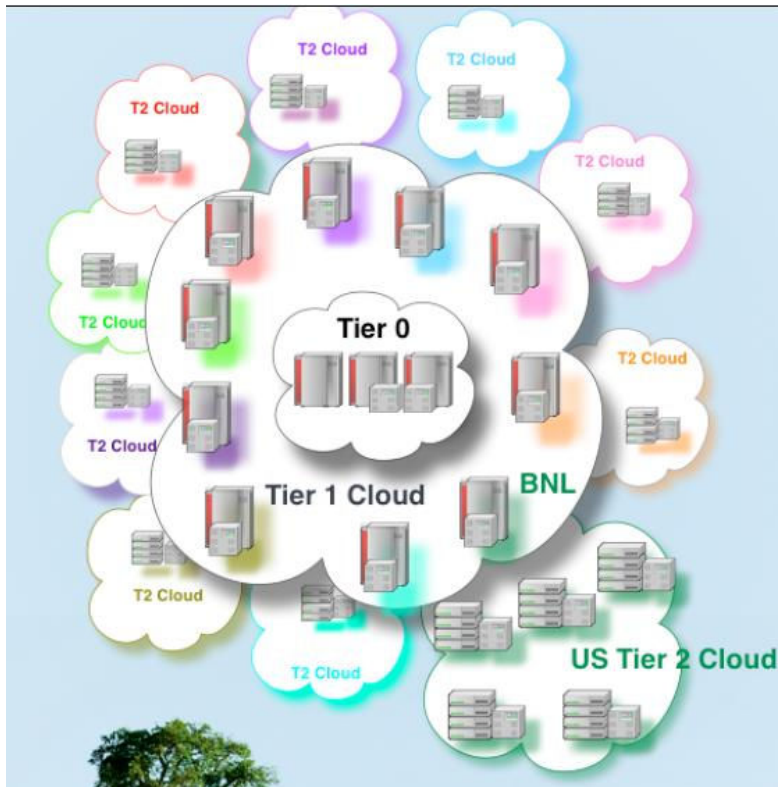# Submitting jobs to the grid

Argonne Jamboree January 2010

R. Yoshida   (revised March 2010)

Esteban Fullana

Grid is an operating system

• You should only be sure that the data that you want are on the grid

• Don't need to know where the data is.

• Don't need to know where your job is going to be running.

that is in principle…

- And to an extent in practice as well..
- In this introduction, we'll mostly stick to the cases where it's true.
- After some experience, you'll probably want (and need) to look under the hood to see what is happening.

Before you can start:

- First you will need to get a Grid certificate and join the ATLAS VO (virtual organization). This is like a "user account" for the Grid.
- How you do this is explained at http://atlaswww.hep.anl.gov/asc/working.php
- At on ascint0y or int1y, set up your account as instructed in https://atlaswww.hep.anl.gov/twiki/bin/view/UsAtlasTier3/Tier3gUsersGuide

Step 1: Get an Athena job running interactively

- The start point is an athena package that makes what you want

- If your **job options** runs, you are basically ready.

- You don't really need to change anything with the Options.py file—but you probably want the messages to be at the "ERROR" level rather than "INFO" to prevent the log files from getting too large.

- Note: input file specified inside the Options.py, number of events, etc. does NOT need to be changed.

Setting up athena release + grid

```
export ATLAS_LOCAL_ROOT_BASE=/export/share/atlas/ATLASLocalRootBase

alias setupATLAS='source${ATLAS_LOCAL_ROOT_BASE}/user/atlasLocalSetup.sh'

setupATLAS

localSetupGcc --gccVersion=gcc432_x86_64_slc5

export ATLAS_TEST_AREA=~/testarea/15.6.6

source /export/home/atlasadmin/temp/setupScripts/setupAtlasProduction_15.6.6.sh

localSetupPandaClient
```

Do localSetupDQ2Client
in a separate session (Python clash with Athena setup)

Step 2: Find the input data files in the Grid

- I've found that finding useful data or MC in ATLAS is one of the most confusing things.

