

# CORAL and COOL news

**Andrea Valassi (IT-ES)**

On behalf of the Persistency Framework team

*Distributed Database Operations Workshop,  
16<sup>th</sup> November 2010*

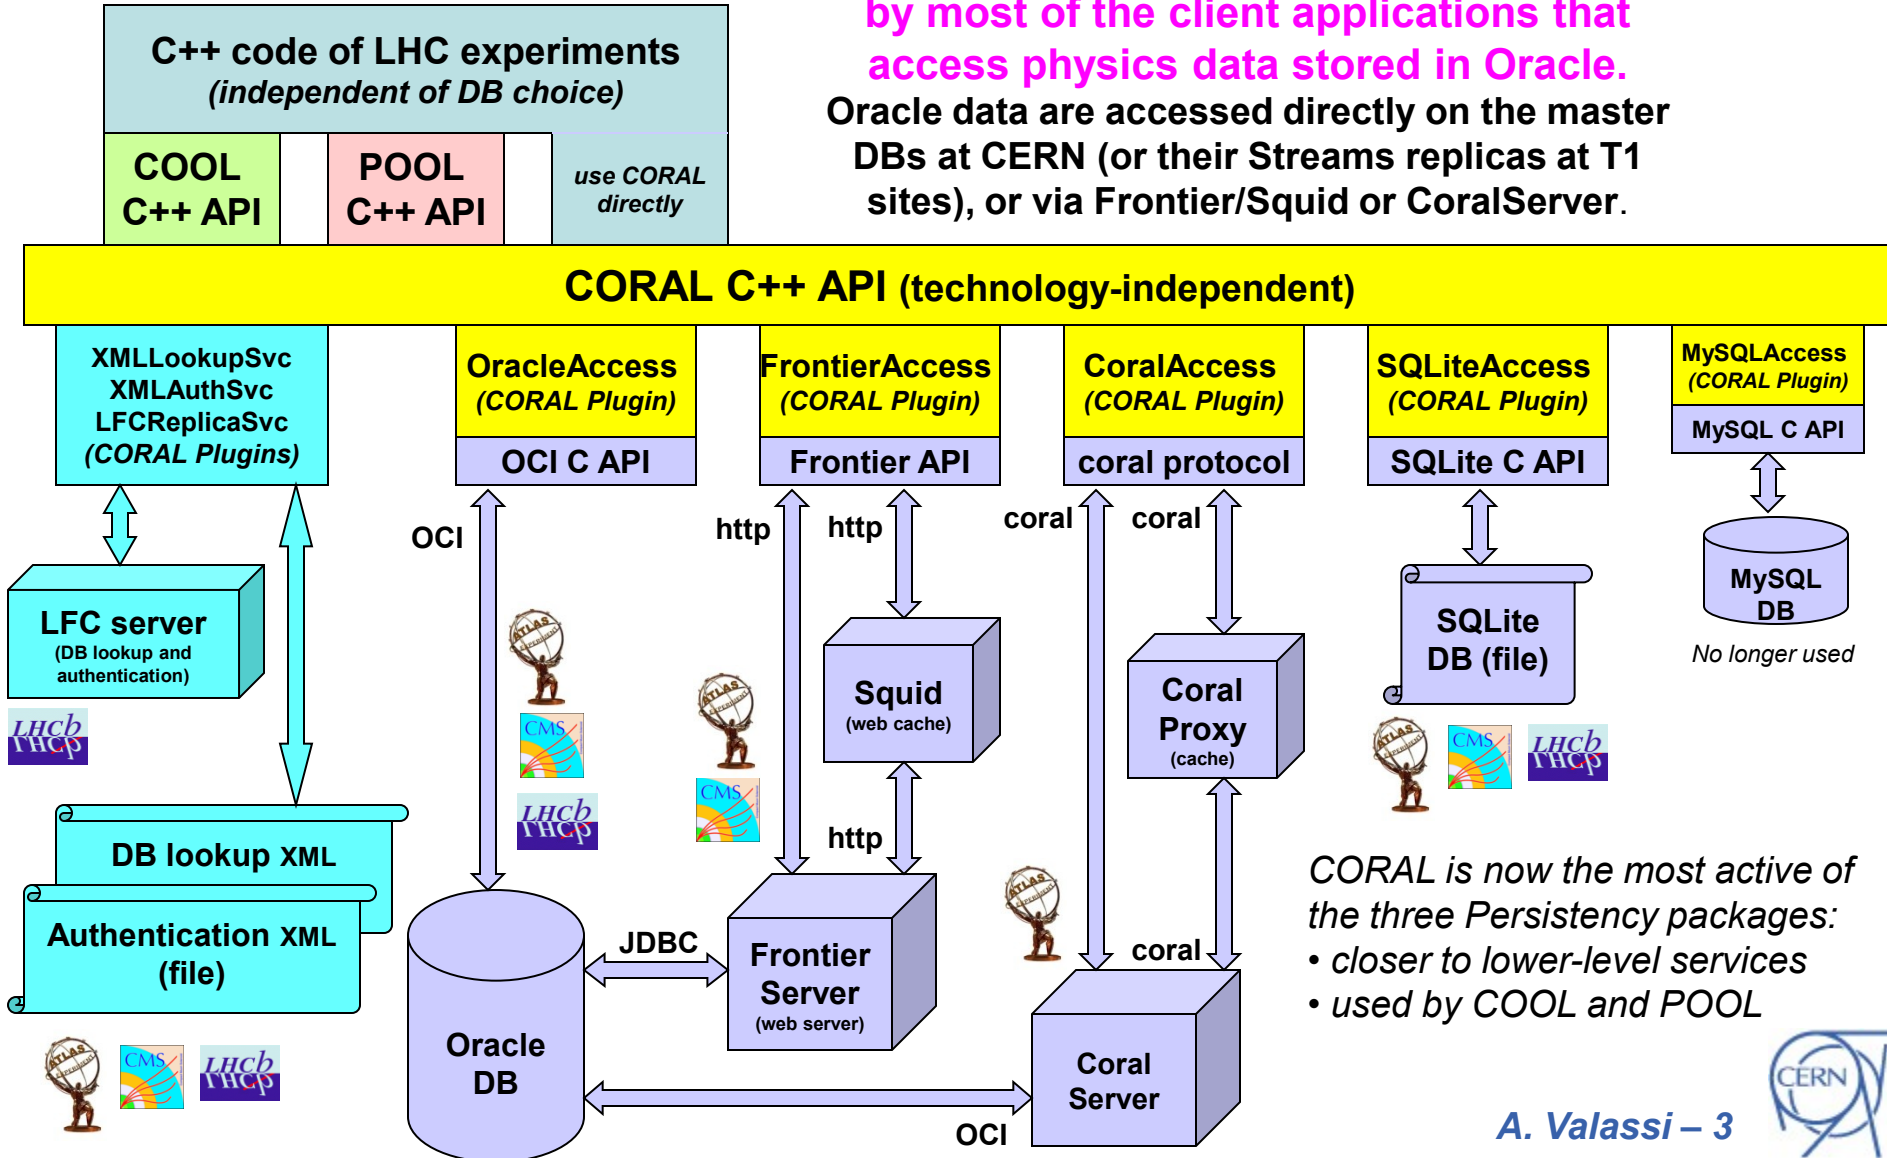
- **Overview of CORAL architecture and usage**
- **Recent activities and work in progress**
- **A few details/comments about selected subjects**
  - ORA-07445 server crashes
  - Oracle client for CORAL
  - Network glitch handling
  - FroNTier
  - CoralServer for ATLAS High Level Trigger



