LCG Persistency Framework

CORAL, POOL, COOL – Status and Outlook

A. Valassi, **R. Basset**, **R. Chytracek**, **D. Duellmann**, **A. Kalkhof**, **I. Papadopoulos**, **W. Pokorski**, **Z. Molnar** ***(CERN IT-DM)***

M. Nowak ***(BNL / Atlas)***
S. A. Schmidt, **M. Wache** ***(University of Mainz / Atlas)***

D. Dykstra ***(FNAL / CMS)***
G. Govi ***(Northeastern University / CMS)***
Z. Xie ***(Princeton University / CMS)***

M. Clemencic, **M. Frank** ***(CERN PH / LHCb)***

CHEP2009 (Prague), 23rd March 2009

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Outline

- **Introduction**
  - PF components
  - Usage in the experiments
- **Recent achievements**
- **Work in progress and outlook**

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF components

CERN IT Department

- **CORAL**
  - Abstraction of access to relational databases
    - Support for Oracle, MySQL, SQLite, FroNtier
  - Used directly or indirectly via COOL/POOL
- **POOL**
  - Technologically-neutral hybrid data storage
    - Streaming of objects (e.g. to ROOT or relational DBs)
    - Object metadata catalogs (e.g. in XML or relational DBs)
- **COOL**
  - Conditions data management
    - Conditions object metadata (interval of validity, version)
    - Conditions object data payload (user-defined attributes)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF architecture

