


Automating Real-time Seismic Analysis

Through Streaming and High Throughput Workflows

Rafael Ferreira da Silva, Ph.D.

USC Viterbi
School of Engineering
Information Sciences Institute



<http://pegasus.isi.edu>

Do we need seismic analysis?

San Francisco

Tokyo

Manila




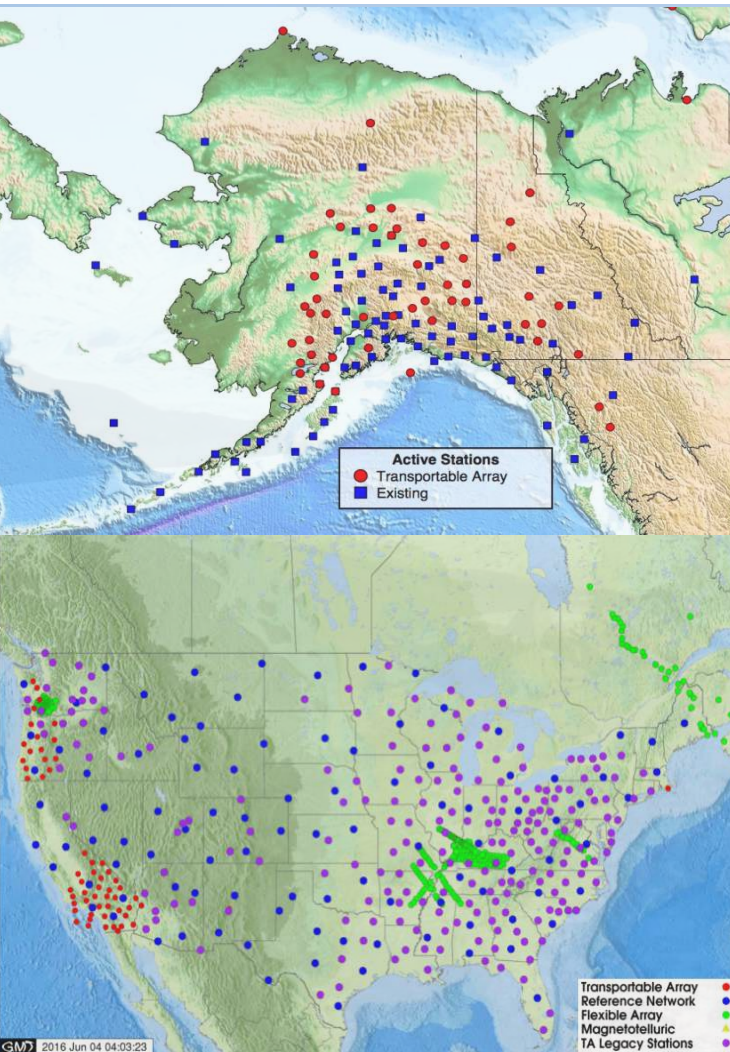