# Overview of CORAL architecture & usage

**CORAL is used in ATLAS, CMS and LHCb by most of the client applications that access physics data stored in Oracle.**

**Oracle data are accessed directly on the master DBs at CERN (or their Streams replicas at T1 sites), or via Frontier/Squid or CoralServer.**



*CORAL is now the most active of the three Persistency packages:*

- closer to lower-level services
- used by COOL and POOL

## Recent news/achievements (last 12 months):

- **High load from service and software support**
  - Oracle: server crashes (ORA-07445), 11g server (Desy), client-side issues (selinux/SLC5, AMD quadcore, light client)
  - Integration with Grid software: gssapi version in Globus, libexpat version in Frontier, fixes in CORAL LFCReplicaSvc
  - Releases for ATLAS/LHCb, tags for CMS
- **CORAL server used by ATLAS HLT in data-taking**
  - New features in progress (slower due to other priorities)
- **Frontier/Squid adopted for ATLAS offline analysis**

## In progress:

- **Reconnection after network or database glitches**
  - Several problems reported, highest priority at the moment
- **Redesign and extend coverage of CORAL tests**



## Recent news/achievements (last 12 months):

- **New relational schema options**
  - Can store payload in a separate table from IOV metadata
  - Can associate more than one payload record to each IOV
- **Faster retrieval of CLOB values**
- **Fixes in the handling of NULL and NaN values**

## In progress:

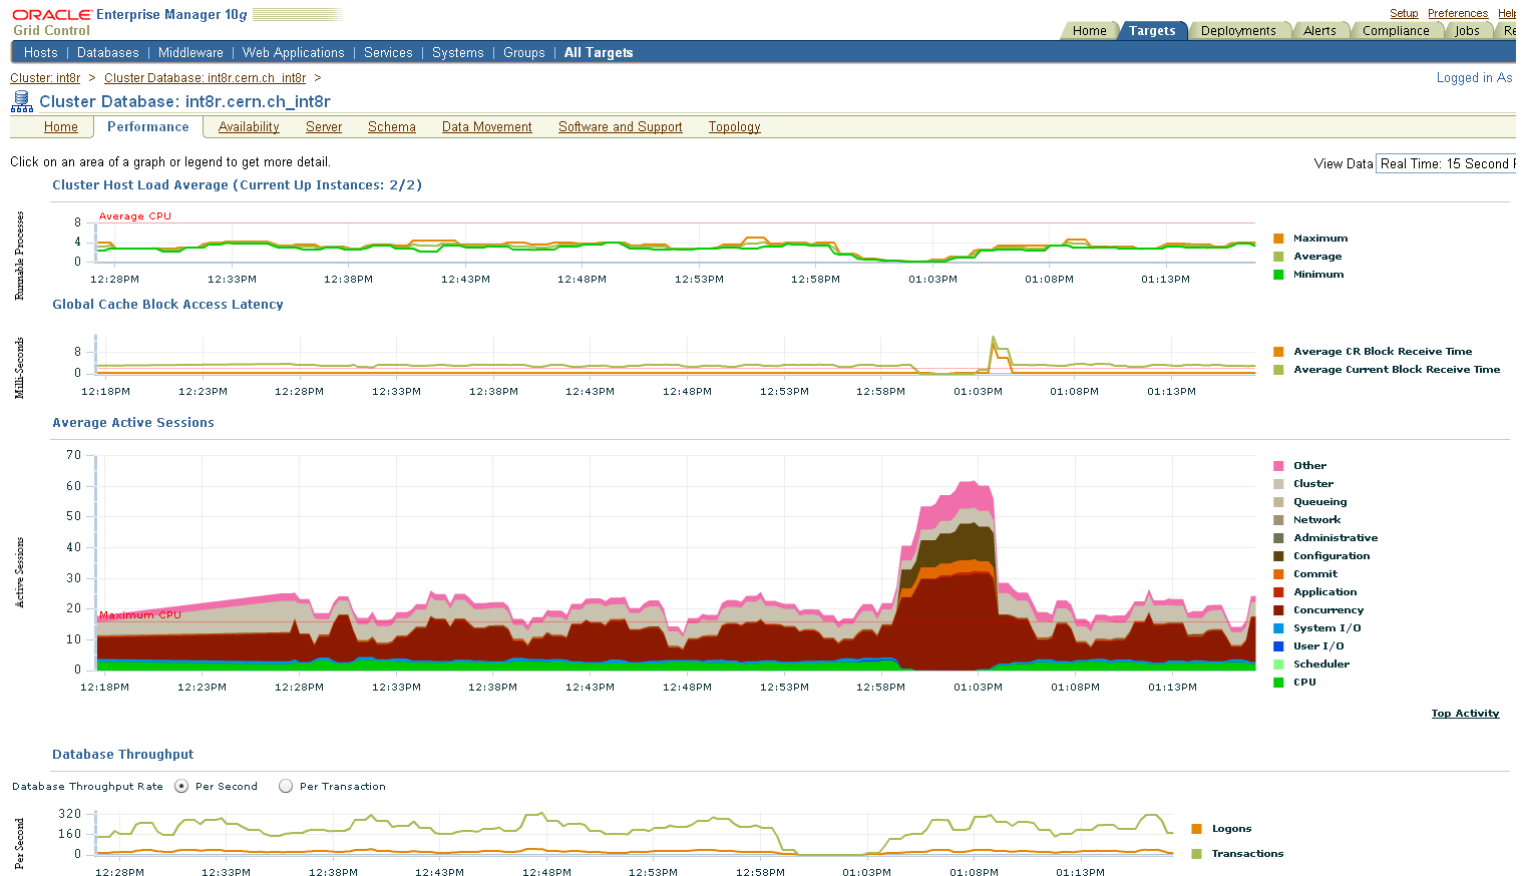
- **Internal cleanup of transactions**
  - Prepare the way for session sharing in COOL



