




Data Store and Query System

- *Data Organisation*
- *Import Status*
- *Data Verification & Correction*
- *Import Status*
- *Servers & Services*
- *Web GUI*
- *Journal*
- *Problems & Tasks*

latest version: 1.17.4

*J. Hrivnac, ET WS, 22 June 15, Valencia
+ Justin, Fedor, Rainer, Andrea,...*

- *XX.0, XX.ES not created because too slow*
- *may be reactivated later*



- re-sort in GUID to provide fast EventService
- executed immediately after import for Run2
 - processing Run1 in parallel

XX.2/<project>.<runNumber>.<streamName>.<prodStep>.<dataType>.<version>/<guid> <transid> <consumerid> <pandataskid> <pandauserid>

Import Status



IMPORTED		IMPORTED	
=====		=====	
EI 2009 2	6471	MC 2009 2	0
EI 2009 1	525	MC 2009 1	0
EI 2009 1ES	525	MC 2009 1ES	0
EI 2010 2	173296	MC 2010 2	0
EI 2010 1	3798	MC 2010 1	0
EI 2010 1ES	1157	MC 2010 1ES	0
EI 2011 2	267518	MC 2011 2	0
EI 2011 1	3642	MC 2011 1	0
EI 2011 1ES	880	MC 2011 1ES	0
EI 2012 2	291021	MC 2012 2	40
EI 2012 1	4398	MC 2012 1	2
EI 2012 1ES	867	MC 2012 1ES	0
EI 2013 2	184553	MC 2013 2	0
EI 2013 1	523	MC 2013 1	0
EI 2013 1ES	525	MC 2013 1ES	0
EI 2014 2	32168	MC 2014 2	0
EI 2014 1	492	MC 2014 1	0
EI 2014 1ES	492	MC 2014 1ES	0
EI 2015 2	87783	MC 2015 2	0
EI 2015 1	2078	MC 2015 1	0
EI 2015 1ES	2078	MC 2015 1ES	0

done TagFiles

```
WAITING
=====
EI 2009 0
EI 2010 728
EI 2011 347
EI 2012 0
EI 2013 0
EI 2014 0
EI 2015 0
MC 2009 0
MC 2010 0
MC 2011 0
MC 2012 0
MC 2013 0
MC 2014 0
MC 2015 0
```

waiting conf files

- Run 2 data are imported and checked immediately (within hours after Consumer VALIDation)
- Import of remaining Run 1 almost stopped because of HBase timeout problem
 - 10s of 10000 TagSets imported last week
 - after full import, data will have to be verified and corrected
 - it would be nice to have a complete & final list of correctly consumed data to check imported data with

Data Verification & Correction



- Some problems during import:
 - Some TagFiles should be re-imported because modified/corrected
 - HBase timeouts (very frequent)
 - HDFS timeouts (rare)
 - AFS ticket expires
 - Machine reboots
 - ...
- Not fatal, but create important delays
 - detailed verification is slow

*other checks will be added
(like comparing *.1 with *.1ES)*

```
1) verify internal Catalog consistency and existence of referenced TagFiles:
$ catalog -query <path> -verify
fix:
$ catalog -query <path> -verify fix
3) find non-cataloged files:
$ verifactor files [MC|EI] [09|10|11|12|13|14|15|all]
fix: copy conf file from *_OLD
$ importer <conf file> <dest dir> <src prefix> <TagFiles to be re-imported>
re-run import to re-create *.1 files
3) verify that all files were imported:
compareImport.sh <conf file> <dest dir>
fix: like 2)
4) verify that all TagFiles have the same key and schema:
$ verifactor schema [MC|EI] [09|10|11|12|13|14|15|all]
5) verify that TagFiles are correctly merged:
$ verifactor merge [MC|EI] [09|10|11|12|13|14|15|all]
6) verify that merged TagFiles have the correct number of lines (slow):
$ verifactor mergelines [MC|EI] [09|10|11|12|13|14|15|all]
7) set *.1 attributes:
$ verifactor attributes [MC|EI] [09|10|11|12|13|14|15|all]
8) verify Catalog entry of all TagFiles covered by the import config files (slow):
$ importer <conf file> <dest dir> <src file prefix> verif
```

Servers & Services



CLI

Simplified

Web GUI

local commands:

\$ catalog

\$ ei

\$ inspector

\$ importer

\$ verifier

AFS
CVMFS
TagConverter.tar.gz

\$ catalog

\$ ei

\$ runEI.py

\$ EventLookup.exe.jar

\$ runEL.py

\$ EventServer.exe.jar

\$ runES.py

```
EventLookup.exe.jar [-e <runs-events>] [-f <runs-events filename>] [-s <stream name>] [-p <amitag>] [<ei  
args>]  
EventServer.exe.jar [-guid <guid>] [<ei args>]
```

atlas-event-index

EIHadoop

atlas-service-EIES

EI

~~ES~~

~~EP~~

Tomcat
visible in CERN

J2EE
visible WorldWide

CLI
anywhere



selections derived automatically (at build time) from data

Event Index

- [Catalog](#)
- [Event Index \(Expert Mode\)](#)
- [Bookmarks](#)
- [System Journal \(for admins\)](#)



