Status of tests in the LCG 3D database testbed

LCG Database Deployment and Persistency Workshop

Eva Dafonte Pérez

Agenda

- STREAMS Overview
- TESTBED Configuration
- TESTS
- Future TESTS

STREAMS Overview

- Flexible feature for information sharing
- Basic elements:
  - Capture
  - Staging
  - Consumption
- Replicate data from one database to one or more databases
- Databases can be non identical copies

STREAMS Architecture

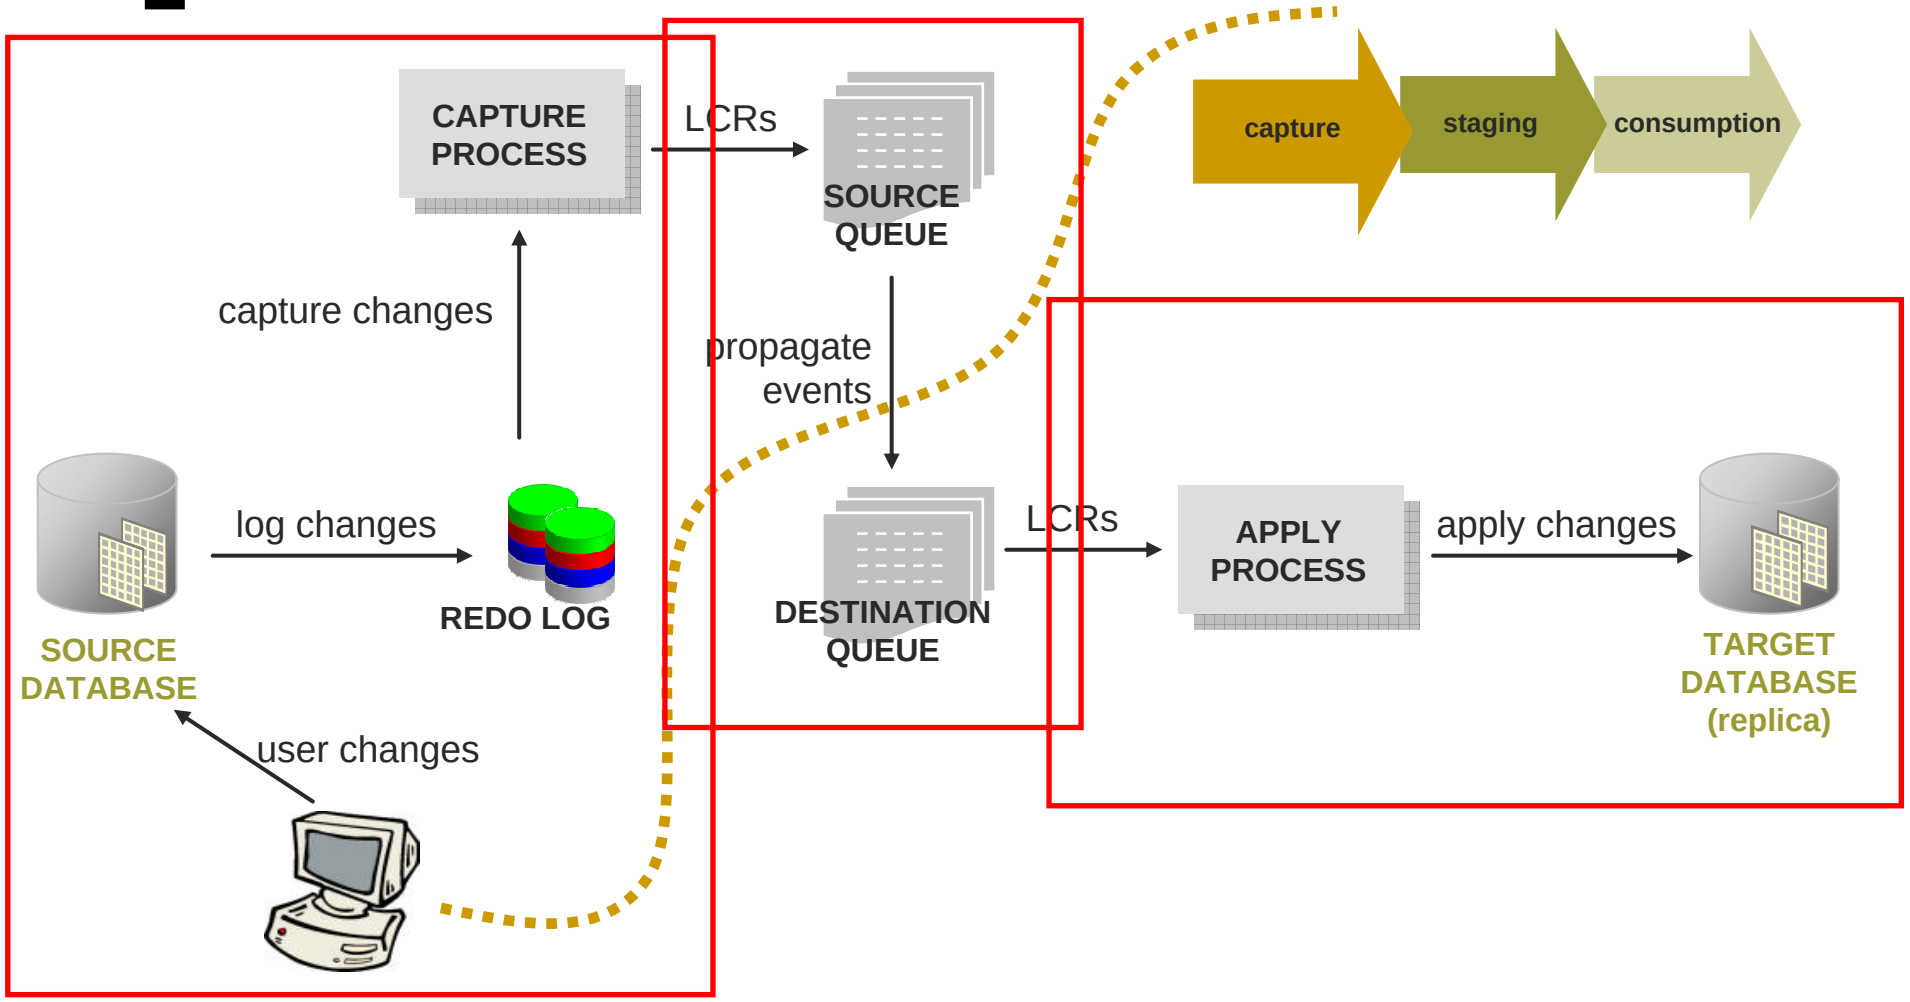

TESTBED Configuration

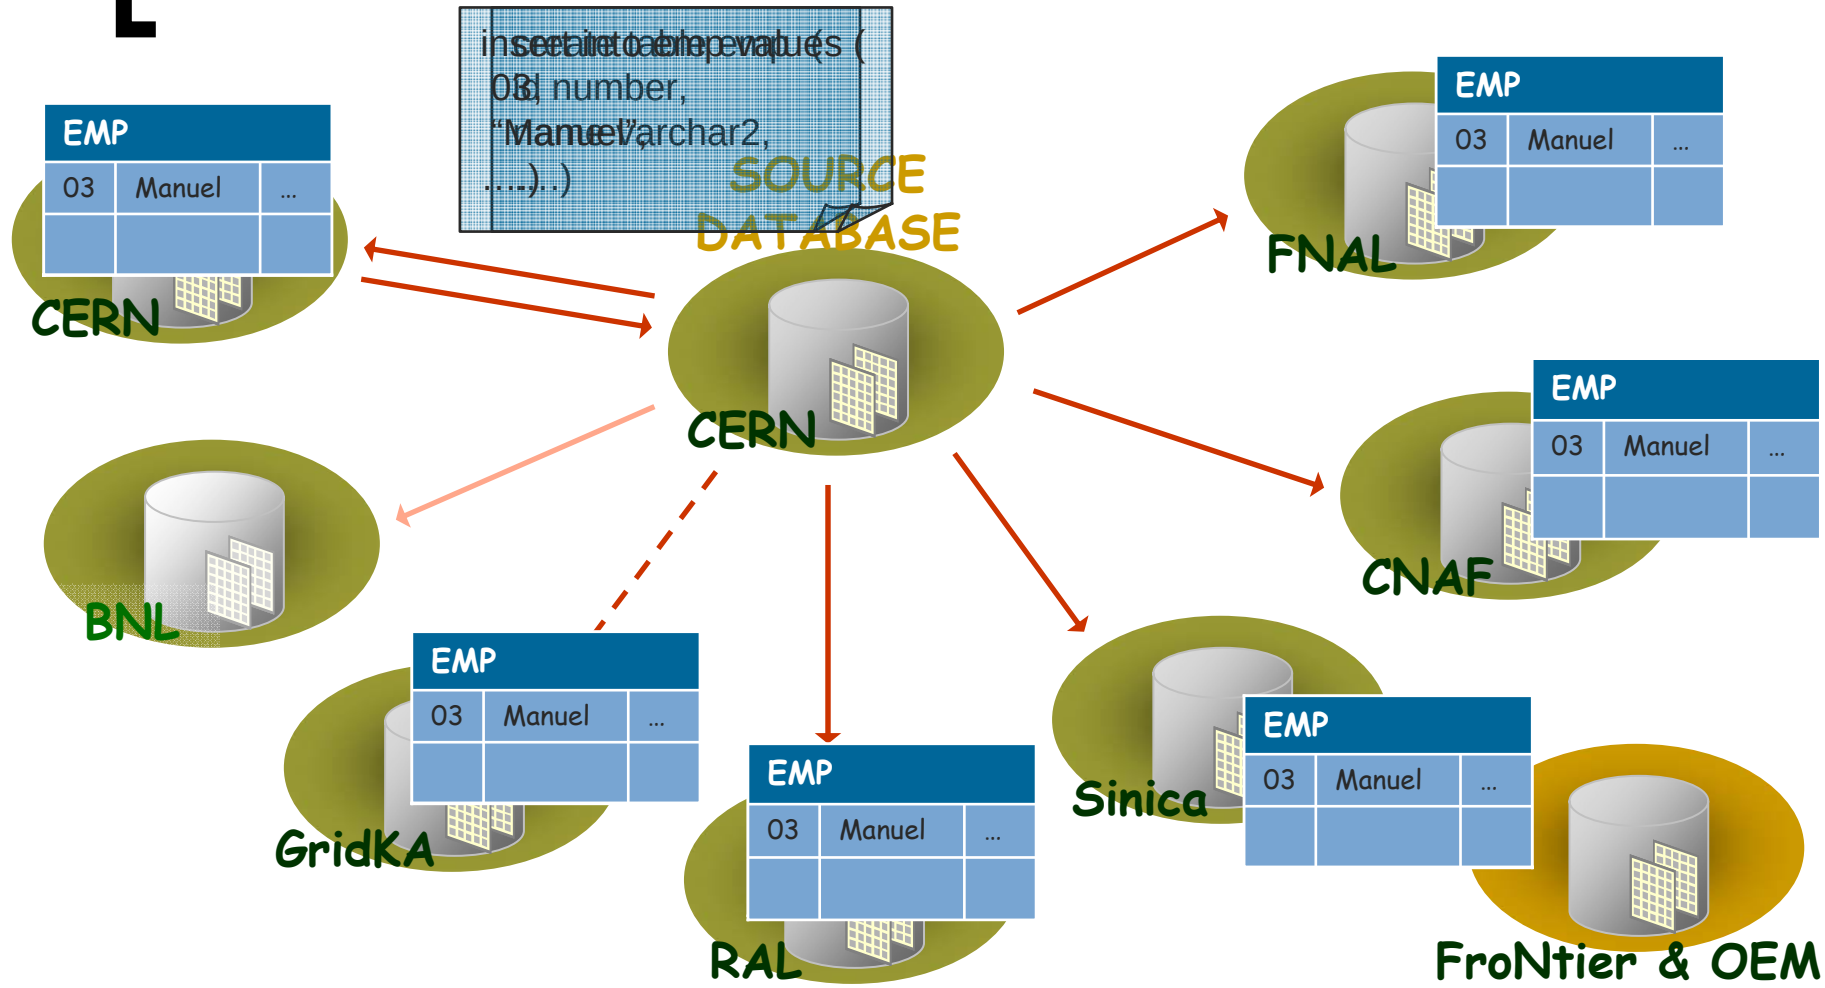

TESTBED Configuration

- Database version 10.1.0.4
  - Apply patch for bug 4291110 (ORA - 25228)
- Database version 10.1.0.3 not usable at source database due to bug 3840917
- Connection problems with GridKA

Network people working on

- Propagation to target disabled because of:
  - Listener down
  - Database not available
  - …

ORA – 04031
Unable to allocate %s bytes of shared memory

- Resources limited at CERN

WIKI and SAVANNAH

- LCG 3D Project - Distributed Deployment Of Databases For LCG wiki page
  - General information, minutes from meetings, configuration descriptions and scripts, …
  - 3D users should register as wiki users
  - Redistribution on validation and production system

https://uimon.cern.ch/twiki/bin/view/ADCgroup/LCG3DWiki

- Project tracker: Savannah
  - Log operations by each site

    https://savannah.cern.ch/projects/lcg3d/

3D OEM

- Installation script available in the LCG 3D wiki
  - Every site should join
  - Access to the OEM (account required):

    http://lxfs6043.cern.ch:7777/em/
- 10gR1: no Streams monitoring
  - own monitoring scripts used
- 10gR2: Look for Streams on Maintenance page of EM

