BOX
    - Job submission to RB is possible
    - Still some infrastructure missing (env variables, lcg-infosites...)
    - Completed for next release (LCG2.7.0 in October)

ALICE Feedback

- LFC available in all sites
  - Considered the local Alice catalog
  - Central Alien storage index
  - Perl API implemented in the AliEn Framework
    - More than 10000 entries (LFC as unique catalog)
    - Too many authentications slow the process
- FTS available in all sites
  - Perl API almost implemented in the Alice framework
  - Tests among T0-T1 performed this summer
  - FTS will be used as FTD plugin
  - FTD was tested between native AliEn sites
- SRM
  - Single implementation for all SEs types provided by LCG

ATLAS: Summary of GDB (July 05)

This is what we said:

Requirements:

- Deployment of FTS
  - Comparison of FTS performance with Don Quijote RFT
  - Following the throughput phase and investigating the integration of FTS in Don Quijote
- For the moment RLS entries are being migrated to global LFC
  - Using their pilot as a global copy of the RLS to run some analysis
  - Cleaning the entries in the catalog before splitting it in many local catalogs

ATLAS Feedback

2 months later…

- ATLAS is joining the Production in October
- Production split in two parts:

Phases and Principal Goals:

- Part 1: During the SC3 Throughput phase
  - LFC testing
    - Migration of RLS to LFC (pilot) entries
    - Cleanup of corrupted entries in the catalogs
    - Use of the LFC pilot as a RLS copy to run some analysis

ATLAS Feedback

- FTS testing
  - Comparison of FTS with Don Quijote RFT
  - Integration of Don Quijote with FTS and distribution to sites
  - First FTS exercises and tests at CERN, CNAF, Pisa and Milano

Part 2: Two major exercises during their Production

1. Tier-0 exercise
2. MC production in Tier2 and reprocessing in Tier 1

ATLAS Feedback

Principal Requirements

- VO Boxes in all Tier1 and Tier2
  - Already available at CERN, Pisa and Milano
- FTS deployed in all sites
  - Servers at CERN, deployment ongoing at CNAF
- LFC is considered local catalog and mandatory in all sites
  - Available at CERN, soon at CNAF and Milano

CMS: Summary of the GDB (July 05)

This is what we said:

Requirements:

Transfers: Done with PhEDEx

- FTS is not currently used
- Integration with PhEDEx likely to happen later

Catalog: Local file catalog is needed by PhEDEx

- Only POOL MySQL catalog in use now
- POOL-LFC interface is being tested by EIS
  - Functionality tests (done)
  - Performance tests (partially done)
  - Catalog migration from XML, MySQL to LFC (partially done)
  - Some issues discovered and already fixed, some yet to be solved

CMS Feedback

2 months later…

- **CMS is planning to begin the production middle September**

Principal goals:

- Testing of the data transfer and the data serving infrastructure known to work for realistic use
- Testing the Workload Management components including the Resource Broker and the Computing Elements

CMS Feedback

- Status in terms of the Catalogue

  - POOL LFCCatalog plugin released (2.1.0)

  - Functionality tests performed
    - POOL CLI thoroughly tested
    - Some bugs found, fixed in later releases (now 2.1.2)

  - Performance tests made
    - Test suite provided by CMS, same used to validate any POOL catalog no changes needed to use it with LFC
      - Performance looks reasonable, but no clear indication on whether they are up to production requirements

Configuration of the Test Suite

Site	LFC client	LFC server	Backend
CERN	PIII 1 GHz	dual Xeon 2.4 GHz	Oracle
Bari	PIII 1.266 GHz	PIV 2.8 GHz	MySQL

- CERN is the LFC pilot installation
  - Database shared among all VOs
- 10000 fake entries inserted

CMS Feedback

	LFC				Oracle(*)		MySQL(*)	
No. of clients	Oracle		MySQL					
	Query by GUID (ms)	**Query by PFN (ms)**	**Query by GUID (ms)**	**Query by PFN (ms)**	**Query by GUID (ms)**	**Query by PFN (ms)**	**Query by GUID (ms)**	**Query by PFN (ms)**
1	**384**	**87**	**264**	**102**	**13**	**7**	**4**	**4**
5	**111**	**18**	**79**	**20**	**5**	**6**	**12**	**13**
20	**159**	**6**	**120**	**6**	**5**	**5**	**15**	**14**
50		**6**		**8**	**7**	**7**	**16**	**14**

CMS Feedback

Results of the tests

Queries by PFN as fast as with other catalogs for parallel queries

Queries by GUID are a lot slower than those by PFN

* Due to the session or the transaction handling
* Some queries require a reconnection to the LFC server outside of the current session or transaction
* The extra authentication causes the overhead

Time to publish an XML fragment of 19 files to DB:

~ 0.5 s Oracle

~ 3 s LFC Oracle

~ 2.2 s LFC MySQL

CMS Feedback

ToDo

- Short-term (2-3 weeks)
  - Configure PhEDEx with POOL-LFC interface
  - Run realistic transfers from remote sites
  - Test migration from existing POOL catalogs to LFC
  - More extensive scalability tests?
- Medium term (1-2 months)
  - Test LFC as a POOL file catalog for CMS analysis jobs
- Long term (3-4 months)
  - Stronger integration with LCG SRM SEs (DPM, dCache, CASTOR)?

LHCb: Summary of the GDB (July 05)

This is what we said:

Requirements:

- **LFC central catalog**
  - **Needed a conversion of their replicas PFN format**
  - **Incompatibilities among both catalogs**
  - **Inclusion in Dirac (Python APIS have been required)**
- **Data Transfer: FTS**
  - **Inclusion of the FTS Client in Dirac and first tests:**

    CERN-CERN CERN-GRIDKA
    CERN-INFN CERN-IN2P3
    CERN-PIC CERN-SARA

LHCb Feedback

2 months later…

- LHCb is planning to begin in October

Principal Goals:

- Evaluation of the services during the throughput phase
- Demonstrate the Data Management to meet the requirements of the computing model
- Test the full data processing
- Full integration of the Data and the Workload Management systems

LHCb Feedback

Current Status in terms of FTS

- Working in the integration of FTS client (still not tested 1.3 version) with the LHCb Data Management System
  - LFC client, lcg-*, FTS, Bookkeeping File Catalog, etc
- Implementation of the FTS in their Transfer Agents
- Their (only) tests of FTS have been done on July (mid of July)
  - There is success on the CERN-CERN, CERN-GRIDKA, CERN-INFN channels

LHCb Feedback

Current Status in terms of LFC

- Implementation of the different catalogs (LFC, Bookkeeping File Catalog, AliEn) in the general Data Management System
- They are considered central services (right now placed at CERN) and all of them will be synchronized
- LHCb T1 hold local file catalogs (read only. Replica of the central catalog) not to collapse the central service and redundancy
- Very good experience using the LFC