



PanDA Summary


Kaushik De

Univ. of Texas at Arlington

ADC Retreat, Naples


Feb 4, 2011

Conclusions from Torre's Talk




- Highest priority – DB scalability, Oracle and noSQL
- Efficient Tape access/data carousel
 - Start with PRODDISK for Tier 1 (next slide)
 - Add special tape queue (in 2 slides)
 - Mix in PD2P for Tier 1 (previous discussion)
 - If we need further improvement, go to fixed 'carousel' scheme
- PanDA in the cloud – important work, need convergence
- glexec – try soon or give up!
- Longer term
 - Need ProdDb/PandaDB optimization
 - Dynamically defining jobs

PRODDISK at Tier 1



- **New idea from brain-storming yesterday**
 - Setup PRODDISK at Tier 1
 - If source is *TAPE, request subscription to local PRODDISK
 - Put jobs in assigned state, wait for callback, use _dis blocks etc (same workflow as used now for Tier 2)
 - Activate jobs when callback comes
 - Increase priority when job is activated, so jobs run quickly
 - Panda will set lifetime for _dis datasets on PRODDISK
 - Lifetime is increased when files are reused from PRODDISK
 - Cleaning via Victor: but different algorithm for PRODDISK (should make this change even for Tier 2 PRODDISK)
 - Clean expired files, or when 80% full (dedicated high priority agent?)

Special Panda Queue for Tape




- Need to regulate flow if we use tapes more
- Setup separate queue in Panda
 - Simplest implementation – use new jobtype, if source is *TAPE
 - Jobs pulled by priority as usual
 - Different base priorities – highest for production, medium for group production, lowest for user event picking (with fair share)
 - We can control this per Tier 1, throttle to balance load, use large queue depth to reduce number of tape mounts

Removing Cloud Boundaries



- **Multi-cloud Tier 2**
 - This will balance load, reduce stuck tasks...
 - Allow some large Tier 2's to take jobs from many clouds
 - Only after setting up FTS channels and testing
 - Put a list of clouds in schedconfigDB for a site (based on DQ2 topology – manual at first, automated later)
 - PanDA will broker as usual – some jobs will come from each Tier 1
 - May need to add weight (fraction per cloud) if we see imbalance
- **May even be possible to have Tier 1's set up like this, to help other clouds by using PRODDISK!**
- **Make CERN a Tier 1 with multi-cloud Tier 2's**

Other Panda Conclusions



- Need to reduce looping job limit to 3 hours for analysis pilots
- PD2P – already discussed this morning