

# Multi-process Production

---

Douglas Smith  
*SLAC National Accelerator Lab.*

Paolo Calafiura  
*Lawrence Berkeley Lab.*

# Production testing of Athena MP

---

- Athena MP has been in testing for awhile, but just as singular test jobs from patched releases.
- Stable releases to support AthenaMP have been around for a few months.
- Moving tests into production system, where stable tags are created, and tasks requested that turn into sets of jobs.
- At meeting in Napoli start of Feb., in discussion it was desired to see the ability of this to work for data production.
- A new task force has been created in the ADC group for Athena MP, and results presented earlier this week (and today).
- Goal is to show data production fully proven by the summer.

# Use in Panda

---

- Changes to the Panda system to support multi-process running.
  - Added param to the queue def. “corecount”, defines the number of processes for athena to use.
  - Pilots use this value to set an env. Variable before job start, “ATHENA\_PROC\_NUMBER”.
  - The athena job then uses this variable to set the number of processes.
  - If the corecount value doesn't exist, the env. Variable doesn't get set.
  - If the release doesn't support multiprocess running, the job ignores the variable, still runs with one process.
- Queue can then run any job, and number of processes for athena are defined at run time, for the given number of cores on the machines in the queue.
- Changes were made to the log view, to see the mutli-process logs. (Each process has a separate log in a sub-dir.)

# Queues setup

---

- First testing queue setup at CERN, CERN\_8CORE
  - Defined as whole node machines, dedicated machines scheduling only these multiprocess jobs.
- Others queues are either setup, or getting setup
  - OU\_OSCER\_ATLAS\_MPI
    - Also now defined as a whole node queue.
  - UKI-SCOTGRID-ECDF\_8CORE
  - A new one at RAL, and Glasgow, not defined yet.
    - These three queues setup as shared resource scheduling, need to carve out a whole node to schedule jobs as submitted.
- We probably need more than just the CERN queue for testing. But at this time, probably don't need more queues than this.

# What is working...

---

- Production tests have been done with release 16.5.0.3
  - Multiprocess use of athena works, with caveats...
  - Reco jobs have been tested, from RAW to ESD and this works.
  - Except, because of a bug in the metadata, the validation step fails at the end, so you can only run with “omitvalid” flag on.
  - Also, there is bug in the production of ESD to AOD, and this fails.
  - So, normal production reco jobs actually don't work.
  - Also simu jobs don't work in 16.5.x.y. Releases.
- Confined to one specific type of job that works at this time.
- Bug fixes exist for these problems, and are going into new releases.

# This week's testing on new Release

---

- New release ready this week, 16.5.0.4:
  - This include bug fixes for the meta data issues
  - Should allow running with the validation step working.
  - Also allow creations of the AOD's and HIST's.
- Tests done on this over Mon-Tues:
  - OmitValid still on, and now including HIST output:
    - Jobs failed:
      - [http://panda.cern.ch/server/pandamon/query?job=\\*&taskID=292840&detail=no](http://panda.cern.ch/server/pandamon/query?job=*&taskID=292840&detail=no)
      - <http://panda.cern.ch/server/pandamon/query?job=1216080028>
  - Only ESD output, and OmitValid set to none:
    - Jobs failed:
      - [http://panda.cern.ch/server/pandamon/query?job=\\*&taskID=292689&detail=no](http://panda.cern.ch/server/pandamon/query?job=*&taskID=292689&detail=no)
      - <http://panda.cern.ch/server/pandamon/query?job=1216542810>
- Although now reports of stand alone tests that do work...

# This week's testing on new queues

---

- Still testing queue at OU – OU\_OSCER\_ATLAS\_MPI
  - Tests have been failing on mc10 reco tests, changes were made but:
    - [http://panda.cern.ch:25980/server/pandamon/query?job=\\*&taskID=293026&detail=no](http://panda.cern.ch:25980/server/pandamon/query?job=*&taskID=293026&detail=no)
    - <http://panda.cern.ch/server/pandamon/query?job=1216650060>
  - So, moving tests to data10 reco tests:
    - [http://panda.cern.ch/server/pandamon/query?job=\\*&taskID=294294&detail=no](http://panda.cern.ch/server/pandamon/query?job=*&taskID=294294&detail=no)
    - <http://panda.cern.ch/server/pandamon/query?job=1217104965>
  - Data10 reco failures seem to be pool file catalog access issues, Horst knows...
- Testing of queue in UK cloud – UKI-SCOTGRID-ECDF\_8CORE
  - Had problems with pool file access at site for reco jobs.
  - This was fixed, and test jobs were shown to work there:
    - [http://panda.cern.ch/server/pandamon/query?job=\\*&taskID=292691&detail=no](http://panda.cern.ch/server/pandamon/query?job=*&taskID=292691&detail=no)
  - Jobs now fixed to the queue only, with and without nprocs value
  - Jobs aren't starting there...
    - Problem with pilot submission, and fixed, but still stuck in activated at this time...

# Current plans...

---

- Waiting for a number new releases to use for testing.
  - New 16.6.x.y and 16.7.x.y. Releases:
    - Needed to test MC production.
    - Problems with job starting here still, to be fixed soon.
    - Might have difficulties in the 16.6.x.y branch, and will need 16.7.x.y. For testing.
- Should be ready very soon...
- Next steps –
  - show “real” production jobs work, and on all currently defined queues.
  - Create tables of memory use, time per event, and so on, for each queue.
  - Get experience with use of “whole-node” and “shared scheduling” queues.
  - Define some data to be used for analysis, and validate this production.
- People really do expect some demonstration of data production and validation by the summer.