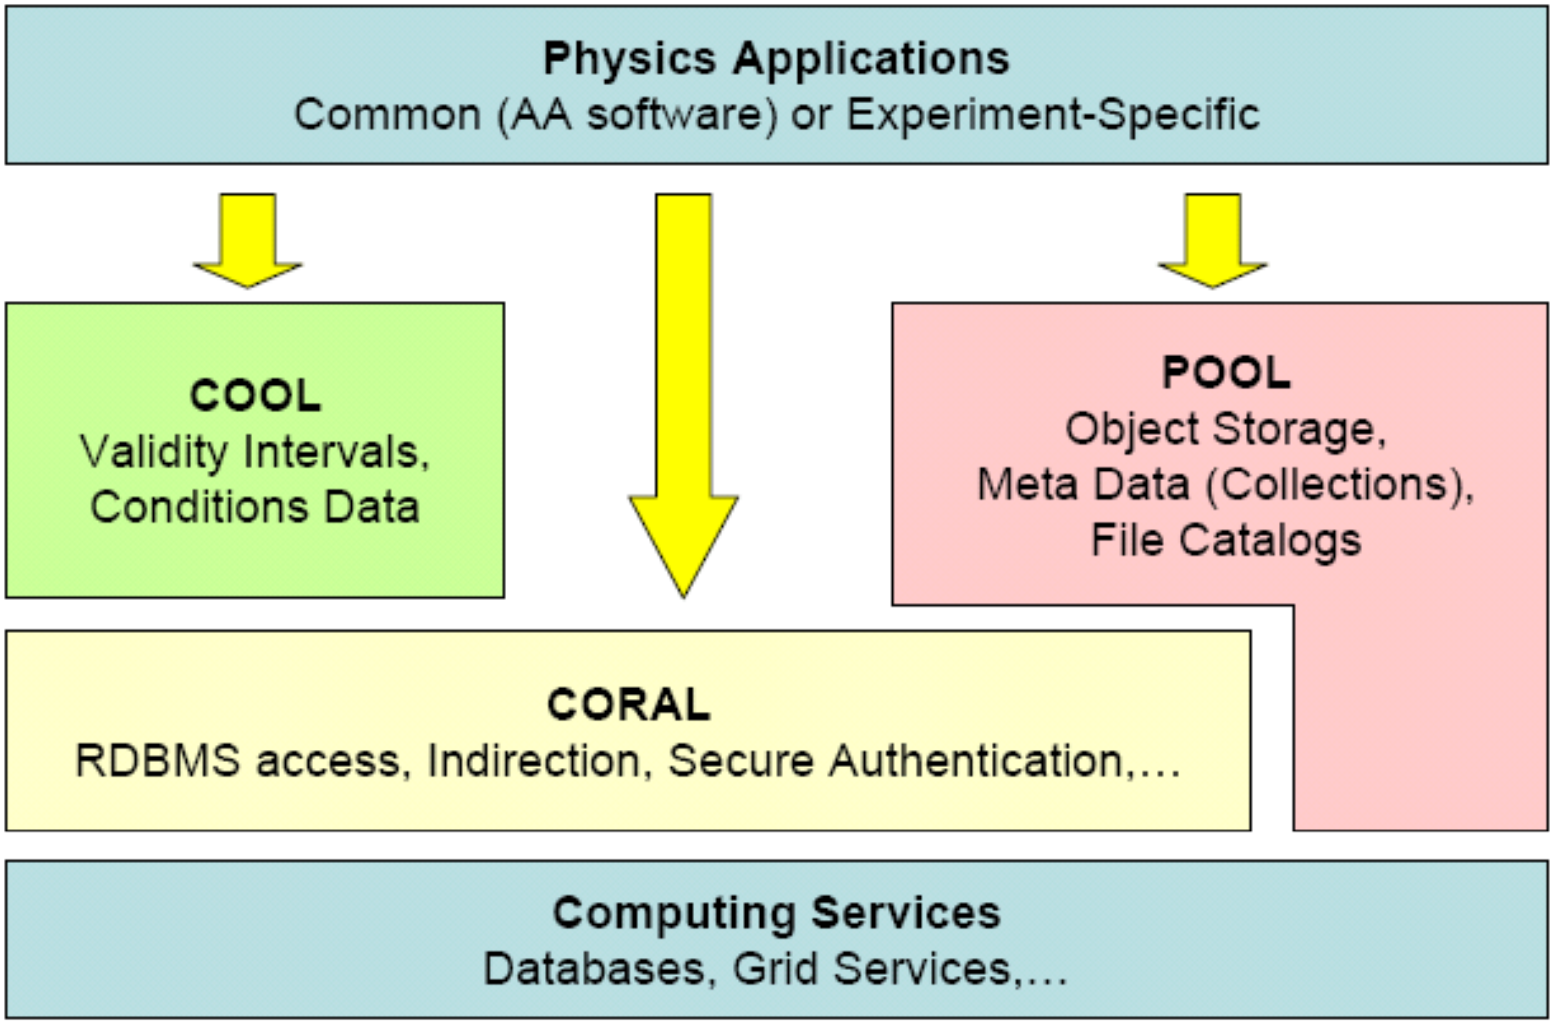

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF usage in the experiments

CERN IT Department

- **ATLAS**
  - POOL: event data (ROOT), event collections
  - CORAL: geometry, trigger (Oracle, SQLite)
    - Migration from MySQL to Oracle has been completed
  - COOL: conditions data (Oracle, SQLite)
    - Evaluating Frontier access (talk by R. Walker)
- **CMS**
  - POOL: conditions data (Oracle/FroNtier, SQLite)
  - CORAL: online conditions data (Oracle)
- **LHCb**
  - POOL: event data (ROOT with XML catalogs)
  - COOL: conditions data (Oracle, SQLite)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF and AA

CERN IT Department

- **Interaction with other projects**
  - SPI (external software, CMT config, nightly builds)
  - ROOT (object streaming in POOL, PyCool)
  - GRID middleware (for CORAL authentication)
  - 3D (relational data deployment and distribution)
- **Release plan agreed with the experiments**
  - Bi-weekly Architect Forum meetings
    - Including PF and all other LCG AA projects
    - Following bi-weekly meetings with SPI and librarians
  - One or two major releases per year
    - Releases are now built by SPI team (talk by S. Roiser)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF and AA – new in 2008

CERN IT Department

- **Dropped SEAL dependency**
  - Main motivation for the LCG55 release (June 2008)
  - Main functionalities relevant to PF were moved to CORAL
    - Simplified component model and runtime loading of plugins
    - Other functionalities are taken from ROOT or Boost
- **Several new platforms and external s/w upgrades**
  - In total, 20 platforms are now supported (only 7 previously)!
    - Add gcc43 on Linux, add SLC5 (32/64 bit, gcc34/gcc43)
    - Add VC9 (Visual Studio Express 2008 – free) on Windows
    - Move from OSX 10.4 to10.5, add Oracle support on OSX
  - Major upgrades of external software versions
    - Upgrade from ROOT 5.18 to 5.22 (talk by F. Rademakers)
    - Other major upgrades: Boost, CMT...
  - Most changes were in the LCG56 release (January 2009)
    - LCG56 will be the baseline for LHC data taking in 2009

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF – activities in 2008

CERN IT Department

- **CORAL**
  - Start of developments for the CORAL server
  - FroNtier enhancements for CMS by the FroNtier team
- **COOL**
  - Performance optimizations and scalability tests
  - New features: payload queries, partial tag locks, tag cloning
- **POOL**
  - Largely in maintenance mode since summer 2008
    - Code changes largely driven by new ROOT releases
  - Review of component usage by experiments
    - Most components are used and maintained by one experiment
- **Several changes in the team**
  - Many departures (especially in CORAL), some arrivals
  - New PF leader (A.V. replacing D. Duellmann)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PF plans for 2009

