ATLAS Analysis Pilots

Peter Love and Graeme Stewart

2008-11-06

University of Glasgow | Department of Physics & Astronomy

LANCASTER UNIVERSITY

ATLAS User Analysis

- ATLAS computing model sees user analysis as the dominant activity at Tier-2s
- Currently a lot of user analysis happens at CERN or at T3
  - This will not scale when real data arrives
  - Some non-grid user activities already have caused significant problems for T0→T1 data replication
- Distributed analysis campaigns being mounted

Analysis Job Submission

- Two front ends:
  - ganga - generic job submission tool
  - pathena - athena on the grid front end
- Two back ends:
  - gLite WMS - standard EGEE resource broker
  - panda - ATLAS resource broker based on a pilot job model
- Ganga can submit to both backends, pathena only to panda
  - In the past only the WMS backend was used for ganga in EGEE

Typical Day at the WMS

WMS ate my job(s)

ATLAS UK Grid Summary

The results of the last 10 jobs are shown, most recent to the right. S denotes Success, F Failure, A Aborted and X Cancelled. C denotes Current jobs that are still running or queued.

Institute	Release	Replica	Hello World	New Package	User Analysis	Last 24 hrs	Status
UKI-IRELAND-TRINITY	14.2.24	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S F	97%	
UKI-LT2-Brunel-40	14.2.22	DPM	C C C C C C C C C C	C C C C C C C C C	C C C C C F C C C C	0%	
UKI-LT2-Brunel-44	14.2.23	DPM	S S S S S S S S S C	S S S S S S S S S	F F F F F F F F F C	65%	
UKI-LT2-IC-HEP-HEP	14.2.24	dCache	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S C	98%	
UKI-LT2-IC-HEP-HPC	14.2.10	dCache	C C C C C C C C C C	C C C C C C C C C	C C C C C F C C C C	0%	
UKI-LT2-IC-LeSC	14.2.10	dCache	S S S S S S S S S C	F F F F F F F F F	F F F F F F F F F F	31%	
UKI-LT2-QMUL-01	14.2.24	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S S	98%	
UKI-LT2-QMUL-02	14.2.23	Yes	C C C C C C C C C C	C C C C C C C C C	C C C C C F C C C C	0%	
UKI-LT2-RHUL-01	13.0.40	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S C	98%	
UKI-LT2-RHUL-02	14.2.24	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S C	97%	
UKI-LT2-UCL-CENTRAL	Unknown	DPM					
UKI-LT2-UCL-HEP	14.2.24	DPM	S S S S S C S S S C	S S S S S C S S S	F F F F F F F F F C	65%	
UKI-NORTHGRID-LANCS-HEP	14.2.23	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S C	98%	
UKI-NORTHGRID-LIV-HEP	14.2.24	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S S	98%	
UKI-NORTHGRID-MAN-HEP-01	14.2.23	dCache	C C S S S S S S S C	S C S S S S S S S	S S S S S F S S S C	60%	
UKI-NORTHGRID-MAN-HEP-02	14.2.24	dCache	S S S S S S S S C C	S S S S S S S C C	S S S S F F S S S C	96%	
UKI-NORTHGRID-SHEF-HEP	14.2.24	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S C	98%	
UKI-SCOTGRID-DURHAM	14.2.23	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S C	98%	
UKI-SCOTGRID-ECDF	14.2.24	DPM	S S S F F C S S S C	S S S C F C S S S	S S C C C F S S S S	92%	
UKI-SCOTGRID-GLASGOW	14.2.24	DPM	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S S	98%	
UKI-SOUTHGRID-BHAM-ESCI	14.2.23	DPM	S S S S S S S S S C	S S S S S C S S S	S S S S S F S S S C	96%	
UKI-SOUTHGRID-BHAM-HEP	14.2.23	DPM	S S S S S S S S S C	S S S S S C S S S	S S S S C C S S S C	96%	
UKI-SOUTHGRID-BRIS-HEP-01	14.2.23	DPM	S S S S C C S S S C	S S S S C S C S S	S S S C S F S C S C	98%	
UKI-SOUTHGRID-BRIS-HEP-02	Unknown	DPM					
UKI-SOUTHGRID-CAM-HEP	14.2.24	DPM	X C C C C C S S S C	C C C C C C S S S	C C C C C F S S S C	54%	
UKI-SOUTHGRID-JET	14.2.20	DPM	C C C C C C C C C C	C C C C C C C C C	C C C C C F C C C C	0%	
UKI-SOUTHGRID-OX-HEP	14.2.24	DPM	C C C C C C C C C C	C C C C C C C C C	X C C C C F C C C C	0%	
UKI-SOUTHGRID-RALPP	14.2.10	dCache	S S S S S S S S S C	S S S S S S S S S	S S S S S F S S S S	98%	
UKI-TIER-1-RAL-LCG2	14.2.23	Castor	C C C C C C C C C C	C C C C C C C C C	C C C C C F C C C C	0%	
Average						76%	

