Oracle technical update for CERN physics services

CERN openlab II quarterly review
31 January 2007

Dawid Wojcik

Quarterly review

- Database recovery tests in the streams environment
- Streams optimizations
- Streams monitoring
- Physics services expansion
- Storage configuration
- Data Guard
- RAC monitoring
- Quarterly Summary
- Future Plans (PoW 2007)

Database recovery in the streams environment

Eva Dafonte Pérez

- Objectives
  - simulate real scenarios of failure at one or more sites
  - perform recoveries on source or destination database
  - document all tests results, report feedback, resolve issues

- Streams test environment

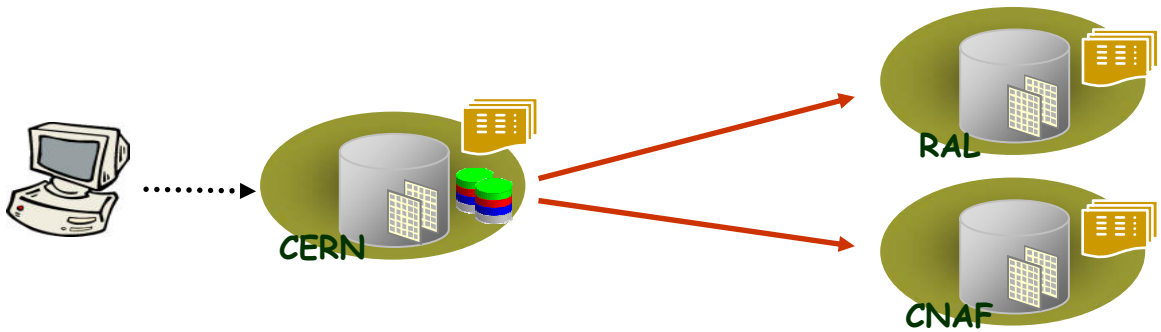

Recovery scenarios

- Point-in-time recovery on destination database

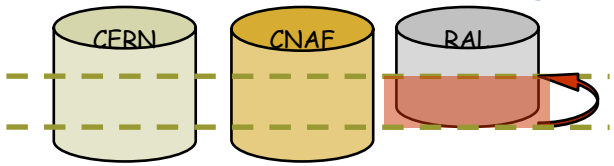

- Point-in-time recovery on source database
  - Latest data on source was not applied on destination

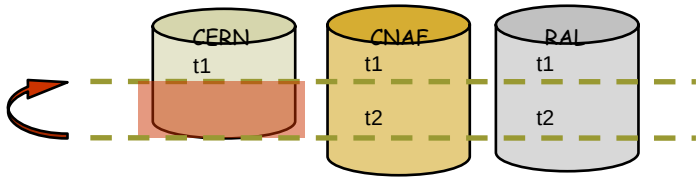

  - Apply process has applied some transactions from the source database **after point-in-time recovery**

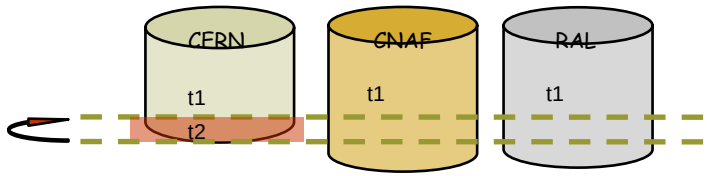

- Recovery on source and/or destination
  - Tablespace point-in-time recovery

Recovery scenarios - results

- All scenarios have been tested ✓
  - Oracle Recovery Manager (RMAN) does the job well
  - More tests to come on production streams setup
- Some problems encountered:
  - Issues resolved ✓
  - Bug fixes done ✓

