

Table of Contents

LHCb Bookkeeping Working Group Meeting.....	1
Date and Location.....	1
Attendees.....	1
Apologies.....	1
Subjects.....	1
Presentation of the existing system.....	2
Matters arising during the presentation.....	3
File ordering from queries.....	3
Quality groups.....	3
Raw data organisation.....	3
Schema extension in roottree.....	3
Query capabilities.....	3
Meeting closed at 15:00.....	4

LHCb Bookkeeping Working Group Meeting

Date and Location

23 March 2007
13:30 - 15:00

Attendees

Andrew, Carmine, Markus, Michael (partly, phone), Niko, Olivier C. (phone), Olivier D., Philippe (chair, notes)

Apologies

None

Subjects

Presentation of the existing system

Carmine presented the existing BK system, consisting of two separate sets of tables: the warehouse DB (contains the job, files and quality tables with their links), the views tables (contains a flat table *roottree* with one row for each combination of the query parameters: configuration, event type, file type, processing versions...). *roottree_'s* rows point to a list of associated files (one table per *_roottree* row). Currently there are 3000+ rows in *roottree*.

Matters arising during the presentation

File ordering from queries

At each query the whole file list is ordered by alphabetic order of the file name. This was considered acceptable by the WG as for the raw data it represents an approximate chronological order. It however poses a problem for users performing analysis "on the fly" and willing to update with new data as files are not necessarily entered in chronological order. A tool should be provided (command line and possibly in Ganga) in order to be able to "subtract" a former dataset from the result of the query.

Quality groups

Each file may have a set of quality groups with a quality flag each. For example: Tracking=Good Calorimeter=Bad Muon=Absent. Currently there is only one group (production) and no possibility to query on this parameter (the views are only created for files that are of good quality). This must be added to the views and the search.

Raw data organisation

As a run consists of a set of files taken with the same running conditions, **the WG recommends that data from different runs are not mixed** down to the level of streamed DSTs. This puts constraints on the data distribution as one should optimise the size of the final DSTs.

Considering an average fill of 6 hours, it consists of 720 2.1GBytes files of raw data (60000 events per file, 30s). The reconstruction will produce 720 rDSTs of approximately 1.2 GBytes. Assuming DST streams will contain 0.5% of the events each, with a size of 100 kBytes, each stream would have a size of 21 GBytes. One should evaluate in due time if data of a given run can be distributed to all Tier1 sites or only to a limited number in order to avoid too small DST files. The quoted figures would possibly be OK as smaller sites would have 10% of the data (i.e. 2.1 GBytes). File sharing should take carefully into account this reduction factor for grouping files to a given site (in the quoted example 70 raw data files are a minimum for a site). There will always anyhow be a remainder...

It is accepted that for raw data the initial "job" represents a run. Its *CONFIGNAME* should consist of the DAQ partition and the *CONFIGVERSION* of the actual DAQ configuration. Additional job/run parameters need to be better defined, but the first that come to mind are: fill number, start/end date-time, HLT version.... These parameters at least should be all queryable, hence be used in a *roottree* table of the views. As the parameters for real data and MC data are quite different, **the WG recommends that real data and MC data are in two different *roottree* tables**. Users would select initially which kind of data they are interested in.

Schema extension in *roottree*

Currently *roottree* allows only 3 parent jobs, which is clearly insufficient. This should be extended rapidly to cope with stripping of MC data.

Query capabilities

Query capabilities should be enhanced with the (non-exhaustive) additional features:

- Query on quality flags
- Query on ranges of parameters rather than single value
- Selection of a number of events rather than a number of files

- For dataset queries at a given site, the server should internally loop until it provides the required number of eligible PFNs of replicas rather than selecting the number of LFNs and give only those PFNs that match (urgent)

The need for querying a file from the event number (3*32-bit number) was not considered a must as at any level of processing one can get the GUID of the raw data file and use it as input selector option if one cannot navigate to it because it is not present at the site.

Meeting closed at 15:00

Next meeting Friday 30 March at 9:30, same place, same code. Agenda to come...

-- Main.phicharp - 23 Mar 2007

This topic: LHCb > BKWGMinutes070323

Topic revision: r1 - 2007-03-23 - PhilippeCharpentierSecondary



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback