First Internal Pool Release 0.1

Dirk Duellmann

Release Platform

- POOL V0.1 has been released on October 3rd
  - the cvs release tag is POOL_0_1_0
- Single supported platform so far
  - RedHat 7.2 using gcc 2.95.2
    - Code should also work on rh61 and (partially) on win32/VC6
- Require several external packages hosted by SPI
  - see /afs/cern.ch/sw/lcg/external
  - MySQL (4.0.1-alpha)
  - MySQL++ (1.7.9)
  - Root (3.03)
  - Xerces-C (1.6.0)
- Supported build systems
  - scram and (less complete) cmt

Release Content

- Storage Service
  - two operation modes objects in a tree or objects named into directory structure
  - provides coherent access to the two very different root optimisation models with the user same code
- File Catalog
  - two implementation based on mysql and xml
  - provide api & admin tools
- Refs and Cache
  - simplistic cache implementation
  - testbed for Ref classes – example for integration with experiment caches
- Reflection & Conversion
  - Full reflection interface
  - some refactoring of the interface going on (not yet in the tagged release)

Release Content II

- Release Target:
  - Provide component implementations for pool core component and show basic navigation between persistent objects
- Test Programs – eg /pool/test/Pool_1
  - write a few interrelated objects into a root file
  - read them back using pool object references
  - keep objects available for repeated local access in pool object cache
  - register and lookup files as required by the navigation from either of the two available catalog implementations
- But…
  - rather bare bones component integration so far!
  - Pool_1 test shows **not** the expected end user code
  - in many cases still calling low level component api which is to be hidden behind a proper end user facade (Refs/PersistencySvc)

Building POOL Libraries and Executables

- … is not as easy yet as it should be. Because of
  - bugs…
  - weaknesses in our integration testing (no automised test yet)
  - too little time for thorough integration testing in this cycle
  - not yet existing documentation
- Plan to tag a bug fix release
- Complicated setup because conflicting build system conventions
  - include directory names
  - path prefixing conventions
  - CVS tag conventions
  - compiler settings (GNU_SOURCE)
- All minor issues
  - but hard to resolve in a single development system

POOL development on CMT and SCRAM?

- Was this useful?
  - yes, at least during the project start-up phase
  - it allowed developers to work in their "home development environment"
- Should we keep **release** support for both or even extend this to other build mechanisms
  - definitely not!
  - this is convenient for single component development
  - but would likely turn into a nightmare for more serious integration testing and release QA

Next Steps

- Code review starts now
  - in 3 teams of 2 pool developers for 2 half days
  - SPI provides one external person for the review
  - Tagged reflection code will be send to Tyng-Ruey Chuang
  - To be finished by end of next week (also ROOT ws week!)
- Standarise and simplify cvs structure
  - Flatten sub directory structure to /pool/component
  - Use “pool” for component interface directories and #include “pool/ISomeThing.h” for inclusion
- Working with SPI project
  - On code documentation (cvsview)
  - Move asap to automatised unit and integration testing (cpp_unit/oval?)

Communication

- Pool work package meetings will be announced on our web and email list
- Everybody contributing to pool development is welcome to attend
  - WARNING: Don't expect to leave without a development task assigned to you
- The pool list became after the release quite active
  - Propose to setup an dedicated developer list for bug reporting traffic and internal developer discussions
  - Spare the traffic on the main list for architectural discussions interesting to a wider community

Next release

- In only less than a month …
  - Planning discussions for the release content started
- Collection support is coming
  - Commit of first prototype is expected soon
    - Steve Eckman replaces Chris Lain as of now
  - One day planning & integration meeting with Collection WP will be held tomorrow in Orsay
  - Expect Storage Svc requirement for container and file level iterators
- Object Conversion and Foreign Class Support
  - gcc-xml, ADL, XML based reflection & conversions prototype?
  - Need to improve the manpower problems in this work package
- Move to MySQL 4.0.3-beta
  - Alexandre Vaniachine has taken over from Zhen the responsibility for MySQL and MySQL++ (external packages used by pool)

Next release - Ideas

- Data extraction and publishing tools
  - take a few file home
  - publish file generated in a decoupled setup
- Catalog browser & admin tool?
  - complement command line tools developed by Zhen
    - eg Python based
  - rather decoupled - could be done eg outside CERN
- Move release platform to RedHat 7.3?
  - CERN will be moving soon – we should as well either for V0.2 or V0.3
- Provide additional support for gcc 3.2?
  - Only if we can get rid of one build system by end of October

Summary

- The internal POOL V0.1 release took place with planned content and only minor delay
  - Thanks to the full commitment of all pool developers and significant help from outside!
- This is the FIRST DEVELOPER release
  - report bugs/problems to the developer list – we left enough ☹
    - we'll create a pool developer list if traffic becomes unbearable
  - even better: fix bugs – most of them are still easy to fix ☺
- Many lessons learned from the very first cycle
  - We will try to fix the observed setup problems
    - The CVS structure will be simplified and standardised
    - Plan to drop release support either scram or cmt as soon as possible
- LCG POOL as a project is up an running
  - a growing fraction of the pool software is as well …