Streams throughput optimizations

- Streams performance

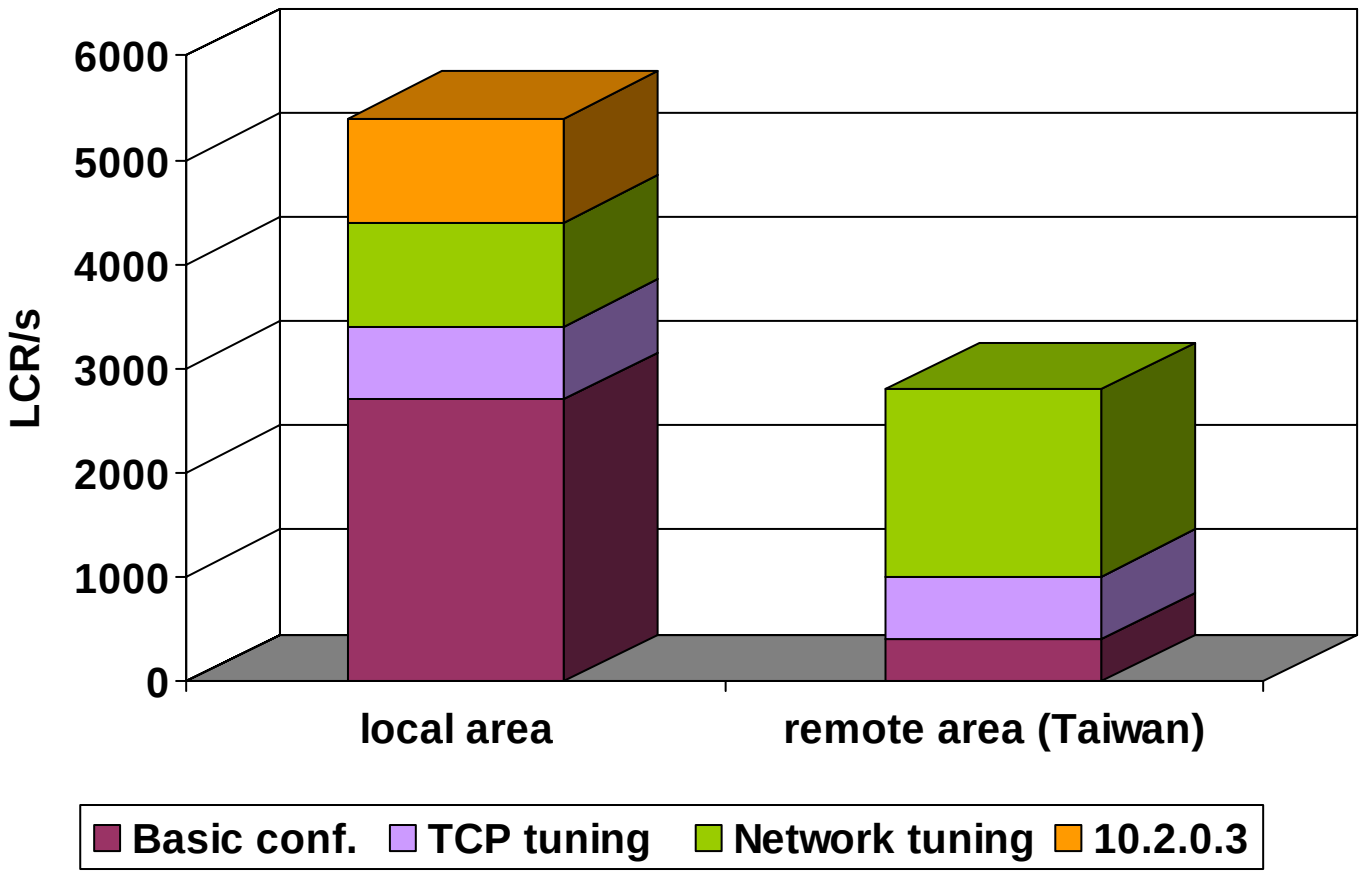

Streams monitoring

Zbigniew Baranowski

- Requested features:
  - Visualization of streams topology
  - Status of streams connections
  - Error notifications
  - Monitoring of streams performance (latency, throughput, etc.)
  - Monitoring of resources related to the streams performance (streams pool memory, redo generation)
- Architecture:
  - ‘strmmon’ daemon written in Phython
    - collects streams and instances info
    - sends errors and warnings
  - End-user web application