- **Signature: ORA-07445 server process crash**
  - “core dump [ksxpmprrp()+267] [SIGSEGV]” in trace files
  - Effect on end users: high load spikes, DB server ‘freezes’
- **Observed in COOL (ATLAS & LHCb) after April PSU in June**
  - Rolled back the April PSU until problem better understood
  - *Tried to reproduce it on test DB to validate possible patches*
- **Used COOL nightly tests to try and reproduce it**
  - June PSU was not rolled back on test1 (lcg\_cool\_nightly)
  - “ORA-07445 [ksxpmprrp()+267]” seen 15x in 2 months (June-July)
    - Always associated to “update sys.aud\$” on disconnecting
    - Always the same test “test\_RelationalCool\_RelationalFolder”
    - Client apps succeed: crash is on server side when disconnecting
  - Developed a script to maximize chances of reproducing the issue
    - Successfully managed to reproduce ORA-07445 (~25 per hour)
    - Also observed freezing of DB server (and of client at the same time)



- Observed freezing of database server during COOL tests



- Observed freezing of client CPU at same time...

- **Which applications are hit by ORA-07445?**
  - COOL applications with connection sharing
    - Clearly demonstrated using the COOL based test script
    - No errors observed if connection sharing is disabled
  - Issue was also observed in applications other than COOL
    - e.g. on int8r: ATLAS tags (CORAL-based POOL collections)
- **Validation of patches on int8r (thanks to Marcin!)**
  - **No April PSU: OK...** (no ORA-07445)
  - **April PSU: NOT OK** (~90 ORA-07445 on 900 COOL tests)
  - **April + July PSU: NOT OK** (ORA-07445 still there)
  - **April + July PSU + 6196748 patch: OK!** (no ORA-07445 on 300 tests)
- **COOL test script can now be used by PDB team**
  - Further tests of ORA-07445 and similar issues
  - Generate some COOL load with/without connection sharing
  - Warning: not fully representative of production-like activities
    - Much more DDL (create/drop tables) than in production!



- **Bugs relevant to CORAL fixed in 11.2.0.1.0p2 (May 2010)**
  - Crash on AMD Opteron quadcore ([bug #62194](#))
    - ATLAS had downgraded back to 10g in April 2010 because of this bug, then moved back to the 11g client when it was fixed
    - Patch tested in Ljubljana – thanks to many site admins for their help!
  - SELinux issues on SLC5 ([bug #45238](#))
    - Now fixed for all four combinations 32/64bit OCI/OCCI
- **No showstopper in 10g, but 11g client better for the future**
  - Several bugs (SELinux and others) are only fixed in 11g client
  - Light instant client (smaller memory) available on both 10g/11g
  - Servers will eventually move to Oracle 11g...
- **One new bug recently reported to Oracle Support**
  - Oracle client uses its own version of gssapi ([bug #71416](#))
    - No problem observed in production, but a potential hazard
    - Similar to the problem reported by LHCb in Globus ([bug #70641](#))



