

NGE Project Report

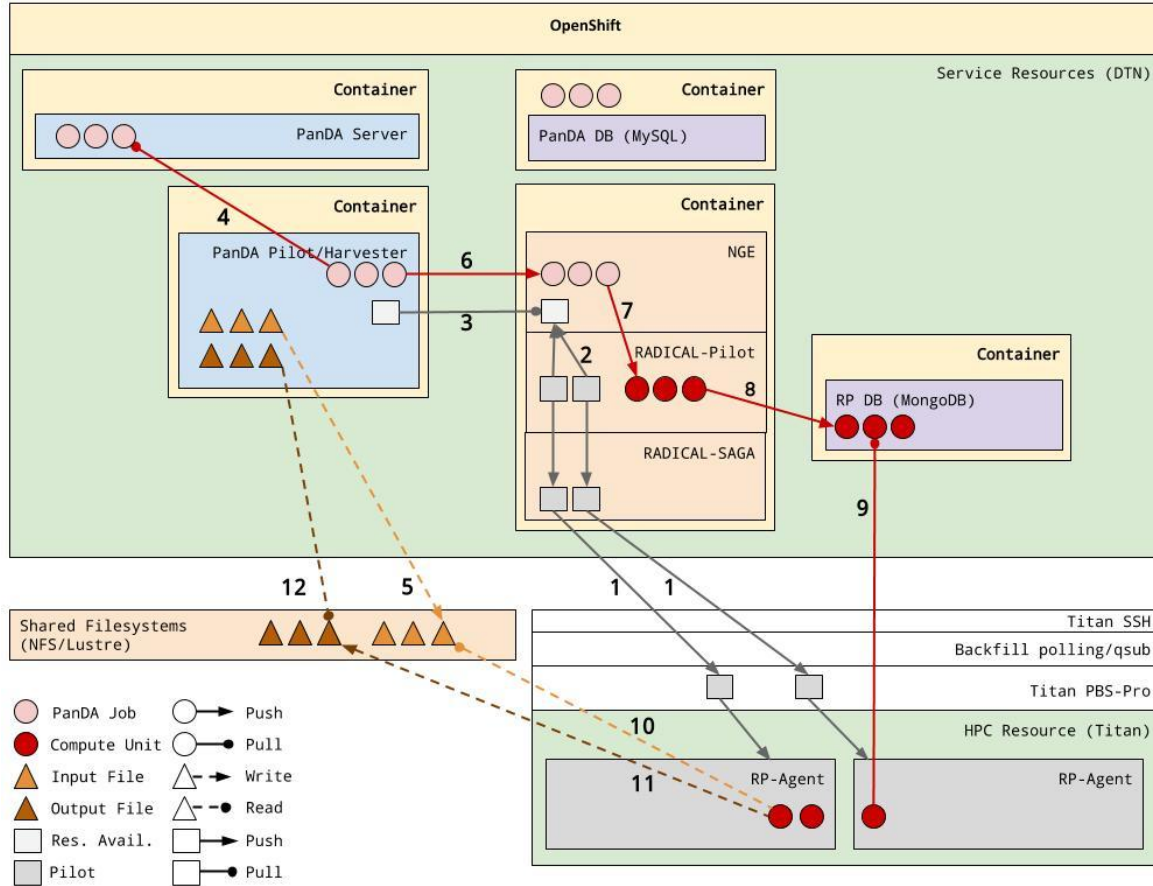
July 2018

Matteo Turilli, Andre Merzky, Pavlo Svirin

Progress of NGE/Harvester Integration Since TIM

- Design:
 - Architecture: direct communication between Harvester and NGE. No more dedicated MongoDB instance.
 - Task submission: Harvester submits tasks to NGE in bulk instead of one by one. This reduces latency and number of REST calls.
- Deployment: <https://github.com/radical-cybertools/radical.nge>
 - Harvester deployed on an OpenShift container.
 - NGE still deployed on Titan (tunneled to DTN): We need dedicated help/tutorial from ORNL.
- Experiments:
 - Compare weak/strong scaling performance among: Harvester alone; NGE alone; NGE/Harvester integration.
 - Submit pilots either directly to a queue or by using the backfill functionality.
 - Execute Molecular Dynamics (MD); ATLAS; mixed MD/ATLAS workloads.

