



Open-ended production

Andrea Favareto (INFN-Genova)

Fedor Prokshin (FSTU)

Santiago González de la Hoz (IFIC)

(EIPM team)

Event Index workshop at IFIC-Valencia

22nd-23rd June 2015



Goal

- Implement automatic system to launch Grid tasks using ProdSys2 (PS2) to index the relevant datasets as soon as they appear in the AMI catalogue
- Agreements:
 - Collect from AMI the list of datasets to be indexed that were completed (“all events available”) the previous day.
 - Save locally the list of datasets and metadata
 - Number of files
 - Number of events
 -
 - Launch PS2 tasks to index them
 - Check the completion consistency
 - Provide input to monitoring of data/dataset/job flow, success and error, etc.



Dataset selection from AMI

(see Fedor's slides)

- Information collected by python script using PyAMI and Rucio APIs
Selection parameters:
 - from AMI:
 - project: mc15% or data15%
 - dataType= EVNT or AOD or DAOD
 - prodsysStatus= "ALL EVENTS AVAILABLE"
 - lastModificationTime = (yesterday) 2015-06-14 00:00:00 - 2015-06-15 00:00:00
 - from Rucio:
 - scope: mc15 or data15
 - dataset name (from AMI)
 - Intended to run as a cron job under “atlevind” account. We had some problems with VOMS registration of it; with RUCIO rights to attach TIDS to technical containers. Now SOLVED!!
 - Need setup atlas software on aiatlas016 (ask central service to install cvmfs)
Better to run cron job here than lxplus, because proxys goes to /tmp, so will be different from different nodes. Ask request to atlas-adc-central-services through a Jira ticket : <https://its.cern.ch/jira/browse/CSOPS-725>
- Retrieve for each dataset:
 - Dataset name, Number of files, Number of events, List of files.
- **Store this info in EventIndex space**
 - **Hbase looks like the simplest solution**



AMI TAG

- We created the AMI tag for the EventIndexProducer transformation POOLtoEI_tf.py
 - AtlasProduction 20.1.5.4 and later contain EventIndexProducer-00-02-09 as AMI tag i2
- go to https://ami.in2p3.fr/new/?subapp=amiTags_add
 - AMI-Tag type: select "i"
 - Select the proper ATLAS release. E.g. AtlasProduction_20.1.5.4
 - Transform name is POOLtoEI_tf.py
 - click "Add" button. A pop-up window will appear asking for a comment, then you will see a green box at the top of the page telling that a new AMI Tag is inserted.

Input

`inputAODFile` `inputESDFile` `inputEVNTFile` `inputHITSFile` `inputPOOLFile` `inputRD0File`

Output

`outputEIFile`



Creation Technical containers

- We created 5 containers to be used for the open-ended production
 - One for each type of data to be indexed (MC EVNT, MC AOD, MC DAOD, Data AOD and Data DAOD)
- Add the list of “new” datasets to the right technical container every day
 - With Fedor’s script
 - Now tested it by hand



Creation Technical containers

- A VOMS group was created “group.proj-evind”, attaching the necessary people who will need to manage the containers.
- We created with rucio 5 technical containers:
 - group.proj-evind:mc15.EventIndex.OpenEnded.EVNT.i2
 - group.proj-evind:mc15.EventIndex.OpenEnded.AOD.i2
 - group.proj-evind:mc15.EventIndex.OpenEnded.DAOD.i2
 - group.proj-evind:data15.EventIndex.OpenEnded.AOD.i2
 - group.proj-evind:data15.EventIndex.OpenEnded.DAOD.i2
- For instance,
 - rucio add-container group.proj-evind:mc15.EventIndex.OpenEnded.DAOD.i2 (linked to AMI TAG i2)
- We attached by hand TIDS to those containers
 - rucio attach group.proj-evind:mc15.EventIndex.OpenEnded.EVNT.i2 mc15_13TeV.361102.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Wplustaunu.evgen.EVNT.e3601_tid04972712_00
- We went to AMI and we selected these criteria:
 - Simulated Data, mc15 and dataType EVNT
 - we got the selection (Selected datasets: 2544 (events: 2133266608 , files: 1687068))

If new AMI TAG is created we must set up new technical containers

PS2 defines tasks to run EI jobs



- PS2 is creating a task for each dataset that is added to the technical containers automatically
 - This is done every ~2 hours (run the job when new TIDs are attached to the containers).
- PS2 people created an interface for us:
 - https://prodtask-dev.cern.ch/prodtask/eventindex_request_create/
- With this interface we created new tasks for each container.