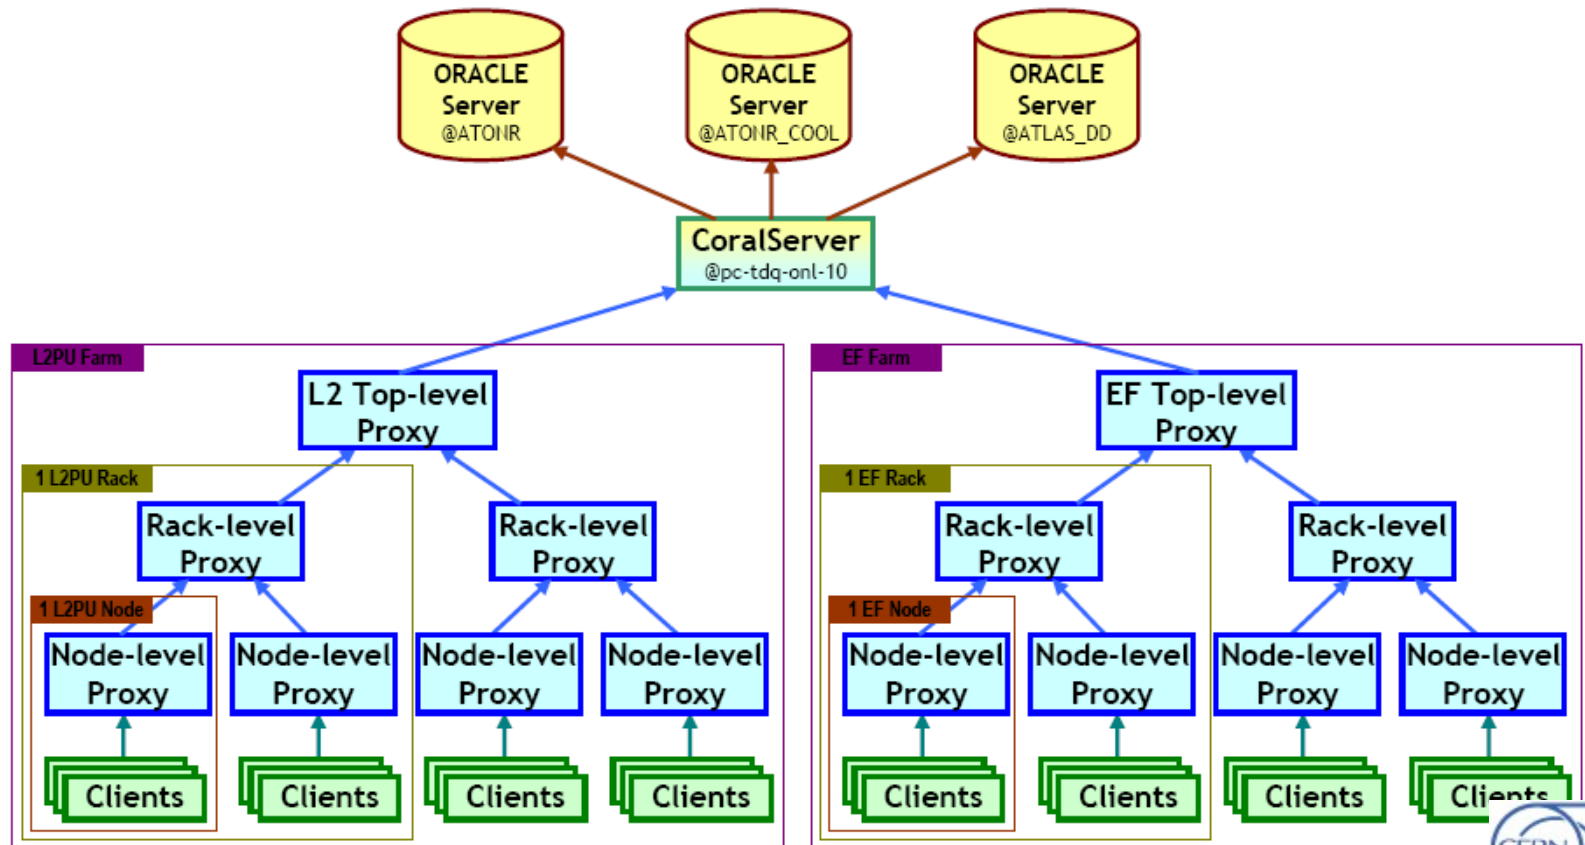
- **Reported by all experiments in various cases**
  - “A transaction is not active” in CORAL server ([bug #65597](#))
  - ORA-24327 “need explicit attach” in ATLAS/CMS ([bug #24327](#))
  - OracleAccess crash after losing session in LHCb ([bug #73334](#))
- **What should CORAL do? Many different scenarios**
  - e.g. non serializable R/O transaction: should reconnect and restart it
  - e.g. DDL not committed in update transaction: cannot do anything
- **What is CORAL doing now?**
  - Correctly reconnecting in some cases (*existing useful features*)
  - Not doing anything in other cases (*missing useful features*)
  - Reconnecting in the wrong way in other cases (*bugs!*)
- **Work in progress** (A. Kalkhof, R. Trentadue, A.V.)
  - Catalog the different scenarios and prepare tests for each of them
  - Prototype the implementation changes (ConnectionSvc and plugins)



