Worldwide distribution of databases in the LHC Computing Grid (LCG)

Eva Dafonte Pérez

Overview

- Provide access to relational database data at CERN tier 0 and collaborating LCG tier sites
- Database replication via **Oracle Streams**
- Initial setup includes 6 tier sites
  - CERN as source database
  - 4 additional sites to be joined now
- Streams **log mining configuration** survey
  - Downstream capture
  - Split & Merge solution

Downstream Capture

- **Real-Time** downstream capture setup tests
  - using single database as source ✓
  - using RAC database as source: issue being solved
- Oracle Support
  - reviewed Downstream setup and logs
  - Streams setup problem discarded
  - Focus on the standby configuration
    - increase Redo log files size
    - stress tests ongoing

Streams Performance Tests

- Between CERN and T1 sites
- Script written in python, based on condition data, insert only workload
- Preliminary numbers: 10 - 100 MB/min reached
  - typical 30 MB/min
- WAN replication running at ~50% of LAN rates
- Sufficient for planned use with conditions data
  - Working with **Oracle Development** on rate improvement
- Experiments now taking over T1 setups for their Replication tests

Experiment and Grid Activities

- ATLAS
  - Online → Offline → T1 sites (GridKA, BNL)
  - Throughput tests ongoing:
    - replication rate: 16 MB/min
    - performance problems on BNL
- LHCb
  - Offline → T1 sites (RAL, GridKA, IN2P3)
  - Online → Offline: preparation
- Grid File Catalog (LFC)
  - CERN → T1 site (CNAF)
  - sustained rate achieved: 33 replicas per second

Backup and Recovery Tests

- Objective
  - simulate real scenarios of failure
  - perform tests to gain experience and document Streams synchronization steps
- Scenarios and tests
  - Point-in-time recovery on the destination database ✓
  - Point-in-time recovery on the source database
    - 'before' latest data sent to destinations
    - destinations beyond source ✓
  - Point-in-time recovery on the source and destination databases
  - Tablespace point-in-time recovery ✓

Streams Monitoring

- Previous: Status scripts, Streams Monitoring Tool (STRMMON) and OEM
- Problems:
  - limited access
  - impossible to monitor entire streams setup
  - no central repository for streams setup status
- Working on an extended tool for streams monitoring (together with a CERN technical student)
  - Daemon Script written in python, getting streams activity from database, archiving logs to the files repository
  - End user application – available and still extending
- Feedback to OEM development

Future Work

- Completion of downstream capture setup for October production
- Integration of 4 additional sites (reaching the final 10 site setup)
  - Completing Experiments setups
- Completing streams monitoring

Programme's Feedback

The feedback is circulated between the people involved.

Oracle EMEA

Monica Marinucci Lopez	Management of the programme
June Farmer	Management of the programme
Graeme Kerr	Technical liaison

Oracle Development

Patricia McElroy	Principal Product Manager Distributed Systems/Replication

CERN Openlab

Sverre Jarp	Chief Technologist Officer
Jürgen Knobloch	IT-PSS Group leader
Dirk Düllmann	IT-PSS-DP Section Leader