

Persistency framework summary and input for phase-2 planning

Dirk Duellmann, CERN IT

<http://pool.cern.ch> and <http://lcgapp.cern.ch/project/CondDB/>



LCG Application Area Internal Review,
March 31, 2005

POOL Phase 1 Highlights



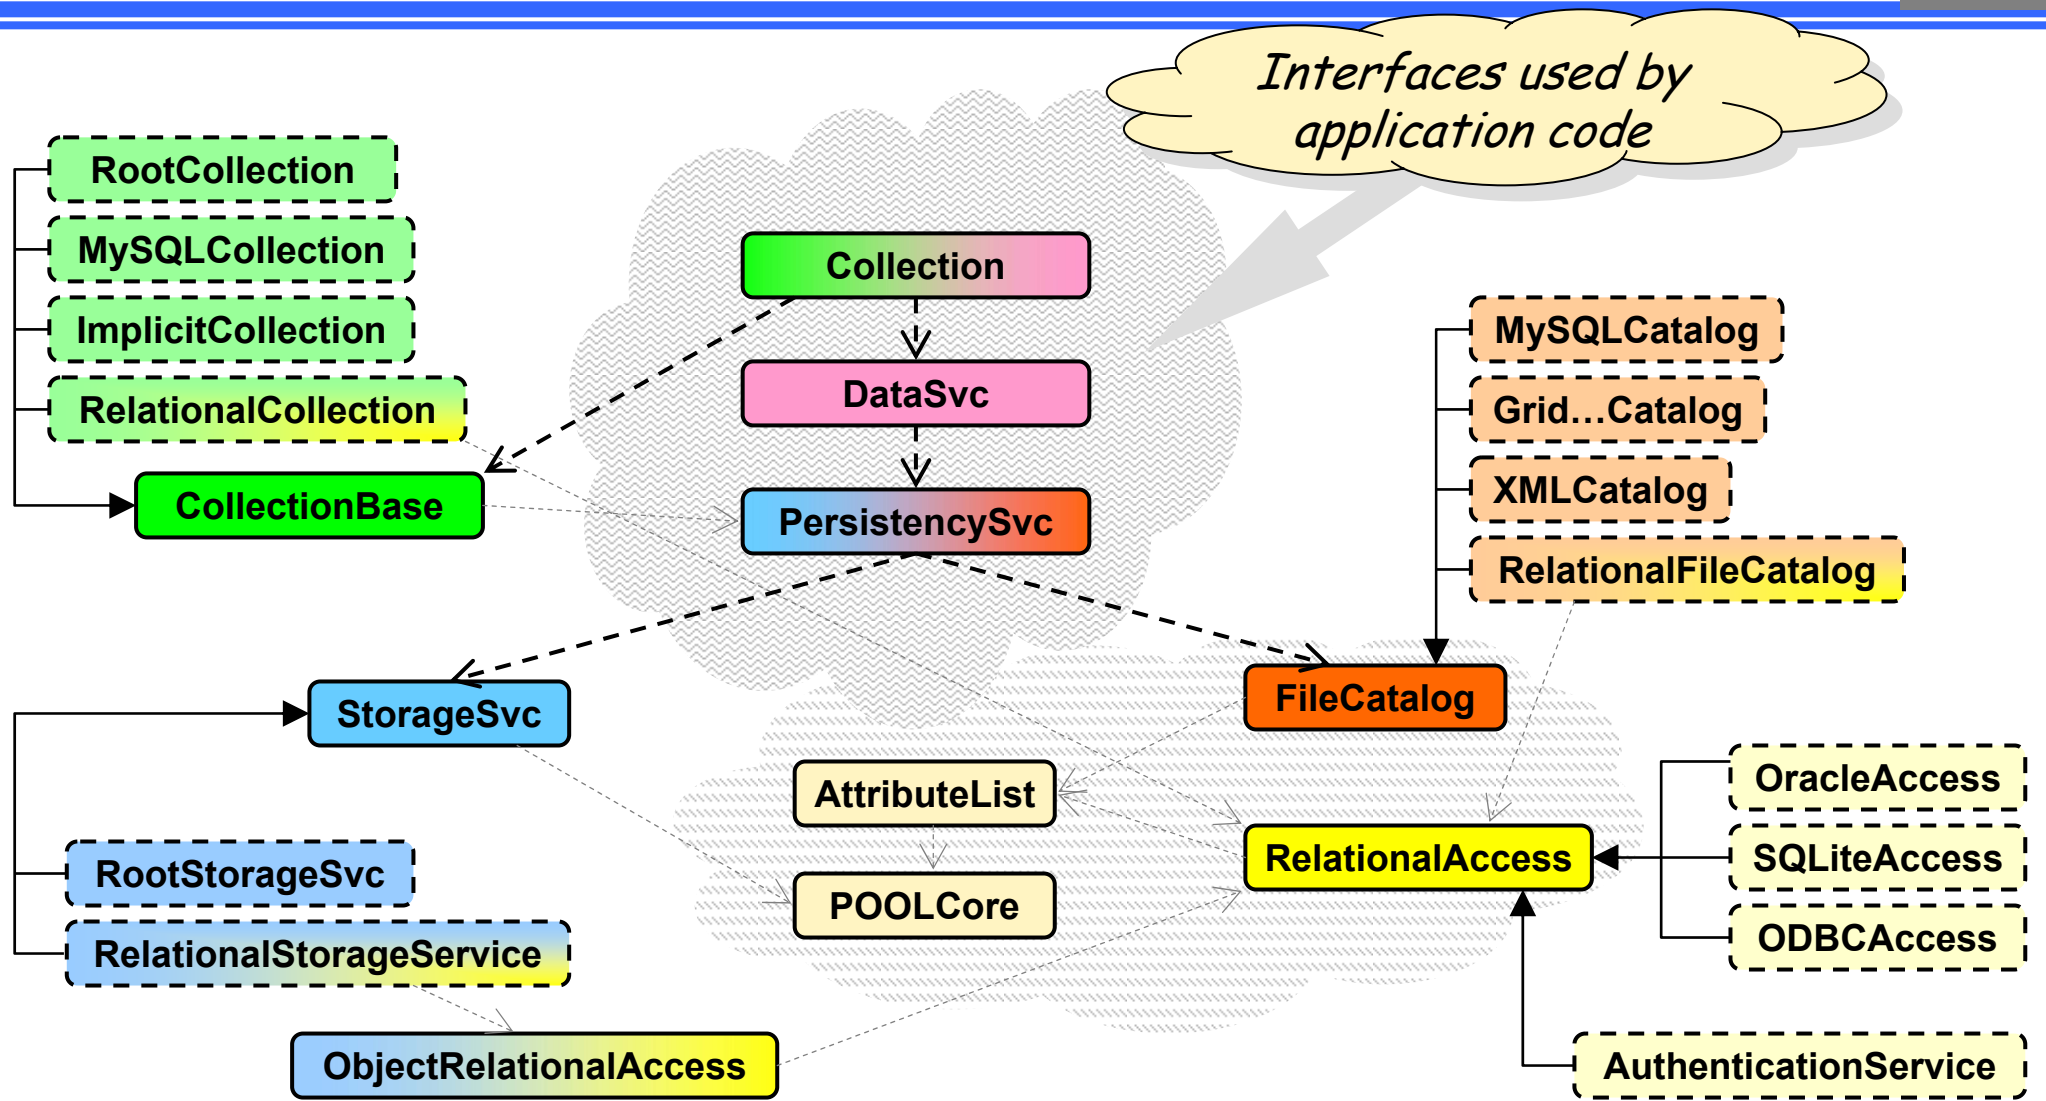
- **Schema extraction, representation and ROOT connection proven**
 - Prototype started in POOL and migrated to SEAL
 - gccxml and SEAL/Reflection are a valid approach
 - Interface to cint implemented in POOL and now moving to cintex
- **Complete STL handling in POOL**
 - Now migrated into ROOT 4 and used from there
- **POOL File Catalog deployment model picked up by all experiments**
 - consistent s/w for very different environments (development, non-grid and grid connected cases)
 - Catalog fragment extract/publish model is baseline for most productions
- **Technology abstraction has largely been achieved**
 - POOL RAL and RDBMS Storage Manager completed this goal and will allow for a consistent database foundation (monitoring, security, error handling)
 - POOL stayed open eg wrt to middleware developments (File Catalogs)
 - Blueprint plug-in model was essential to achieve this
- **Integration in experiment offline frameworks has been achieved**
 - Strong experiment involvement was in design and implementation was key
- **POOL successfully deployed close to PB scale**
 - Re-use of existing components (ROOT I/O, RDBMS) helped a lot!

