




Event Index Core Schema

(Hadoop & HBase)

- Global Schema
- Catalog
 - Catalog Attributes
- Dataset TagFiles
- EL HBase Table



It may be useful for future evolution.

[Julius Hrivnac](#)

EI WS, 12-13 Feb 2018, Valencia



Global Schema

HBase

Catalog

Contains entries for all TagFiles,
primary (datasets) or derived,
permanent or temporary,
full or index-like.

**Indexed by TagFile id (= dataset name for primary
dataset, with RunNumber at the end for optimisation
wrt to most frequent prefix searches).**

EL Table

Contains all events.
Optimised for EL searches.
Contains pointers to events in TagFiles (to get all
available info).

Indexed by global RunNumber_EventNumber.

HDFS

TagFile

Contains actual tag data.
May have different schemas.
Described in Catalog.
**Indexed by RunNumber_EventNumber for
primary datasets.**

upload

*There is redundant information (RunNumber,
streamName,...) to allow faster searches and merging.*



Global Schema (for dataset TagFiles)

Catalog entry

id = year.project.**streamName**.prodStep.dataType.version.**runNumber**

name
path
type
format
info
key
schema
before
after
master
slaves
indexes
imported
nevents
amiNevents
status
trigger
project
prodStep
dataType
runNumber
version
streamName
idver
multi
missing
consumer
... freely extensible with other attributes

path points to TagFile
(= dataset)

redundant info
to allow merging and
to speed up

TagFile entry

key = **RunNumber**_EventNumber

LumiblockN
BunchId
EventTime
...
Trigger Masks, etc
...
nam0
db0
clid0
tech0
oid0
... the same for 1,2,3
project
streamName
prodStep
dataType
version
...
L1/L2/EF Trigger Chains PH/PT/RS

EL Table entry

key = **RunNumber**_EventNumber

ref_project
ref_stream
ref_step
ref_type
ref_amitag
[stream_id
guid]
...
[Projects,**Streams**,Steps,**Datasets**,Types,Amitags referenced above]

refs to keys
in the same table

dataset contains pointer
to TagFile entry
(= event)



Catalog (1)

Kudu Metadata ?

- **id:** The unique identifier, HBase key, chosen to use prefix search for most frequent queries
 - EI17.1.data17_13TeV.physics_Main.deriv.DAOD_HIGG1D1.f889_m1902_p3402.00339500 for dataset (RunNumber @end)
 - D0overlap/EI15.1/00271388 for dataset overlap table
 - T0overlap/EI15.1/data15_13TeV.physics_ZeroBias.merge.AOD.r7600_p2634.00271744@S1000 for trigger overlap table
 - @S1000 = based on selection of 1000 pseudo-random events
 - TStat/EI15.1/data15_13TeV.physics_ZeroBias.merge.AOD.r7600_p2634.00271744@00 for trigger statistics table
 - @00 = -tlevel/-tdecision options of generating ti command
 - EICache/2017.12.18.10.13.56.688/atlevind for CLI command result
 - Cleaned periodically
 - EIWSCache/EICache/2017.12.18.06.30.56.096/tomcat for WS command result
 - Cleaned periodically
 - Anything else chosen by -outname option to ei command
- **name**
- **path:** HDFS path
- **type:** tag, table,...
- **format:** map,kudu,...
- **info**
- **before:** for chained TagFiles forming TagSet
- **after:** for chained TagFiles forming TagSet
- **master:** for derived TagFiles
- **slaves:** derived TagFiles
- **indexes:** index TagFiles (can be either full subsets or pure indexes)



Catalog (2)

- **key:** HDFS file key (RunNumber_EventNumber for datasets)
- **schema:** HDFS file schema (different TagFiles can have different schemas)
- **attributes:** anything else, can be freely attached



Catalog Attributes

- **imported**: when imported
 - There is also **timestamp** of last Catalog modification of an entry
- **nevents**: number of events
- **amiNevents**: number of events in AMI, if different
- **status**: good (=null), bad, modified (from AMI), invalid (from AMI), trashed
- **trigger**: bad for bad triggers
- **project**: data15_13TeV,...
- **prodStep**: merge,...
- **dataType**: AOD, ESD,...
- **runNumber**
- **version**: AMI tag
- **streamName**: physics_Main
- **ldver**: catalog schema version, currently = 2
- **multi**: if duplicated events, <multiplicity>x<number of events>
- **missing**: missing TagFiles in a TagSet
- **consumer**: consumer version, currently os

open - can be freely added/updated



Dataset TagFiles (1)

➤ **key:** RunNumber_EventNumber=String

➤ **schema:**

- LumiBlockN=int
- BunchId=int
- EventTime=int
- EventTimeNanoSec=int
- EventWeight=float
- McChannelNumber=int
- Lvl1ID=String
- IsSimulation=int
- IsCalibration=int
- IsTestBeam=int
- L1trigMask=String
- L2trigMask=String
- EFtrigMask=String
- SMK=int
- HLTPSK=int

- Different for other types of TagFiles (like Overlap tables).
- May be even different for different dataset TagFiles.



Dataset TagFiles (2)



schema (cont):

- nam0=String
- db0=String
- cnt0=String
- clid0=String
- tech0=String
- oid0=String
- nam1=String d
- ... (till nam3)
- project=String
- streamName=String
- prodStep=String
- dataType=String
- version=String



Dataset TagFiles (3)



schema (cont):

- L1trigChainsTAV=String
- L1trigChainsTAP=String
- L1trigChainsTBP=String
- L2trigChainsPH=String
- L2trigChainsPT=String
- L2trigChainsRS=String
- EFtrigChainsPH=String
- EFtrigChainsPT=String
- EFtrigChainsRS=String



ELHBase Table

- `<run_number>_<event_number>` ->
 - `[r:<ref_dataset>:<pointer_event_in_dataset> : <ref_project>,<ref_stream>,<ref_step>,<ref_type>,<ref_amitag>,<stream_id>,<guid>]`
 - `<ref_dataset>` ->
 - `r:ec : <pointer?>`
 - `r:i : <ref_status>`
 - `<ref_status>` ->
 - `r:n : <dataset>`
 - `r:s : <status>`
 - `<ref_project>` -> `r:n : <project>`
 - `<ref_stream>` -> `rn:n : <stream>`
 - `<ref_step>` -> `r:n : <step>`
 - `<ref_type>` -> `r:n : <type>`
 - `<ref_amitag>` -> `r:n : <amitag>`
 - `<stream_id> = [A|D|E|G|H|R]`
 - `<status> = [delete[d]]`
-
- Several logical tables within one physical table.
 - The basic structure is:
 - `RunNumber_EventNumber` ->
 - `[GUID]`
 - link to TagFile
 - Aux
 - Rainer or Grigori for details.