Jakarta

Osaka

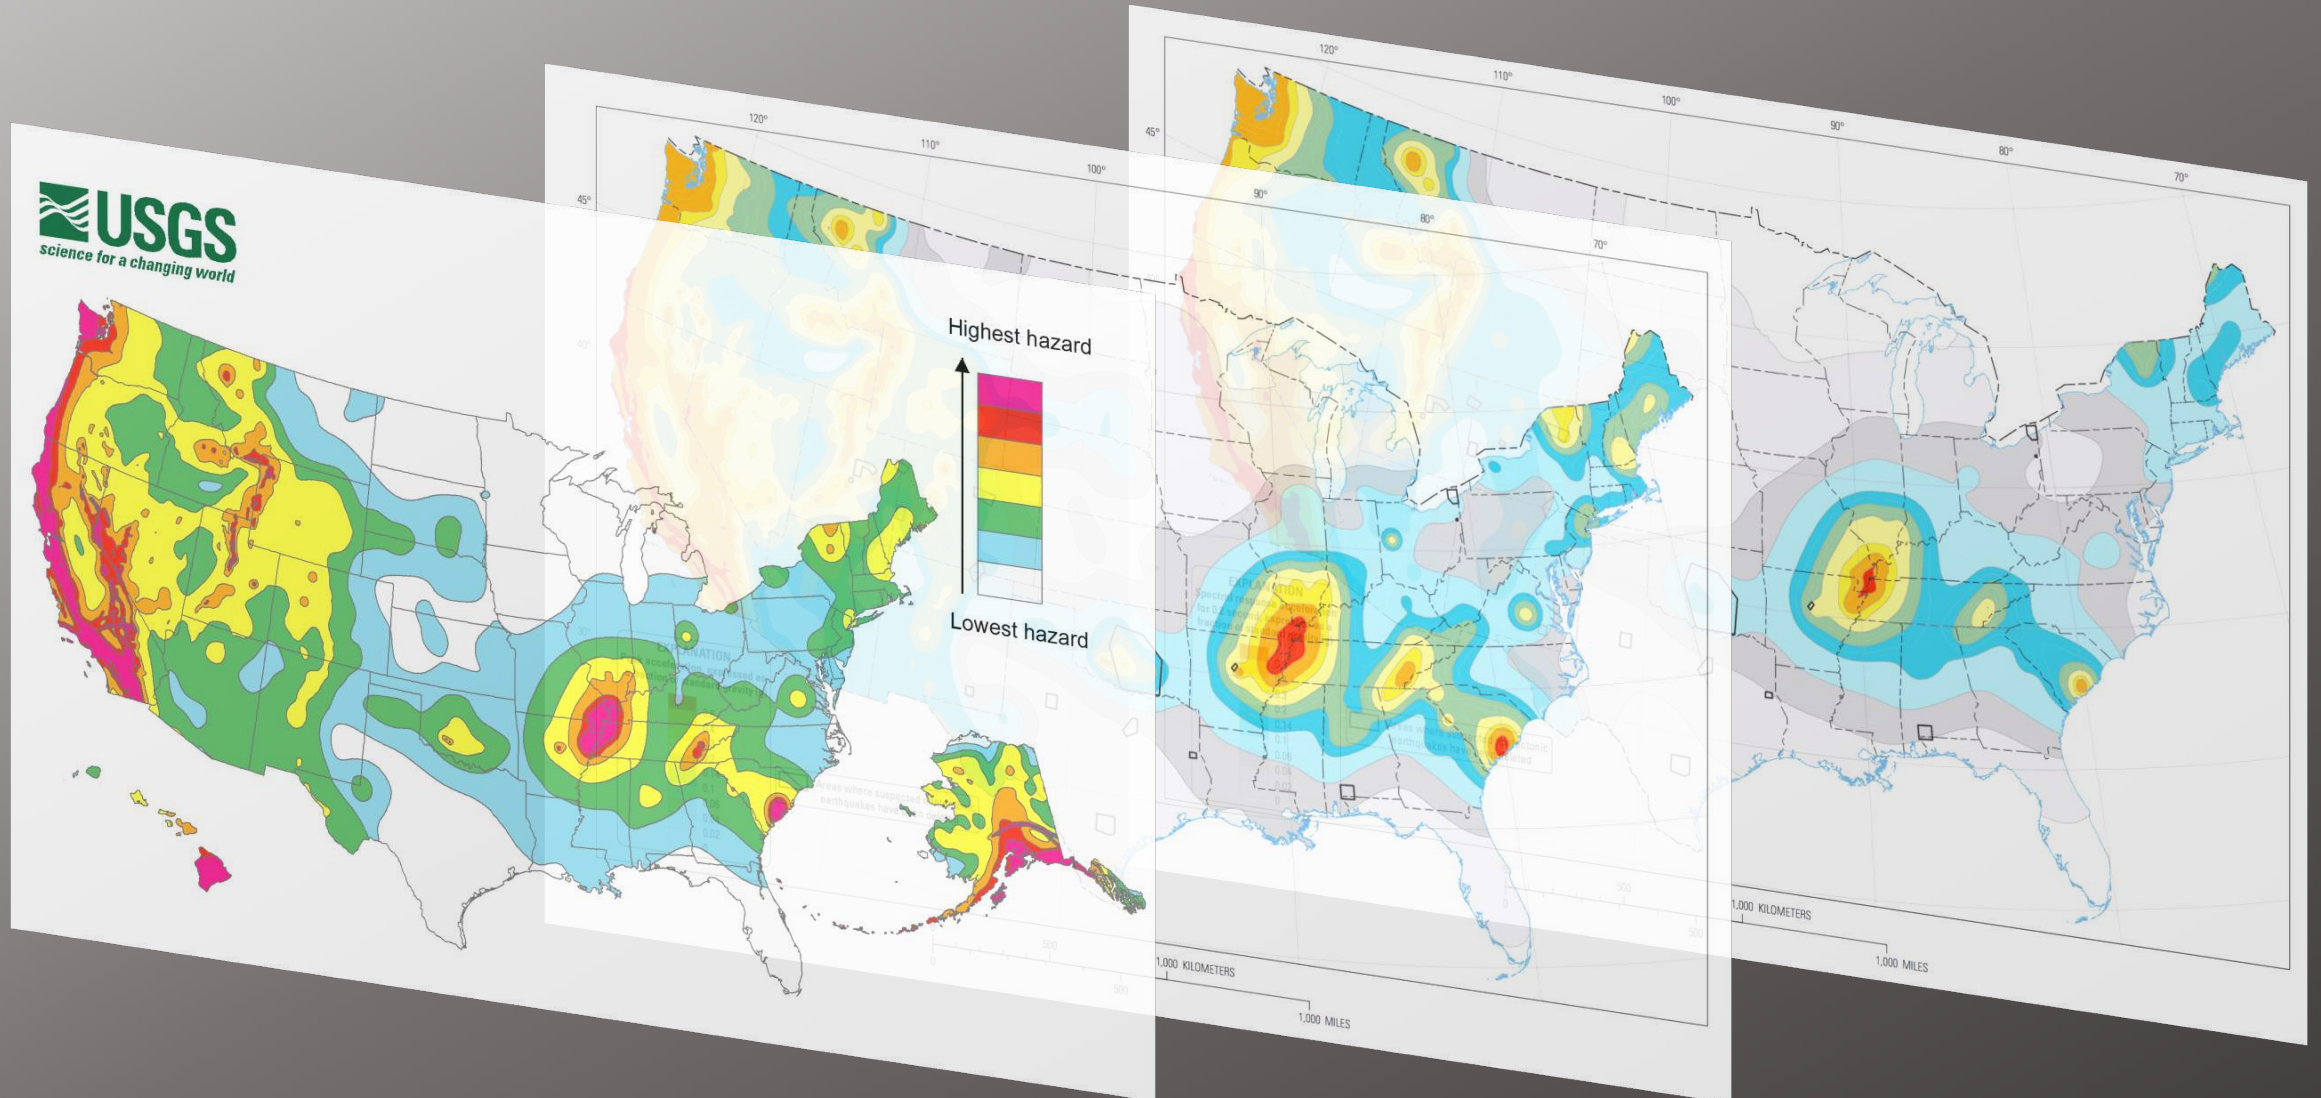
Los Angeles

USArray

A continental-scale Seismic Observatory



The development of reliable risk assessment methods for these hazards requires real-time analysis of seismic data



So, how to efficiently process these data?

Experiment Timeline

Scientific Problem

Earth Science, Astronomy,
Neuroinformatics,
Bioinformatics, etc.



Computational Scripts

Shell scripts, Python, Matlab, etc.

```
#!/usr/bin/perl
use strict;
use warnings;

my $script = "results_log.pl";
my $script_dir = "scripts";
my $script_path = "$script_dir/$script";

my $server = "10.1.1.1";
my $port = 8080;

my $url = "http://$server:$port/$script";

my $content = "Hello, World!";

my $response = get($url, $content);

print "Response: $response\n";
```


Distributed Computing

Clusters, HPC,
Cloud, Grid, etc.



Scientific Result

Models, Quality Control,
Image Analysis, etc.



$$U \frac{d^2}{dx^2} U^{-1} g = \left(U \frac{d}{dx} U^{-1} \right) \times \left(U \frac{d}{dx} U^{-1} \right) g$$

$$= \frac{d}{dx} \left[g' \psi' + \frac{1}{2} g \frac{\psi''}{\psi'} \right] \cdot \psi' + \frac{1}{2} \left[g' \psi' + \frac{1}{2} g \frac{\psi''}{\psi'} \right] \times \frac{\psi''}{\psi'}$$

$$= g'' \psi'^2 + 2g' \psi'' + \frac{1}{2} g \times \left[\frac{\psi'''}{\psi'} + \frac{\psi''^2}{\psi'^2} \right]$$

Analytical Solution


Automation

Workflows, MapReduce, etc.



Monitoring and Debug

Fault-tolerance, Provenance, etc.



What is involved in an experiment execution?



Why Scientific Workflows?

Automates complex, multi-stage processing pipelines

Enables parallel, distributed computations


Automatically executes data transfers

Reusable, aids reproducibility

Records how data was produced (provenance)

Handles failures with to provide reliability

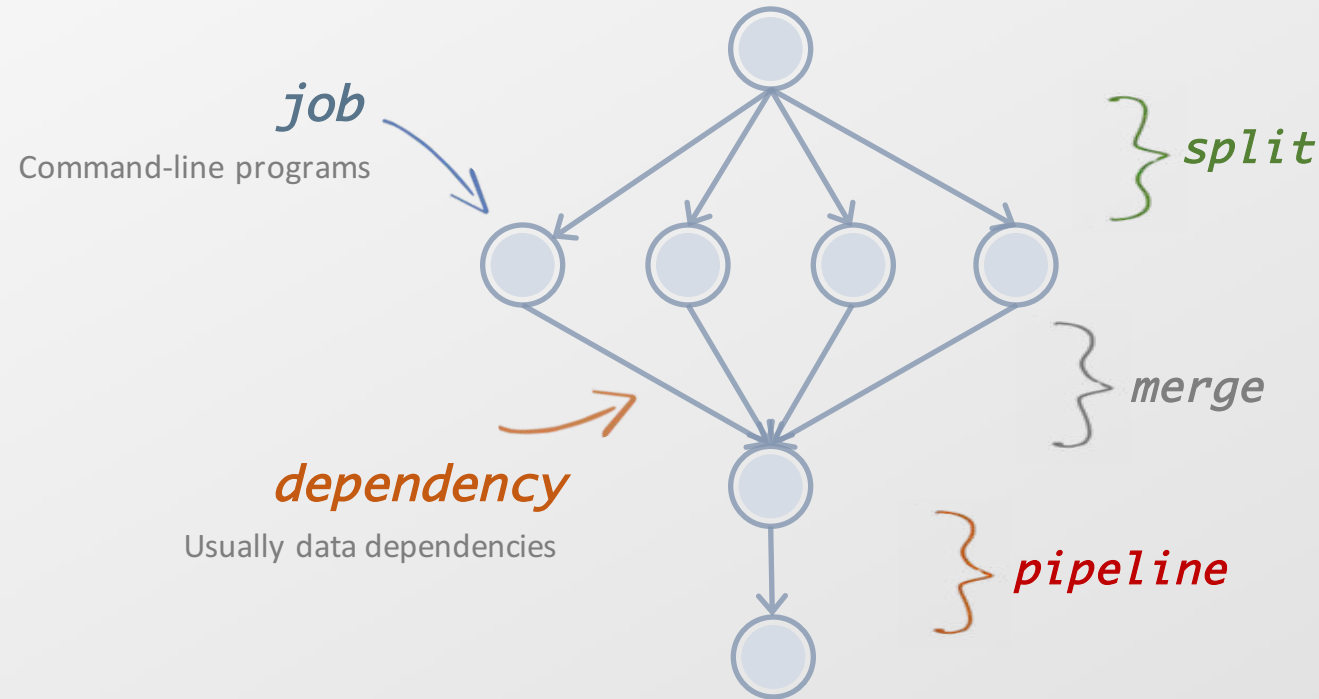
Keeps track of data and files



Taking a closer look into a workflow...

