

# *Event Index Core - Quick Summary* *(Hadoop & HBase)*



- Import & Data Flow
- Commands
- Performance
- Journal
- Missing Pieces & Planned Improvements
- Info
- **What's New**
- **To Do**

*Andrea Favareto (testing)*

*Julius Hrivnac*

*Fedor Prokoshin (trigger)*

*Rainer Toebicke (infrastructure, HBase EL)*

*Ruijun Yuan (graphics)*

*EI WS, 10 May 2016*



# Import

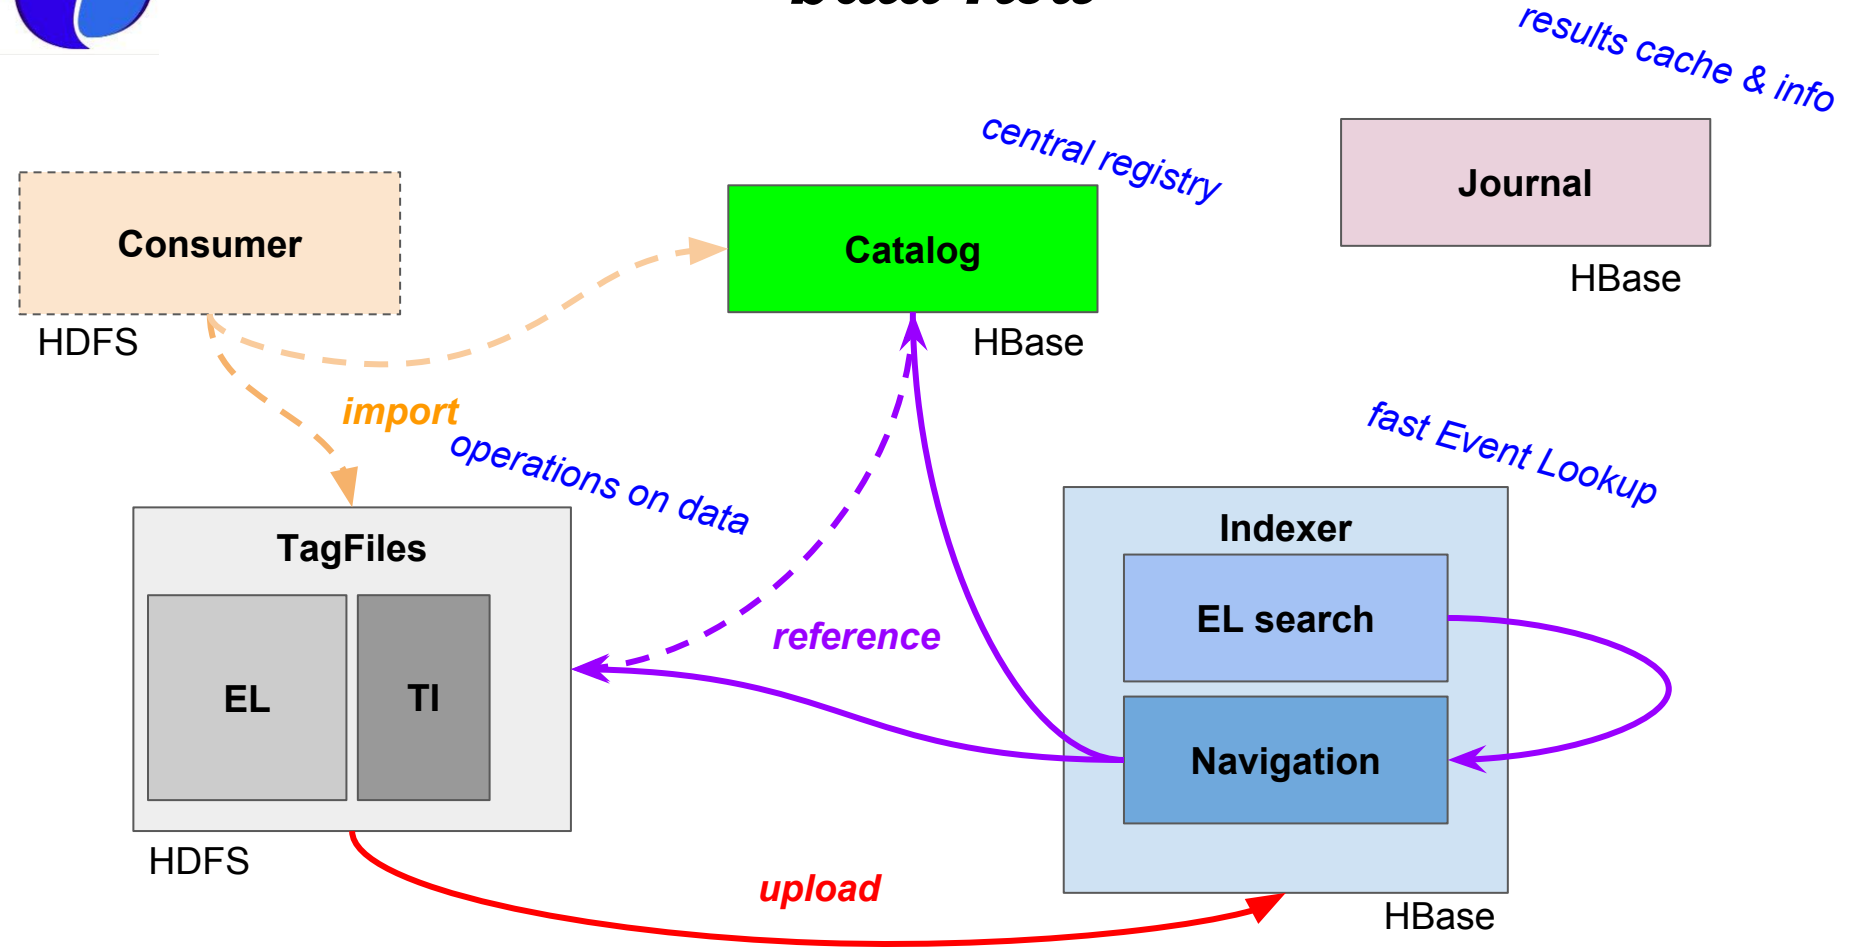
All imported events:  
EI-AOD: 14 670 609 118  
EI-ESD: 716 471 588  
EI-DAOD: 10 086 567 398  
EI-DESD: 164 011 085  
MC-AOD: 9 978 796 397  
MC-EVNT: 4 997 862 966  
all: 47 304 982 753