PS2 defines tasks to run EI jobs



Manager:

Short description:

Ref link:

Campaign:

Subcampaign:

Phys group:

Energy gev:

Project:

- Group set to: SOFT
- Project set to: data_evind or mc_evind

Prev Slice: 0 Next

Dataset:

group.proj-evind:data15.EventIndex.OpenEnded.DAOD.i2/

job options:

Input events:

-1

Comment:

0 1 2 3 4 5 6 7 8 9

AMI tag

i2

Events per Input file

50

Events per job

50

Total events

-1

Files per job

50

nGB per job

50

max failure attempt

10

Output formats

EI

Input format

project mode

spacetoken=ATLASDATADISK;skipScout=yes

Priority

880

Destination

- AMI tag is i2
- 50 files/job
- High priority (880)

PS2 defines tasks to run EI jobs



- Existing tasks:
 - group.proj-evind:mc15.EventIndex.OpenEnded.EVNT.i2/
(https://prodtask-dev.cern.ch/prodtask/inputlist_with_request/2879/)
 - * group.proj-evind:mc15.EventIndex.OpenEnded.AOD.i2/
(https://prodtask-dev.cern.ch/prodtask/inputlist_with_request/2914/)
 - group.proj-evind:mc15.EventIndex.OpenEnded.DAOD.i2/
(https://prodtask-dev.cern.ch/prodtask/inputlist_with_request/2936/)
 - group.proj-evind:data15.EventIndex.OpenEnded.AOD.i2/
(https://prodtask-dev.cern.ch/prodtask/inputlist_with_request/2905/)
 - group.proj-evind:data15.EventIndex.OpenEnded.DAOD.i2/
(https://prodtask-dev.cern.ch/prodtask/inputlist_with_request/2935/)
- To monitor all EVENTINDEX existing tasks: https://prodtask-dev.cern.ch/prodtask/request_table/#/?type=EVENTINDEX
- All tests ~OK so far with some exceptions:
 - A few retried jobs but finally all worked
 - 1 broken job, we had to clone it
 - 5 failed tasks, 0 events in all the files belong to the same TID

PS2 defines tasks to run EI jobs



Request ID	Link	Group	Description	Campaign	Manager	Type	Approval status	Created	Approved
2936		SOFT	Test open ended container for EI DAOD mc15	mc15	sgonzale	EVENTINDEX	processed 2 2	2015-06-10	
2935		SOFT	Test open ended container for EI DAOD data15	data15	sgonzale	EVENTINDEX	processed 140 5 1 134	2015-06-10	
2933		SOFT		mc15	sgonzale	EVENTINDEX	processed 1 1	2015-06-10	
2914		SOFT	Test open ended container for EI AOD mc15	mc15	sgonzale	EVENTINDEX	processed 17 1 16	2015-06-09	
2905		SOFT	Test open ended container for EI AOD data15	data15	sgonzale	EVENTINDEX	processed 2 1 1	2015-06-08	
2879		SOFT	Test open ended container for EI	test	sgonzale	EVENTINDEX	processed 3 3	2015-06-04	

Showing 1 to 6 of 6 entries

- To resubmit failed jobs, we had to select the slice with broken/failed tasks, **clone it**, select new slice and approve it.
 - For instance, for task 2914, the PS failed for one TID. We clone only this TID in task 2933 and later ran ok.
 - **Important to monitor because human action is needed.**

PS2 defines tasks to run EI jobs



- PS2 jobs running on DAOD for run 266919 requested by Jame Catmore have finished OK,
- In AMI you can find 79 dataset (events 31328705, files; 818); Now by hand attached to the data15 DAOD technical container. In the near future will be automatic.
- But 5 datasets failed:

- if the files belong to a TIDS have all of them 0 events (only metadata), our EI jobs are failing
- For instance:

```
bash-4.1$ rucio list-files
data15_13TeV:data15_13TeV.00266919.physics_Main.merge.DAOD_STDM2.f594_m1435_p2361_tid05608882_00
```

