PanDA Pilot Submission using Condor-G: Experience and Improvements

Xin Zhao, John Hover, Tomasz Wlodek, Torre Wenaus
Brookhaven National Laboratory
Jaime Frey, Todd Tannenbaum, Miron Livny
University of Wisconsin-Madison

PanDA (Production and Distributed Analysis) is the workload management system of the ATLAS experiment, used to run production and user analysis jobs on the grid. As a late-binding, pilot-based system, to maintain a stable and scalable pilot submission system is critical for PanDA operation.

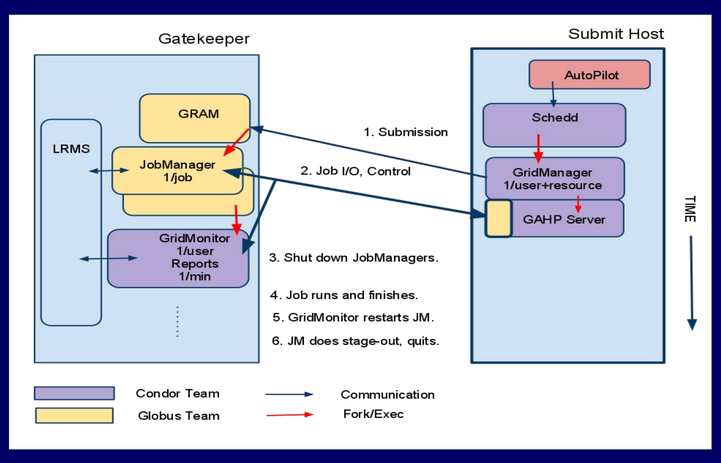

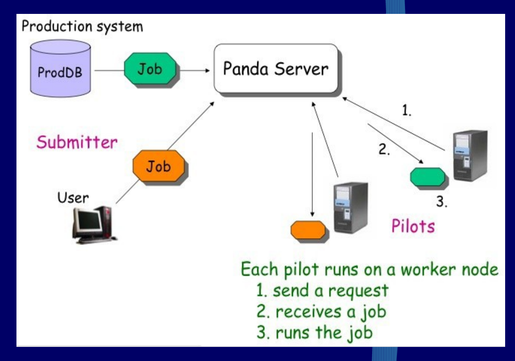

The ATLAS Computing Facility (ACF) at BNL, as the ATLAS Tier1 center in the US, operates the pilot submission systems for the US cloud. This is done using the PanDA "AutoPilot" scheduler components which submits pilot jobs via Condor-G, a grid job scheduling system developed at the University of Wisconsin-Madison.

Improving Scalability and Stability of Condor-G and AutoPilot

- Isolation of site specific issues from affecting submission to other sites
  - Separate GridManager process for each remote CE
  - GRIDMANAGER_SELECTION_EXPR = GridResource
- Introduction of the "Nonessential" job attribute
  - Unconditional removal and cleanup of stale pilots to prevent clogging the whole submission process
- Role-based throttle on limiting jobmanager processes on remote CE
  - Half for job submission, half for job completion/cleanup
- Grid Monitor restart behavior adjustable
  - GRID_MONITOR_DISABLE_TIME
  - Refined GridManager error handling to avoid flooding sites with the Grid Monitor jobs
- Better scheduling in AutoPilot to avoid overloading remote CE
  - Nqueue adjustable based on real job status
  - Multi-job pilots
- Bug fixes in Condor and underlying Globus libraries
  - Internal inefficient loops fixed
  - Integrated updtream Globus bug fixes into GAHP server
- Improved grid job status monitoring
  - condor_status -grid

PanDA AutoPilot Operation at BNL ACF

- Pilot submission in the US cloud
  - 92 PanDA queues at 43 gatekeepers, with HEPSPEC-06 ~ 111,000
  - Full scale production has >10,000 real jobs running all the time in the US cloud, daily pilots peak reached ~ 480,000 pilots
  - 5 Condor-G submit hosts (3 primary) at BNL Tier1
  - Stress test results: 50,000 jobs managed on one submit host; 30,000 jobs submitted to one remote CE (GT2 gatekeeper)

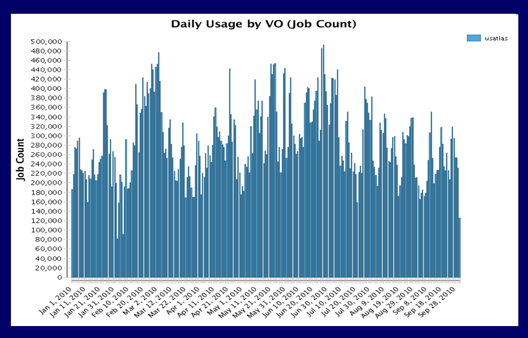

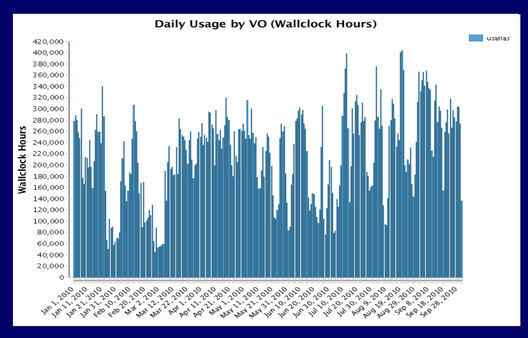

- Some best practices in Condor-G submission
  - Reduce frequency of proxy renewal
  - Avoid hard-kill (-forcex) jobs from client side
  - Optimize PanDA queues setup among several submit hosts to avoid starvation of pilots for any queues.
  - Multiple AutoPilot instances to increase injection rate for large PanDA queue
  - Caching condor_q command results to reduce load of condor Schedd on the submit host
- Monitoring of pilot submission
  - Dashboard & Nagios alerts
  - https://nagios.racf.bnl.gov/nagios/cgi-bin/prod/dashboard.php

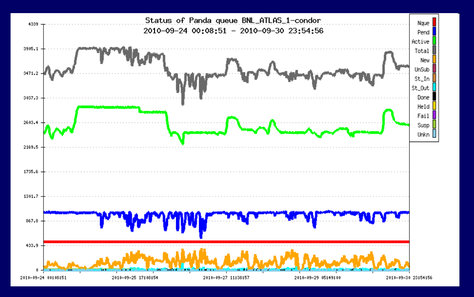

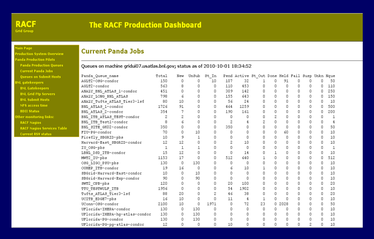

BROOKHAVEN NATIONAL LABORATORY

DEPARTMENT OF ENERGY • UNITED STATES OF AMERICA •