3D OEM

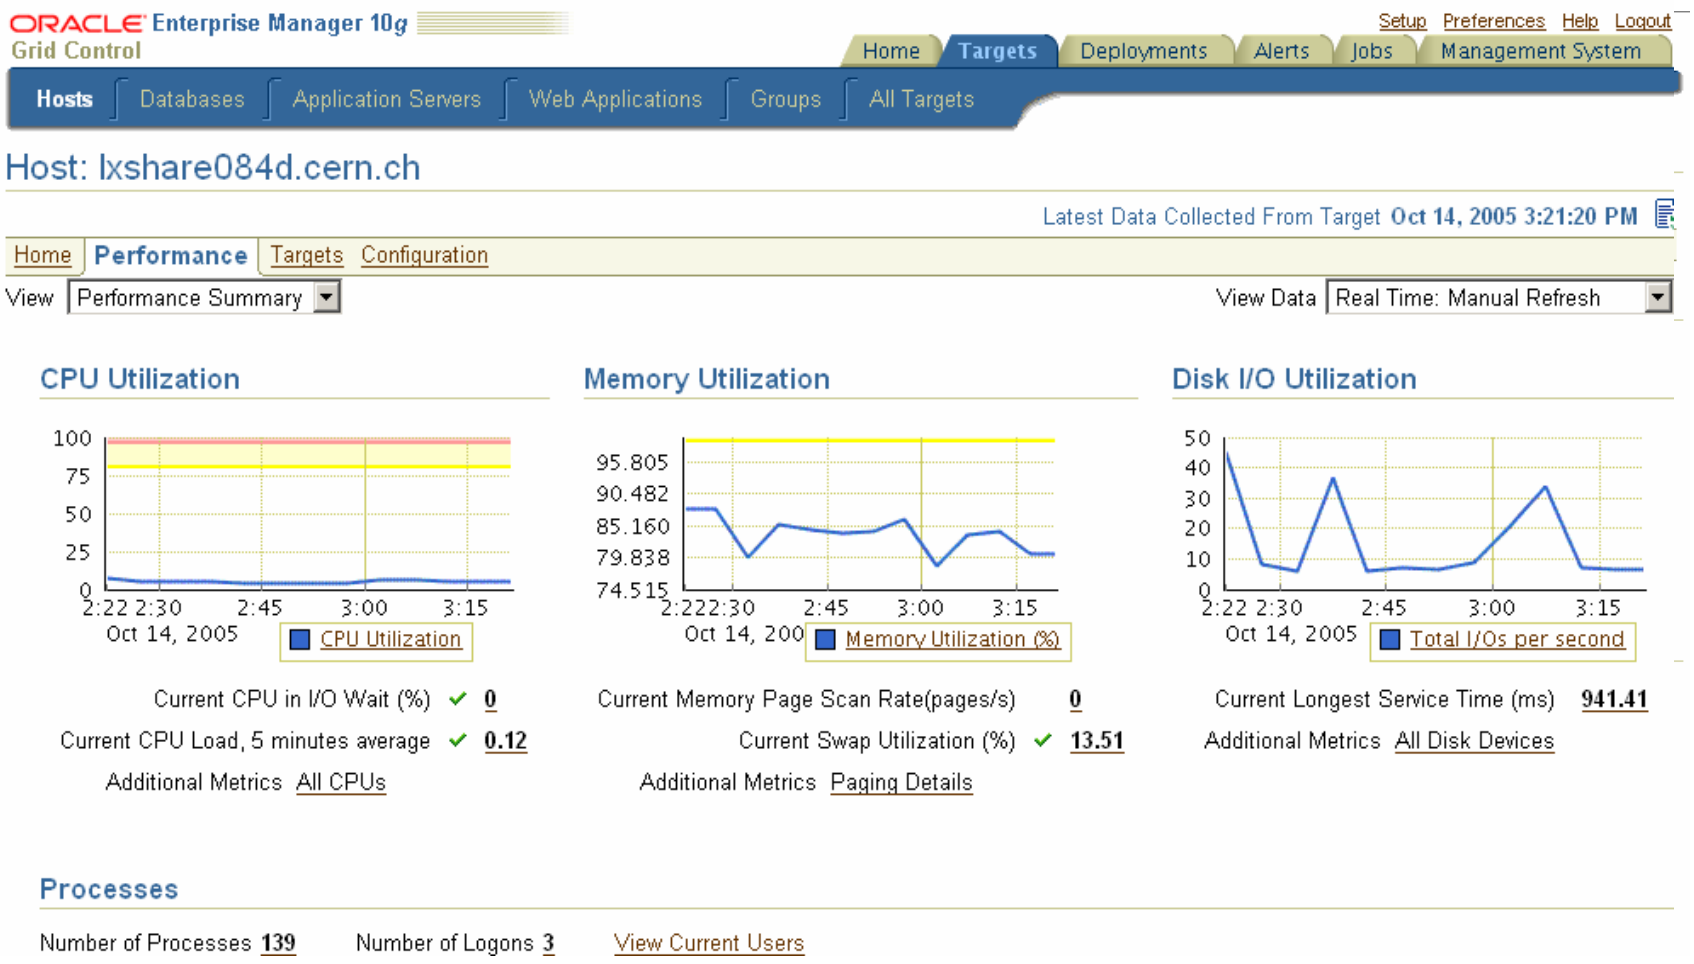
ORACLE Enterprise Manager 10g
Grid Control
Setup Preferences Help Logout
Home Targets Deployments Alerts Jobs Management System
Hosts Databases Application Servers Web Applications Groups All Targets
Host: lxshare084d.cern.ch
Latest Data Collected From Target Oct 14, 2005 3:21:20 PM
Home Performance Targets Configuration
View Performance Summary
View Data Real Time: Manual Refresh
CPU Utilization
100
75
50
25
0
2:22 2:30 2:45 3:00 3:15
Oct 14, 2005
CPU Utilization
Current CPU in I/O Wait (%) 0
Current CPU Load, 5 minutes average 0.12
Additional Metrics All CPUs
Memory Utilization
95.805
90.482
85.160
79.838
74.515
2:222:30 2:45 3:00 3:15
Oct 14, 200
Memory Utilization (%)
Current Memory Page Scan Rate(pages/s) 0
Current Swap Utilization (%) 13.51
Additional Metrics Paging Details
Disk I/O Utilization
50
40
30
20
10
0
2:22 2:30 2:45 3:00 3:15
Oct 14, 2005
Total I/Os per second
Current Longest Service Time (ms) 941.41
Additional Metrics All Disk Devices
Processes
Number of Processes 139
Number of Logons 3
View Current Users

