CHEP – Kraków, October 19-25th 2024

Optimizing the database architecture for metadata in preparation for ATLAS Run 4

ATLAS Metadata Interface (AMI) is a generic ecosystem for metadata aggregation, transformation and cataloging, benefiting from more than 20 years of feedback in the LHC context. This poster shows how it is planned to optimize the database architecture for metadata in preparation for ATLAS Run 4.

Migrating the ATLAS metadata to a single catalog

→ AMI is the official metadata catalog of the ATLAS experiment. It allows searching for data based on metadata criteria.
→ Around 100 M data and MC datasets distributed within 28 ORACLE databases (a.k.a. catalogs).
→ Datasets are derived from a root dataset. A dedicated database records the parentage tree.

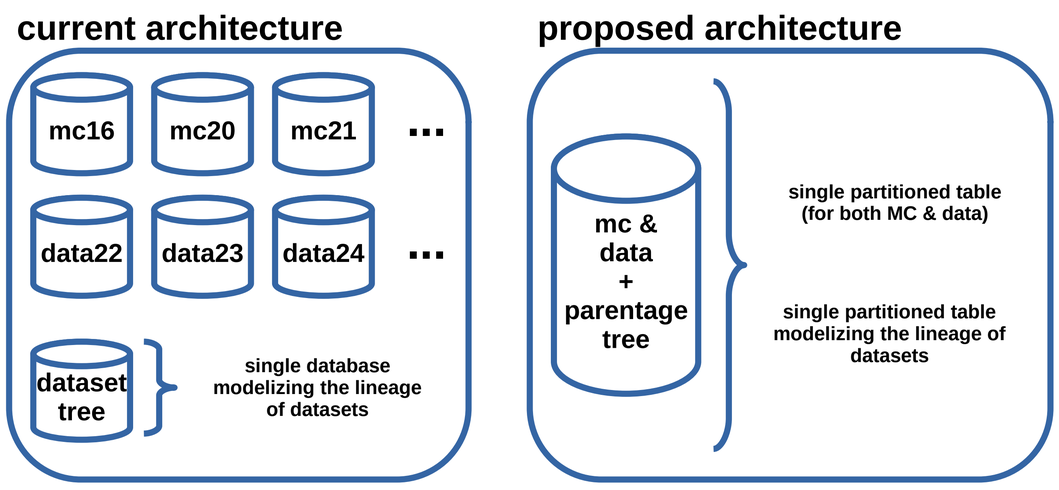

- Current architecture (20 years old):
  - 1 database per year for data and per production campaign for MC.
  - Single database with a table recording the parentage tree.
  - Suboptimal cross-catalog queries (slow).
- Proposed architecture:
  - Single database for MC, data and parentage tree.
  - MC and data in a single table + views to emulate the old arch.
  - DB partitioning using scopes (e.g. “mc23_valid”).

→ Performance tests ongoing for both inserting and searching.

Efficient propagation of physics parameters

→ For MC, in a given temporal interval of validity, dataset may have physics parameters (cross section, generator efficiency, …) inherited from a parent dataset.
→ Currently, ~500 M physics parameters.

- Current architecture:
  - Each physics parameter is propagated, across catalog, and duplicated for each children dataset.
  - High redundancy (up to 25x) → coherency errors.
  - Time consuming propagation task.
- Proposed architecture:
  - Physics parameters attached to a single dataset (often the root one) and not duplicated.
  - Physics parameter values are retrieved by a global SQL view crossing information from:
    - the unique dataset table,
    - the unique parentage tree table,
    - the unique table storing the parameter values.
  - Partitioning and indexing to guarantee ~constant performance over time.

Physics parameters propagation rules

- For a given interval of validity, a physics parameter can be attached to the root dataset or on any of its children.
- A physics parameter entry has 2 propagation rule flags:
  - override: the new value overrides the value of children,
  - propagate: indicate whether the value apply for children.

Example of dataset lineage

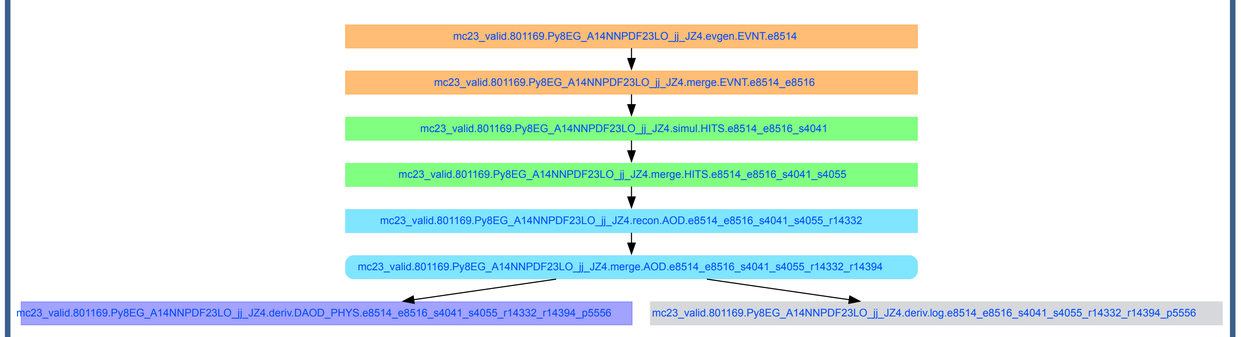

Automatic scoping with the Metadata Query Language (MQL)

- MQL is the AMI ecosystem's domain-specific language for writing data queries without having to specify joins → end users only need to know the metadata names.
- To guarantee simple, high-performance queries, the MQL engine has to automatically specify the partitioning scope values in the generated SQL (ongoing development).

Performance

- Performing a dataset selection in the new single catalog architecture, regrouping ~7.5 M entries, gives similar performance (~100ms) as querying a test catalog of the old architecture regrouping ~300 k entries.
- Physics parameter retrieval (full history or latest values) takes less than 200ms → encouraging result.

J Odier, F Lambert, J Fulachier, P-A Delsart (LPSC – Grenoble),
on behalf of the ATLAS Computing Activity.

Refs.: [1] Jérôme Odier, Fabian Lambert and Jérôme Fulachier.
The ATLAS Metadata Interface (AMI) 2.0 metadata ecosystem: new design principles and features.
J. Phys.: Conf. Ser., vol. 214, year 2019.
https://doi.org/10.1051/epjconf/201921405046