



Software Infrastructure

(Data Management L3 Area, WBS 1.09.03)

Ted Kisner & Debbie Bard

Overview

- Establish key interfaces used by Simulation and Analysis L3 areas
- Enable Simulation and Analysis workflows to run on a wide range of computing platforms used by the project
- Indexing of data / metadata and provenance tracking of generated products
- Automation of both real-time processing and bulk reprocessing

Data Interfaces

- Formats on disk and in memory
 - SPT3G hand-off from DAQ, used for on-site computing and HTC platforms
 - TOAST used at HPC centers
- Indexing of data
 - Starting point is metadata from DAQ
 - Interfaces to Data Movement L3
 - Used for data selection and updated by data processing

Processing Interfaces

- API for simulation and analysis operations (G3Module / toast.Operator)
 - Minimize interface differences between frameworks (underway)
 - Copy / redistribute only when necessary
- Tools for constructing workflows for all our compute needs on HTC / HPC resources
 - Rapid data quality checks, transient detections, bulk reprocessing for data releases, Monte Carlo simulations, etc
- Optimizations sufficient to meet requirements on available compute resources
 - GPUs on Perlmutter (NVIDIA) and Aurora (Intel)

Deployment & Execution

- Software Deployment
 - Range of solutions, depending on target compute platform: containers (docker / singularity), Python packages, CVMFS
 - Most of the release process can be automated using existing practices
- Job Management
 - Need to run workflows on local, HTC, and HPC systems
 - Job tracking across all resources- separate but related to data provenance tracking
 - Automate as much as possible; Dedicated personnel resources for managing the execution of these systems
- Leverage upcoming Superfacility infrastructure for HPC work (next slide)

Superfacility Model

- Growing need to support workflows that can run across DOE/ASCR facilities (and institutional computing, cloud...)
 - Focus on connecting compute, network and experiment facilities - “superfacility” model
- A suite of tools for integrated use at multiple centers is being developed (at NERSC, and also at other ASCR/DOE facilities). This includes:
 - Federated Identity (i.e. using home institution as best source of identity)
 - Services gateway (typically K8-based, eg Spin/Slate) for databases, web portals, workflow management..
 - REST APIs for: Data movement / staging; System status; System reservations; Submitting / managing compute jobs...
- This is being discussed/coordinated at an ASCR level - aim is to have a coherent set of best practices for running cross-facility workflows.

Summary

- Currently making progress on specific tasks that are needed immediately on the schedule (prior to CD-1 / PDR)
 - Data and processing formats and interfaces
 - Deployment to upcoming NERSC system (Perlmutter)
- In the design phase of the project (now through CD-2 / FDR) need to:
 - Create and refine system for data indexing
 - Create draft system for job / data provenance tracking
- Continue tracking progress on Superfacility tools and how they can guide some of our automation efforts which are needed later during construction