EI

-legend

Year	Projets	run Number	Stream Name	Data Step	Data Type	Version
E115.1 ▼	data15_13TeV ▼	00267069	physics_Main ▼	merge ▼	AOD ▼	594_m143

-query

☒ key

-key/scan/mr

☐ scan

☐ mr

00267069-00000008169 00267069-00000013077

-filter

ID
RunNumber_EventNumber
LumiBlockN
BunchId
EventTime
EventTimeNanoSec
EventWeight
McChannelNumber

-email

-name

-info

Search

Reset

Journal

```
new Record(id, action, args, rc, time, id_before, id_after, comment).register();
```



Event Index



- [Catalog](#)
- [Event Index \(Expert Mode\)](#)
- [Bookmarks](#)
- [System Journal \(for admins\)](#)

in HBase

Journal

end
period *in hours*

*more selective searching
can be added*

*anyone can register in Journal
so far in Journal:*

- *import*
- *merging*
- *requests via WebService*

Entries between -0.5 and 0.0:

```
x atlevind/2015.06.21@23.49.16.494: rc = 0, time = 12, action = {-outname=EI11.1ES/data11_7TeV.001
x atlevind/2015.06.21@23.50.21.806: rc = 0, time = 23, action = {-outname=EI10.1ES/data10_7TeV.001
x atlevind/2015.06.21@23.50.38.848: rc = 0, time = 13, action = {-outname=EI12.1ES/data12_8TeV.002
x atlevind/2015.06.21@23.50.49.048: rc = 0, time = 9, action = {-outname=EI11.1ES/data11_7TeV.0018
x atlevind/2015.06.21@23.52.35.354: rc = 0, time = 10, action = {-outname=EI11.1ES/data11_7TeV.001
x atlevind/2015.06.21@23.52.37.703: rc = 0, time = 14, action = {-outname=EI12.1ES/data12_8TeV.002
x atlevind/2015.06.21@23.53.06.973: rc = 0, time = 23, action = {-outname=EI10.1ES/data10_7TeV.001
x atlevind/2015.06.21@23.53.59.237: rc = 0, time = 21, action = {-outname=EI11.1ES/data11_7TeV.001
x atlevind/2015.06.21@23.54.15.489: rc = 0, time = 18, action = {-outname=EI12.1ES/data12_8TeV.002
x atlevind/2015.06.21@23.54.43.969: rc = 0, time = 11, action = {-outname=EI10.1ES/data10_7TeV.001
x atlevind/2015.06.21@23.55.45.915: rc = 0, time = 13, action = {-outname=EI11.1ES/data11_7TeV.001
x atlevind/2015.06.21@23.56.00.787: rc = 0, time = 14, action = {-outname=EI10.1ES/data10_7TeV.001
x atlevind/2015.06.21@23.57.14.038: rc = 0, time = 6, action = {-outname=EI11.1ES/data11_7TeV.0018
x atlevind/2015.06.21@23.57.26.068: rc = 0, time = 14, action = {-outname=EI10.1ES/data10_7TeV.001
```

Problems & Tasks



- **Serious** problem with HBase timeout
 - some days, almost all imports have failed !
 - hurts also queries
- Merging to *.0 and *.ES is very slow (days)
 - HDFS is fast in reading, but slow in writing !
 - reading profits from parallelism
 - writing suffers from parallelism
 - as it sorts billions of entries (by a long key - guid)
 - growing probability of a problem during the job
 - gave up so far, accessing *.1 and *.1ES
 - will try to improve the sorting algorithm
 - I think: we have the same problem in HBase
 - but HBase clients are not isolated, so one heavy client spoils others), unlike direct HDFS
- HBase Catalog is backedup (nightly), but what about the data ?
 - one stupid command can destroy everything !!!
 - should find a way to “lock” data
- Should think about multiple servers (fallback, load balancing)

- Home:
 - <http://atlas-event-index.cern.ch/doc> (*atlas-event-index = aiatlas016*)
 - <http://cern.ch/hrivnac/Activities/Packages/TagConvertor>
- CLI:
 - **catalog, ei, inspector, importer, verifactor**
- 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*
 - `java -jar EventLookup.exe.jar` *for remote Event Lookup CLI*
 - `runEI.py, runES.py, runEL.py`
- SVN:
 - `svn+ssh://svn.cern.ch/repos/atlasoff/Database/TAGHadoop/TagConvertor`
- Ant targets:
 - `ant -p` *to get help on available targets*