Some views on metadata from Data Preparation

Beate Heinemann & Andreas Hoecker

ATLAS metadata workshop, Grenoble 30−31 August, 2010

Metadata in ATLAS

Distinguish metadata levels (see metadata task force report)

- Event
- File
- Luminosity Block (LB) [~2 mins]
- Stream
- Dataset
- Run

Metadata storage

- Event level metadata are written into the bytestream
- File level metadata are stored in the SFO/Tier0/DDM/AMI DBs, also in COOL
- LB metadata are stored in bytestream, COOL and the SFO DB
- Stream-level metadata are stored similarly to file-level metadata. Streams are configured via Trigger configuration DB (filled from python configuration)
- Dataset metadata are stored in Tier0/DDM/AMI DBs, and in COOL
- Run metadata are stored everywhere

Metadata relevant for Data Preparation – **Data Periods**

Runs are organised along "Data Periods"

- Correspond to Year plus Major and Minor period initials, eg, **2010-periodE6**
- Somewhat arbitrary, but cut along coherent detector, trigger and software configuration.
- Minor periods important for scheduling before conferences
- Stream unaware (ie, same for all physics streams)

Data periods are defined in text files located on Data Preparation AFS area

The "RunQuery" and DQ tools read this area as information for users

The information is book-kept on DataPeriods TWiki

"Data Periods" are commonly used as unit for communication between physics analysts, and also as input to the creation of "Physics Containers"

Metadata relevant for Data Preparation – **Physics Containers**

Data periods are used to create corresponding "Physics Containers"

- Physics containers are "super-datasets" in the nomenclature document
- A physics container is collection of *coherent* (and stream-wise) datasets
- Coherence means in particular identical reconstruction software for all runs belonging to a given physics container

Description on TWiki: always one physics container per stream and data period

Physics containers are defined via formatted text files on DP AFS (web), with comment field, data types, list of datasets. They are fully referenced in AMI

Naming convention:

[projectTag].[period].[streamName]. PhysCont .[dataType].[version]

- project tag, stream name and data types are just as for datasets
- period is data period, eg, "periodE6"
- version tag is, eg, "t0pro04_v01"

Physics containers are DDM containers just as any other "dataset" (aka, container) and can be used in same way by DDM tools (eg, as input to jobs)

Metadata relevant for Data Preparation – **Data Quality**

The data quality status is encoded in LB-wise "Data quality flags"

- Various sources: online and offline
- Various levels: DCS, DQMF (automatic monitoring), SHIFT, Final ("LBSUMM")
- Stream unaware (ie, same for all streams)
- Overall "ReadyForPhysics" flag indicates warmstart status (not a DQ flag)

The flags are written to COOL (DCS flags obtained via *DCS Calculator* from DCS status information written to COOL via PVSS2COOL copy process)

The DQ flags are NOT anymore written into the Detector store
(or *should not* be written anymore → use GRL mechanism to select good runs/LBs – see next slide)

Preliminary (Tier-0 processing) "LBSUMM" DQ flags for given run period frozen in COOL Tags. Preliminary "pass1-analysis-2010XY" tags available 4 days after run ends (or earlier in rush periods). Refined tags made after reprocessing.

DQ flags are set and visualised in "DQ status browser" and also visualised in Run-Query and other DQ tools

Metadata relevant for Data Preparation – **Good Run Lists**

DQ flags used to derive Good run lists (GRL)

- Correspond to combination of DQ flag and ReadyForPhysics requirements applied to set of runs
- Output is XML file with list of runs and selected LBs per run
  (some additional metadata given)

Official GRL maker web page

Tools exist to apply GRL selection of LBs at Athena and ROOT levels, and for luminosity calculation

Metadata relevant for Data Preparation – **Luminosity**

Online and refined offline integrated luminosity per LB is stored in COOL

- Luminosity of several detectors and algorithms stored
- Large amount of LHC and luminosity related metadata in COOL (incl. DQ)

Integrated luminosity per run, LHC fill, day and corresponding cumulative luminosities are displayed by the Data-Summary web

Run-wise luminosity information is also provided by Run-Query

Users compute luminosity for their analysis via their specific GRL using the iLumiCalc.exe tool, either in the release or (recommended) via a web interface

- Inputs are:
  - **GRL XML file**
  - Trigger to compute **live fraction** (currently: L1_MBTS_2, information from COOL)
  - Physics trigger used (EF) to compute **prescale corrections** (information from COOL)

The computation of integrated luminosity is corrected for TDAQ deadtime, but not for any data losses downstream…

Metadata relevant for Data Preparation – **Luminosity** *contd…*

During the user data processing, processed LBs must be book-kept

- Athena "metadata store" keeps track of analysed LBs, even if no event has passed selection
- GRL tools allow user to recreate GRL after private data processing and data analysis, taking into account possible data losses on the way

Complete LBs that are lost will hence not bias integrated luminosity calculation (and hence, eg, cross section measurement)

A problem is if files within an LB have been lost …

Can be circumvented by merging RAW files before any data processing to a single file per data stream and LB

- Mostly done at Tier-0
- However, not possible if file size exceeds 5 GB (not happening with well balanced streams and 200 Hz output rate)

Metadata relevant for Data Preparation – **File-level data**

Complete file-level information is available in the SFO/Tier0/DDM/AMI DBs

- The SFO is the master DB for RAW files (however, not covering recovered files from the debug stream)
- Files and numbers of events per file of reconstruction products are available in the Tier0 and higher-level DBs
- Event losses due to, eg, job crashes in the RAW→ESD reconstruction step of the Tier-0 are book-kept in COOL (/GLOBAL/FILECOUNT/PROMPT stream-wise folder), but not downstream losses
- No other bookkeeping in COOL so far (required in case of luminosity corrections)

Datasets and file-level information is displayed by the Tier-0 monitoring and Run-Query (up to Tier-0 only), AMI (all levels)

Metadata relevant for Data Preparation – **Other data**

Large amount of other DP and analysis related metadata available in COOL

- Detector mask and DAQ partition
- Trigger configuration information (tagged via “keys”) and trigger rates (all levels, but not sure it is fully debugged for HLT)
- Beamspot fit results
- Magnetic field values
- *… as well as the huge set of detector conditions (which are also metadata)*

Referenced versus time or run/LB numbers

All of this information is accessible in Athena (via COOL IovDbService), and many are browsable with Run-Query

Conclusions / open areas

Metadata in ATLAS in rather good shape from DP point of view

Open tasks:

- All production file/event losses should be stored in COOL, and be easily tracked/displayed (probably Run-Query)
- Need a tool that checks integrity of all metadata in system (started for files by reprocessing coordinators)

Need to clarify use of GRL in production tools: DP recommendation is *not* to apply GRL selection during Athena production, but at ROOT (ntuple) level

- Reason is that DQ flags may get updated
- New "StandBy" stream reduces fraction of useless (ie, non-stable-beams and non-ReadyForPhysics) data to minimum

→ Run grid Athena jobs over Physics Container, and apply GRL selection at final stage of analysis ☺

Question to ADC: would it be useful to replicate ESD along data periods?