NGE-PanDA and NGE-Harvester Integrations

Jan 17, 2018

Motivations

- Investigate the advantages of scale, flexible execution and interoperability outside production constraints.
- Support both backfill and regular queueing capabilities at the same time on Titan. Executions at 3750 nodes/12 hours and then at 11250 nodes/24 hours.
- Support both Molecular Dynamics and ATLAS-like workloads.
- Support concurrent execution of multiple workloads on the same pilot.
- Enable concurrent distribution of jobs on multiple OLCF resources, e.g., Titan and Summit.

Integrations: Design & Implementation

- Phase 1: loose coupling and minimal development effort
- Phase 2? On the base of the results, tight coupling (more development effort).
- Implement NGE as a persistent service with a REST interface, hosted on an OpenShift container on a Titan's DTN.
- NGE uses a database to hold resource and task states: #cores and walltime; submitted, executing done.
- Implement a coordination protocol between PanDA Pilot and NGE: NGE advertises resources; PanDA Pilot pushes jobs; NGE updates jobs' states.
- Use Gluster and NFS shared filesystems for input/output data staging.

PanDA-NGE Integration: Proposed Architecture

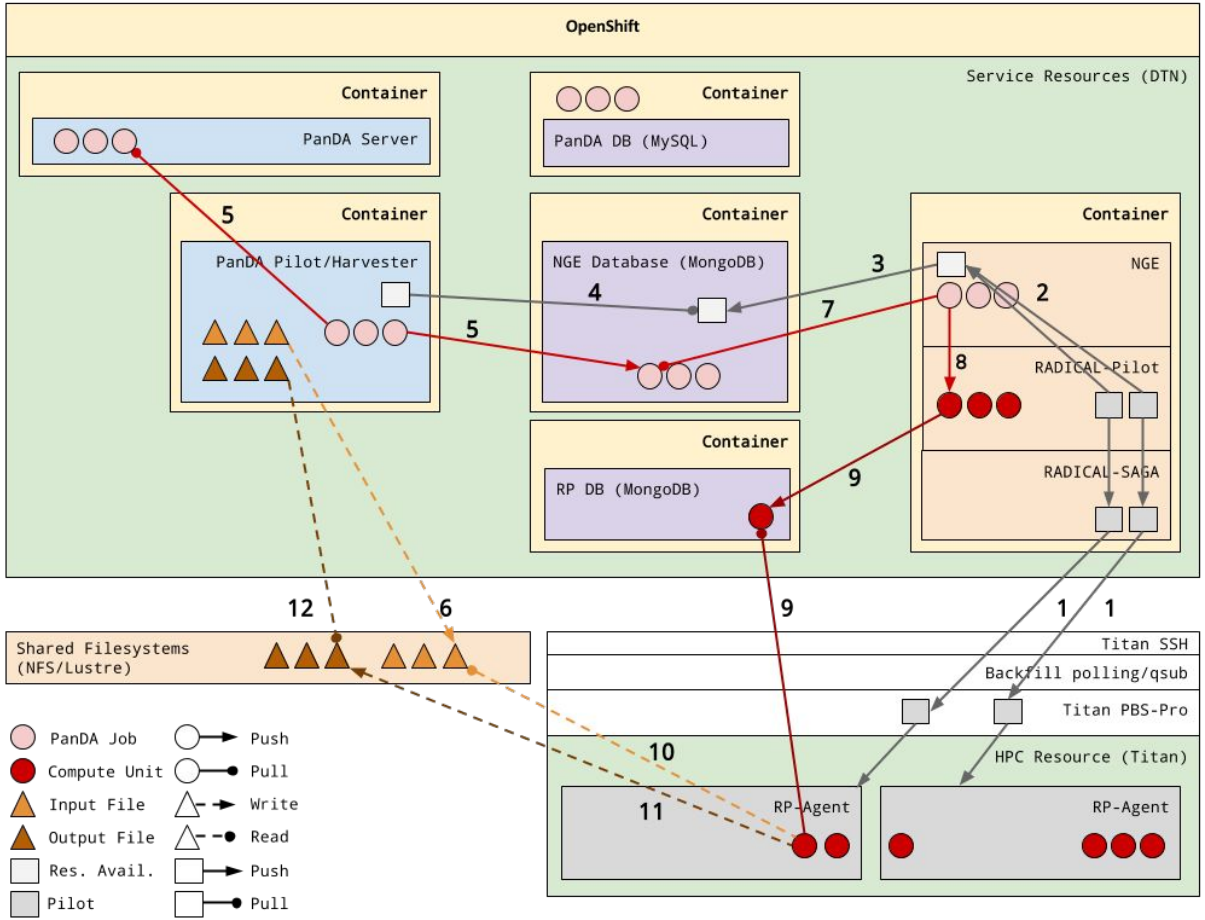

Development Status

- NGE:
  - Executions up to 131K cores on Titan.
  - Support for both CPU and GPU, exclusive or concurrent.
  - Optimization and dedicated scheduler.
- PanDA-NGE integration:
  - Phase 1 completed, demo at: https://drive.google.com/file/...
  - Code: https://github.com/ATLAS-Titan/PanDA-NGE
  - Design document: https://docs.google.com/document/…
- Harvester-NGE integration:
  - Deployed Harvester and NGE on Titan
  - Developed calls to NGE from Harvester

Development Roadmap

- NGE
  - Deployment of both NGE and RP on containers
  - Iterating/extending API for Harvester
- PanDA-NGE integration:
  - Phase 2?
  - Support of actual use cases/workload?
- Harvester-NGE integration:
  - Complete NGE API support in Harvester
  - End-to-end run of null workload and real-life workload as with PanDA-NGE.
  - Phase 2?