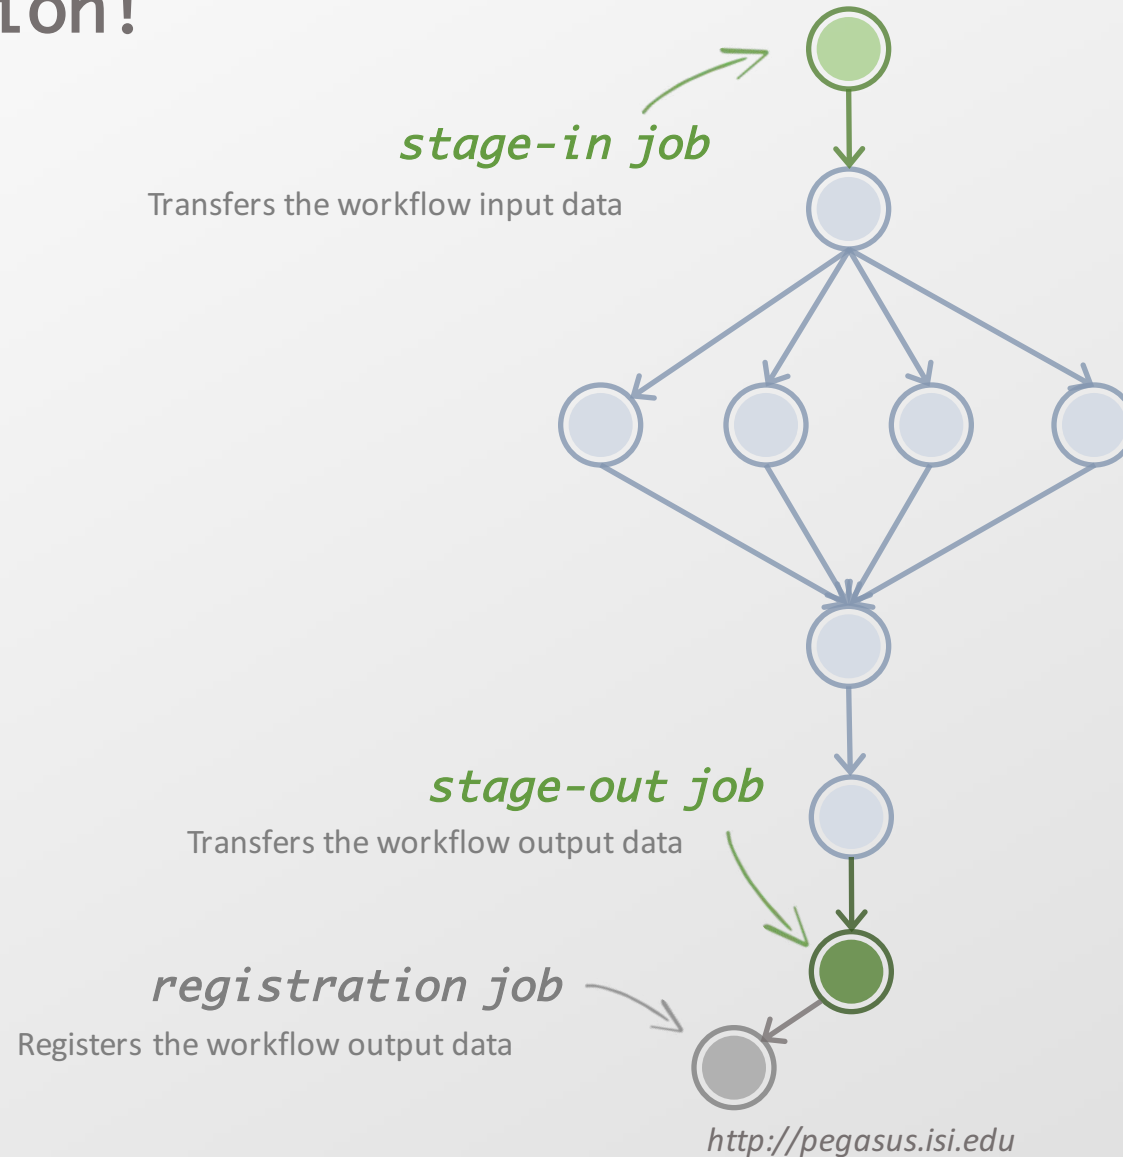
abstract workflow
executable workflow
optimizations
storage constraints

directed-acyclic graphs




From the abstraction to execution!

abstract workflow
executable workflow
optimizations
storage constraints

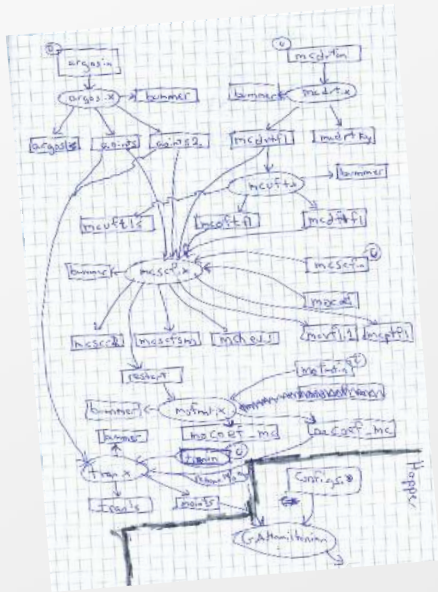


Optimizing storage usage...