Architecture




Novelties:

- Eliminated dedicated MongoDB instance between Harvester and NGE.
- **#6:** communication via REST interface.
- **#6:** task bulk submission from Harvester to NGE

NGE Performance: Null Workload Execution

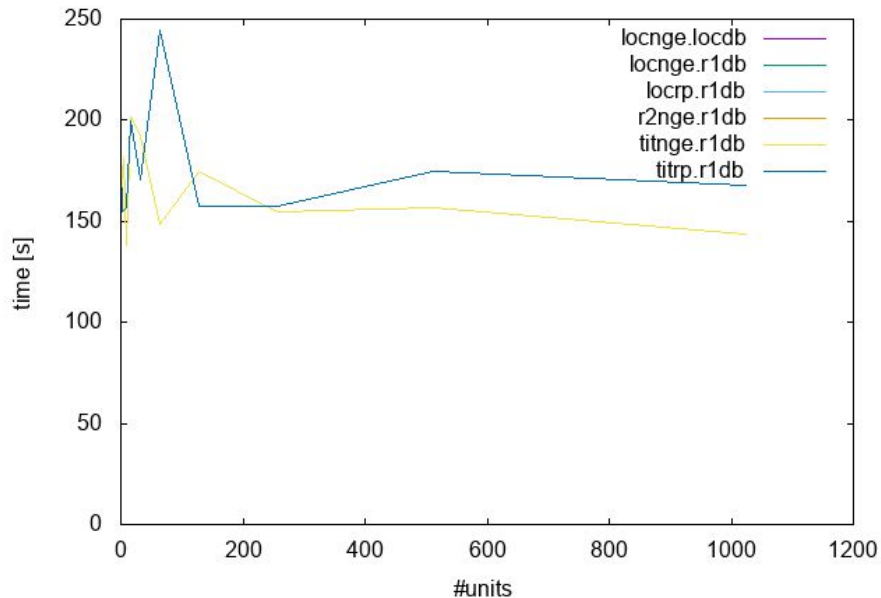
Workload Execution Time (NEW to DONE)




- Characterization and comparison of null workload execution time with NGE and RP in different deployment scenarios.
- NGE and RP have analogous performance.

NGE Performance: Resource Acquisition

Pilot Queue + Bootstrap Time (PENDING to ACTIVE)

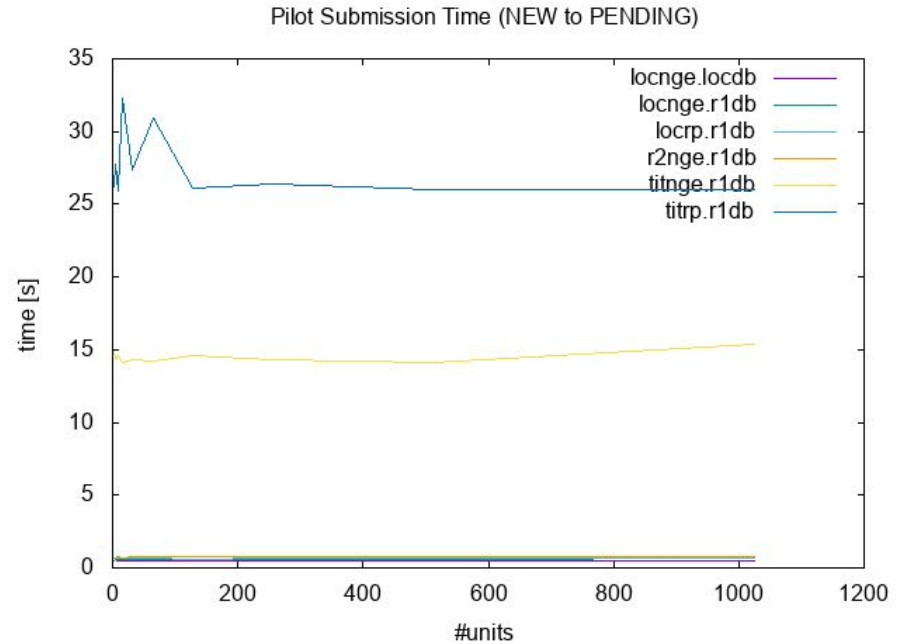
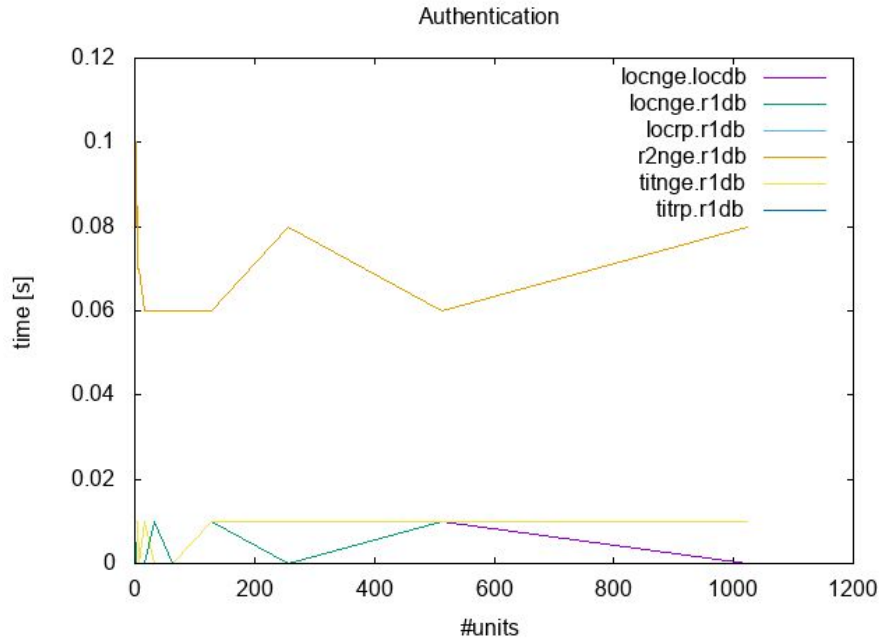


