BigPanDA@OLCF: Research Track 4 Update

Jack Wells and Sarp Oral

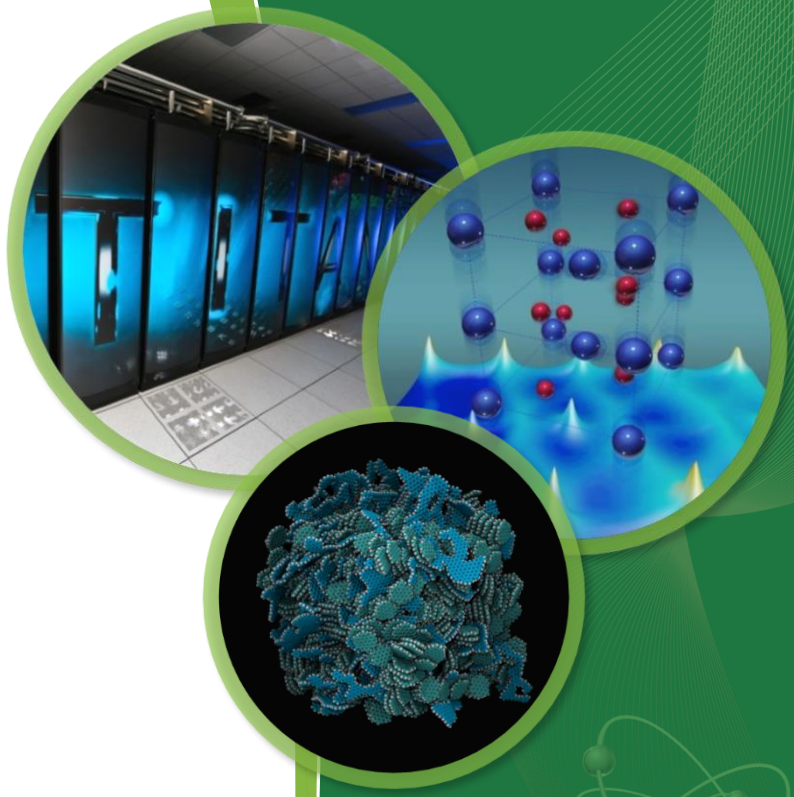

Scaling PanDA Operations for ATLAS

Explore the variations along the following four dimensions in the scheduling of PanDA jobs:

- Increase number of PanDA pilots/brokers running continuously at OLCF
  - **Currently at 20 pilots/brokers**
  - **Can we move forward with a study of preemption that would include an increase in pilot/broker count?**
- Increase the size range of compute nodes launched by PanDA
  - Varies between 7 and 300 nodes. No sense to increase with serial data management
- Increase time-duration flexibility for jobs launched by the PanDA pilot
  - Should be performed in systematic manner
- Increase diversity of HENP applications running at OLCF through PanDA.
  - TBD

Monitoring Shared Resources

- Quantify compute hours delivered through PanDA that would not otherwise been possible to deliver
  - Done
- Monitor the scheduling and delivery of leadership-class jobs on Titan during the operation of the PanDA workflow.
  - Done, but not documented. Opportunity for more structured study.
- Monitor the impact of the I/O associated with the PanDA-enabled payloads on the wider Titan I/O performance.
  - Finish analysis on ATLAS/Athena IO footprint and publish
  - Perform IO analysis on another payload for which we can recompile with Darshan (e.g., molecular dynamics)
  - Perform study on potential IO interference between large Capability jobs (e.g., XGC) and PanDA-enabled workload
- Quantify the marginal increase in network contention on Titan‘s interconnection network.
- Monitor the marginal increase in contention for access to the data transfer nodes (DTN). TBD
- ~~Monitor the increased load on the service-node infrastructure due to the PanDA Pilot jobs.~~
- Quantify the marginal increase in WAN traffic into and out of the OLCF data center. TBD

Monitoring Shared Resources, continued

- “Explore the tradeoffs between data pre-fetch and parallel file-system usage.”

- We expect the nature of the operational impacts on the broader system performance to depend rather directly on certain elements of the PanDA payloads deployed,
  - Payload check-pointing and/or **preemption**,
    - Can we envision a study of preemption with the current PanDA/ATLAS deployment?
  - Destination of the data offload,
  - Ratio between input and output data,
  - Dependencies between and among tasks,
  - Existence of "real time" access requirements.

PanDA server deployment

- “We expect that the experience gained as a result of performing this operational demonstration will be sufficient to position OLCF to deploy, toward the end of our two-year project, a PanDA server instance operated locally and capable of serving the broader OLCF user community.”
  - Started this activity in year 1.
  - Need more discussion and planning for operational model.
  - More discussion of monitoring tasks that are enabled through openshift containers.
- We also envision plug-in interfaces to other workload execution systems, e.g. Globus, or Condor.
  - What is the opportunity space with this?