Streamlining Production Processing

David Jaffe & Brett Viren

Physics Department

BROOKHAVEN NATIONAL LABORATORY

091219 Offline Workshop

Outline

What we do now

Doing Better

Panda

Going forward

MDCXX

How "mock data challenge" simulation production happens now :

- David and/or Cheng-Ju collect requirements.
- One or other enumerates the jobs and parcels our run assignments to the contributing clusters (eg DocDB 4299)
- Each cluster contact must interpret the run assignment and implement local batch submission scripts.

This last step leads to a lot of unnecessarily duplicated effort.

Once small step further

One simple idea is instead of producing a wiki/spreadsheet table with job parcels, produce one text file for each job assignment that has one `nuwa.py` command line per line.

Once big step further

Next, it was recognized that we only had two batch systems (PBS-derived and Condor, others?) so submitting jobs shouldn't be **that** different. $\Rightarrow$

- Unify all the batch submission scripts into a single Python application.
- Factor out any cluster-specifics
- Keep central database of run assignments, jobs submissions and status.
- Allow failed jobs to be identified and rerun.
- Put up a fancy web interface.

Batch submission becomes trivial.

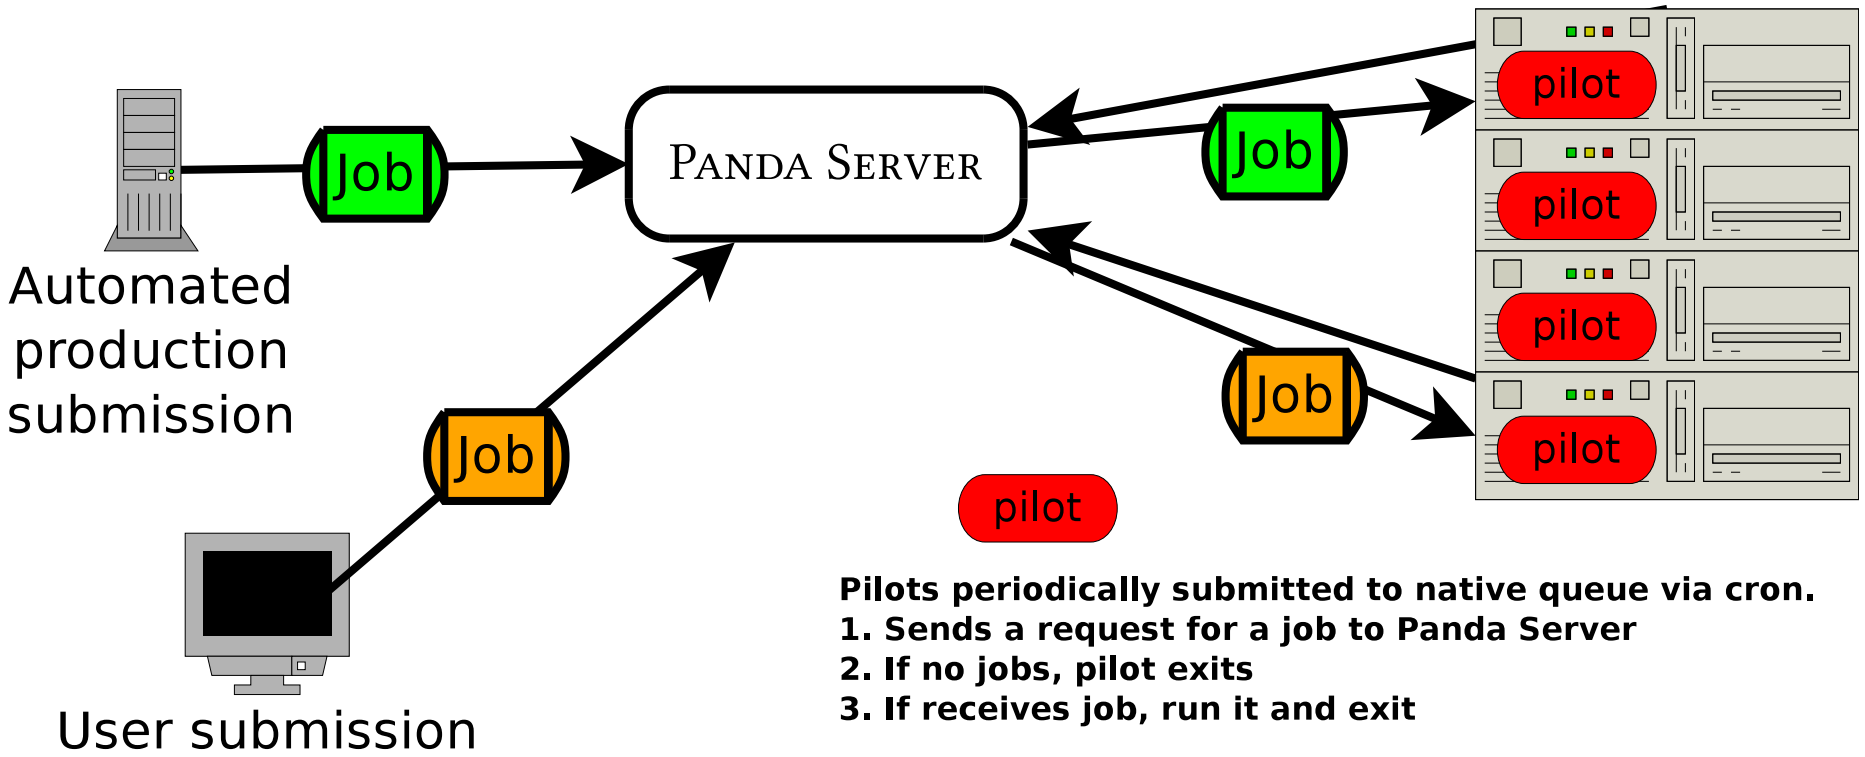
Wait, this sounds familiar

Of course, with such a good idea, it is likely already done.
Enter **Panda**:

- PanDA = **P**roduction **an**d **D**istributed **A**nalysis
- Developed for ATLAS by Physics Application Software (PAS) group at BNL.
  - They are very eager to support us.
- Single, unified batch system on top of multiple, native batch systems.
  - Specifically, PBS and Condor are supported.
- Opens up larger computing resources via Open Science Grid.
- Status and history visible on web application.

