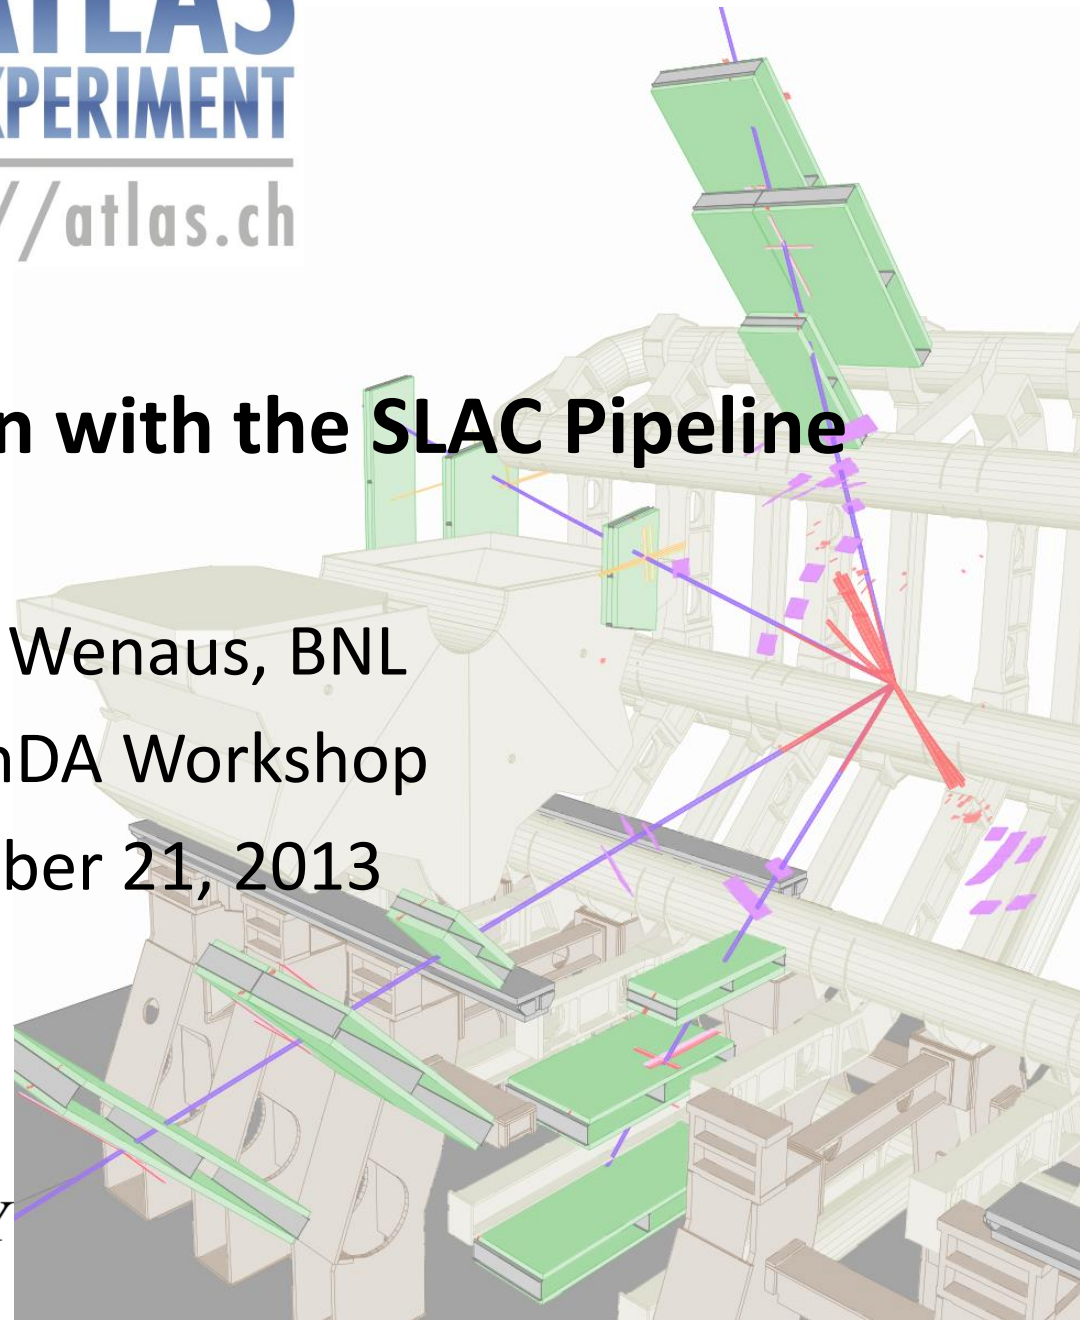


PanDA Integration with the SLAC Pipeline

Torre Wenaus, BNL
BigPanDA Workshop
October 21, 2013

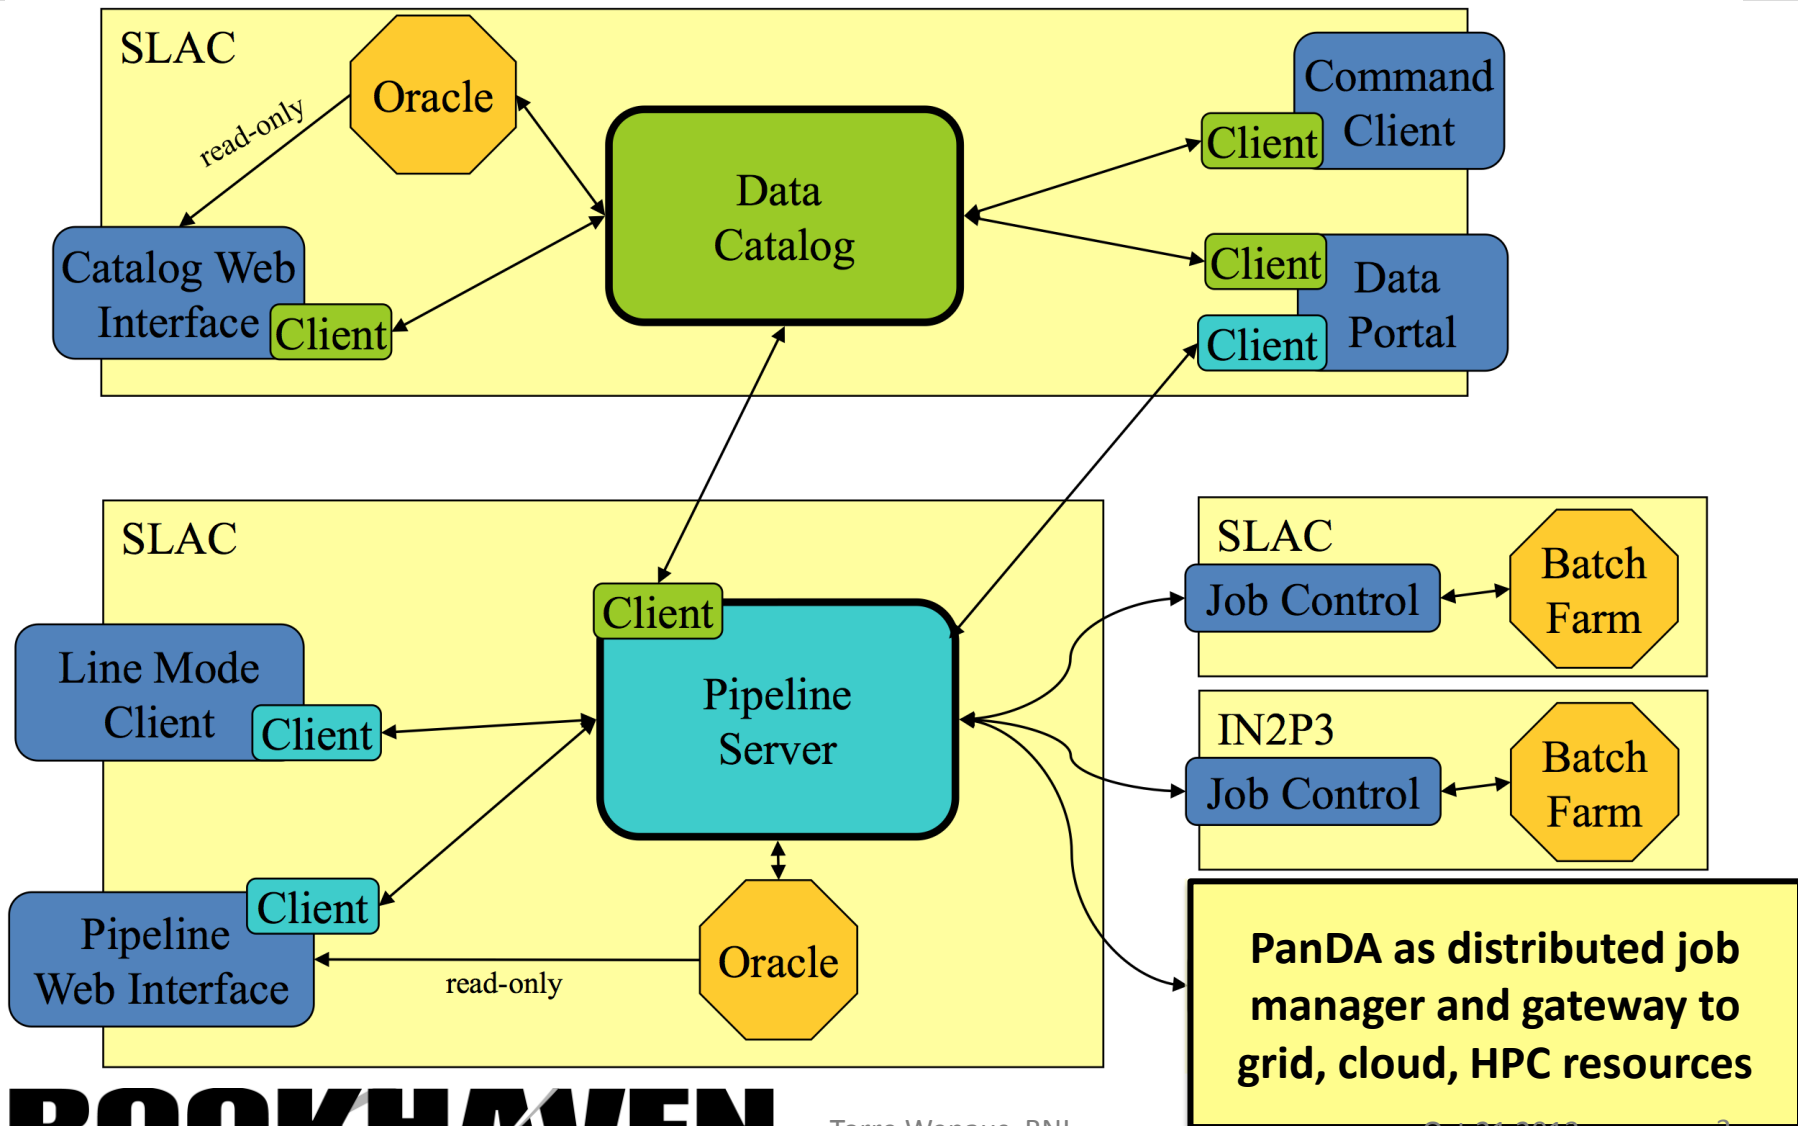