Other WMS Issues

- The WMS knows nothing about ATLAS dataset model so it does not do any real brokering anyway
  - Extra components always introduce new failure modes and obscure lower layers
- Job priorities are extremely difficult to manage with WMS - lots of work, crude and inflexible

So in a nutshell we'd like to...

- Be able to submit user analysis jobs through panda as well as the WMS

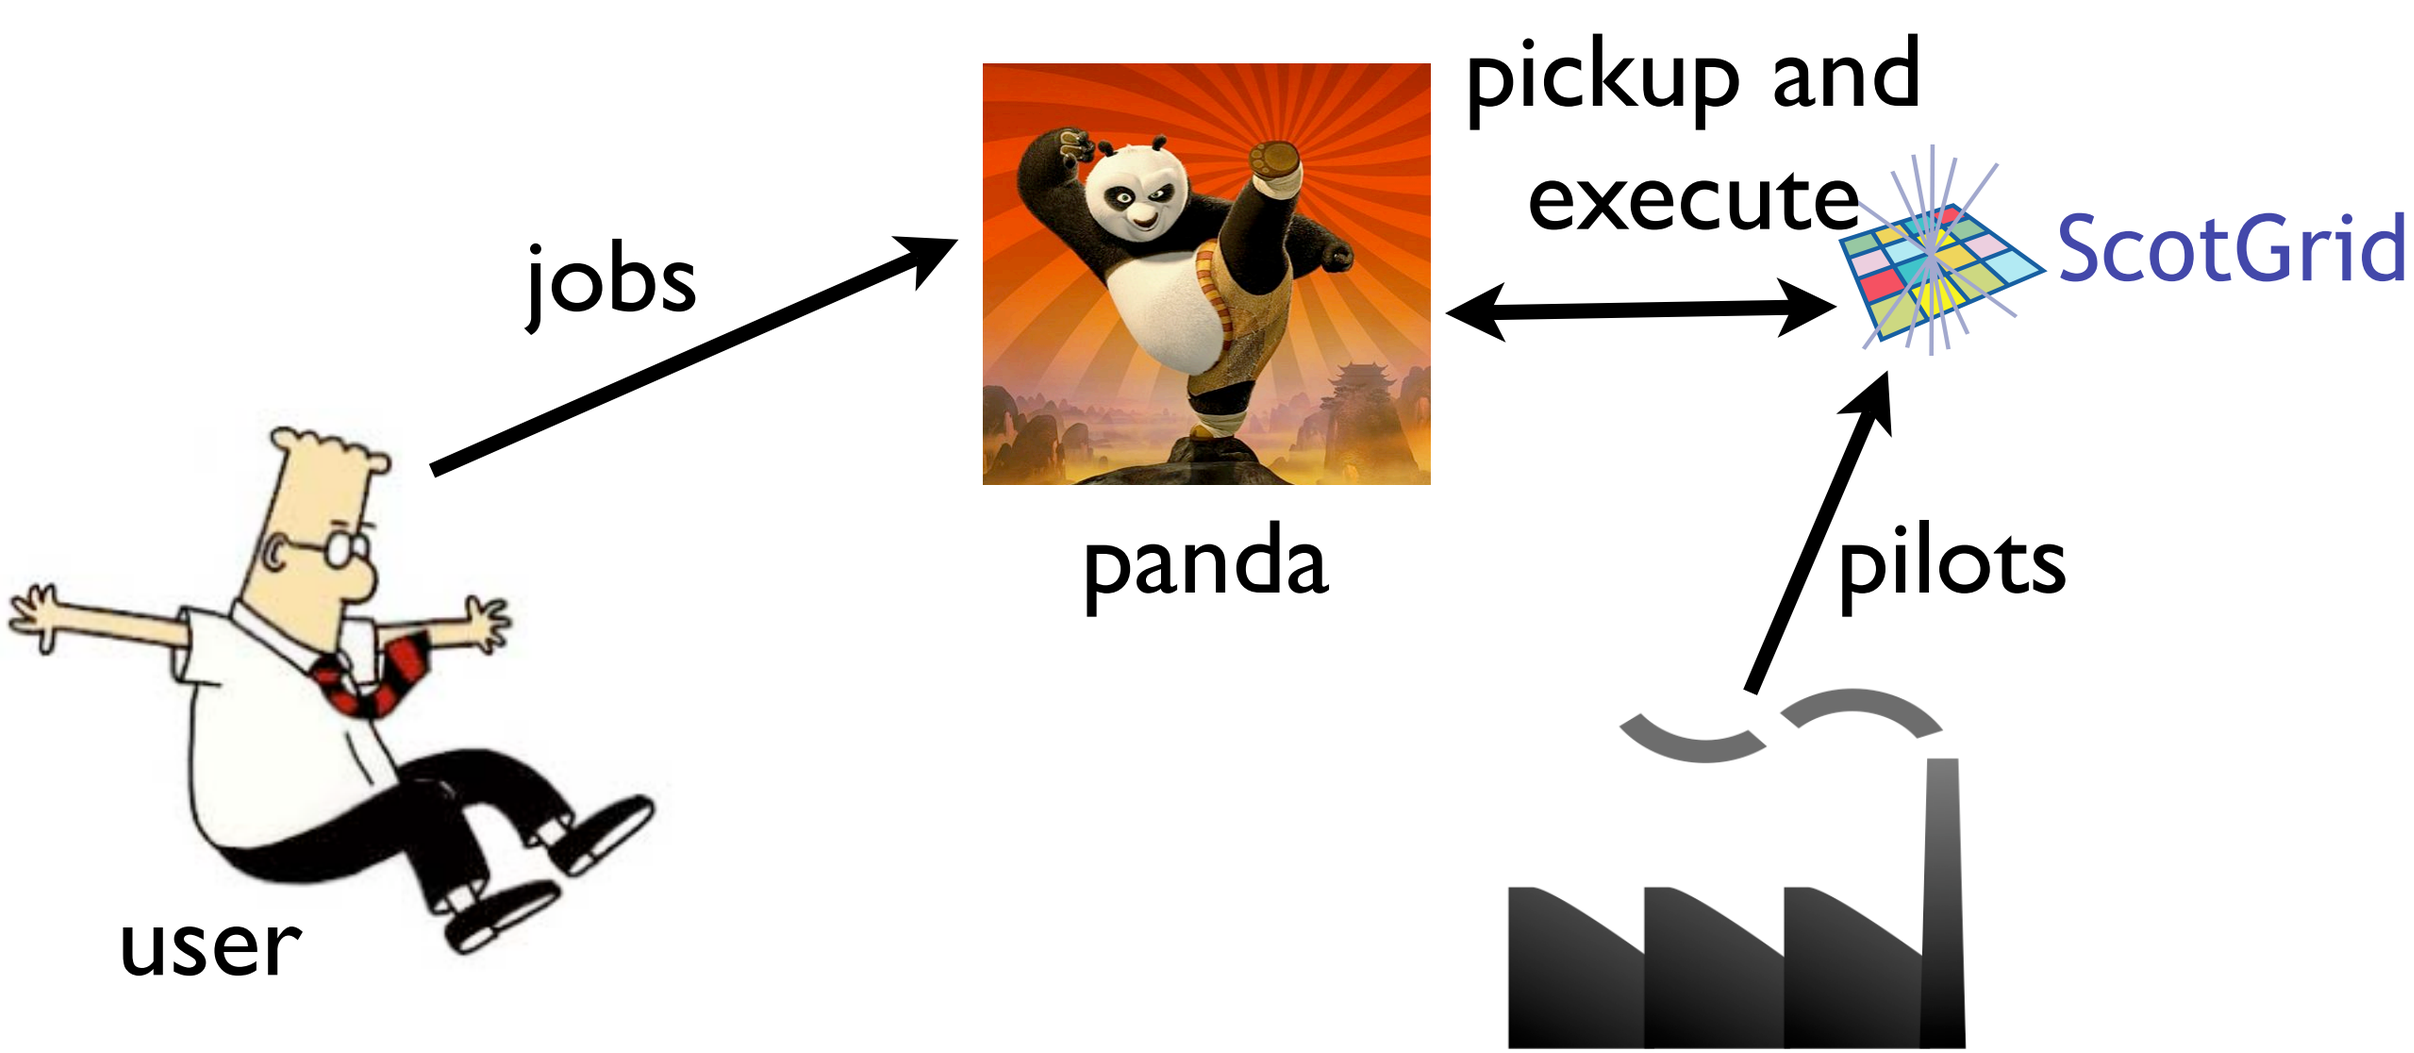

DN Problem

- Pilots are submitted from one DN
- Jobs are submitted from another
- glexec should be used to change the UID of the job to the pool account of the submitter
  - But it is not yet released to production
    - Should have been delivered in July
  - So we would like to ask that sites accept these jobs even without UID switching