abstract workflow
executable workflow
optimizations
storage constraints




Workflow systems provide tools to generate the abstract workflow



```
dax = ADAG("test_dax")
firstJob = Job(name="first_job")
firstInputFile = File("input.txt")
firstOutputFile = File("tmp.txt")
firstJob.addArgument("input=input.txt", "output=tmp.txt")
firstJob.uses(firstInputFile, link=Link.INPUT)
firstJob.uses(firstOutputFile, link=Link.OUTPUT)
dax.addJob(firstJob)
for i in range(0, 5):
    simulJob = Job(id="%s" % (i+1), name="simul_job")
    simulInputFile = File("tmp.txt")
    simulOutputFile = File("output.%d.dat" % i)
    simulJob.addArgument("parameter=%d" % i, "input=tmp.txt",
        output="%s" % simulOutputFile.getName())
    simulJob.uses(simulInputFile, link=Link.INPUT)
    simulJob.uses(simulOutputFile, link=Link.OUTPUT)
dax.addJob(simulJob)
dax.depends(parent=firstJob, child=simulJob)
fp = open("test.dax", "w")
dax.writeXML(fp)
fp.close()
```



abstract workflow



nextflow



Nimrod



Which Workflow
Management System?




swift

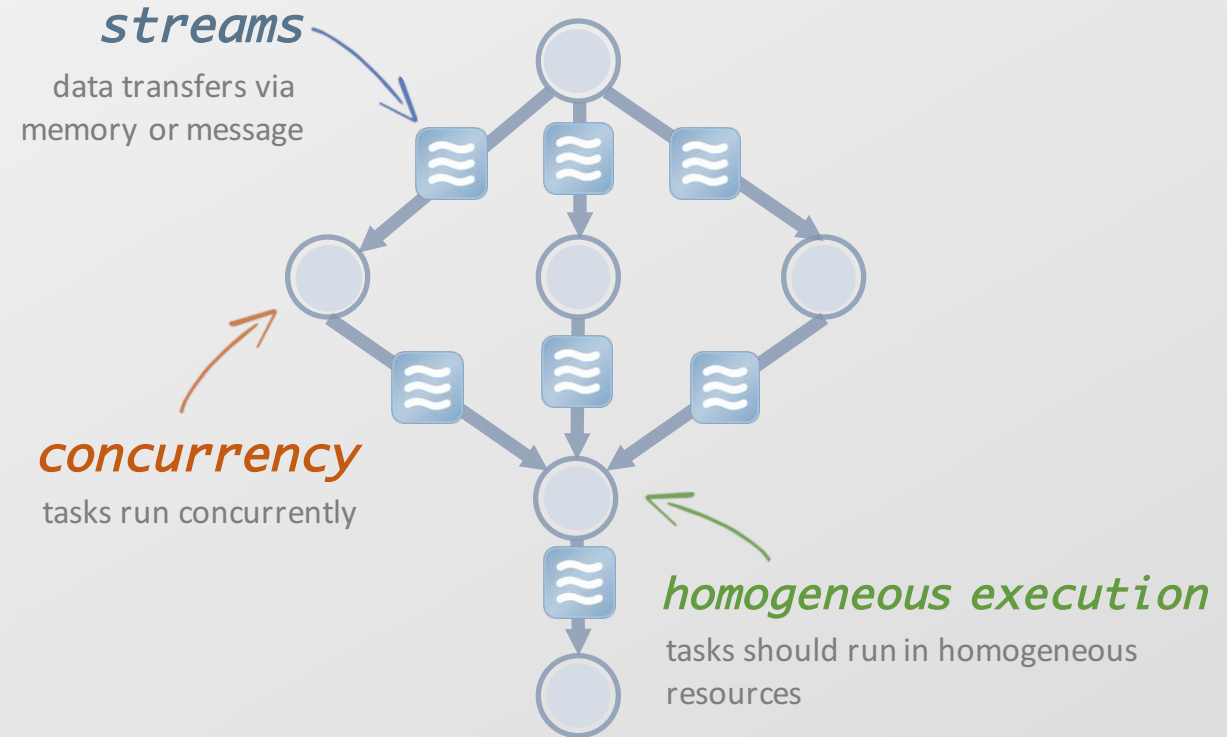


...and which model to use?