    http://oms3d.cern.ch:4889/streams/main.php

Streams monitoring

Monitor view

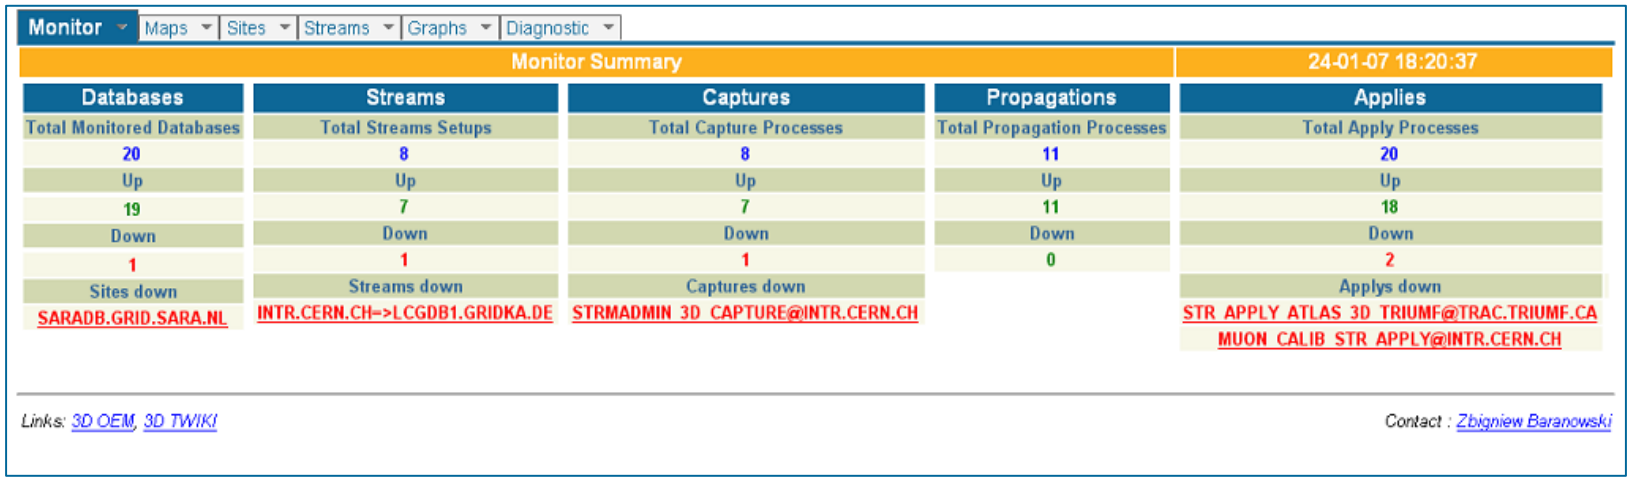

Connection view

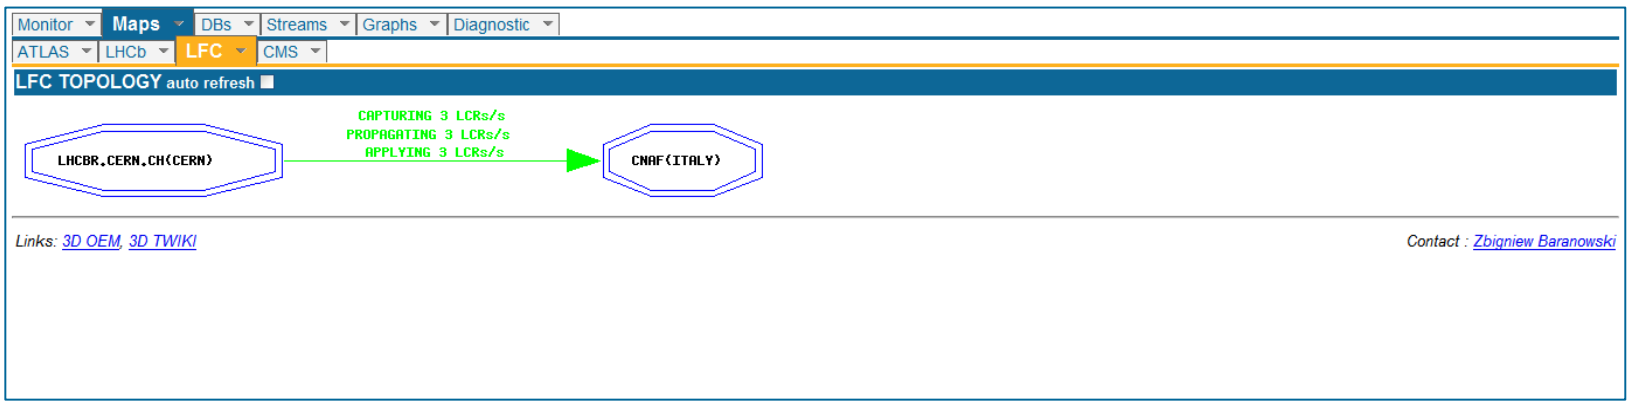

Physics services expansion

- Hardware and software upgrade of physics services at CERN
  - hardware resources has been doubled
  - move from RHEL3 to RHEL4 (32 bit)
- Good opportunity to review
  - Installation procedures & scripts
    - storage configuration
    - network configuration
  - HA solutions
    - online RAC migration using Oracle Data Guard (described later)

Storage configuration

- Oracle ORION tests with two different multipathing solutions
  - RHEL3 – Qlogic
  - RHEL4 – device mapper
  - Equal performance observed, no problems with multipathing failover
- Storage configuration scripts prepared for fast ASM deployment with multipathing failover
  - storage hardware configuration scripts
  - device mapper configuration scripts

