Data Store and Query System

- Data Organisation
- News
- Problems

latest version: ***1.14.0***

J.Hrivnac, SW WS, 3Feb15, CERN
+ Justin, Fedor, Rainer, Andrea,...

Data Organisation - Global

- **EI14**: for “raw” 2014 data (details next slide)
- **EI15**: for “raw” 2015 data (details next slide)
- …
- **EICache**: for results, periodically purged, automatically assigned filename
- **EIWSCache**: for results from WebService, periodically purged, a user can specify filename
- **EI14Derived**: for derived/rearranged data (EventServer, EventPicking,...)

Data Organisation - Detailed

EI14/
<project>.<runNumber>.<streamName>.<prodStep>.<dataType>.<version>/
<guid>_<transid>_<consumerid>_<pandataskid>_<pandauserid>

- Dir/filename fields are also available from Catalog, so:
  - They can be searched at the Catalog stage (fast):
    - **-query path:EI14/<project>.<runNumber>...** (not good way)
    - **-query guid:<guid>**
  - They can be moved somewhere else
- Problem: When a user searches the whole set (-query path:EI14), she gets over 40000 TagFiles, which is not good (at all) for Hadoop - remember: Hadoop likes small number of big files
- Proposal: Merge TagFiles and move dir/filename fields into TagFile schema, so we will have three copies of data (depending on usage and performance, we may later decide to delete some):
  - **EI14.2**/<project>.<runNumber>.<streamName>.<prodStep>.<dataType>.<version>/<guid>_<transid>_<consumerid>_<pandataskid>_<pandauserid>
  - **EI14.1**/<project>.<runNumber>.<streamName>.<prodStep>.<dataType>.<version>
  - **EI14.0**
  - Depending of query granularity, request may go to appropriate copy (transparent to user)
  - Files will be bigger, but compression should be very efficient because we will add just repeated strings
- HowTo proceed:
  - I will convert existing (EI14, EI12) data
  - The future data from Consumer will contain extended Schema (all data should have the same Schema), will go into many small files and will be (re)assembled as needed
    - once partitioned and global files have the same Schema, they can be used together