- In general, you need to sort of know the name of the data you are looking for before you start.

- Rely on your physics group, colleagues, etc..

```
[ryoshida@atlas16 run]$ ls -l AOD.pool.root
lrwxrwxrwx  1 ryoshida atlas 122 Sep  6 18:11 AOD.pool.root -> /data/nas2/users/ryo
shida/sep_jamboree_mc/mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/AOD.023534
. 00001.pool.root.2
```

```
[ryoshida@atlas16 run]$ dq2-ls mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/
mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/
```

dq2-ls command lets me know that the "**container**" of this data exists. (..../ means container)
    files are in **datasets** which are in **containers**

Finding the data (cont)

- Fancier commands (**dq2-list-dataset-replicas-container**) give you more information.

```
[ryoshida@atlas16 run]$ dq2-list-dataset-replicas-container mc08.005144.PythiaZee.r
econ.AOD.e323_s400_d99_r474/
mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474_tid023534:
        INCOMPLETE:
        COMPLETE: AGLT2_MCDISK,BNL-OSG2_MCDISK,FZK-LCG2_MCDISK,IN2P3-CC_MCDISK,INF
N-MILANO-ATLASC_MCDISK,INFN-MILANO_MCDISK,INFN-NAPOLI-ATLAS_MCDISK,MWT2_DATADISK,NE
T2_MCDISK,SLACXRD_MCDISK,SWT2_CPB_MCDISK,WISC

Container name: mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/
Total  datasets: 1
Summary:
        SITE                        / # COMPLETE / # INCOMPLETE /  TOTAL
        ------------------------------------------------------------------
        MWT2_DATADISK                   1              0              1
        WISC                            1              0              1
        NET2_MCDISK                     1              0              1
        AGLT2_MCDISK                    1              0              1
        INFN-MILANO_MCDISK              1              0              1
        FZK-LCG2_MCDISK                 1              0              1
        SWT2_CPB_MCDISK                 1              0              1
        INFN-NAPOLI-ATLAS_MCDISK        1              0              1
        IN2P3-CC_MCDISK                 1              0              1
        SLACXRD_MCDISK                  1              0              1
        INFN-MILANO-ATLASC_MCDISK       1              0              1
        BNL-OSG2_MCDISK                 1              0              1
```

Finding the data (cont)

- Now you can find out what files are in the dataset(s). **dq2-ls -f**

```
[ryoshida@atlas16 run]$ dq2-ls -f mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474_tid023534 | less

mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474_tid023534
[X]     AOD.023534._00240.pool.root.2   5CDC9202-2150-DD11-90DE-0015C5E3F815   md5:1996248a7ad39ad5a9d63a6817d0b18c   36160335
[X]     AOD.023534._00207.pool.root.2   E8761C00-2350-DD11-A09C-00145EDD7B05   md5:78be2488769725dc4429ccf007441079   37263185
[X]     AOD.023534._00333.pool.root.2   9A49B93B-2A50-DD11-BD56-00145ED6DEC8   md5:f44e6d9e564440e764d05b66e9df73ed   36585137
[X]     AOD.023534._00205.pool.root.6   D2ADCE6E-815A-DD11-AFC4-001D0967D085   md5:12fce35d4c7493875d4a78aa371e877b   37136889
[X]     AOD.023534._00141.pool.root.2   B47E3000-1350-DD11-A9B9-0019B9E48FFC   md5:39e79bb4c9b1cd0d250c07c6fc35f398   36914052
[X]     AOD.023534._00377.pool.root.2   8C47FE71-1E50-DD11-B786-000D604E6524   md5:a2f704d265a284adcc0bd19d0819f155   37467876
[X]     AOD.023534._00398.pool.root.3   4A4D1692-2250-DD11-B445-00145EDD74C1   md5:cf4d716b141c91902111431bc9523e27   35950078
[X]     AOD.023534._00097.pool.root.3   F679B168-1750-DD11-B3FA-0015C5E44BF3   md5:02149077189704771944eu24564196    37458078
[X]     AOD.023534._00083.pool.root.3   AA04C16F-2350-DD11-9FC9-00145E6D496A   md5:45b04fe0ebe22da996e2058fa6e1a428   37589924
[X]     AOD.023534._00004.pool.root.2   40FA2AA1-3D50-DD11-8E48-0019BB360420   md5:86b08acd4c85d5e70dfc1e58ad5dc5c1   34465908
[X]     AOD.023534._00104.pool.root.6   F84F1A0F-E65A-DD11-AE08-001D0967C9AA   md5:599f802d2624faf0601402fba21ede14   37024242
[X]     AOD.023534._00314.pool.root.3   5293CD24-1750-DD11-AC20-0015C5E42D3F   md5:119b6f7db49e7b5188a473ba86a729be   37253767
[X]     AOD.023534._00200.pool.root.3   CC78B32A-1750-DD11-AEF7-0015C5F735B4   md5:701cdd6f33ca0f4600b7bdd01e8b185e   38167796
[X]     AOD.023534._00081.pool.root.3   E681A90B-2750-DD11-8B7D-00145ED6DE0E   md5:53ca876f2c069af0ba8ccb338b43809b   34848042
[X]     AOD.023534._00224.pool.root.6   92B2CC10-7C5A-DD11-B3C5-001D0968C880   md5:bb44ae23a5662470502b2b9d5ae40d    35914757
total files: 400
local files: 400
total size: 14866439667
date: 2008-07-30 04:18:11
```