Data occupy about 70 TB,  
which can be compressed 5x

- Runs smoothly
  - Single EI dataset import takes minutes and happens within one hour
    - New validated conf files checked every 20 min
  - Some MC datasets take longer as they can be very big and come in big chunks
  - Parallelism: up to 10+10 imports on aiatlas016 (EI) + aiatlas054 (MC)
  - The “full speed” imports about 1000 datasets per day
- No further development
  - Will be re-factored with Consumer
- Upload to HBase not yet fully automated
- System handles disabling, removing and replacing of bad data
- All data are backuped, archived, can be re-created or are not essential



# Data Flow





# Commands

- All non-local commands pass via Web Service (load-balanced & fallback aiatlas016,054)
- Some (destructive) options are only available locally
- All commands accept -email option, which triggers async. execution
- Different command incarnations keep the same options (if possible)
- Further functionality (like duplicate events detection & correction, statistical gathering,...) is available from **ei** command - it can be made available as more user-friendly single-purpose commands if requested

available

available / undocumented

can be quickly implemented

<u>Command</u>	<u>Local &amp; API</u>	<u>Web GUI</u>	<u>Remote Full CLI</u>	<u>Direct URL</u>	<u>Native Python</u>	<u>NOTE</u>
Catalog						
Event Index						<i>rich functionality</i>
Event Lookup						<i>fast</i>
Trigger Info						<i>new</i>
Inspector						
Importer						
Journal						<i>access to all results</i>
<b>NOTE</b>			<i>richest</i>	<i>fastest</i>		



# Event Lookup - Nightly Benchmarks - Remote Full CLI

- Single database request takes x10ms
- Single remote request with Full CLI takes < 1s
- Simultaneous 10x access from the same Ixplus machine (asking for the same data or different data) takes 6-7s for 10000 ev. sample
- There is still room for at least 2x speed up
- Direct URL and Native Python clients are faster