PanDA and LSST/DESC

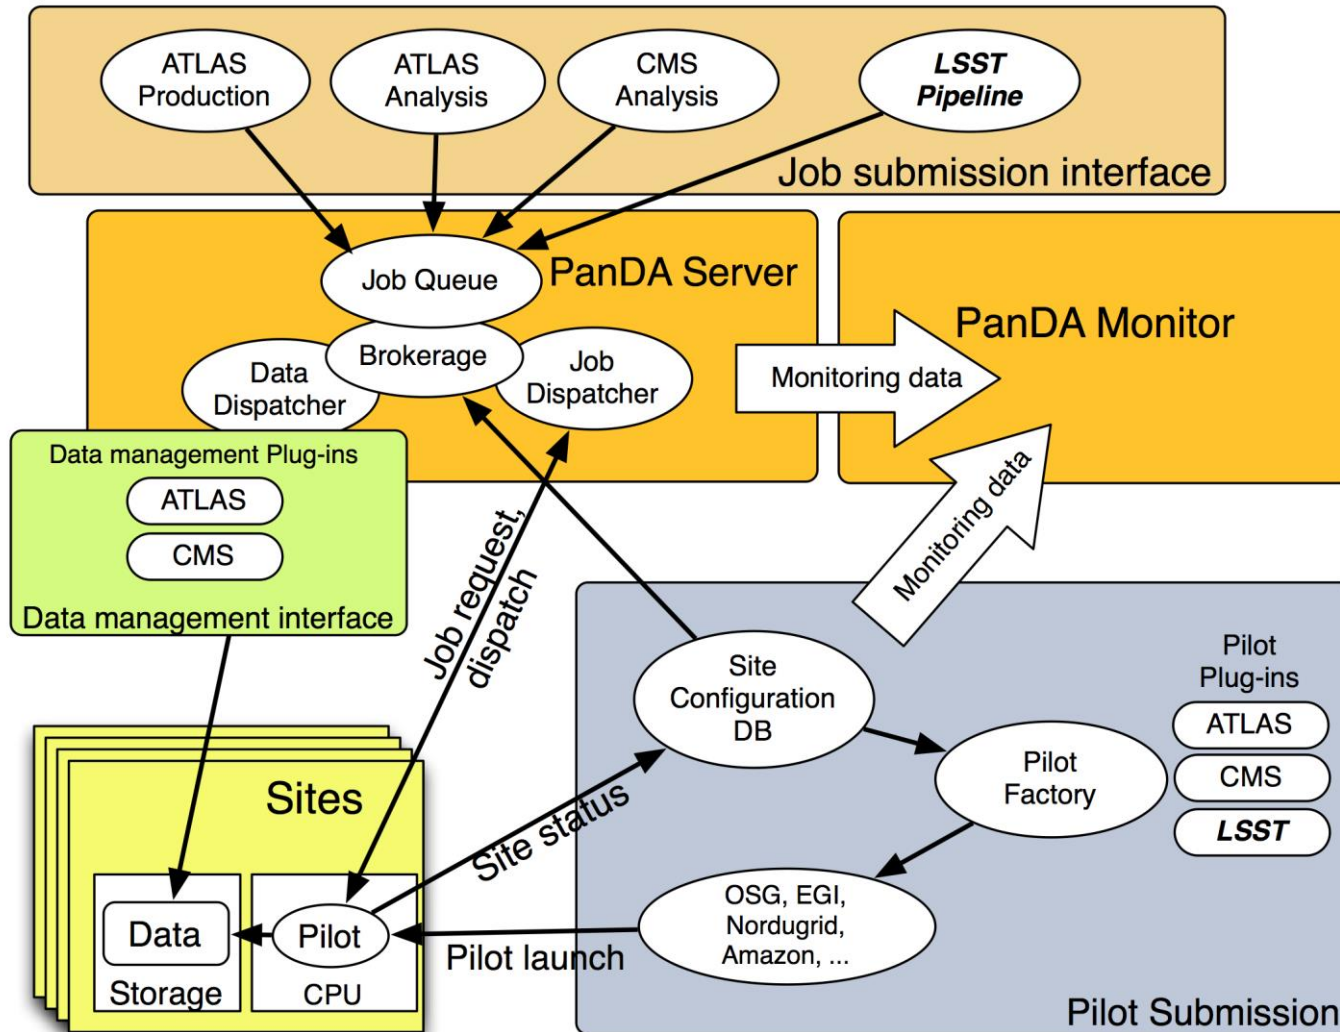


- Gave a PanDA overview at the January 2013 LSST Dark Energy Sciences Collaboration (DESC) meeting
- Decision made to prototype PanDA as (another) submission engine for Tony Johnson's pipeline workload manager used by Fermi and now DESC (among others)
 - Gain experience in both directions
 - Gain DESC access to more resources via PanDA
 - Evaluate PanDA's usefulness
- In February Tony provided a Java interface to the pipeline for submission engines

