



Tier3 and Panda

A Primer

Alden Stradling
12 October 2010

Tier3g and Panda

- **Major differences from T2:**
 - Lower maintenance manpower
 - No DDM Site Services
- **Addressing those differences:**
 - No SRM requirements
 - Simplify, simplify, simplify



Panda, the System

- **Databases:**
 - Jobs DB
 - Metadata DB (*schedconfig* and etc.)
- **Pilots:**
 - Submitter sends pilots to the site queue... somehow (options)
 - Pilots get jobs based on info from the DB, send updates during run and upon completion



Devil in the Setup

- The DB must be told about the queue (78 params!)
 - Most of them have no bearing on simple T3s
 - *Working on a standard T3 template*
- Getting the pilots into the local T3 queue
 - Can be complex... or as simple as a couple of command lines and a cron job
 - *In keeping with KISS, opt for the cron job*



Inserting a Queue

- Really, really simple (be sure to ask for permissions)

<https://twiki.cern.ch/twiki/bin/view/Atlas/SchedConfigNewController>

- Take the example queue parameters, copy them to a new config file named <my queue nickname>.py
- Take the example jdl parameters, copy them to a new config file named <jdl>.py in pandaconf/JDLConfigs (the jdl name should match the *jdl* field in the queue config file)
- Modify the file to reflect the differences on your site (if you are creating a **Standard** T3 site, there should be very little to change)
- `svn add; svn ci -m "Added my queue"`
- Wait (at most) 1/2 hour for cron insertion



Turn On the Queue

- https://twiki.cern.ch/twiki/bin/view/Atlas/ADCoS#HOWTO_change_queue_status
- Change your queue status to “brokeroff”
 - “Brokeroff” tells pathena not to auto-submit to the queue via brokerage



Queue

- You'll need to have a condor batch system already set up across the cluster (this is an assumed precondition)
- Set up condor with a default queue that you want to use for panda



Getting Pilots In

- On the head node:
 - Check out SVN version of autopilot (will become obsolete in the near future, but that's the procedure for now)
 - Create the scripts shown on the documentation Twiki
 - Create cron jobs to run those scripts
 - ~~Profit!~~ I have gone this far, but have run into authentication problems. I will make sure the docs are exhaustive when I figure out the solution.



The “Other” Way

- One can also set up a full Production-style pilot submitter:
 - CondorG
 - Adding to another pandameta table
 - Showing up on the Autopilot page
 - etc...
- I recommend AGAINST it as a general way forward
 - Too many moving parts, for unimpressive gains



Security

- For now, we match a regex in schedconfig to your pattern of compute nodes (sad panda)
 - Best to control user access:
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/PandaUserInfoTool>
- With the advent of Condor 7.4, we can use proxy sending. It's not a part of VDT right now
 - Works with AutoPilot, will continue to work with autopyfactory



Data Management

- This part is in flux
 - Present schema: *Year/Username/Dataset/File*
 - Requires pathena submissions with a text file, listing all the file paths! (not such fun)
 - Will be moving to a much more DDM-style schema
 - Allows the user to just specify a DQ2 “dataset”, not a list of files



Data Affinity

- Charles Waldman and Tadashi Maeno are working on it
- Xrootd (as presently constituted) does not apply, really
- Ways of modifying pcache to use inotify to call back and inform the brokerage
- Charles' email will be attached to the agenda as well



Existing Docs

- <https://twiki.cern.ch/twiki/bin/view/Atlas/Tierg3InstallingT3Panda>
 - Doug's recent documentation
- <https://twiki.cern.ch/twiki/bin/view/Atlas/PandaTier3>
 - Torre's original doc (small Twiki problems are being addressed)
- <https://twiki.cern.ch/twiki/bin/view/Atlas/SchedConfigNewController>
 - Modifying your own Panda entry



Forthcoming Docs

- Detailed discussion of schedconfig fields – with versions for T1/T2 and T3, to keep from scaring people
- Subversions for Panda configs, cron jobs and execution scripts
- Testing methods and troubleshooting, twiki FAQ
- User-directed submission documentation (made to be adapted for your site!)



Conclusions

- This is not yet ready for prime time (he says reluctantly) – it will need another few weeks in the oven
 - The Tier3 definition is in flux
 - The advantages of standardization will save substantial time
 - DDM simulation is evolving in the near term
 - In the meantime – make your cluster as standard as possible before your users get used to another way of doing things!
- ***Once the recipes are ready (and your cluster is standard), Panda installation will be a snap***