*just a snapshot of nightly tests & benchmarks*

```
sample: TagConvertor/test/data/EL/ev10000.txt
real    0m32.711s
sample: TagConvertor/test/data/EL/ev1000.txt
real    0m1.559s
sample: TagConvertor/test/data/EL/ev100.txt
real    0m1.168s
sample: TagConvertor/test/data/EL/ev10.txt
real    0m1.103s
sample: TagConvertor/test/data/EL/ev_1_1000.txt
real    0m1.544s
sample: TagConvertor/test/data/EL/ev1.txt
real    0m0.938s
sample: TagConvertor/test/data/EL/ev_2_1000.txt
real    0m2.011s
sample: TagConvertor/test/data/EL/ev_3_1000.txt
real    0m2.357s
sample: TagConvertor/test/data/EL/ev_4_1000.txt
real    0m2.433s
sample: TagConvertor/test/data/EL/ev5000.txt
real    0m2.838s
sample: TagConvertor/test/data/EL/ev500.txt
real    0m1.553s
sample: TagConvertor/test/data/EL/ev50.txt
real    0m1.100s
sample: TagConvertor/test/data/EL/ev_5_1000.txt
real    0m1.933s
sample: TagConvertor/test/data/EL/ev_6_1000.txt
real    0m2.064s
sample: TagConvertor/test/data/EL/ev_7_1000.txt
real    0m2.572s
sample: TagConvertor/test/data/EL/ev_8_1000.txt
real    0m1.950s
sample: TagConvertor/test/data/EL/picklist_full.txt
real    0m3.694s
```

*The first request after period of sleep is always slower because the system must wake up.*

*second slowest (8879 events)*



# Event Index

- The most generic access
- Two versions with different richness and user-friendliness
- Allows, in principle, to run any job against stored data

## Event Index

- [Global Help](#)
- [Catalog](#)
- [Event Index \(Expert Mode\)](#)
- [Event Lookup](#)
- [Trigger Info](#)
- [Bookmarks](#)
- [System Journal \(for admins\)](#)



## Event Index

-legend	Year	Projets	Stream Name	Prod Step	Data Type	Version	Run Number
-query	EI15.1	data15_13TeV	physics_Main	merge	AOD	f594_m143f	00267069
-key/mr	<input type="radio"/> key <input type="radio"/> mr 00267069-00000008169 00267069-00000013077						
-filter	RunNumber_EventNumber LumiBlockN BunchId EventTime EventTimeNanoSec EventWeight McChannelNumber Lvl1ID						

```
-limit 0
-outname EIWSCache/2016.05.09.16.11.54.429/tomcat
-query id:EI15.1.data15_13TeV.physics_Main.merge.AOD.f594_m143f.00267069 status:good
-name 2016.05.09.16.11.54.429/tomcat
-show 20
-filter BunchId EventTime Lvl1ID
-key 00267069-00000008169 00267069-00000013077
```

```
2 events found, result 2016.05.09.16.12.00.239/tomcat stored in EIWSCache/2016.05.09.16.11.54.429/tomcat, 5s spend
SimpleResult(KeyQuery(00267069-00000008169..00267069-00000008169 00267069-00000013077..00267069-00000013077 , null, null)):
>>>
BunchId = 2757
EventTime = 1433423693
Lvl1ID = 16785385
>>>
BunchId = 2340
EventTime = 1433423695
Lvl1ID = 33557215
2 results found
```



# Event Lookup

- Several apis, HBase based index is the fastest and gives access to all data
- Physical Containers selection under way

## Event Index

- [Global Help](#)
- [Catalog](#)
- [Event Index \(Expert Mode\)](#)
- [Event Lookup](#)
- [Trigger Info](#)
- [Bookmarks](#)
- [System Journal \(for admins\)](#)



## Event Lookup

List of 'runnumber evtnumber' (-e)

AMI tag (-p) (substring match)

Stream name (-s)

Data type (-d)

GUID type (-t)  AOD  ESD  RAW  all

-api  simple  rich  indexer  mc (indexer)