SLAC Workflow Pipeline and PanDA Prototype



PanDA System Overview



PanDA and LSST



- Pipeline => PanDA submission working, tested with HelloPanda and a standard DESC simulation workflow (phosim, DM stack)
 - Anyone who can submit a pipeline task can submit to PanDA
- Initial tests used ATLAS PanDA queue at BNL (epsilon level) and CERN based standalone pilot submission (no APF)
 - BNL PanDA instance being set up for non-ATLAS use in US
 - Discussions started with OSG (Miron so far) re: an OSG supported PanDA instance once PanDA is cleanly packaged and (re)ported to MySQL
- Currently borrowing BNL ATLAS resources during testing phase
 - Can submit jobs (pilots) as LSST VO as soon as desirable
- Plenty of US ATLAS resources to be had opportunistically right now; a quiet period (there will regularly be quiet periods like this until 2015)

Pipeline/PanDA integration: Pipeline interface

- Java daemon JobControlClient accepts 3 control commands from the pipeline: submit, status, kill
 - Submit
 - Receive command and arguments
 - Submit the job
 - Job emails the pipeline server on job start and finish
 - Status
 - Report job status to pipeline server
 - Kill
 - Kill job
- PanDA implementation extends the JobControlClient class
 - Communicates with pipeline server via Java RMI
- Tested first with an emulator for the pipeline server (provided so a JobControlClient can be tested standalone) then PanDA daemon was connected to the real pipeline service at SLAC

Pipeline/PanDA integration

- PandaJobControlService implemented supporting the submit, status, cancel functions
- Submit:
 - Service creates and submits a PanDA job configured with the specified command/arguments
 - Job is configured to invoke a DESC-specific transformation (payload script)
 - Following pipeline conventions, the script emails the pipeline server on payload job start/finish. Depends on emailing always working from WNs.
 - Returns PandaID as the ID identifying the job
 - LSST jobs are picked up from the PanDA job queue by standard PanDA production pilots
 - Amazon S3 used as staging area for job files
 - No special provision for data management at present; jobs are on their own in accessing data

Pipeline/PanDA integration (2)

- Status:
 - Service queries PanDA for status of current pipeline jobs
 - Programmatic API to the PanDA monitor used to obtain job status (curl query returns json result)
 - Service maps PanDA job states, etc. onto pipeline conventions and returns pipeline JobStatus
- Cancel:
 - Service directs PanDA to cancel the job (request to PanDA server)
 - The PanDA server immediately sets the job state to 'cancelled'
 - The job itself will initiate self-destruct when it gets its next heartbeat reply (every 30min) that will contain the cancel command

LSST PanDA jobs in pipeline and PanDA monitors (June)



← → ↻ 🏠 panda.cern.ch:25880/server/pandamon/query?job=*&ui=user&name=Torre+Wenaus

Showing 40 jobs modified from 2013-06-05 15:04 to 2013-06-05 20:00
add &nojobs=yes parameter to your URL to hide this table

Jobs:

