

Persistence planning

David Malon

malon@anl.gov

Argonne National Laboratory

Joint ATLAS, LHCb, LCG/IT meeting

1 December 2010

Introduction

- ATLAS has been informally reviewing its I/O and persistence architecture
 - And design and implementation and deployment and performance and ...
- Motivated by many factors:
 - Improving performance across a variety of storage and access platforms, data products, and use cases
 - Support for increasingly-many-core architectures
 - Refactoring to make implementations cleaner, more consistent, and more maintainable
 - Feature wish lists that did not make it into production before current LHC running
 - ...
 - But also by longer-term planning, particularly in the context of the computing side of ATLAS upgrade planning
- One part of this process is reconsideration of our strategy regarding persistence technologies and interfaces thereto
 - Motivation for today's meeting
- POOL is an essential element of that strategy today

POOL—and Gaudi

- What role does POOL play, what benefits do we derive from it, and at what cost?
- What should our long-term strategy be toward POOL and toward the functionality POOL currently provides?
- For functionality we wish to retain, what should our strategy be?
 - Moderated by the realities of expected support levels
- How should we approach functionality that we would like to have in the long term, but currently lack?
 - And where should it go?
- Important for ATLAS to understand better what LHCb is doing (and planning to do), and persistence-related services in Gaudi
 - ATLAS has not taken significant advantage of the evolution of Gaudi persistence-related services after adoption of POOL, for a variety of reasons
- Is there functionality in POOL (or missing from POOL) that both experiments might care about, that might better be supported in Gaudi?
- Are there current or planned Gaudi persistence capabilities that ATLAS should consider adopting or exploiting?

Time scale? Not yet

- **IMPORTANT:** No precipitous change will be made before the start of 2011 running
 - Expect that we can rely upon the current level of POOL support for 2011 data taking (essentially maintenance)
- Propose not to talk about COOL, or about CORAL very much today
 - Insufficient time, and different stories regarding use across experiments
- Purpose today is to improve the technical foundation for longer-term planning