-details  event  type  id  dataset  rich output

-email (implies asynchronous execution)

```
-e 00278880 558085589, 00278880 210318172
```

```
-details typenulldataseteventnull
```

```
-p all
```

```
-api indexer
```

```
12 guids found for 1 runs with 2 events, 0 guids missing, 0s spend
00278880 00210318172 6E501FA2-174B-3548-A890-CC381DF11538 StreamESD data15_13TeV.00278880.physics_Main.merge.AOD.f628_m1497
00278880 00210318172 8E053F00-EBAE-9E4F-8425-2BF57C68E368 StreamESD data15_13TeV.00278880.physics_Main.merge.AOD.r7562_p2521
00278880 00210318172 AAB98A83-7374-694C-A44D-38FD1F285832 StreamAOD data15_13TeV.00278880.physics_Main.merge.AOD.f628_m1497
00278880 00210318172 B7FBAFCA-206B-8F44-A08B-371CEAEC1EA7 StreamAOD data15_13TeV.00278880.physics_Main.merge.AOD.r7562_p2521
00278880 00210318172 FC28C1E6-1D56-E511-AEDE-44A8420A8576 StreamRAW data15_13TeV.00278880.physics_Main.merge.AOD.f628_m1497
00278880 00210318172 FC28C1E6-1D56-E511-AEDE-44A8420A8576 StreamRAW data15_13TeV.00278880.physics_Main.merge.AOD.r7562_p2521
00278880 00558085589 75915627-0977-D648-8C77-D1B690F086C7 StreamESD data15_13TeV.00278880.physics_Main.merge.AOD.r7562_p2521
00278880 00558085589 886C7D3A-3E56-E511-9D00-44A8420A5EB7 StreamRAW data15_13TeV.00278880.physics_Main.merge.AOD.f628_m1497
00278880 00558085589 886C7D3A-3E56-E511-9D00-44A8420A5EB7 StreamRAW data15_13TeV.00278880.physics_Main.merge.AOD.r7562_p2521
00278880 00558085589 C1AD55A4-19E9-C649-9832-DB0AB2D0EACB StreamAOD data15_13TeV.00278880.physics_Main.merge.AOD.r7562_p2521
00278880 00558085589 D3D51279-40A3-1E40-B4DA-E0EA4D6BE21F StreamESD data15_13TeV.00278880.physics_Main.merge.AOD.f628_m1497
00278880 00558085589 E6CC54DB-32FB-054E-AC05-C3CEDA07550C StreamAOD data15_13TeV.00278880.physics_Main.merge.AOD.f628_m1497
```



# Trigger Info

- Sometimes very long jobs, async mode is useful
- Graphical output on the way

## Event Index

- [Global Help](#)
- [Catalog](#)
- [Event Index \(Expert Mode\)](#)
- [Event Lookup](#)
- [Trigger Info](#)
- [Bookmarks](#)
- [System Journal \(for admins\)](#)

## Trigger Info

-legend	Year	Projets	Stream Name	Prod Step	Data Type	Version	Run Number
-query	E115.1	data15_13TeV	physics_Main	merge	AOD	f641_m151	00284213
-tlevel	<input type="radio"/> L1 <input type="radio"/> L2 <input checked="" type="radio"/> EF or HLT <input type="radio"/> all						
-tdecision	<input checked="" type="radio"/> L1 AfterVeto or L2/EF/HLT Physics <input type="radio"/> L1 AfterPrescale or L2/EF/HLT PassThrough <input type="radio"/> L1 BeforePrescale or L2/EF/HLT Resurrected						
-toverlap	HLT_e120_loose,HLT_2e17_vloose,HLT_2g50_loose,HLT_2j55_bloose_L14J20.0ETA4 (disables -tlevel and -tdecision)						
-email	<input type="text"/> (implies asynchronous execution)						
<input type="button" value="Search"/> <input type="button" value="Reset"/>							

```
-tdecision 0
-query id:E115.1.data15_13TeV.physics_Main.merge.AOD.f641_m1511.00284213
-tlevel 3
-toverlap HLT_e120_loose,HLT_2e17_vloose,HLT_2g50_loose,HLT_2j55_bloose_L14J20.0ETA49,HLT_g120_loose
```

```
Result 2016.05.09.16.48.31.399/tomcat stored in EICache/2016.05.09.16.48.31.399/tomcat, 1350s spend
Processed 126364190 events
Found 10 overlaps of trigger chains
HLT_e120_loose:HLT_g120_loose,107482 (136718:883752)
HLT_2g50_loose:HLT_g120_loose,17743 (71274:883752)
HLT_2j55_bloose_L14J20.0ETA49:HLT_g120_loose,7010 (829505:883752)
HLT_2e17_vloose:HLT_2g50_loose,4624 (196517:71274)
HLT_e120_loose:HLT_2g50_loose,3299 (136718:71274)
HLT_e120_loose:HLT_2e17_vloose,2286 (136718:196517)
HLT_2e17_vloose:HLT_g120_loose,2120 (196517:883752)
HLT_e120_loose:HLT_2j55_bloose_L14J20.0ETA49,1376 (136718:829505)
HLT_2e17_vloose:HLT_2j55_bloose_L14J20.0ETA49,628 (196517:829505)
HLT_2g50_loose:HLT_2j55_bloose_L14J20.0ETA49,589 (71274:829505)
!Trigger info not found,63324101
```