SCOPE:NAME GUID	ADLER32	FILESIZE	EVENTS
data15_13TeV:DAOD_STDM2.05608882._000001.pool.root.1	12BE91EE-D46C-9040-88AB-3E4DB7D83301	ad:6ea4bf35	122594 0
data15_13TeV:DAOD_STDM2.05608882._000002.pool.root.1	4433FBE6-4B0A-9445-BCFC-FBC91EEF9C8C	ad:c8ba0fab	120688 0
data15_13TeV:DAOD_STDM2.05608882._000003.pool.root.1	567D7749-2584-7945-985E-C290D80F0A9D	ad:47f14eba	135667 0
data15_13TeV:DAOD_STDM2.05608882._000004.pool.root.1	29554500-2B47-4542-8BFB-14083D1D9080	ad:08985eed	142831 0
data15_13TeV:DAOD_STDM2.05608882._000005.pool.root.1	20A2E128-C10E-4648-AC8D-90BD9F4E6393	ad:32e79c86	120702 0
data15_13TeV:DAOD_STDM2.05608882._000006.pool.root.1	7B375855-F0A4-5240-B25D-D42EFA125AEE	ad:770ab86e	122305 0
data15_13TeV:DAOD_STDM2.05608882._000007.pool.root.1	43355E3E-0779-CE4E-8E08-903A0BDFC5E4	ad:7426f55e	122098 0
data15_13TeV:DAOD_STDM2.05608882._000008.pool.root.1	9377690C-B123-C641-97B6-C62991EE7B6C	ad:b6391363	122799 0
data15_13TeV:DAOD_STDM2.05608882._000009.pool.root.1	C02D2990-38DF-EC49-A13C-9ADA72C51E6C	ad:1780bfca	131423 0

- Javier worked to taking into account this error inside the EI transformation (New Tag has been created EventIndexProducer-00-02-10, now EI transformation is able to process empty DAOD).
- For this particular data taking we know that there is a problem in the HLT trigger containers. It should be fixed for the next data taking (e.g.; <https://its.cern.ch/jira/browse/ATR-11367>)
- This issue is going to be solved in the next ATLAS production releases, according to Sasha Vanyashin, and it was already reported in a Jira ticket (<https://its.cern.ch/jira/browse/ATEAM-120>)

PS2 defines tasks to run EI jobs



- We checked that if there is at least one file with some events, the jobs ran ok:

```
-bash-4.1$ rucio list-files
```

```
data15_13TeV:data15_13TeV.00266904.physics_Main.merge.DAOD_EGAM3.f594_m1435_p2361_tid05608858_00
```

SCOPE:NAME GUID	ADLER32	FILESIZE	EVENTS
data15_13TeV:DAOD_EGAM3.05608858._000001.pool.root.1	ED065310-3A5E-FA4F-9D5E-A83E9613D9A4	ad:2544a681	124848 0
data15_13TeV:DAOD_EGAM3.05608858._000002.pool.root.1	F2A79298-870B-914C-B015-490AD630A889	ad:76e666bc	1120328 17
data15_13TeV:DAOD_EGAM3.05608858._000003.pool.root.1	5E67E7F3-3DC4-454B-9283-FC7F1D9E4557	ad:d9b511f6	695528 7
data15_13TeV:DAOD_EGAM3.05608858._000004.pool.root.1	ECF8F293-141B-6C44-ACFD-2815E4EA1834	ad:e9d48acd	1266359 20
data15_13TeV:DAOD_EGAM3.05608858._000005.pool.root.1	6DF8170D-011F-674A-A2DB-5EED48435EFE	ad:28826d5e	1634364 29
data15_13TeV:DAOD_EGAM3.05608858._000006.pool.root.1	F106C345-E421-C241-8D3C-40FC9516A75A	ad:3fed4490	810791 10
data15_13TeV:DAOD_EGAM3.05608858._000007.pool.root.1	5CA1A104-B93D-2340-88C3-1E014C26242B	ad:a9bdf58b	118170 0
data15_13TeV:DAOD_EGAM3.05608858._000008.pool.root.1	DACF26C1-4B8D-0A46-A44F-FE7CEE178EFB	ad:79250538	122175 0
data15_13TeV:DAOD_EGAM3.05608858._000009.pool.root.1	31BC32D1-DA22-EA4E-AA1F-86D63A207085	ad:c7f6e2f8	118209 0



More items to discuss

- When a new tag transformation has been created, how to know in which ATLAS release is going to be included?
 - <http://atlas-computing.web.cern.ch/atlas-computing/projects/releases/recentReleases.html>
 - Send email to atlas-dp-proc@cern.ch to ask for?
- If AMI status of a Dataset (TID) changes (e.g. form All Events Available to Obsolete), we should flag properly the already imported files in Hadoop
- Validation:
 - Which info should come back from PS2 when tasks are submitted and then when completed? In order to be used to check the completeness of EI production against the original information from AMI
- Clean up of the technical containers from time to time.
- DAOD containing files with 0 events (only metada), do we have to store this information in our system?