OPTIMIZING ACCESS TO CONDITIONS DATA IN ATLAS EVENT DATA PROCESSING

A Formica[1], E Gallas[2] , L Rinaldi[3]

CHEP 2018 - Sofia (Bulgaria)

The ATLAS Conditions DB

The COOL DB architecture

Folders represent Conditions DB tables
- Each **Folder** is owned by a specific **Schema**
  - Each has subsystem, instance, and if used offline or strictly online

Multi-version **Folders** have one/more **FolderTags**
- For Conditions that allow different versions over time intervals

FolderTags may be included in one/more **GlobalTags** (when designated to be used in event-wise processing)

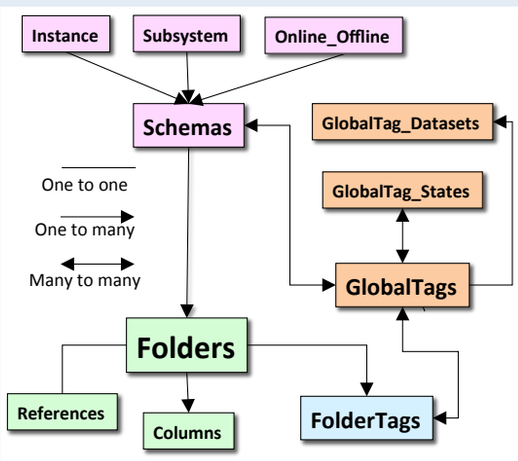

What is stored in the Conditions DB

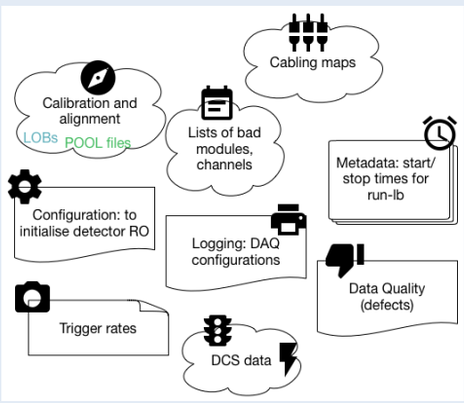

Condititions data are in general not event-wise, that characterize the states of all ATLAS sub-systems, organized in specific Intervals of Validities (IoV).

They contain a large variety of information such as Detector Control System (DCS), trigger and DAQ, Data Quality, information from LHC and sub-detectors

At present time Conditions data are stored in a COOL DB with more than 17000 Folders, distrubuted over 30 Schemas

How Conditions-data are used in event-data processing

Conditions data are accessed during almost all ATLAS data processing activities.

Online applications save conditions which have been observed during a run such as configuration, run parameters, measured physical parameters, etc. Such data are stored into single version folders and tags are never used; their payload is stored into columns of relational tables in accordance with data model provided by COOL (numeric values and strings).

Offline Distributed Computing: the Athena jobs access concurrently conditions data via COOL API using the intermediate caching system provided by the Squid/ Frontier services. Payload POOL ROOT files are distributed over the CVMFS file system.
In some specific cases, a subset of conditions data is exported to SQLite files and accessed directly.

Gathering DB usage information

New information added upon an existing collection of metadata about Conditions database folder and global tags:
- Conditions DB Metrics: database table row counts, average row volumes and counts of distinct references to external data by folder and folder tag (red columns)
- Conditions Usage patterns of existing representative event data processing jobs (yellow tables)