Progress map : 100%





# Journal

- All results, incl. timing and eventual errors are available
- Useful for async operations and problem hunting
- Could be used as a cache

### Event Index

- [Global Help](#)
- [Catalog](#)
- [Event Index \(Expert Mode\)](#)
- [Event Lookup](#)
- [Trigger Info](#)
- [Bookmarks](#)
- [System Journal \(for admins\)](#)

### Journal

start: 16:12:26 09/May/2016  
stop: 17:12:26 09/May/2016  
periodic refresh:

actor:  Catalog (modifying)  Event Index  Eventy Lookup  Trigger Info  Importer  \*

rc:  0  -1  \*

Entries of \* with rc = \* between 16:12:26 09/May/2016 and 17:12:26 09/May/2016:

actor ▲▼	rc ▲▼	time ▲▼	timestamp ▲	action ▲▼	result ▲▼
TICLI	0	1350	17:11:01.266 09/May/2016	{-show=0, -aux=net hep.atlas.Database.EIHadoop.Acc... more	Info: Result 20
Importer	0	339	17:09:41.078 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
TICLI	0	3102	17:09:26.308 09/May/2016	{-show=0, -aux=net hep.atlas.Database.EIHadoop.Acc... more	Info: Result 20
Importer	0	326	17:08:53.633 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
CatalogCLI	0	89	17:08:39.932 09/May/2016	{-modify=projectmc15_13TeV.runNumber428359.strea... more	Info: 89s spent
Importer	0	380	17:07:46.592 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	361	17:07:45.688 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	372	17:07:38.330 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	374	17:07:10.418 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
EICLI	0	550	17:07:07.056 09/May/2016	{-outname=MC15_1/mc15_13TeV.428359.ParticleGun_pid... more	Info: 1000000 events found, result MC15.1.mc15_13T... more
Importer	0	13	17:04:31.778 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	20	17:04:18.429 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	396	17:04:01.431 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	18	17:03:57.368 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	131	17:03:37.124 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	375	17:03:26.675 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	3	17:02:17.490 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20
Importer	0	5	17:02:13.626 09/May/2016	/afs/cern.ch/user/a/atlevind/work/DatasetsConf/mc1... more	Info: Result 20

```
Info: Result 2016.05.09.16.48.31.399/tomcat stored in EICache
Result:
Processed 126364190 events
Found 10 overlaps of trigger chains
HLT_e120_loose:HLT_g120_loose,107482 (136718:883752)
HLT_2g50_loose:HLT_g120_loose,17743 (71274:883752)
HLT_2j55_bloose_L14120.0ETA49:HLT_g120_loose,7010 (8294:196517)
HLT_2e17_vloose:HLT_2g50_loose,4624 (196517:71274)
HLT_e120_loose:HLT_2g50_loose,3299 (136718:71274)
HLT_e120_loose:HLT_2e17_vloose,2286 (136718:196517)
HLT_2e17_vloose:HLT_g120_loose,2120 (196517:883752)
HLT_e120_loose:HLT_2j55_bloose_L14120.0ETA49,1376 (136718:883752)
HLT_2e17_vloose:HLT_2j55_bloose_L14120.0ETA49,628 (196517:883752)
HLT_2g50_loose:HLT_2j55_bloose_L14120.0ETA49,569 (71274:883752)
!Trigger info not found,63324101
```



# *Missing Pieces*

- Full set of commands with Shibboleth protection
  - Currently only Graphical GUI and direct URL are protected and available outside CERN
- More interactive & graphical GUI
- Event Lookup with Physics Containers
- Better Event Lookup for MC
- Data compression on HDFS: 5x, lazy transition
- More automated data flow & connection with AMI
- New / better data formats
- ...