We propose having all Daya Bay clusters allow Panda jobs.

How does Panda work?



- Uniform layer on top of disparate native batch systems.
- So called “pilot” jobs submitted to native queue, run real jobs
- Pilots can be specialized if needed.
- Could consider automating submissions or have single user submit all production jobs on all participating clusters from a single location.

Panda web: job sets at a site

Panda job information

Click for help

Summary of all **jobs for the last** 3 **days, jobID** any **in** any **state at** TEST2 **site** Go Retrieve All

19 jobs. Click job number to see details.
States: activated:6 finished:13
Users: Brett Michael Viren:1 Jose Caballero:18
Processing type: ITB_Integration:3
Sites: TEST2:19

Hide Job Sets

Showing 1 jobsets modified from 2009-12-11 19:46 to 2009-12-08 23:13

Job Sets:

User:jobID	Created	Latest	Jobs	Pre-run	Running	Holding	Finished	Failed	buildJob	Site
Jose Caballero:10030	2009-12-11 16:55	2009-12-11 16:55	1	1						TEST2
	Out: None									
Jose Caballero:10029	2009-12-10 21:55	2009-12-10 21:55	1				1			TEST2
	Out: None									

- Status of all job sets at a site.

Panda web: individual jobs at a site

Showing 19 jobs modified from 2009-12-11 19:46 to 2009-12-08 23:13

Jobs:

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1035120271 Brett Michael Viren	fake	activated	2009-12-11 19:46	0:00:02		12-11 19:46	US/TEST2, analysis-run	1000
1035089876 Jose Caballero	pathena jobID=10030 trivial	activated	2009-12-11 16:55	0:00:02		12-11 16:55	US/TEST2, analysis-run	1000
1034903404 Jose Caballero	pathena jobID=10029 trivial	finished	2009-12-10 21:55	0:18:08	0:00:02	12-10 22:13	US/TEST2, analysis-run	1000
1034850572 Jose Caballero	pathena jobID=10028 fake	finished	2009-12-10 15:14	3:40:01	0:00:02	12-10 18:54	US/TEST2, analysis-run	1000
1034850571 Jose Caballero	pathena jobID=10028 fake	finished	2009-12-10 15:14	3:40:01	0:00:02	12-10 18:54	US/TEST2, analysis-run	1000
1034850570 Jose Caballero	pathena jobID=10028 fake	finished	2009-12-10 15:14	3:39:59	0:00:02	12-10 18:54	US/TEST2, analysis-run	1000
1034850569 Jose Caballero	pathena jobID=10028 fake	finished	2009-12-10 15:14	0:11:56	0:00:02	12-10 15:26	US/TEST2, analysis-run	1000

- Status of individual jobs.
- Per-job view (not shown) gives access to `stdout` and `stderr`

Panda web: per-user view

Production Clouds Incidents DDM PandaMover AutoPilot Sites Analysis Stats Physics data Usage ProdDash DDMDash

Show **my page** users groups Brett Michael Viren Log out

Brett Michael Viren

Click for help

Jobs: 1 total jobs, last at 2009-12-11 20:57
Sites used: Panda (1)
Job types run: analysis (1)
Groups:

Usage	1 day (quota)	7 day (quota)	30 day (quota)
Analysis	0 (500)	0 (3500)	0 (15000)
User production	0 (30)	0 (210)	0 (900)
Express	0 (150)	0 (1050)	0 (4500)

Summary of all **jobs for the last** 5 **days, jobID** any **in** any **state at** any **site** Go Retrieve All

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1035120271 Brett Michael Viren	fake	finished	2009-12-11 19:46	1:04:39	0:00:02	12-11 20:51	US/TEST2, analysis-run	1000

What is needed to use Panda

To make use of Panda need:

- Download Python code for the pilots to each cluster.
- Set up `crontab` to submit pilots to native queue.
- Download Python code for submitting.
- Write the actual job scripts.

Job scripts:

- Job script downloaded by pilot via HTTPS
- Give job script command line when submitting job
- Script responsible for setting environment, staging files, and running `nuwa.py`

Security

Certificates are used to authorize all access:

- Request Daya Bay Virtual Organization (VO) certificate.
  - BNL is setup and ready to request and manage the Daya Bay VO.
  - Will only allow Daya Bay jobs to run on participating Daya Bay nodes
- Job and pilot submitters need a personal certificate
  - Request certificates through DOE Certificate Authority (CA)[1]
  - Brett went through this process and it was simple.
  - Requires a "sponsor" known to DOE CA to vouch for an applicant. (Craig is right person?)

[1] `http://www.doegrids.org/pages/cert-request.html`

Going forward

Our plans going forward:

- BNL will request a Daya Bay VO.
  - Objections? Anyone else want to do it?
- Will run Panda pilots on RACF nodes and gain experience.
- Assist other sites that want to join in.

Information is being collected in the wiki[2]

[2] `https://wiki.bnl.gov/dayabay/index.php?title=Batch_Processing`