POOL Project Evolution



- **POOL entered its third year of active development**
 - Joint development between CERN and experiments
 - During the last years we followed the proposed work plan and met the rather aggressive schedule to move POOL into the experiment production
 - Last year POOL has been proven in the LCG data challenges with volumes ~400TB
- **Changing from pure development mode to deployment, support and maintenance**
 - Several developers moved their effort into experiment integration or provision of back-end services
 - This is a healthy move and insures proper coupling between software and deployment!
 - But it affects the available development manpower
 - Task profile changing from design and debugging to user support and service optimisation
- **Need to maintain stable and focused manpower from CERN and the experiments**
 - Only this close contact has made POOL a successful project
 - Experiments and CERN have confirmed their commitment to the project
- **Manpower available to project developments now rather limited**
 - LCG Side - funding now insured
 - (re-)hiring round for phase 2 has started
 - Experiment Side - contributions are in several cases still without commitment
 - Valuable contributions received (new Tree Branch code from Bill)
- **Reliable planning needs committed manpower for the duration of the plan**

POOL Architecture



Main Phase 2 Topics



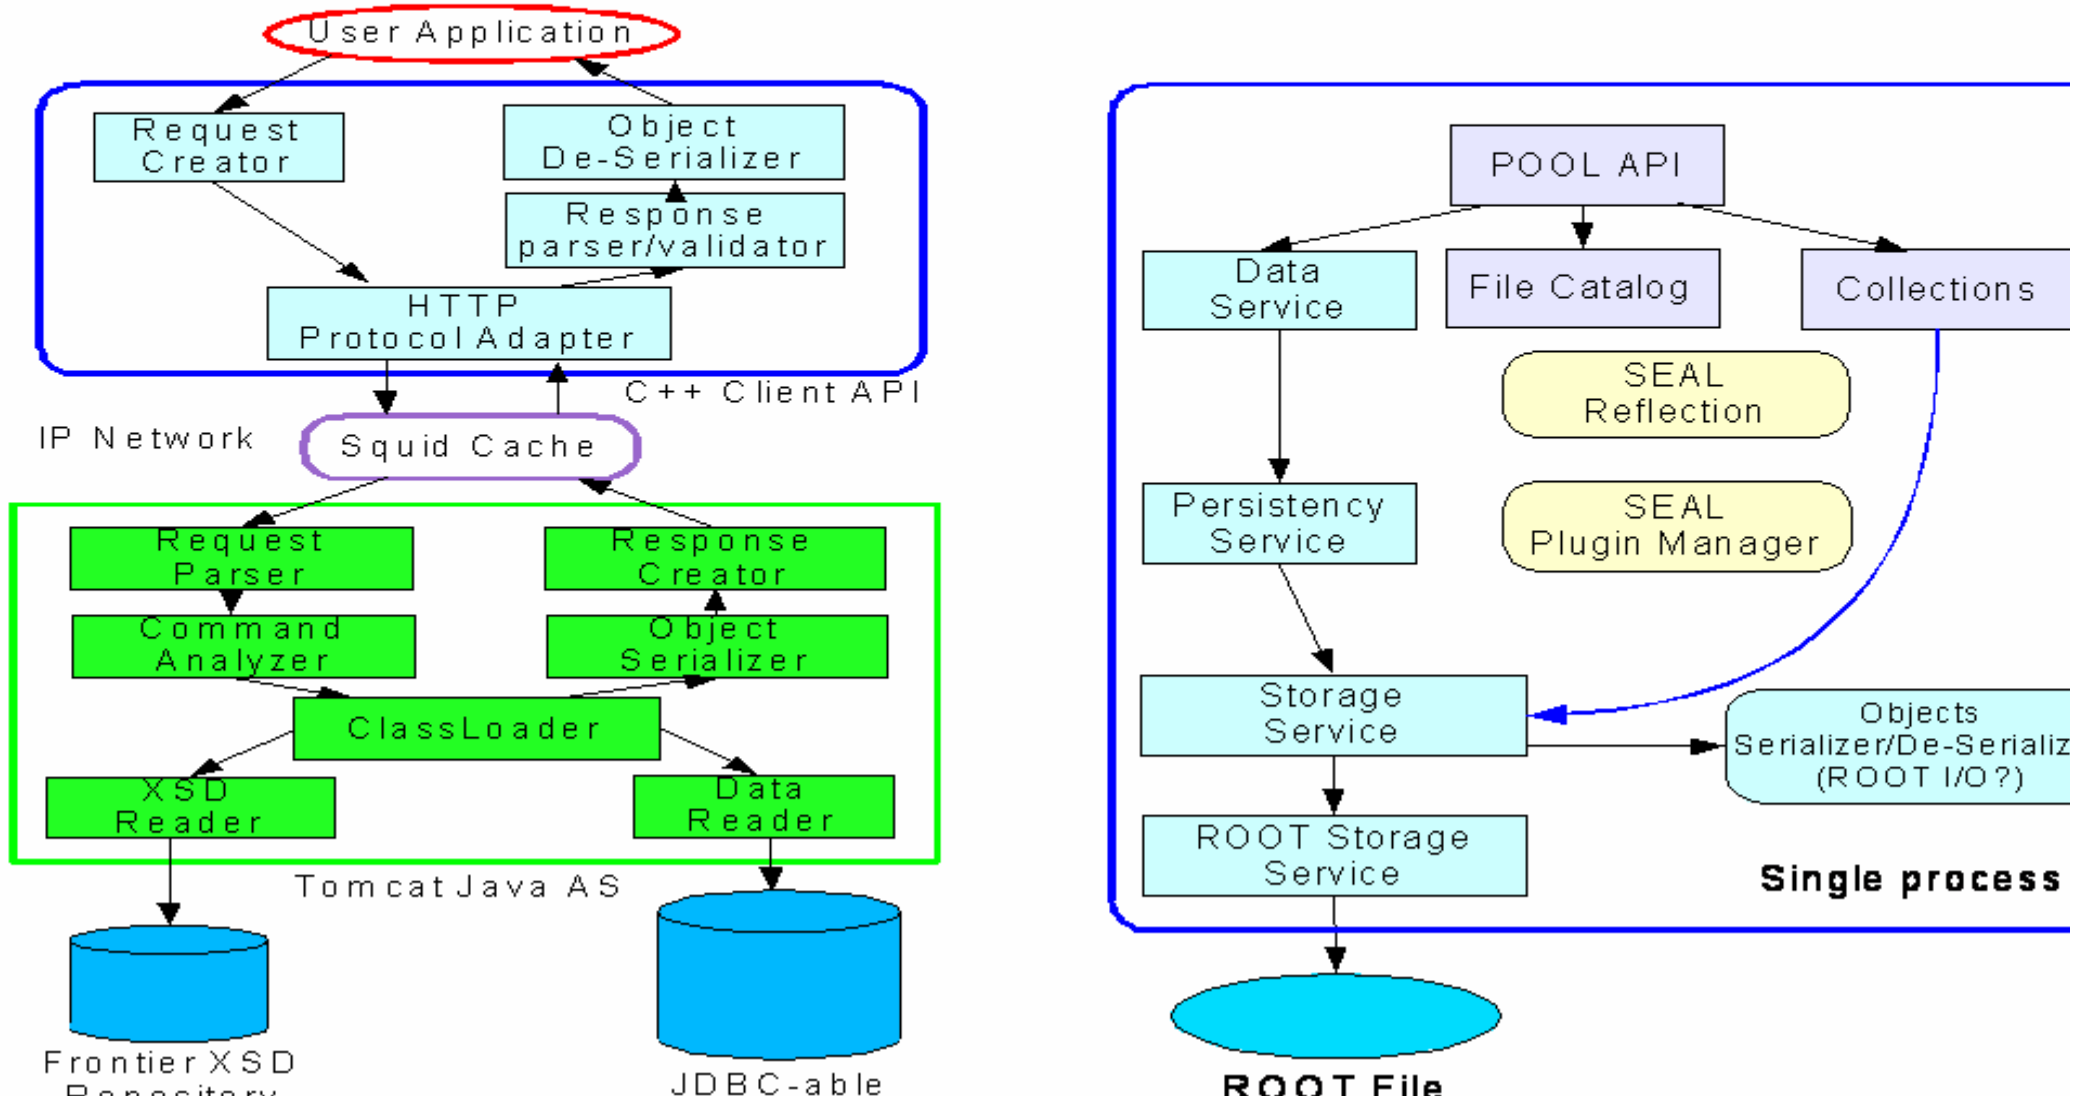
- Getting close to end of the first design & deployment phase
 - Main functionality has been provided
 - Some RTAG features have not been provided yet..
 - .. and not asked for anymore
 - Experiment requirements and focus change in some areas after more deployment experience
 - Need to insure that POOL experiment communication stays effective
 - POOL link people attend relevant experiment meetings (shared with IT service representation)
- AA Phase 2: Maintenance, Consolidation and Service Integration
 - Faster, simpler, more reliable
 - Better service integration with low level services
 - (Mass)storage and database services at Tier 0
 - LCG StorageElement, file catalogs and distributed database services
 - Need to establish closer technical links to grid deployment groups to insure that requirements from AA are communicated