Data Organisation - Schema

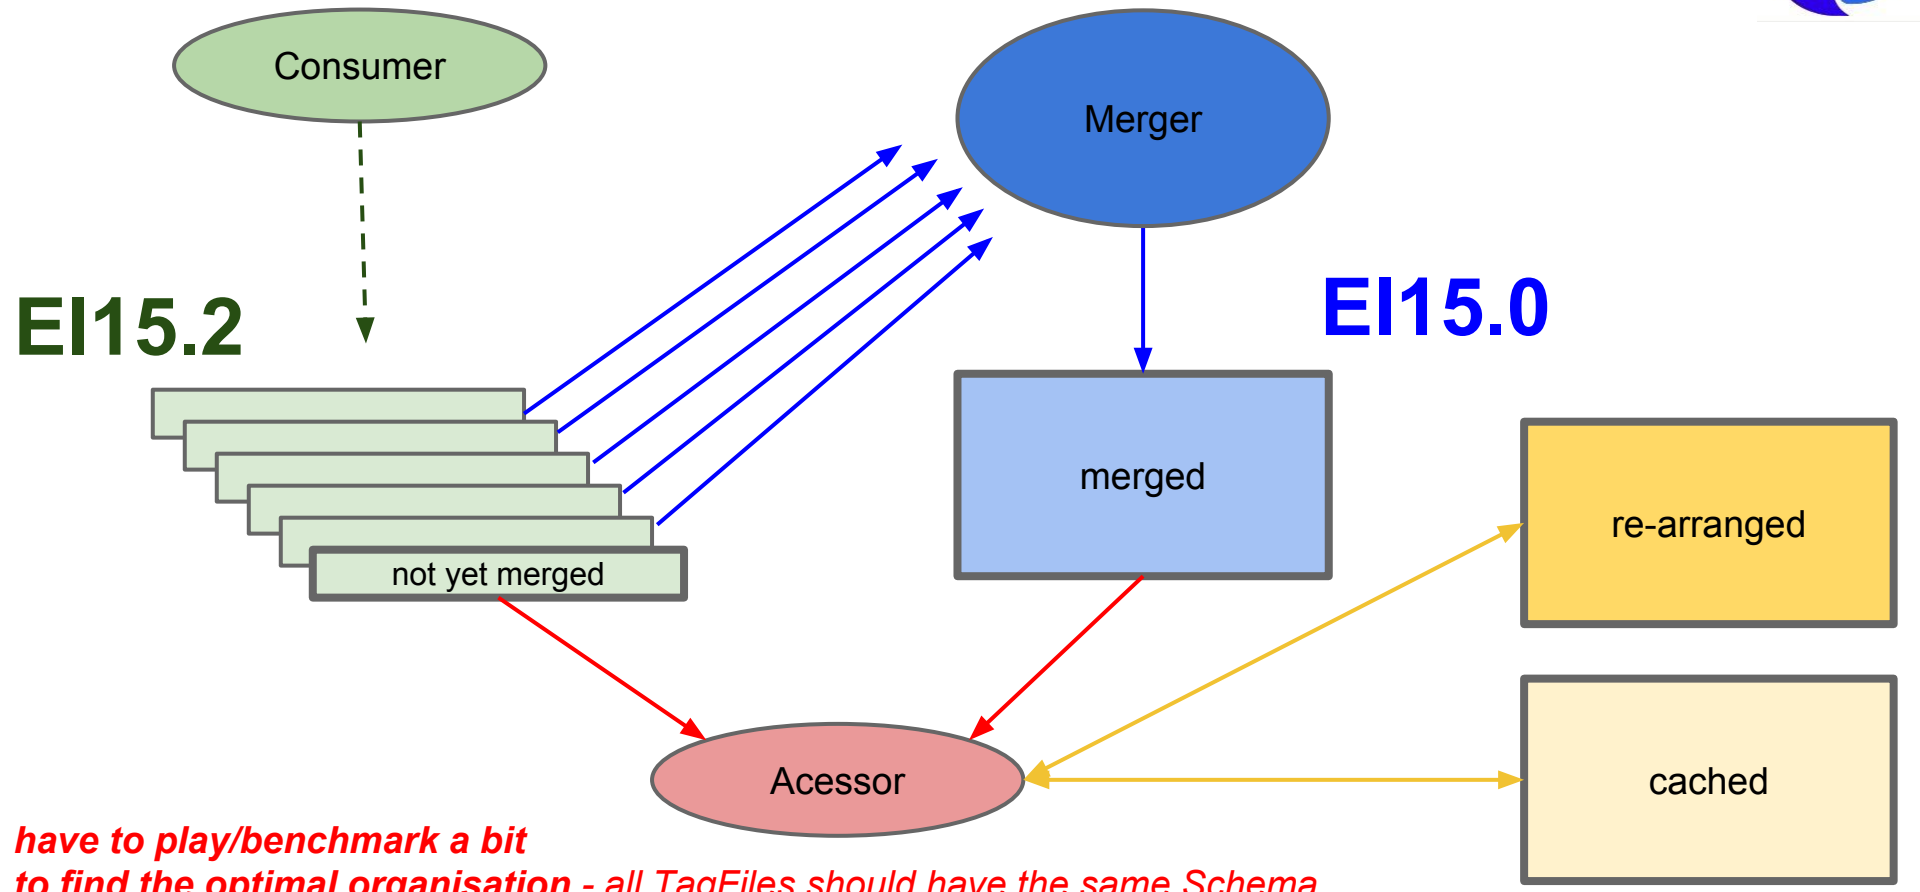

have to play/benchmark a bit
to find the optimal organisation *- all TagFiles should have the same Schema*

Other News

- User can specify directory for results (otherwise unique name is generated) via **-outname**
  - Useful for long-living cache
- Web Service writes big results into local disc, following **-filter** option

