

NGE-Harvester Integration

Phase 2

Planning NGE-Harvester Integration Phase 2

- Keep loosely-coupled design for integration between NGE and Harvester:
 - Complete deployment via OpenShift
 - Iterate on API, focusing on coordination with Harvester.
 - Deploy NGE as a continuously running, standalone service.
- Execute non-production MD and ATLAS-like workloads:
 - Which MD workload?
- Workload constraints:
 - Task runtime: 10m-12h; $\leq 65K$ tasks; $\leq 6K$ nodes;
 - compute-dominated tasks; pre-production reliability.
- Start characterizing capabilities that could be transferred/implemented in Yoda:
 - concurrent use of CPU/GPU;
 - different MPI implementations;
 - OpenMPI ORTE/ORTElib;
 - application-containers
- Investigate porting of the integration to Summit-dev.

Effort and Deliverables

- Duration:
 - Phase 2: Feb-Jun 2018
 - Phase 3 (upon review): July-Dec 2018
- Effort:
 - PanDA:
 - Pavlo, Ruslan: Development of coordination protocol between NGE and Harvester
 - Danila: Harvester, Pilot2
 - RADICAL:
 - Andre, Matteo: NGE API, RADICAL-Pilot
 - OLCF:
 - Arjun Shankar?
- Deliverables:
 - Demo integration for two types of workloads;
 - Characterization of capabilities: Experiments and measures on Titan
 - Contribute to the porting of Harvester and NGE to Summit-dev