There are 400 files in this data set.

Finding the data (cont)

- Alternately you can use AMI at http://ami.in2p3.fr/opencms/opencms/AMI/www/

- Eventually you get to a page like this…

Finding the data (cont)

- With all sorts of information

Using Pathena to submit to the Grid

- Now we're ready to submit to the Grid

Output dataset I name

```
[ryoshida@atlas16 run]$ pathena --outDS user09.RikutaroYoshida.8SeptTest.V1
--inDS mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/ --nFiles=50 Ana
lysisSkeleton_topOptions.py
```

Input container name

run over 50 files (not 400)

Do
➤pathena ─-help
to get many options you might use..

Using Pathena to submit to the Grid

- Pathena command submits to the Grid and gives you some info about the job.

```
INFO : use US as default cloud
INFO : extracting run configuration
INFO : ConfigExtractor > Input=POOL
INFO : ConfigExtractor > Output=AANT AANTupleStream AANT
INFO : archiving source files
INFO : archiving InstallArea
INFO : checking symbolic links
INFO : uploading source/job0 files
INFO : query files in mc08.005144.PythiaZee_recon.AOD.e323_s400_d99_r474/
INFO : scanning LFC heroatlas.fas.harvard.edu for ANALY_NET2
INFO : 300 files are missing or skipped at ANALY_NET2
INFO : use 50 files
INFO : submit to ANALY_NET2
==================
 JobID  : 294
 Status : 0
  > build
    PandaID=1021050513
  > run
    PandaID=1021050514-1021050516
```

Running at Northeast Tier 2

Setup, compile etc…

Split into 3 subjobs

Keeping track of your jobs

- Command pbook shows you what is happening with the job

```
[ryoshida@atlas16 run]$ pbook
INFO : Synchronizing local repository ...
INFO : Got 2 jobs to be updated
INFO : Updating JobID=293 ...
INFO : Updating JobID=294 ...
INFO : Synchronization Completed

Start pBook 0.1.81
>>> show(294)
INFO : Getting status for JobID=294 ...
INFO : Updated JobID=294
====================================
         JobID : 294
          type : pathena
       PandaID : 1021050513-1021050516
         nJobs : 3 + 1(build)
          site : ANALY_NET2
         cloud : US
          inDS : mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/
         outDS : user09.RikutaroYoshida.8SeptTest.V1
         libDS : user09.RikutaroYoshida.0909040325.820836.lib._000294.lib.
tgz
       retryID : 0
   provenanceID : 0
  creationTime : 2009-09-09 04:03:32
    lastUpdate : 2009-09-09 04:11:52
        params : --outDS user09.RikutaroYoshida.8SeptTest.V1 --inDS mc08.0
05144.PythiaZee.recon.AOD.e323_s400_d99_r474/ "--nFiles=50" AnalysisSkeleto
n_topOptions.py
     jobStatus : running
       defined : 3
       running : 1
```