*Should prioritise due to limited/disperse manpower.*

**General strategy:**

1. *Functionality*
2. *Performance (speed)*
3. *Size*



# Info

**Web Service:** <https://atlas-event-index.cern.ch/ElHadoop>

**Documentation & Distribution:** <https://atlas-event-index.cern.ch/doc>

**Twiki:** <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/EventIndex>

**Twiki Tutorial:** <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/EventIndexTutorial>

**Sources:** `svn+ssh://svn.cern.ch/repos/atlasoff/Database/TAGHadoop/TagConvertor`

**AFS:** `/afs/cern.ch/sw/lcg/external/Java/TagConvertor`

**CVMFS:** `$ lsetup eiclient`



# *What's New*

- Catalog query supports NOT argument
  - ... -query 'id:xxx status:~bad ...'
- SimpleCatalog.java added
- ELdirect.jsp added
  - Like EL.jsp, but without decorations
  - Will make also Tldirect.jsp
- el can give dataset name
  - ... -details dataset ...
- Emailing results generalised to all commands
  - ... -email aaa@bbb.ccc
- Remote command with email will be run asynchronously
  - Results by email and in Journal
- Event Lookup works for MC (with -api mc)
  - Should be improved
- Python client to Event Lookup
  - Will make also for Trigger Info



# To Do

- Direct URL & Python client for TI
- Duplicate detection
  - Who does it ?
- Shibboleth for Java & Python clients
  - Precondition for new CVMFS release
  - Help ! (in contact with AMI & ConditionsDB)
- Physical containers
  - TagSets @ Catalog (ready)
  - Include in HBase (Rainer)
  - Fill (Fedor)
- EL defaults (the same for all implementations)
  - If no ami-tag or ph.container & not asking for RAW => the highest ph.container, then no result (with error message)
- MC in HBase
  - And in API/CLI
- Verify import/upload/removal/correction data flow



# URL & Shibboleth

```
### in CERN:
curl --data e=00278880+558085589 http://aiatlas016.cern.ch/EIHadoop/ELdirect.jsp
curl --data e=`cat MyEventFile.txt | sed ':a;N;$!ba;s/\n/%2C/g'| tr ' ' '+'` http://aiatlas016.cern.ch/EIHadoop/ELdirect.jsp
### on any SLC machine:
# first get a kerberos cookie:
cern-get-sso-cookie --krb -r -u https://atlas-event-index.cern.ch -o ~/private/ssocookie.txt
# or personal cookie:
# get myCert.p12 from http://cern.ch/ca
openssl pkcs12 -clcerts -nokeys -in myCert.p12 -out ~/private/myCert.pem
openssl pkcs12 -nocerts -in myCert.p12 -out ~/private/myCert.tmp.key
openssl rsa -in ~/private/myCert.tmp.key -out ~/private/myCert.key
rm ~/private/myCert.tmp.key
chmod 644 ~/private/myCert.pem
chmod 400 ~/private/myCert.key
cern-get-sso-cookie --cert ~/private/myCert.pem --key ~/private/myCert.key -r -u https://atlas-event-index.cern.ch -o ~/private/ssocookie.txt
# then access the service:
curl -L -k --cookie ~/private/ssocookie.txt --cookie-jar ~/private/ssocookie.txt --data e=00278880+558085589
https://atlas-event-index.cern.ch/EIHadoop/ELdirect.jsp
curl -L -k --cookie ~/private/ssocookie.txt --cookie-jar ~/private/ssocookie.txt --data e=`cat MyEventFile.txt | sed ':a;N;$!ba;s/\n/%2C/g'| tr ' ' '+'`
https://atlas-event-index.cern.ch/EIHadoop/ELdirect.jsp
```