POOL RAL and LCG 3D



- Propose POOL RAL as reference implementation interfacing to a distributed database service (LCG 3D)
 - Location independent connection to a database in the grid
 - Service catalog using the POOL (File) catalog components underway
 - Mapping of grid identity to local database/file user and role
 - VOMS implementation for DB (and files)
 - Database connection pooling and consistent DB error logging/handling
 - Consistent database client and database server monitoring
 - Eg query hitslist and database time
 - For single application or aggregated for a database domain
 - Relational data distribution
- Provide RAL independent from other POOL components
 - Packaged separately for applications which have no other POOL dependency but would profit from a consistent database service integration
- RAL as candidate foundation to be used by ROOT Database Trees?

Frontier & POOL (Simplified)



FroNtier Integration



- Access to remote Databases is provided by the FroNtier system developed at FNAL
 - http based data transfer with light-weight, database independent client
 - Full encapsulation of database details (schema, queries, physical storage) behind application server
 - Simplified deployment option for read-only data in particular for higher level tiers with only limited resources to provide a database service
- Proof of concept POOL integration has been done by FroNtier team at FNAL
 - Frontier objects appear as normal POOL objects as described by the LCG Dictionary
 - More work required to setup a test deployment infrastructure (within LCG 3D) but very promising option eg for conditions data
 - Integration with COOL project is being discussed
- POOL Integration of the pre-existing FroNtier system with was possible in relatively short time