Keeping track of your jobs (cont)

- Panda monitor also lets you see the status of your jobs

4 jobs. Click job number to see details.
States: defined:3 running:1
Users: Rikutaro Yoshida:4
Releases: Atlas-15.4.0:4
Processing type: pathena:4
Sites: ANALY_NET2:4

Showing 1 jobsets modified from 2009-09-09 04:03 to 2009-09-09 04:03

Job Sets:

User:jobID	Created	Latest	Jobs	Pre-run	Running	Holding	Finished	Failed	buildJob	Site
Rikutaro Yoshida:294	2009-09-09 04:03	2009-09-09 04:03	4	3	1				1021050513 libDS	ANALY_NET2
	In:mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/ **Out:**user09.RikutaroYoshida.8SeptTest.V1									

Showing 4 jobs modified from 2009-09-09 04:03 to 2009-09-09 04:03

Jobs:

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1021050516 Rikutaro Yoshida	pathena jobID=294 runAthena	defined	2009-09-09 04:03	0:00:10		09-09 04:03	US/ANALY_NET2, analysis-run	1000
	In: mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/ **Out:** user09.RikutaroYoshida.8SeptTest.V1							
1021050515 Rikutaro Yoshida	pathena jobID=294 runAthena	defined	2009-09-09 04:03	0:00:11		09-09 04:03	US/ANALY_NET2, analysis-run	1000
	In: mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/ **Out:** user09.RikutaroYoshida.8SeptTest.V1							
1021050514 Rikutaro Yoshida	pathena jobID=294 runAthena	defined	2009-09-09 04:03	0:00:11		09-09 04:03	US/ANALY_NET2, analysis-run	1000
	In: mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/ **Out:** user09.RikutaroYoshida.8SeptTest.V1							
1021050513 Rikutaro Yoshida	pathena jobID=294 buildJob	running	2009-09-09 04:03	0:00:15	0:01:45	09-09 04:03	US/ANALY_NET2, analysis-build	2000
	libDS: user09.RikutaroYoshida.0909040325.820836.lib._000294							

Notification of finished jobs

- Eventually Panda will email you that your job is finished.

PANDA notification for JobID : 294 Inbox | X

atlpan@cern.ch to Rikutaro.Yoshi.   show details 1

Summary of JobID : 294

Created : 2009-09-09 04:03:31 (UTC)
Ended : 2009-09-09 04:57:20 (UTC)

Site   : ANALY_NET2

Total Number of Jobs : 3
        Succeeded : 3
        Partial   : 0
        Failed   : 0

In : mc08.005144.PythiaZee.recon.AOD.e323_s400_d99_r474/
In : user09.RikutaroYoshida.0909040325.820836.lib._000294
Out : user09.RikutaroYoshida.8SeptTest.V1

