

Panda Meeting 2023-04-05

(back to the [list of all Panda meeting minutes](#))

Zoom Link

Time

8 am PT

Attendees

Richard Dubois Wei Yang Fabio Hernandez Michelle Gower Peter Love Brian Yanny James Chiang Wen Guan Edward Karavakis Mikolaj Kowalik Jen Adelman-McCarthy Zhaoyu Yang Tim Jenness

Regrets

Agenda:

1. Update
 - a. Update from the Panda team on improving performance and reduce stress on various components.
 - b. IDF to DF submission
 - i. issues
 - ii. Richard: requirement
 - c. Site issues
 - i. Squid config at USDF was updated so jobs should be able to reach to the Panda@CERN via Squid instead of NAT
 1. hope it will scale better and address the timeout issue. Can we try it?
 - ii. status of the second ARC CE at USDF (not done yet, interrupted by Wei's travel)
 - d. Panda installation at USDF
 - i. Given the stress on Panda-DOMA, should we speed up the Panda installation at USDF?
2. Next steps
 - a. status of replication and butler ingestion at DFs for the 2.2i/defaults/test-med-1 collections in /sdf/data/rubin/repo/dc2 repo.
 - b. Can we try the full 7 steps on this collection at FrDF via Panda?
 - i. Note that CM team has run full seven steps at USDF (via Panda?) and Zhaoyu also ran a 5-task pipeline at FrDF
3. AOB

Notes:

1. Update:
 - a. Panda team's performance improvement plan: will work on several areas:
 - i. using pull mode for most short jobs, and push mode only for large memory jobs. The goal is to reduce the load and latency on CE, batch and HTCondor (in harvester)
 - ii. Trim non-termination message between Panda and iDDS. non-terminating msg: send, running; terminating: finished/failed
 - iii. use bulk messages to reduce load, and improve message system HW (believe this will help)
 - iv. improve DB backend capacity (how much this will help remind to be seen)
 - v. understand network timeout issue when pilot updates job status with Panda (is this a USDF NAT or Squid capacity issue, or other issue?)
 - vi. request pilot developer to fix out of order job status update (which can cause Panda to kill the wrong job)
 - vii. Panda clustering via M-core jobs (1 pilot, multiple jobs in parallel, sequential, or mixture), believe this will help reducing the load on CE, batch and HTCondor. Michelle thinks that this will not conflict with BPS clustering.
 - viii. Consider using EventService to keep track of M-core jobs.
 - b. User batch (IDF) to USDF [requirements](#):
 - Submit (bps) jobs from cloud to USDF
 - Can include payload different from canned stack at USDF
 - Submitter is identified to system can we use Panda share? Tim pointed out that if 10% batch capacity is allocated to users, but 8% are already used by jupyter, then there won't be much left.
 - Batch allocation set per user Butler will reside at USDF, use client server mode.
 - User can only write to their own repo space at the USDF
 - Register outputs in user accessible butler
 - Will continue discussion in the next meeting (in two weeks, Apr. 19)
 - c. Panda USDF installation update
 - i. Given the limitation of resource at CERN (panda doma) want to speed up USDF Panda installation
 - ii. Eddie: Psql DB schema deployed. Working on other deployment
 - d. Configured Squid to access panda server at CERN for job status update.
2. Next step

- a. CM team ran all 7 steps at USDF by login to USDF to submit panda jobs over a subset of DP0.2
- b. Want them to run the same 7 steps at FrDF (and UKDF)
- c. Zhaoyu has a pull request to enable CM team to do the above by using *bps submit* instead of *prun* from USDF.