Traceability

- You can trace the panda job which was run on your site from the panda monitor
  - http://panda.atlascomp.org/
- Enter a batch system ID into the “job” search field
  - Use batchID@SITE on SGE clusters
- You will recover the panda job, with the DN of the submitter, the transform and the job’s arguments

Here's one I prepared earlier...

- http://gridui07.usatlas.bnl.gov:25880/server/pandamon/query?job=1534575.svr016.gla.scotgrid.ac.uk

BNL monitor | Production | Clouds | DDM | PandaMover | AutoPilot | Sites | Analysis | Physics data | Usage | Plots | ProdDash | DDMDash

0 min old Update

Not logged in. List users

Panda monitor
Times are in UTC

Panda info and help

Panda job information

Jobs: 18667746

Click for help

Jobs - search
Recent running, activated, waiting, assigned, defined, finished, failed jobs
Select analysis, prod, install, test jobs

Quick search
Job
Dataset
Task request
Task status
File

Summaries
Blocks: days
Errors: days
Nodes: days
Daily usage

Tasks - search
Generic Task Req
EvGen Task Req
CTBsim Task Req
Task list
New Tag
Bug Report

Datasets - search
Dataset browser

Showing 1 jobs modified from 2008-11-03 11:14:24 to 2008-11-03 12:55:23

Jobs:

PandaID, Owner	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
18667746 borut.kersevan@ijs.si	valid1.105200.T1_McAtNlo_Jimmy.digit.e357_s462_d137_tid027601._01427.job #1	finished	11-02 20:50	13:56:30	0:27:35	11-03 11:14	UK/UKI-SCOTGRID-GLASGOW, production	960
	In: mc08.105200.T1_McAtNlo_Jimmy.simul.HITS.e357_s462_tid025431 **Out:** valid1.105200.T1_McAtNlo_Jimmy.digit.RDO.e357_s462_d137_tid027601							

Job 18667746 details

20 files for job 18667746:

Filename	Type	Status	Dataset
DBRelease-6.0.1.tar.gz	input	ready	ddo.000001.Atlas.Ideal.DBRelease.v060001
HITS.024528._00040.pool.root	input	ready	mc08.105001.pythia_minbias.simul.HITS.e349_s462_tid024528 (dispatch block: panda.952bf3df-0dc5-4f4b-ae08-a62b50304991_dis18667745)
HITS.024528._01249.pool.root	input	ready	mc08.105001.pythia_minbias.simul.HITS.e349_s462_tid024528
HITS.024528._01495.pool.root	input	ready	mc08.105001.pythia_minbias.simul.HITS.e349_s462_tid024528 (dispatch block: panda.952bf3df-0dc5-4f4b-ae08-a62b50304991_dis18667745)
HITS.024528._01613.pool.root	input	ready	mc08.105001.pythia_minbias.simul.HITS.e349_s462_tid024528 (dispatch block: panda.952bf3df-0dc5-4f4b-ae08-a62b50304991_dis18667745)
HITS.024853._00012.pool.root.1	input	ready	mc08.108863.Hijing_beamgas.simul.HITS.e4_s470_tid024853 (dispatch block: panda.952bf3df-0dc5-4f4b-ae08-a62b50304991_dis18667745)
HITS.024853._00086.pool.root.2	input	ready	mc08.108863.Hijing_beamgas.simul.HITS.e4_s470_tid024853

Tests We'd Like to Do

- Compare the efficiency of ganga analysis done through the WMS and through panda
  - This will be done through a 'personal pilot' from the user who submits the jobs (i.e., pilot submitter and payload submitter are the same)
- Open the system to a small number of users to try more general analysis

Breaking policy

- EGEE Joint Security Policy Group
- Acceptable Use states
- *No intentional sharing of credentials for Grid purposes is permitted.*

Current Status and Future Plans

- Canada, France and US already run analysis via the panda backend
- We will move to using glexec as soon as we can
- WLCG MB to be approached
- This test will provide very important information about which direction the user analysis implementation should go