Next-Generation Navigational Infrastructure and the ATLAS Event Store

Abstract:

The ATLAS event store employs a persistence framework with extensive navigational capabilities. These include real-time back navigation to upstream processing stages, externalizable data object references, navigation from any data object to any other both within a single file and across files, and more. The 2013-2014 shutdown of the Large Hadron Collider provides an opportunity to enhance this infrastructure in several ways that both extend these capabilities and allow the collaboration to better exploit emerging computing platforms. Enhancements include redesign with efficient file merging in mind, content-based indices in optimized reference types, and support for forward references. The latter provide the potential to construct valid references to data before those data are written, a capability that is useful in a variety of multithreading, multiprocessing, distributed processing, and deferred processing scenarios. This paper describes the architecture and design of the next generation of ATLAS navigational infrastructure.

ATLAS Software & Event Data Model:

The ***ATLAS*** experiment at CERN uses the Athena software framework to run **simulation**, **reconstruction**, and **analysis**.

StoreGate allows software modules to **transparently** use a data object that was created by an upstream module or read from disk.

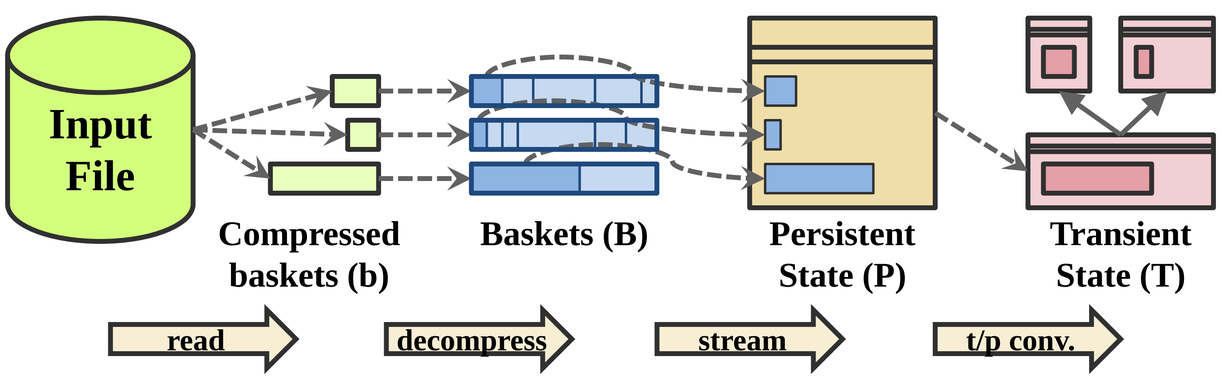

The Athena persistency framework: Reading POOL/ROOT data via the Athena software framework

References / Token:

ATLAS I/O and persistence infrastructure supports ***extensive navigational capabilities***:

- Real-time **back navigation** to **upstream processing stages**.
- **Externalizable** data object references.
- Navigation from any data object to any other both within a single file and across files.

Enhancements to the Navigational Infrastructure:

References to event entry points will be designed to use **unique**, **immutable** {run number, event number} pairs.

ROOT support for content-based indexing:

ROOT supports an in-file indexing capability that provides the foundation for efficient implementation of content-based references to event data in ATLAS. Because ATLAS event numbers will be **64-bit integers**, a number of enhancements are required to ROOT **indexing** and to its underlying **formula evaluation** infrastructure.

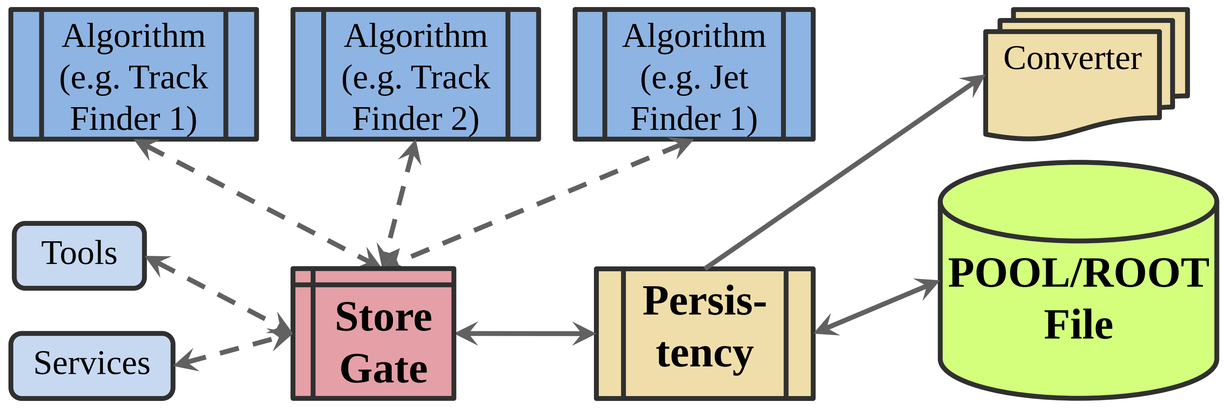

The Athena software framework: Uses a transient store to exchange data and interface to I/O and persistence infrastructure.

POOL / ROOT Persistency Framework:

Athena software uses ***ROOT I/O*** via the ***ATLAS POOL*** persistency framework, which provides **high-performance** and **highly scalable object serialization** to **self-describing**, **schema-evolvable**, **random-access files**.

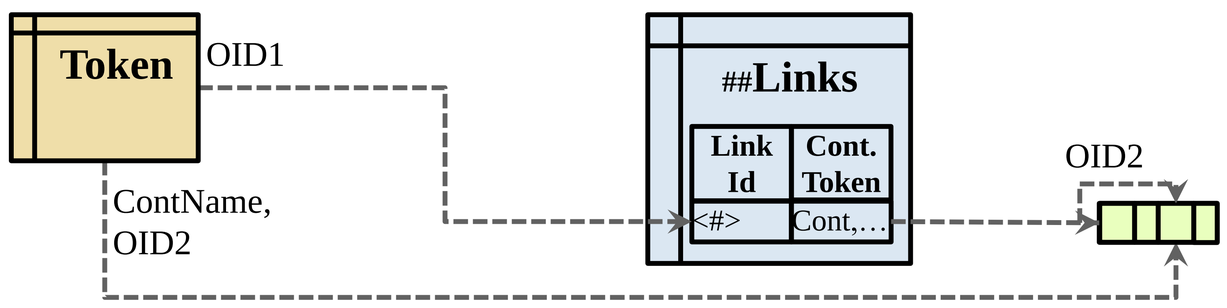

POOL Token references and Link table: Tokens store the storage technology, database identifier, persistent class identifier, container name, and two offsets

Immutable references for merged files:

For event data stored in ROOT, merging files means appending the Tree's entries. But Token references require the offset identifiers to be updated. The use of immutable attributes in the Token will allow **re-indexing** to be handled by the **underlying storage technology**.

Support for forward references:

In practice, to construct an event reference to a standard ATLAS data file one needs only to identify the file and to specify which event, as all other attributes have predefined default values. File identifiers can be pre-generated, allowing ATLAS to define forward references to events in **downstream data products**.

P van Gemmeren[1], D Malon[1] and M Nowak[2] on behalf of the ATLAS Collaboration

[1]Argonne National Laboratory, Argonne, Illinois 60439, USA

[2]Brookhaven National Laboratory, Upton, New York 11973, USA

E-mail: gemmeren@anl.gov