Pegasus WMS and Panorama 360

George Papadimitriou[1], Ewa Deelman[1], Rafael Ferreira da Silva[1], Rosa Filgueira[2], Vickie Lynch[3], Anirban Mandal[4], Cong Wang[4], Jeffrey Vetter[3]
[1]University of Southern California – Information Sciences Institute [2]British Geological Survey [3]Oak Ridge National Laboratory [4]Renaissance Computing Institute

PEGASUS WORKFLOW MANAGEMENT SYSTEM

Overview of the Pegasus WMS

- Pegasus (https://pegasus.isi.edu) is a system for mapping and executing abstract **application workflows** over a range of execution environments
- The same abstract workflow can, at different times, be mapped **different execution environments** such as XSEDE, OSG, commercial and academic clouds, campus grids, and clusters
- Pegasus can easily scale both the size of the workflow, and the resources that the workflow is distributed over. Pegasus runs workflows ranging from just a few computational tasks **up to 1 million**
- Stores static and runtime **metadata** associated with workflow, files and tasks. Accessible via command line tools and **web based dashboard**
- Pegasus-MPI-Cluster enables fine-grained task graphs to be executed **efficiently on HPC** resources

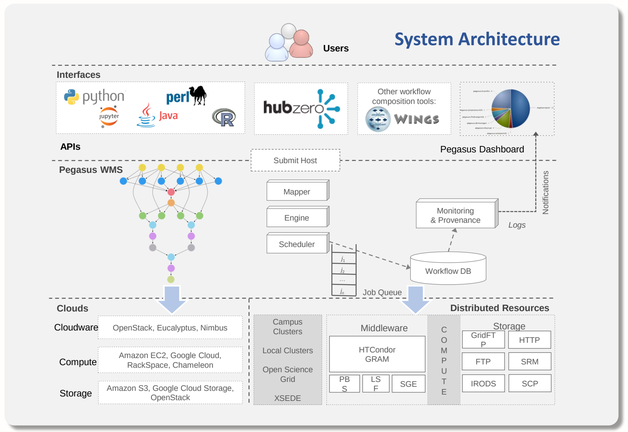

TRANSFERS WITH GLOBUS

Inter-site Data Transfers with Globus Transfer

Globus (https://www.globus.org/) lets you efficiently, securely, and reliably transfer data directly between systems.

We have upgraded support for Globus Transfer in Pegasus, and pegasus-transfer can connect to Globus API in order to submit transfer tasks and facilitate the execution of scientific workflows.

All these, can happen remotely from the machine that as Pegasus submit host

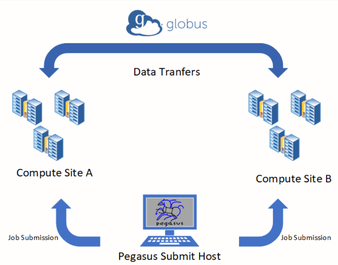

We have been running tests between NERSC and OSG, where workflows run part of their execution on one site and transfer intermediate data to the other site, to complete the run. This scenario can be useful in cases where the existence of a particular resource can significantly expedite the workflow execution.

In Collaboration With

USC Viterbi School of Engineering Information Sciences Institute
OAK RIDGE National Laboratory
renci
Argonne NATIONAL LABORATORY
Rensselaer
BERKELEY LAB

PANORAMA 360

Architecture and Project Goals

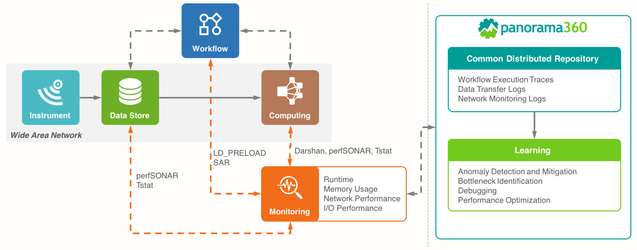

DATA CAPTURE
Characterization of instrument data capture, data summarization, and publication

REPOSITORY
An open access common repository for storing end-to-end workflow performance and resource data captured using a variety of tools

LEARNING
Development of ML techniques for workflow performance analysis and infrastructure troubleshooting

WORKFLOW EXECUTION MONITORING

Towards complete characterization of scientific workflows

With the Panorama 360 we are targeting in capturing data from every aspect of a scientific workflow. Either it is a compute job or a transfer job, we take advantage of a number of monitoring tools in order to capture the statistics we are interested in.