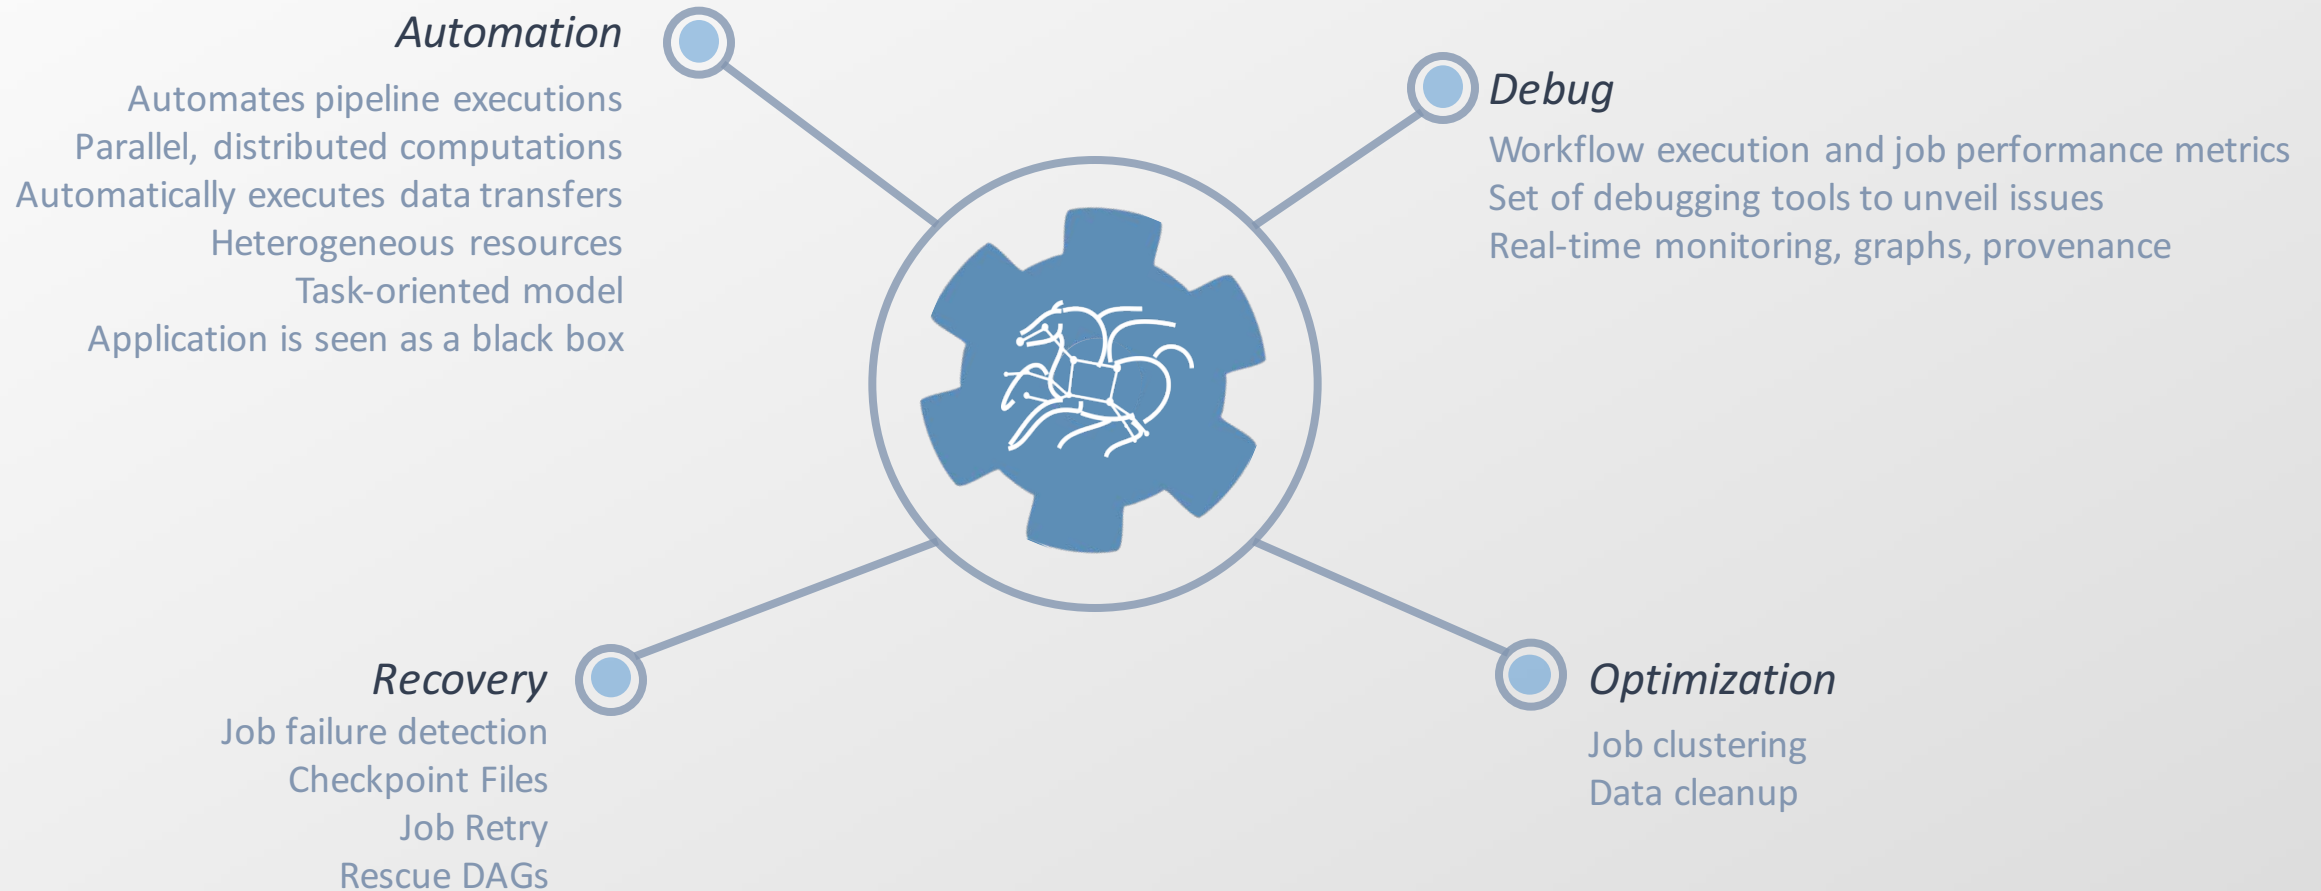
task-oriented



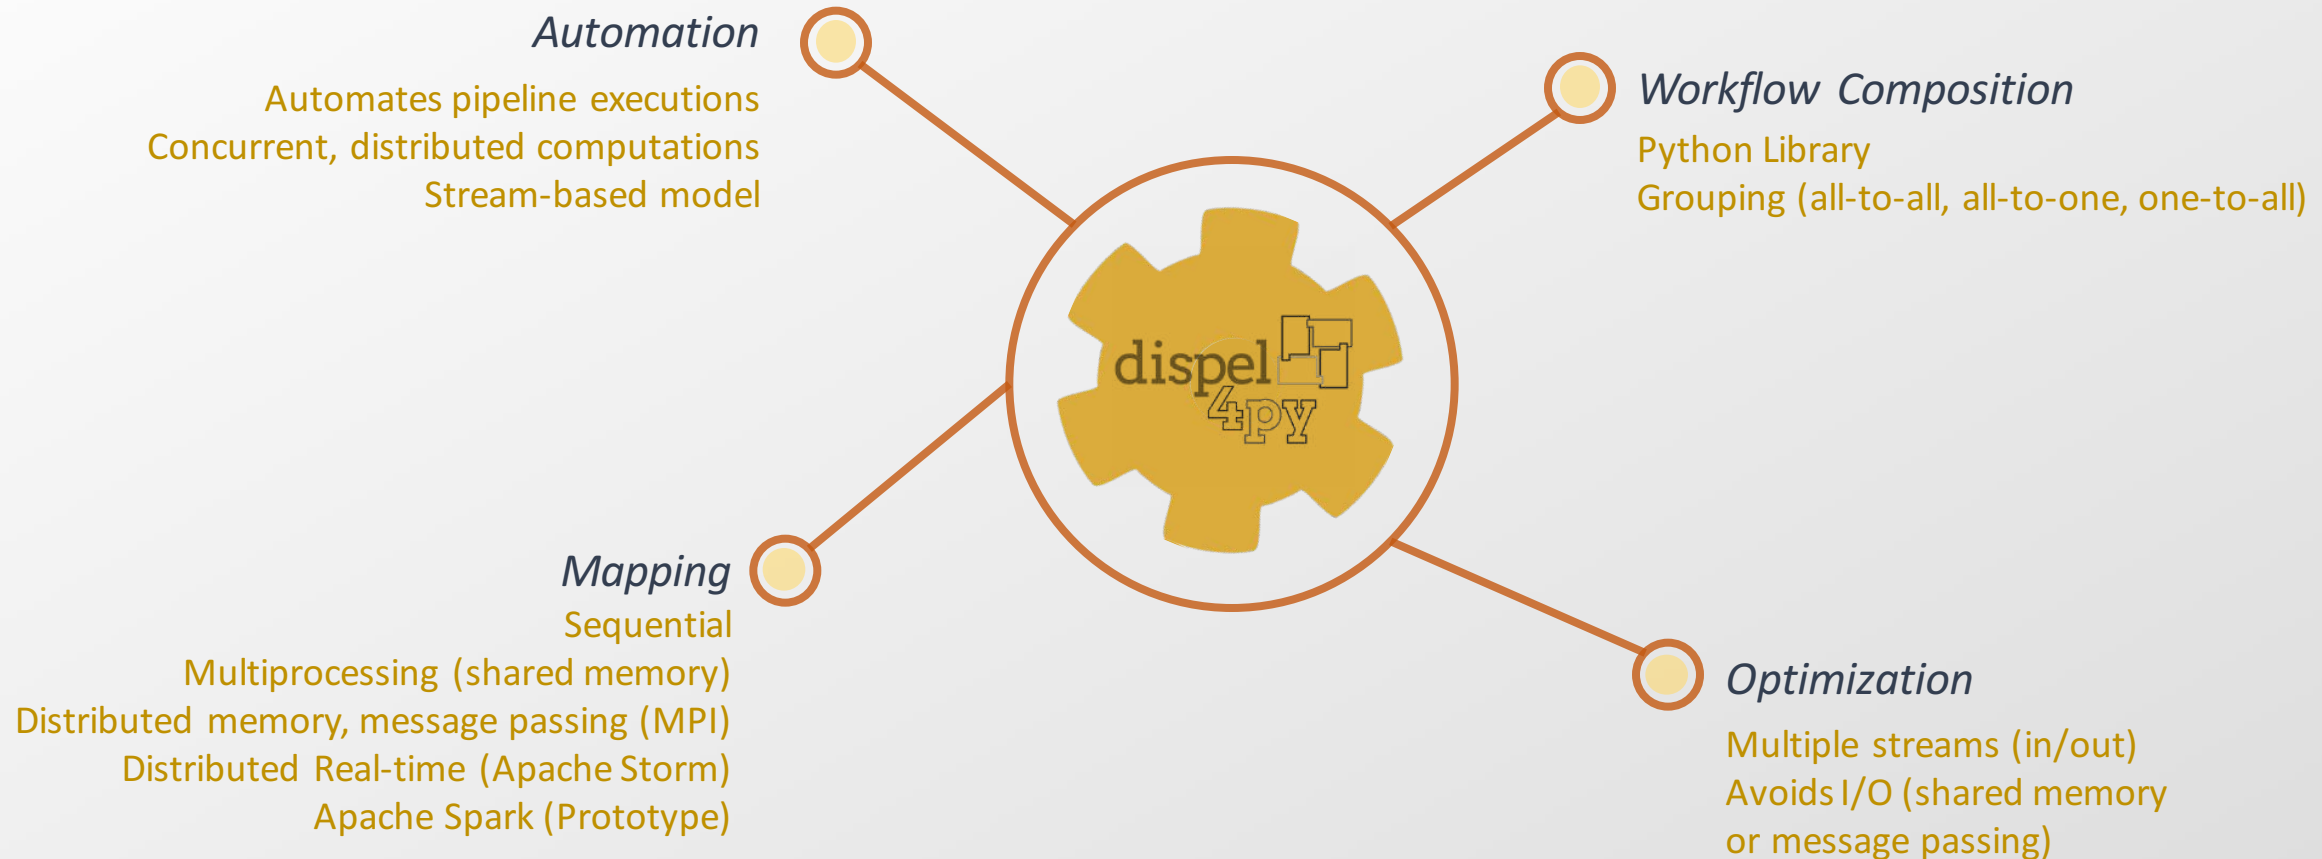
stream-based