Data Guard

- Database migration using Oracle Data Guard (physical standby)
  - Prepared and tested necessary procedures
  - Carried out in production environment
  - Overall downtime minimized to 20 minutes
  - Two RAC databases migrated so far (more to come)
- Future plans
  - More Data Guard tests
  - Oracle Enterprise Manager (OEM) Data Guard tests

RAC monitoring

- New features implemented
  - Monitoring of production services
    - all production clusters' services
  - Monitoring of Oracle Clusterware (CRS)
    - all clusterware services for deployed RACs
  - Unavailability tracing
    - instances unavailability
    - clusterware services unavailability
    - availability plots for specified periods

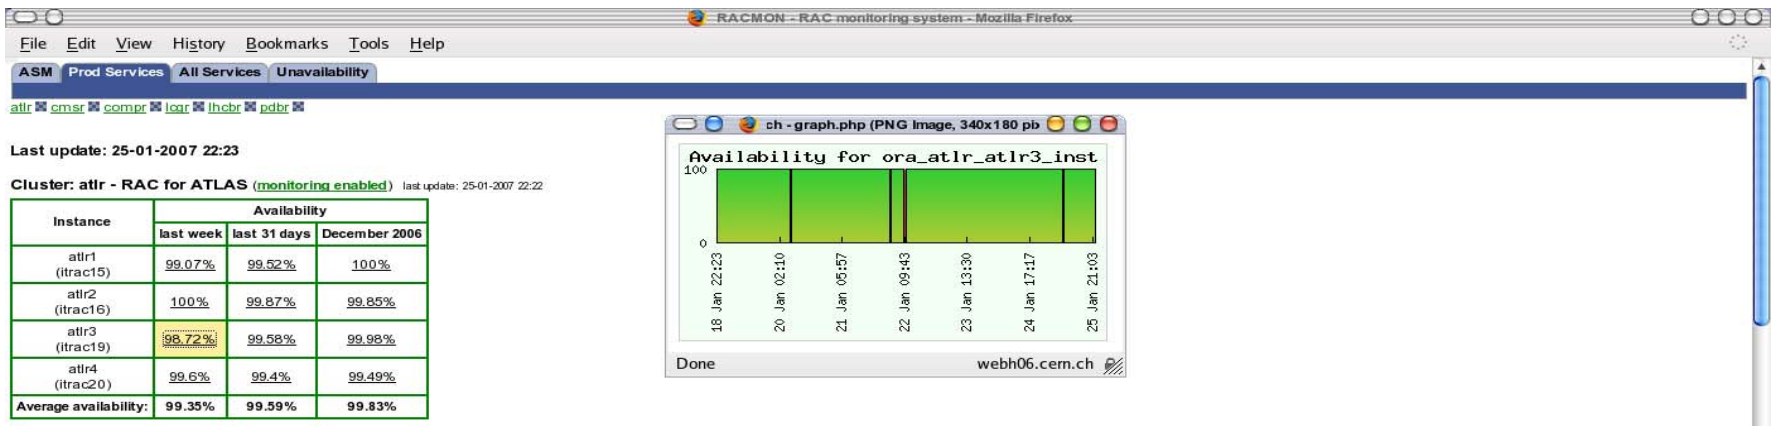

Quarterly Summary

- **Oracle Streams and Data Replication Services**
  - Single Point of Contact: E. Dafonte Pérez (CERN) – G. Kerr (Oracle)
  - Participants: M. Bogusz
- **Streams monitoring – feedback to OEM development**
  - Participants: Z. Baranowski
- **Highly available database services based on RAC/ASM**
  - Procedures review, RAC monitoring, Data Guard tests
  - Single Point of Contact: D. Wojcik (CERN) – G. Kerr (Oracle)

Future Plans (PoW 2007)

- **Oracle Streams and Data Replication Services**
  - Single Point of Contact: E. Dafonte Perez (CERN) – G. Kerr (Oracle)
  - Participants: D. Duellman (CERN) – P.McElroy, A. Downing (Oracle)
- **Oracle Enterprise Manager**
  - Single Point of Contact: C. Lambert (CERN) – A. Bulloch (Oracle)
  - Participants: D. Wojcik, M. Kierebinski, Z. Baranowski, A. Wiecek (CERN) – G. Kerr (Oracle)
- **Oracle Data Guard**
  - Single Point of Contact: A. Topurov (CERN) – G. Kerr (Oracle)
  - Participants: E. Dafonte, M. Girone, E. Grancher, D. Wojcik (CERN)
- **Highly available database services based on RAC/ASM**
  - Single Point of Contact: D. Wojcik (CERN) – G. Kerr (Oracle)
  - Participants: M. Girone (CERN)