FroNtier

- Frontier server available in the testbed
- High performance database access
- Standard web components and additional components
- Build an N-tier architecture for databases

Conclusions

- CERN end of distribution testbed is ready
- Both Streams and Frontier tests are now possible
- OEM based server monitoring and administration tools available
- Frontier client kit and squit needs to be put up for distribution

TESTS: POOL FileCatalog (I)

- written in Python; multi-thread application; producer-consumer pattern
- uses full **RLS application** stack of POOL **FileCatalog**

maintains consistent lists of accessible files (physical and logical names) together with their unique identifier (FileID)

resolves a logical file reference (FileID) to a physical file

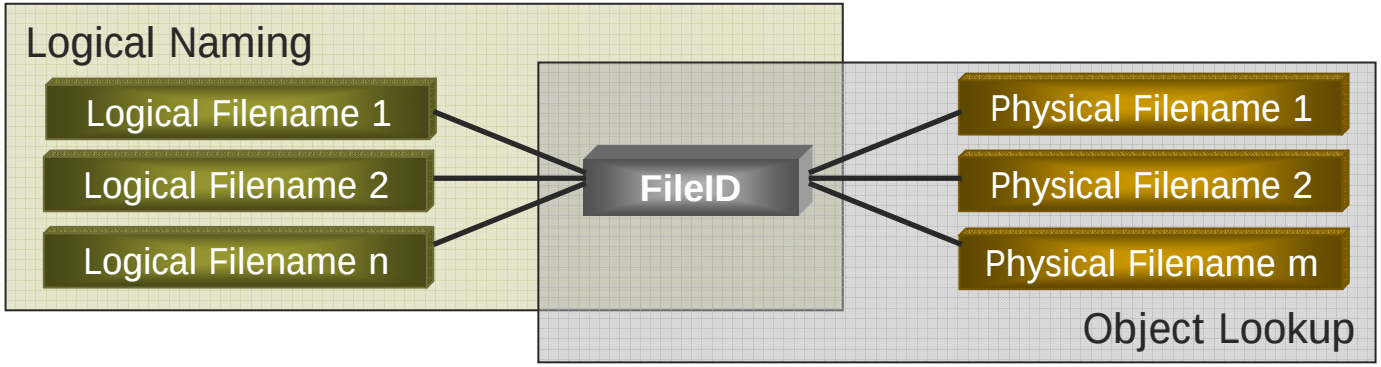

TESTS: POOL FileCatalog (II)

- CERN to CERN
  - Streams setup bi-directionally
  - Tests running during 5 weeks, up to 1.500.000 entries
- Changes to increase performance
- CERN to testbed sites
  - Streams setup un-directionally
  - Tests running during 3 weeks, up to 50.000 entries
- Tests
  - focus on stability and robustness
  - Successful results

TESTS: DDL operations

- Proposal to test DDL operations replication
  - Create and drop table, index, trigger, … , alter tables or index parameters, add columns, rename table, etc.
  - Test automated
- Conclusions
  - Public synonyms not replicated
  - ALTER object RENAME… statements replicated but
  - RENAME… statement not replicated
- Test replication of more than one schema
  - only CERN to CERN

TESTS: COOL and ATLASDD

- COOL is an API for reading and writing conditions data
- Oracle Geometry DB (ATLASDD) with data access through the new POOL/RAL interface
- Replica of COOL and ATLASDD accounts created on CERN testbed databases
- Streams replication setup
- Tests ongoing

Future TESTS

- ATLAS and LHCb condition database
- ATLAS detector description database
- CMS conditions via FroNtier and POOL
- ATLAS condition database
- ATLAS AMI meta data systems
- LCG LFC catalog replication
- ATLAS tag collection database
- ARDA/EGEE meta data catalog

Questions & Answers