Proposed POOL Milestones



- Split File Catalog production release (Q2)
 - Adapters for external FC implementations included in production release (Globus RLS, LFC (GD), gLite (GM), xrootd(?))
- RAL delivered as separately installable component (Q2)
 - SEAL/ROOT dependency will remain
- POOL build and releases passed to SPI (Q2)
- RAL supports logical database lookup, monitoring, connection handling and provides Grid Security integration (Q3)
 - Integration with LCG 3D infrastructure
- POOL FroNtier integration (production version) (?)
 - If requested by the experiments
- POOL Integrated with new reflex/cintex dictionary components (?)
 - Depends on SEAL/ROOT milestones
- POOL schema loading on demand implemented (Q3/Q4)
 - After new dictionary integration?
- POOL files access integrated with experiment analysis frameworks (Q2)
 - Refs and collection support requirements needs clarification

Manpower & Commitments Today



- Contributors with 20% FTE or more in POOL and a CVS commit during the last 6 month
- IT-ADC
 - Ioannis Papadopoulos, Radovan Chytrcek, Giacomo Govi, Maria Girone, {Dirk Duellmann}
- ATLAS
 - {Marcin Nowak}, Kristo Karr, {Yulia Shapiro}
- LHCb
 - Markus Frank(?) -> move to SEAL/ROOT
- CMS
 - Zhen Xie(?), {Bill Tanenbaum}

Proposed COOL Milestones



- **First COOL production release (Q1/Q2)**
 - Followed by a open review (Q2)
- **COOL meets performance requirements (Q2/Q3)**
 - After validation against experiment reference workload
- **COOL deployed in production at Tier 0 (Q3)**
 - Integrated with Physics database services at CERN
- **COOL data distribution deployed according to experiment models (Q3/Q4)**
 - Integrated with LCG 3D service
- **COOL integrated with FroNtier as data distribution mechanism (?)**

Manpower & Commitments Today



- Contributors with 20% FTE or more in COOL
 - and a CVS commit during the last 6 month
- IT-ADC
 - Andrea Valassi
- ATLAS
 - Sven A. Schmidt

Summary



- The LCG POOL project provides a hybrid store integrating object streaming with RDBMS technology
 - POOL has been successfully integrated into three quite different LHC experiments software frameworks
 - Successfully deployed as baseline persistency mechanism for CMS, ATLAS and LHCb at the **scale of ~400TB**
- POOL continues the LCG component approach by abstracting database access in a vendor neutral way
 - A Relational Abstraction Layer has been released and is being picked up by several experiments
 - Minimised risk of vendor binding, simplified maintenance **and data distribution** are the main motivations
- POOL as a project is (slowly) migrating to a support and maintenance phase
 - Need keep remaining manpower focused in order to complete remaining developments and to provide adequate support to user community
 - Maintaining a significant experiment contribution is required insure the the tight feedback loop which made POOL an effective project
- The LCG Conditions DB project has produced several releases of the Oracle and MySQL based implementations within the LCG Application Area
 - After an interface and extension review a concrete plan to consolidate the implementations has been discussed
 - Manpower also from the experiments is now becoming available to the project allowing to re-factor the package based on the Relational Abstraction Layer
- New complementary technologies such as FroNtier are being integrated into the LCG persistency framework as distributed access to database data gets more interest

Main Concerns



- ROOT/SEAL merge promises longer term savings
 - but will come with significant additional effort in the medium term
 - and more compatibility constraints from non-LCG users
- POOL relies on several SEAL components
 - Some of which being “consolidated” with ROOT counterparts
 - Change cascade of this consolidation to POOL and above needs careful evaluation
 - Merged product (ROOT/SEAL/RootStorageService) needs careful dependency check to avoid new or circular dependencies
 - Those can create more maintenance effort than code duplication
- Manpower in POOL and COOL is scarce
 - But proper coupling to experiment users and grid deployment /services will requires significant testing & communication effort
 - COOL project is progressing rapidly, is well aligned with other AA s/w and seems to have many interested users
 - But little commitment in terms of manpower
 - Any unplanned development requests are likely create significant waiting time for experiment users