Workload Submission Time



Pilot queue and bootstrap & workload submission time: NGE and RP have analogous performance.


NGE Performance: Initialization



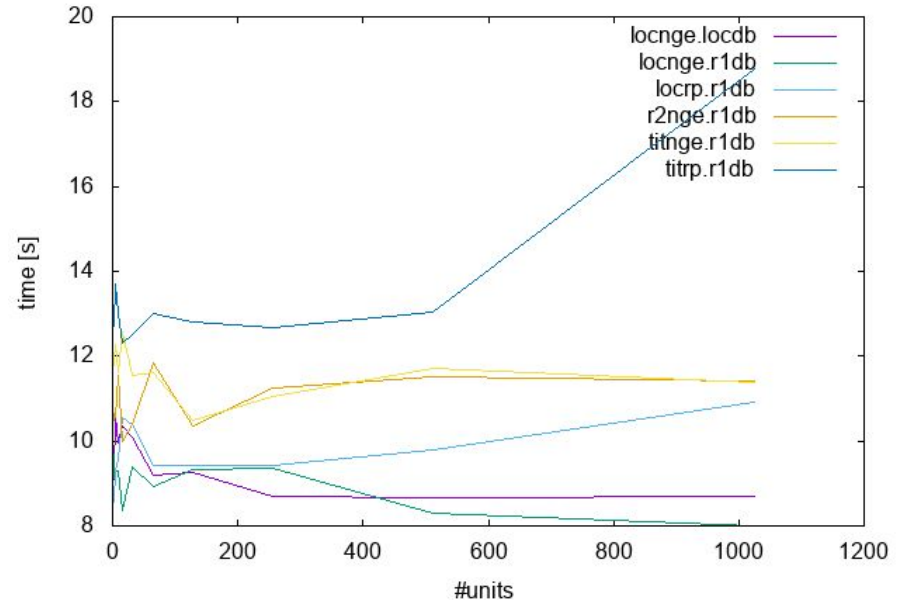
Authentication and pilot submission time: NGE and RP have analogous performance.

NGE Performance: Termination

Pilot Cancellation (ACTIVE to CANCELED)



Stack Termination



Pilot cancellation and stack termination time: NGE and RP have analogous performance.

Experiments: Molecular Dynamics

- <https://github.com/radical-experiments/radical.canon/tree/master/gromacs>
- Machine: Titan
- Queue: backfill, default
- Workload: Gromacs with protein test system (small number of atoms)
homogenous BoT, heterogeneous BoT

Tasks:

- Number of tasks: 1k, 8k, 32k
- cores per task: 1, 4, 8
- Runtime: 1min, 10min, 20min

Pilot:

- number of nodes: 8, 64, 512

NGE Only Experiments

	pilot size	# tasks	task runtime	task cores	ttc ideal	ttc real
small	8 nodes	1k	220s	1	1760s	1790s
medium	64 nodes	4k	300s	4	1200s	
large	512 nodes	16k	600s	8	1200s	

Planned Activities

- Experiments:
 - Run MD experiments with NGE and Harvester/NGE and compare weak and strong scaling performance.
 - Run ATLAS workload via NGE and compare event rate performance with current Harvester deployment.
 - **Run mixed workload experiments.**
 - Repeat most relevant experiments using backfill functionality.

Planned Activities

- Development:
 - Iterate on current integration further developing the NGE connector in Harvester.
 - Extending profiling and analytics capabilities of NGE stack to Harvester stack, enabling fine-grained characterization of MD workload executions.
 - **Support Summit on NGE: no change should be needed in Harvester, showing the potential for the isolation of concerns enabled by the design of the Harvester/NGE integration.**
 - Port design features to Yoda on the base of the results of our characterization.