Improved understanding of Conditions data usage by event-wise processing

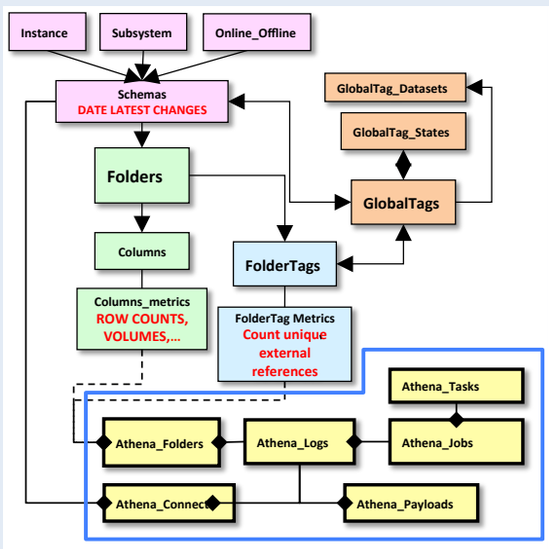

Conditions DB Storage/Insert Rates

Listing ATLAS COOL production database instances

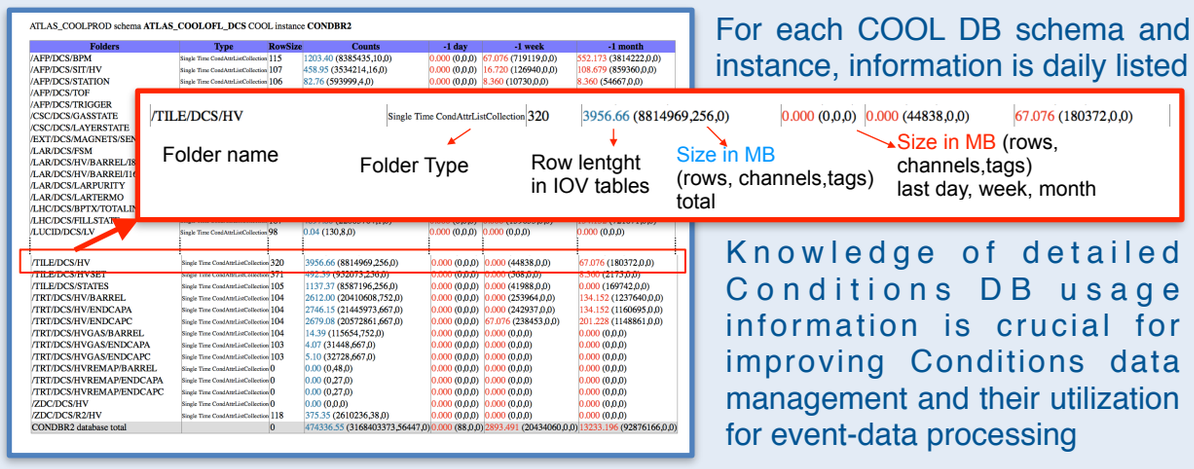

For each COOL DB schema and instance, information is daily listed

Knowledge of detailed Conditions DB usage information is crucial for improving Conditions data management and their utilization for event-data processing

Conditions DB Folder History Report

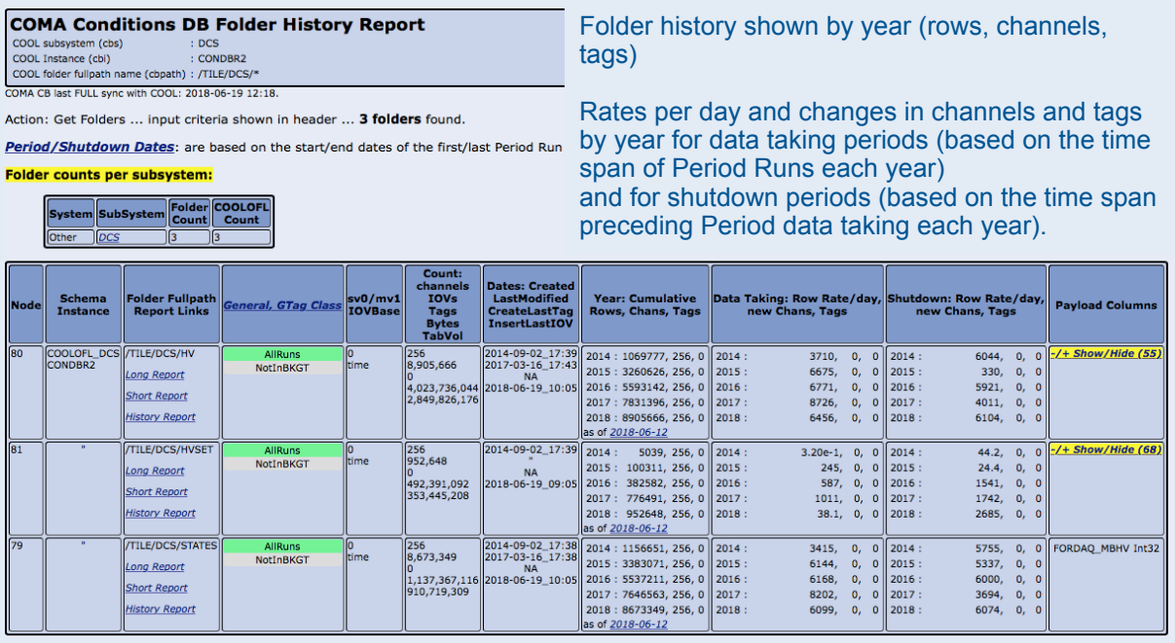

Folder history shown by year (rows, channels, tags)

Rates per day and changes in channels and tags by year for data taking periods (based on the time span of Period Runs each year)
and for shutdown periods (based on the time span preceding Period data taking each year).

The Custom DB Release use case

DB Releases are composed of file-based 'static' SQLite replicas containing relational data from the above databases plus the detectod Geometry and the POOL ROOT Payload data files pointed to by reference in some Conditions Database Folders.
DB Release tarballs are mainly used for Monte Carlo simulation, and are distributed on the grid and on shared file systems.

Initially, DB Releases contained the full MC Conditions dataset, size: O(10 GB)

Simulation @ HPC

Thanks to the the Conditions metadata usage tool, we learned that only a small subset of Conditions data is effectively used for MC Simulation

We are now building and deploying Custom DB Releases, on demand, which size is O(100 MB)

Particularly useful in preocessing environments with limited/low connectivity (Like some HPCs and home/small clusters)

Studies are under way to build Custom DB Releases for specific real data processing tasks (like Overlay Event). Mainly dominated by large sized DCS folder

Benchmark suites

Software tools used for benchmarking should not overload the hosting system, so having light components could enhance the precision of benchmarking measurement.
Recently, custom light DB Release have been implemented into the ATLAS KitValidation benchmarking tool, which is used inside the CERN Cloud Benchmarking suite

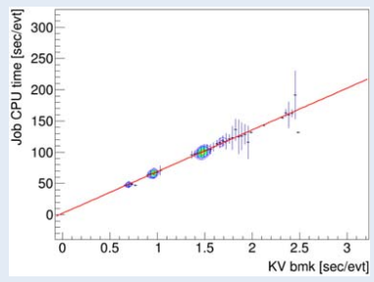

Correlation between KV measurements and average CPU time per event of ATLAS jobs running in several Azure data centres

[1] CEA Saclay IRFU/SEDI, [2] Oxford University, [3] Bologna University and INFN

contact: rinaldi@bo.infn.it