PandalD, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1839075004 Torre Wenaus	jobsetID=4471 lsst-trf	finished	2013-06-05 15:01	0:02:53	0:01:21	06-05 15:05	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.7e1f114f-5312-4d10-931b-d0f5c0685910								
1839074959 Torre Wenaus	jobsetID=4461 lsst-trf	finished	2013-06-05 15:01	0:01:03	0:01:37	06-05 15:03	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.b7762dc5-4392-4451-9ed9-12733675858a								
1839074941 Torre Wenaus	jobsetID=4451 lsst-trf	finished	2013-06-05 15:00	0:01:12	0:01:28	06-05 15:02	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.cf53640d-509b-4970-b3d4-f2								
1839074933 Torre Wenaus	jobsetID=4441 lsst-trf	finished	2013-06-05 15:00					
Out: user.wenaus.lsst.c19a6826-3575-4730-ab57-6								
1839045644 Torre Wenaus	jobsetID=2549 lsst-trf	finished	2013-06-05 14:29					
Out: user.wenaus.lsst.6d07ad4f-a4d8-47c3-b257-1f								
1839045614 Torre Wenaus	jobsetID=2539 lsst-trf	finished	2013-06-05 14:29					
Out: user.wenaus.lsst.edd8055e-c2c0-46fd-ba6c-b								
1839037042 Torre Wenaus	jobsetID=1245 lsst-trf	finished	2013-06-05 14:07					
Out: user.wenaus.lsst.105e35aa-e282-475a-be78-1								
1839035320 Torre Wenaus	jobsetID=1083 lsst-trf	finished	2013-06-05 14:04					
Out: user.wenaus.lsst.3d70ea89-ab79-4720-a065-6								
1839027003 Torre Wenaus	jobsetID=466 lsst-trf	finished	2013-06-05 13:54					
Out: user.wenaus.lsst.4a0d2d9f-5bbf-44a1-b9ab-3c								
1839024031 Torre Wenaus	jobsetID=307 lsst-trf	finished	2013-06-05 13:51					
Out: user.wenaus.lsst.588b6abc-4cdb-4519-8d25-6								
1838995930 Torre Wenaus	jobsetID=8679 lsst-trf	finished	2013-06-05 13:24					
Out: user.wenaus.lsst.c6ab0425-77ff-428c-967f-9f4								
1838012695 Torre Wenaus	jobsetID=6183 lsst-trf	finished	2013-06-04 14:29					
Out: user.wenaus.lsst.28a87efd-62c5-42d8-a1bc-e								
1838004938 Torre Wenaus	jobsetID=5101 lsst-trf	finished	2013-06-04 14:11					
Out: user.wenaus.lsst.337913d9-34b0-4897-819f-e								
1838003396 Torre Wenaus	jobsetID=4929 lsst-trf	finished	2013-06-04 14:08					
Out: user.wenaus.lsst.0236a4a6-aa38-4080-a83a-c								
1837992821 Torre Wenaus	jobsetID=3927 lsst-trf	finished	2013-06-04 13:52	0:04:45	0:01:36	06-04 13:58	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.730de0b2-8d18-4e1e-bf9a-5c0cc0caa7df								
1837977020 Torre Wenaus	jobsetID=2896 lsst-trf	finished	2013-06-04 13:35	0:02:16	0:01:20	06-04 13:38	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.9705cc65-8937-4472-97b8-33380e014e57								
1837933081 Torre Wenaus	jobsetID=8523 lsst-trf	finished	2013-06-04 12:22	0:00:40	0:01:27	06-04 12:24	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.ce0fcda5-0864-4711-ae8-6d8bc2f31268								
1837918391 Torre Wenaus	jobsetID=7451 lsst-trf	finished	2013-06-04 12:04	0:00:28	0:01:21	06-04 12:06	US/US.ANALY_BNL_SHORT, ptest	1000
Out: user.wenaus.lsst.b89c2c4f-2397-444b-958b-37607cee8552								



summary / HelloPanda / HelloWorldJob

Streams for process: HelloWorldJob

Processing plots

Task Summary: Canceled: 0, Canceling: 0, Failed: 0, Queued: 0, Running: 0, Success: 8, Terminated: 15, Terminating: 0, Waiting: 0, Total: 23

Top Level Stream: Min Max Status: All All Not Success Waiting Ready

Stream Filter: Regular Expression (?)