Report Panda problems of any sort to

the eGroup for help request
  hn-atlas-dist-analysis-help@cern.ch

the Savannah for software bug
  https://savannah.cern.ch/projects/panda/

Retrieving your output

- Use dq2-get

```
[ryoshida@atlas16 ryoshida]$ pwd
/data/nas3/users/ryoshida
[ryoshida@atlas16 ryoshida]$ dq2-get user09.RikutaroYoshida.8SeptTest.V1
Group of sites selected, skipping domain name check
Querying DQ2 central catalogues to resolve datasetname user09.RikutaroYoshida.8SeptTest.V1
Datasets found: 1
user09.RikutaroYoshida.8SeptTest.V1: Querying DQ2 central catalogues for replicas...
Your configured site is not recognised. Make sure you use a ToA site name! Close-sites optimisations is not po
ssible.
Querying DQ2 central catalogues for files in dataset...
user09.RikutaroYoshida.8SeptTest.V1: No complete replica available, trying to gather pieces from incomplete so
```

```
user09.RikutaroYoshida.8SeptTest.V1/user09.RikutaroYoshida.8SeptTest.V1.AANT._00002.root: 5242880/9251918 tran
sferred
user09.RikutaroYoshida.8SeptTest.V1/user09.RikutaroYoshida.8SeptTest.V1.AANT._00002.root: 9251918/9251918 tran
sferred
user09.RikutaroYoshida.8SeptTest.V1/user09.RikutaroYoshida.8SeptTest.V1.AANT._00002.root: validated
Finished
[ryoshida@atlas16 ryoshida]$
[ryoshida@atlas16 ryoshida]$ ls user09.RikutaroYoshida.8SeptTest.V1/
user09.RikutaroYoshida.8SeptTest.V1.AANT._00001.root   user09.RikutaroYoshida.8SeptTest.V1._1021050514.log.tgz
user09.RikutaroYoshida.8SeptTest.V1.AANT._00002.root   user09.RikutaroYoshida.8SeptTest.V1._1021050515.log.tgz
user09.RikutaroYoshida.8SeptTest.V1.AANT._00003.root   user09.RikutaroYoshida.8SeptTest.V1._1021050516.log.tgz
[ryoshida@atlas16 ryoshida]$
```

Look at the output
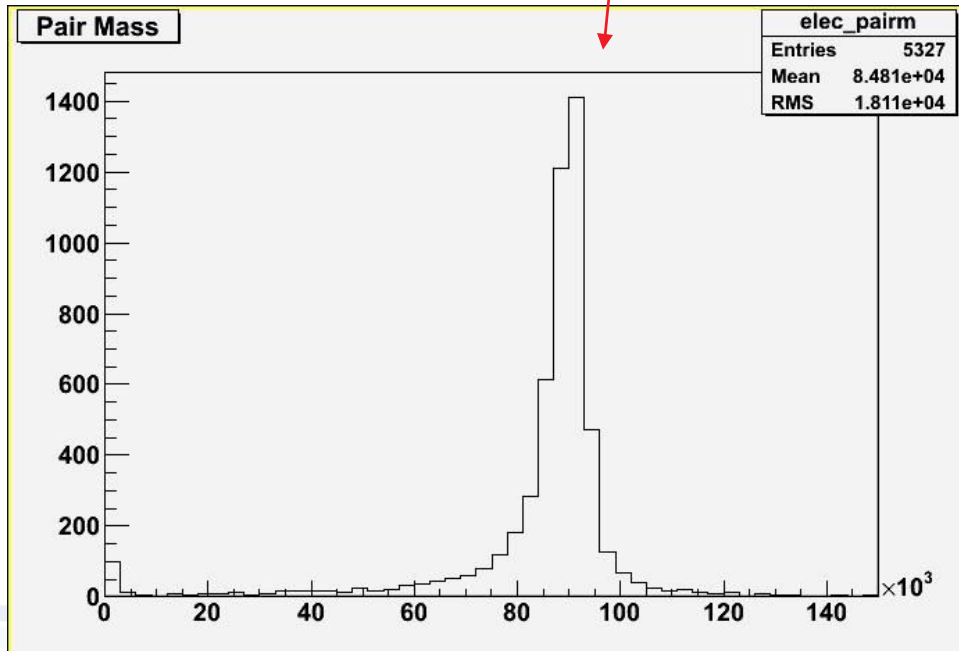
I discovered the Z!

Now you try it.

- If you follow the instructions from
  https://atlaswww.hep.anl.gov/twiki/bin/view/Jamborees/Jamboree2010Mar
- you should be able to start running on the Grid easily from ANL ASC.
- There are many more things you can do with submission and control of your jobs, but this should get you started.