

“Science Automation using Workflows in the Big Data and Extreme Scale Computing Era”

Key idea:


**Need for multiple, customized, collaborating
WMS for
BD (ex-situ) and EC (in-situ)**

Ewa Deelman

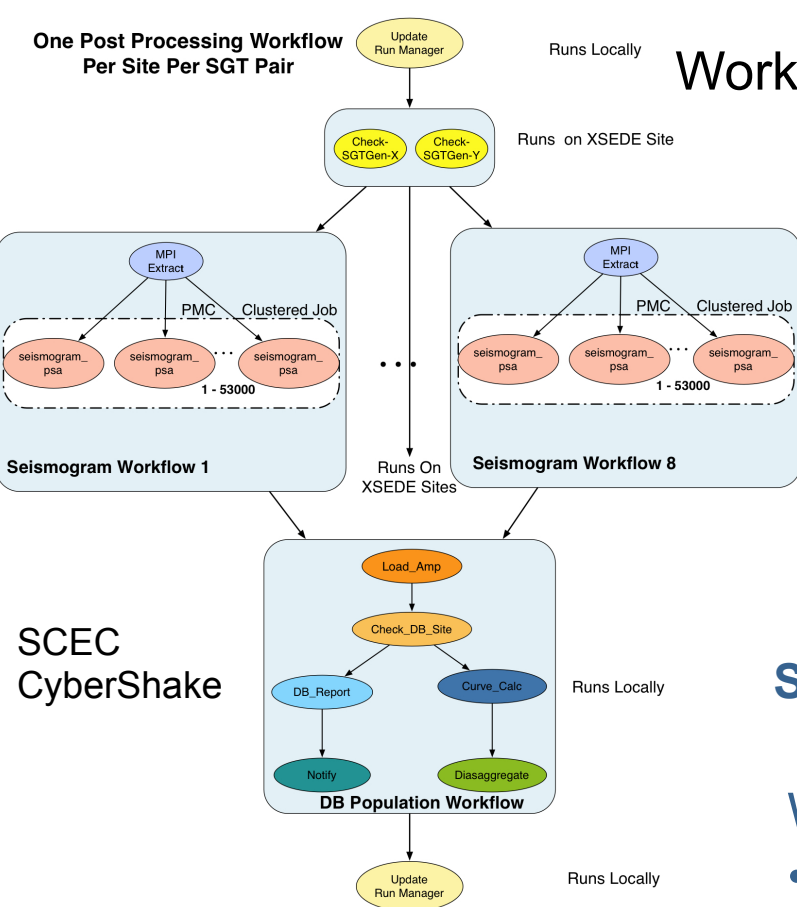
*University of Southern California
Information Sciences Institute*

deelman@isi.edu
<http://www.isi.edu/~deelman>

Pegasus Workflow Management System
<http://pegasus.isi.edu>



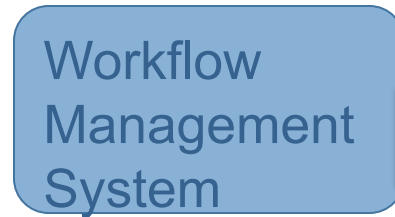
One Post Processing Workflow Per Site Per SGT Pair



Work definition



data



work

Local Resource

- Campus Cluster
- XSEDE
- NERSC
- Open Science Grid
- EGI
- FutureGrid
- Amazon Cloud

Separation of Workflow Definition and Execution

Workflow ensembles:

- Today workflows are managed on an individual basis
- As science is scaling up, it is necessary to manage entire workflow ensembles.
- Opportunity to optimize data transfers, reuse, and storage, across the wide area and inside EC systems.


1,144 geographic locations
 Uses Pegasus with execution on TACC's Stampede
 ~ 470 million tasks total
 Over 739 hours of computing
 ~ 636,000 tasks per hour
 45 TB of data
 12 TB being staged back for archiving

Tom Jordan, USC



Applications will be managed by multiple Workflow Management Systems

- Workflow Management Systems can potentially bridge the gap between big data and extreme scale computing
- Data needs to be staged to the EC resources and staged back
- Computations can involve multiple EC resources
- For efficiency a workflow management system may need to work *in situ* on an EC resource, coordinating fine-grained computation scheduling and data movement across the machine
- There needs to be a delegation of work or collaboration on workload management between BD WMS and EC WMS
- Each WMS needs to tailor and optimize the workflow execution to each specific environment
 - data and computation management decisions that occur inside an EC need to take into account energy efficiency, and thus data locality among others
- Need to worry about reliability and reproducibility
- Need to worry about interactivity with both types of WMSs



Interplay between BD and EC WMS needs to be explored

- Restructuring of the workflows for different environments
- Common capabilities that need to work together:
 - provenance capture (and linking), reliability, and performance
 - need tools for efficient provenance storage and query
- Data management at different scales
 - EC WMS
 - may deal with data in memory
 - potentially streaming data from/to the EC resource
 - Makes use of HPC libraries
 - BD WMS may
 - select to the best replica from a set of possible data repositories, select services, ECs
 - consider the proximity of computing to these storage resources.
 - trigger computations based on the influx of new data products
 - EC WMS may provide hints to BD WMS on how to stage the data into the extreme scale system
 - EC WMS may also give hints about how the output data is structured, or how it is streamed so that the BD WMS can reconstruct the results of the computation.

Workflows need to be easy to compose, reuse, launch, monitor, and interpret --- all from scientist's desktop.