Created Date: None None or last N hours 6

Show all streams on one page

Show only latest execution

Stream	Status	Created	Submitted	Started	Ended	Job Id	CPU	Host	Links
15	Success	05-Jun-2013 07:04:39	05-Jun-2013 07:04:49	05-Jun-2013 07:09:46	05-Jun-2013 07:09:46	1839035320	0	acas1021.usatlas.bnl.gov	Messages : Log : Files
16	Success	05-Jun-2013 07:07:19	05-Jun-2013 07:07:29	05-Jun-2013 07:09:45	05-Jun-2013 07:09:46	1839037042	0	acas1080.usatlas.bnl.gov	Messages : Log : Files
17	Success	05-Jun-2013 07:28:53	05-Jun-2013 07:29:03	05-Jun-2013 07:30:08	05-Jun-2013 07:30:10	1839045614	0	acas0198.usatlas.bnl.gov	Messages : Log : Files
18	Success	05-Jun-2013 07:29:06	05-Jun-2013 07:29:12	05-Jun-2013 07:30:32	05-Jun-2013 07:30:33	1839045644	0	acas1399.usatlas.bnl.gov	Messages : Log : Files
19	Success	05-Jun-2013 08:00:37	05-Jun-2013 08:00:45	05-Jun-2013 08:04:08	05-Jun-2013 08:04:09	1839074933	0	acas1586.usatlas.bnl.gov	Messages : Log : Files
20	Success	05-Jun-2013 08:00:46	05-Jun-2013 08:00:53	05-Jun-2013 08:02:11	05-Jun-2013 08:02:12	1839074941	0	acas1586.usatlas.bnl.gov	Messages : Log : Files
21	Success	05-Jun-2013 08:00:55	05-Jun-2013 08:01:04	05-Jun-2013 08:02:20	05-Jun-2013 08:02:20	1839074959	0	acas1525.usatlas.bnl.gov	Messages : Log : Files
22	Success	05-Jun-2013 08:01:04	05-Jun-2013 08:01:14	05-Jun-2013 08:04:11	05-Jun-2013 08:04:12	1839075004	0	acas1525.usatlas.bnl.gov	Messages : Log : Files

Select all · Deselect all · Toggle selection · Rollback Selected

Activities Since June

- Work program established with Open Science Grid (OSG) to move PanDA towards becoming an OSG service
 - US PanDA instance supporting OSG VOs
- Predicated on generalizing PanDA for non-ATLAS use: modularity, packaging, documentation, no Oracle requirement: the program of the BigPanDA project
 - Jarka in particular is working to take up LSST DESC as the first supported BigPanDA VO using the new BigPanDA infrastructure
- Jarka's work program since June:
 - Implement MySQL based PanDA server as a tandem (with Oracle) implementation for OSG and other non-ATLAS use (done)
 - Establish Amazon EC2 based MySQL PanDA service for OSG pre-production and development (done)
 - to be cloned as stable production instance when need arises
 - Establish a VO-neutral PanDA monitor skeleton operating against the EC2 PanDA service for OSG VO PanDA monitoring (in progress)

Current Activities



- With the infrastructure now in place for OSG VO support, the next step and present work-in-progress is the LSST VO implementation
- PanDA-pipeline daemon (Tony's pipeline interface) moved from a dev EC2 instance to the PanDA OSG VO server
- Jarka and I working on implementing the pipeline-triggered PanDA job submission on the PanDA OSG VO server
- BNL now supports the LSST VO; a new PanDA site BNL-LSST created for LSST jobs
 - Submission to the BNL-LSST site being debugged

Next Steps

- Finish setting up the BNL-LSST PanDA site to receive real jobs from the pipeline
 - Software installed and DESC expert identified to help with simu workload
- Finish setting up LSST PanDA monitoring alpha version
- Enlist and work with users to get their workloads up and running
 - One or two DESCers at BNL are likely first candidates
- Get LSST user input to tailor LSST PanDA monitoring
- Extend to support opportunistic OSG (or non-OSG PanDA) sites – ie sites that don't necessarily have a local LSST person helping out
 - Key issues are software and data availability
 - For software, set up OSG's Oasis service for CVMFS based distribution
 - For data access, xrootd a good candidate
- Integrate the SLAC data catalog
 - The developer presented last week that it is a few months away from the new VO-neutral REST interface being ready

Utility of the pipeline/data catalog for BigPanDA?

- Pipeline provides a user-friendly web based interface to defining potentially very complex DAG workflows, saved and managed as tasks in projects
- Similar, at least superficially, to what we want from DEFT
 - Could it be useful?
- Data catalog similarly user-friendly, supports datasets and folders, dataset and file based metadata, VO-neutral... potentially interesting as a data management offering in association with BigPanDA?