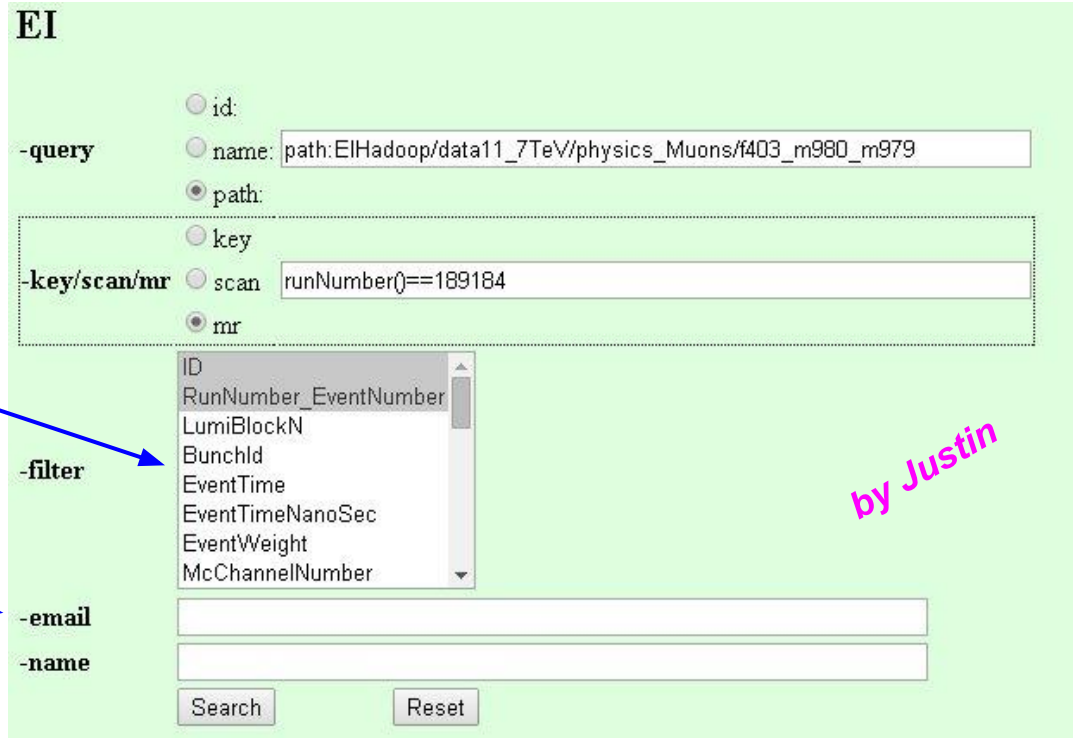

to send notification when job finishes

to have reasonable name in Bookmark, also used as -outname

Going-on & Problems

- **runEventLookup.py** will be updated once definitive directory structure is in place
- Big Catalog result fails via Web Service - will be fixed, but such requests are useless anyway
- CERN J2EE service upgrade - will allow authentication via e-groups (but happened several hours before tutorial and caused problems)
- Some imports have failed (but automatic re-try was successful) - that was due to multiple tasks trying to import the same data, fixed by locking
- Currently importing by copying, will use moving once everything is solid
  - But not a big deal: copying of whole 2014 data takes less than two hours
- Obsolete code removed

Links

- Home:
  - http://atlas-event-index.cern.ch/doc *(atlas-event-index = aiatlas016)*
  - http://cern.ch/hrivnac/Activities/Packages/TagConvertor
- Web Service:
  - http://atlas-event-index.cern.ch/EIHadoop
  - ssh -L 8080:aiatlas016.cern.ch:80
  - java -jar EIHadoopCatalog.exe.jar *for remote Catalog CLI*
  - java -jar EIHadoopEI.exe.jar *for remote Event Index CLI*
  - java -jar EventServer.exe.jar *for remote Event Server CLI (not yet in CVMFS)*
  - java -jar EventLookup.exe.jar *for remote Event Lookup CLI (not yet CVMFS)*
- SVN:
  - svn+ssh://svn.cern.ch/reps/atlasoff/Database/TAGHadoop/TagConvertor
- Ant targets:
  - ant -p *to get help on available targets*