- **Frontier has been used in CMS for many years**
  - Integrated into CORAL (then “RAL”) in 2005
    - Originally developed and used for CDF at FNAL
- **Adopted by ATLAS for offline analysis in 2009**
  - Smooth adoption from the software point of view
    - *Frontier supported and tested in COOL since 2006 (just in case)*
      - Even if it was never used in production before 2009
    - Only minor fixes/optimizations (e.g. [bug #66638](#), [bug #66638](#))
      - Recent patches for ATLASDD too (just in case, again!)
  - Deployment discussions ongoing at the [T1SCM](#)
    - Decision by ATLAS to use own Squids (not shared with CMS)
- **Recent interest also by LHCb**
  - Discussions at CHEP in Taiwan
    - May need COOL patch to retrieve metadata without payload
- **A few performance tests ongoing**
  - e.g. comparison with OracleAccess and CoralAccess



- **CoralServer deployed for HLT in October 2009**
  - Smooth integration, used for LHC data taking since then
  - No problems except for handling of network glitches
  - New developments in progress at low priority (monitoring...)



- **ATLAS, CMS and LHCb access Oracle via CORAL**
  - For conditions data (e.g. COOL in ATLAS/LHCb) and more
- **Support for many backends allows many options**
  - ATLAS switch to Frontier was painless from a software POV
- **Highest load is from software and service support**
  - Releases, Oracle client issues, Oracle server debugging
  - A few new developments too, but with lower priority

***Thanks to the PDB team and the experiment and site DBA's  
for their help and very good collaboration!***



# Reserve slides

- **Based on COOL nightlies, developed a test script to maximize chances of reproducing ORA-07445**
  - See <https://savannah.cern.ch/task/?16836>
  - The script ora07445.csh executes internally several cycles of the test\_RelationalCool\_RelationalFolder executable
    - <http://cool.cvs.cern.ch/cgi-bin/cool.cgi/cool/contrib/ExternalTests/OracleConnectionSharing/ora07445.csh?rev=1.5&content-type=text/vnd.viewcvs-markup>
  - Typically I have run several scripts in parallel
    - <http://cool.cvs.cern.ch/cgi-bin/cool.cgi/cool/contrib/ExternalTests/OracleConnectionSharing/allOr07445.csh?rev=1.4&content-type=text/vnd.viewcvs-markup>
  - Using 30 scripts with 30 cycles each (~4 hours) I observed ~90 occurrences of ORA-07445 in a **reproducible** way!