- Pegasus contains a module called Kickstart, which wraps the execution of jobs and provides execution statistics, such as duration, I/O, memory usage, which are available after the job completion.

Darshan
HPC I/O Characterization Tool

pegasus globus

Tstat

- Additionally there is an extension to Kickstart's functionality in Pegasus Panorama Branch which enables us to collect more refined traces with frequency as low as 1 second.
- Apart from Pegasus statistics we are compiling MPI-Jobs with Darshan, which provides us with POSIX and MPI I/O file access statistics. By default Darshan generates a summary, but by enabling the Darshan-Extended-Module we can collect traces from file accesses.
- On the network front, in order to acquire statistics related to data transfers we are using TSTAT logs and Globus logs. Globus service logs can provide us with transfer summaries (start time, throughput, etc.) and events during data transfers (failures, transfer completion, etc.). On the other hand TSTAT can give us low level network statistics, such as packet loss, which can reveal network related issues that affected the performance of a transfer.

DAKOTA IN PEGASUS WORKFLOWS

Exploring simulation parameter space with Dakota

Dakota (https://dakota.sandia.gov/) can be used for Parameter Studies, Design of Experiments, Uncertainty Quantification, Optimization, Calibration.
It also supports the implementation of Surrogate models and Nested models and has been designed to exploit parallel computing, when possible.

DAKOTA
Explore and predict with confidence.

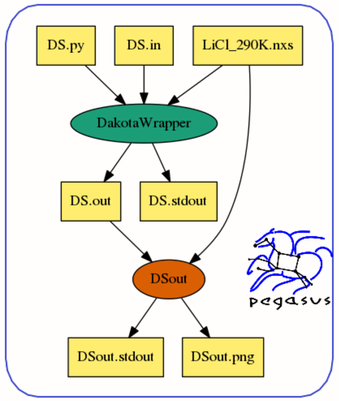

Simulations that rely on Dakota, can benefit from the integration with Pegasus. This integration could be implemented with several ways.

- Pegasus managing and executing Dakota Jobs, that handle the simulations themselves
- Dakota using Pegasus to execute the simulation workflows

In all cases Dakota studies can benefit from the staging capabilities of Pegasus (input files and executables), and from the ability to access supercomputing infrastructure.

IMPACT ON DOE SCIENCE

Diamonds that deliver!

Panorama enabled cutting-edge domain science research and development that has the potential to solve some of the challenges associated with **drug discovery and delivery**:

- The motions of a tRNA (or transfer RNA) model system can be enhanced when coupled with nanodiamonds, or diamond nanoparticles approximately 5 to 10 nanometers in size

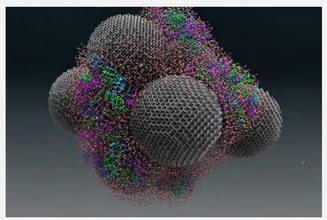

Water is seen as small red and white molecules on large nanodiamonds spheres. The colored tRNA can be seen on the nanodiamond surface. Image :Michael Mattheson, ORNL (https://www.ornl.gov/news/diamonds-deliver).

- We have developed an SNS Pegasus workflow to confirm that nanodiamonds enhance the dynamics of tRNA when in the presence of water. The workflow calculates the epsilon which best matches experimental data. These calculations used almost **400,000 CPU hours on a Cray XE6at NERSC**.
- The workflow runs NAMD parallel simulations, which varies the epsilon between -0.01 and -0.19 for each temperature specified (it requires 800 cores: equilibrium runs take ~1.5hs and production runs 12-16hs). AMBER's cpptraj removes global translation and rotation, and SASSENA calculates neutron scattering intensities from the trajectories (400 cores, 3-6hs). This workflow was used to computer 4 temperatures between 260K and 300K, which generated **~3TB of data**.