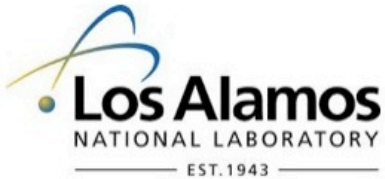

HPCOR 2014

Session Name: Workflows

Co-Chairs: Shreyas Cholia, Kevin Harms

Contributors

- Kevin Harms harms@alcf.anl.gov (Chair)
- Shreyas Cholia scholia@lbl.gov (Chair)
- Yao Zhang yaozhang@alcf.anl.gov
- Laura Monroe lmonroe@lanl.gov
- Sasha Ames ames4@llnl.gov
- Jeff Long long6@llnl.gov
- Norbert Podhorszki pnorbert@ornl.gov
- Andy Wilson atwilso@sandia.gov
- Rudy Garcia rudgarc@sandia.gov
- Laura Biven laura.biven@science.doe.gov
- Stephen Bailey stephenbailey@lbl.gov

Top Findings

Opportunities

- Federated Identity
- Workflows across sites (SDF)
- Documenting WF catalog for potential users

Best Practices

- HPC staff directly involved in WF support
- Documentation
- Virtualization (Future)
- Supporting Backend WF Services (Future)

Challenges

- Security
- Exascale challenges
- Re-using existing tools
- Integration with schedulers

Q.0: What are workflows?

- Workflow, work orchestration: Sequences of compute and data-centric operations
- automate interoperability of applications
  - automate provenance tracking -> enable ability to reproduce results
  - assist with data movement
  - monitor simulation
  - driving/steering simulation run
  - data processing of experimental data (including near-realtime processing)
- HPC batch systems - workflows help work with (around?) batch scheduler and queue policies
- Types of Workflow Tasks:
  - Bag of tasks (DAG)
  - Map-Reduce
  - In-situ
  - Tracking Provenance / Data Movement

Q1. What are your major strategies and initiatives over the next 5 & 10 years? How do they affect staffing levels?

- Move towards formal support for Workflows
  - ALCF – investigating; NERSC, LLNL – formal support; OLCF – limited support between Rhea and Titan.
  - Will require staffing commitment
- Next generation computing systems will impose new constraints.
- Handling the following use cases
  - Designing systems to handle high IO
  - In-situ processing
  - Adaptive analysis
  - Near real-time analysis

Q2. What are your current efforts and/or site configuration in this area?

- Number of active tools being used and in dev
  - OLCF - Kepler - run on Rhea linux cluster and submit jobs to Titan, Hadoop, custom one-off: Dataspaces (in-situ) + Adios + job scripts, Swift
  - PNNL - Velo
  - ALCF - custom one-off tools, allow running script(s) on dedicated script host.
  - SNL - Hadoop / Accumulo / Solr / Pig, custom one-off clusters
  - NERSC - Hadoop, Firework, qdo, custom-off
  - LLNL- UQ pipeline, PSUADE, CRAM (both clusters and sequoia), Hadoop, custom off-one

Q3. What are your mandates and constraints?

- Mandates
  - Support of data-intensive science (leading us to workflows)
  - Connecting experimental facilities with HPC centers
- Constraints
  - Security Policy
  - Integrating with system software
  - Communication between applications

Q4. How do you forecast future needs and requirements?

- Open Question
- What are metrics needed to evaluate workflows?
  - performance, throughput, ability to handle different classes of problems, feature sets, easy of use

Q5. What are the biggest challenges and gaps between what you can do today and what will be required in 5 - 10 years?

- Security
  - Within sites; across sites
- Storing intermediate results on disk not feasible in exascale – in-situ analysis
- User education
  - What exists? How do I pick the right tool for my workflow? How do we prevent people from always writing their own
- Scheduling challenges
  - Batch systems can't handle a million jobs
  - Near-real time analysis
- Typically need a management services outside the batch enviroment
  - Databases, Task Managers, Master Servers, Web Server

Q6. What opportunities exist for productive collaborations among DOE HPC centers?

- BOF at major conference to create a catalog of workflow services and pros/cons
- Federated authentication/authorization between facilities
  - agreement within/between sites
  - existing solutions are possible, mostly limited by policy.
- Share VM images across sites
- SDF for workflows across sites.

Q7. Describe some practices that you think are effective as well as lessons learned that would be helpful to other centers?

- We might be in a pre-best practices phase.
- Enable a virtual machine infrastructure internal to center to spin up supporting services
  - ability to run the same workflow on a laptop as well as the center.
  - support standard VM image "system", docker etc.
- HPC staff need to be more deeply involved in development and deployment of WF tools.
  - There needs to be general recognition that facilities will have to support workflows.
- Workflow tools need to have tighter integration with different batch managers so that it can submit to different job schedulers. Generic job specification language.
- Better Documentation
  - Helping guide users towards the right tools and how to implement workflows
  - List of tools, pros and cons - Feature matrix. How to implement workflows. How to choose between tools. Eg. trade-offs between local vs. remote analysis
- System Configuration
  - Can we run a more full featured linux OS (shared libs etc.)
  - Ability to talk to the network
- Having a common auth infrastructure enables cross-site workflow
- Lesson learned: users still want to write their own workflow engines… why is this happening? Can we guide users to existing tools?