CERN IT Department

- **CORAL**
  - Most active PF sub-project: CORAL server development
    - See details in the upcoming talk in this session
  - Improve coverage and configuration of CORAL test suite
    - Focus on thread-safety (CORAL server is multi-threaded)
- **COOL**
  - Connectivity features (session sharing, transaction control)
  - Partitioning of relational schema
- **POOL**
  - Review and improve DataSvc (used by CMS)
- **All PF projects**
  - Move code repositories from CVS to SVN

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CORAL server (upcoming talk in this session)

CERN IT Department

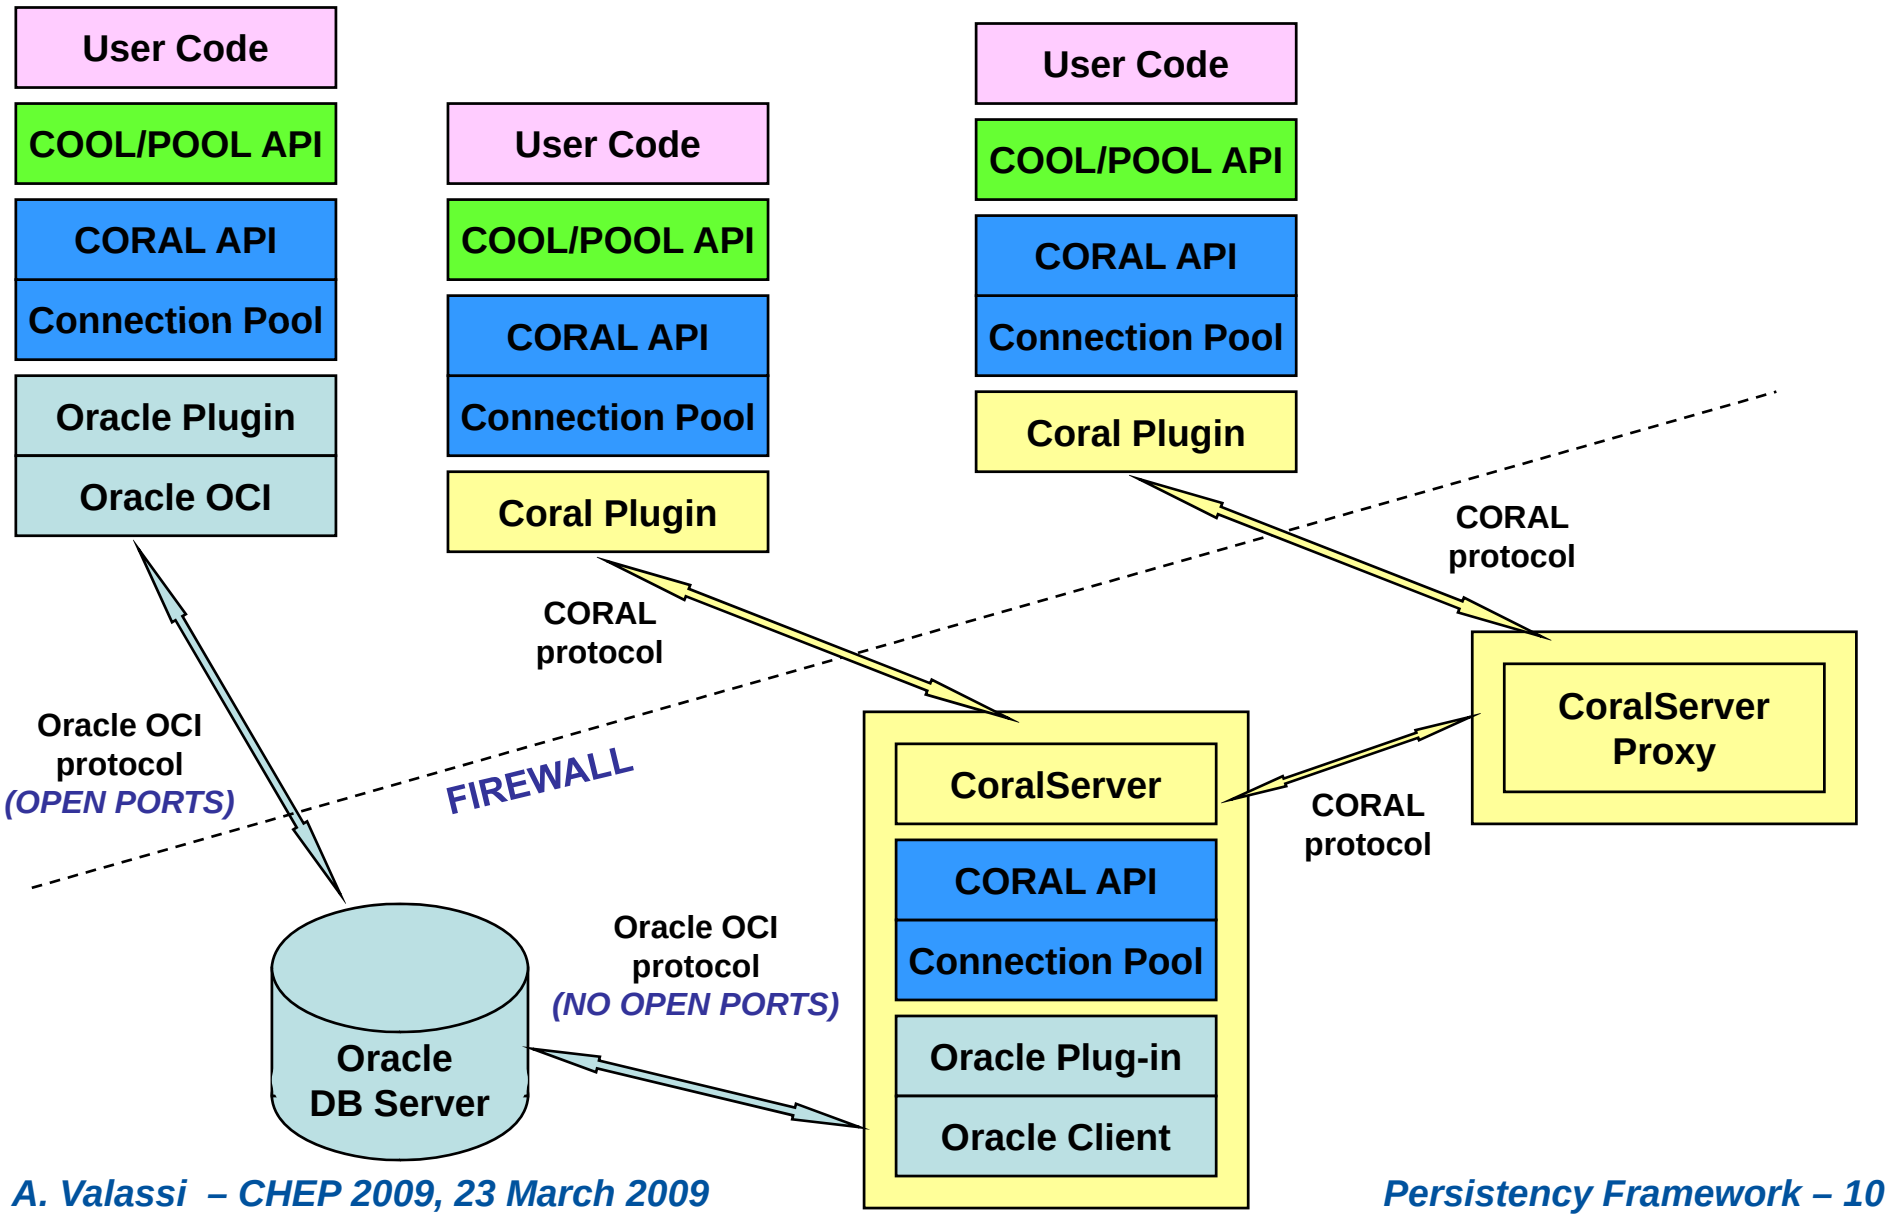

COOL

CERN IT Department

- **Performance optimizations and scalability tests**
  - Completed in 2008 (poster in today's session)
    - Factor out common C++ code for main queries (SV, MV…)
    - Better C++ API – no COUNT(*) queries unless demanded
    - Optimize SQL queries, stabilize Oracle execution plans
    - Scalability tests (insertion/retrieval performance on large tables)
- **Partitioning (in progress)**
  - Atlas expects 300 GB/year COOL data from DCS
  - Main motivation for partitioning: data management issues
    - Data archiving, backup, replication to T1 sites
      - General issues – relevant also for CMS non-COOL conditions
    - COOL performance for data retrieval is OK on large tables
- **Connectivity features (in progress)**
  - Share CORAL session for COOL databases in same schema
  - User control over database transactions

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

POOL

CERN IT Department

- **Review of component usage in the experiments**
  - Completed in the summer 2008 (W. Pokorski)
  - Code repository cleaned up (new CVS)
- **Enhanced collection packages (ATLAS)**
  - Improved performance (Oracle), new features (ROOT)
  - Packages used only by ATLAS and maintained by ATLAS
- **Enhanced RelationalStorageSvc (CMS)**
  - Fixes for schema evolution, schema mapping improvements
  - Package used only by CMS and maintained by CMS
- **DataSvc enhancements for CMS (in progress)**
  - DataSvc used only by CMS, but not adapted to their needs
  - Exact nature of the changes under discussion

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Conclusions

CERN IT Department

- **PF will be essential for LHC data taking**
  - Event data and collections (ATLAS, LHCb)
  - Conditions data (ATLAS, CMS, LHCb)
- **Development plans for 2009**
  - CORAL server development
  - COOL data partitioning
  - POOL is mostly in maintenance mode

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it