What does Pegasus provide?



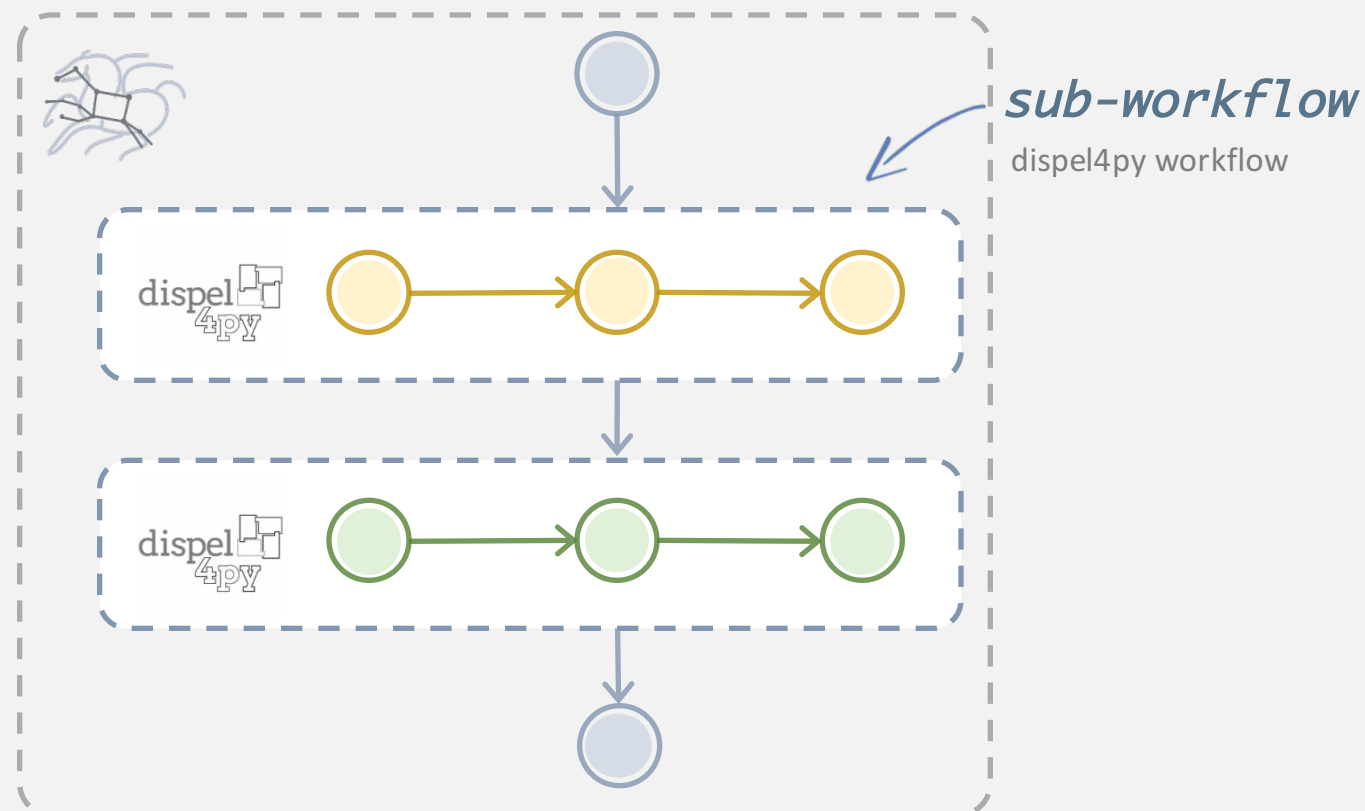
...and dispel4py?





Asterism greatly simplifies the effort required to develop data-intensive applications that run across multiple heterogeneous resources distributed in the wide area

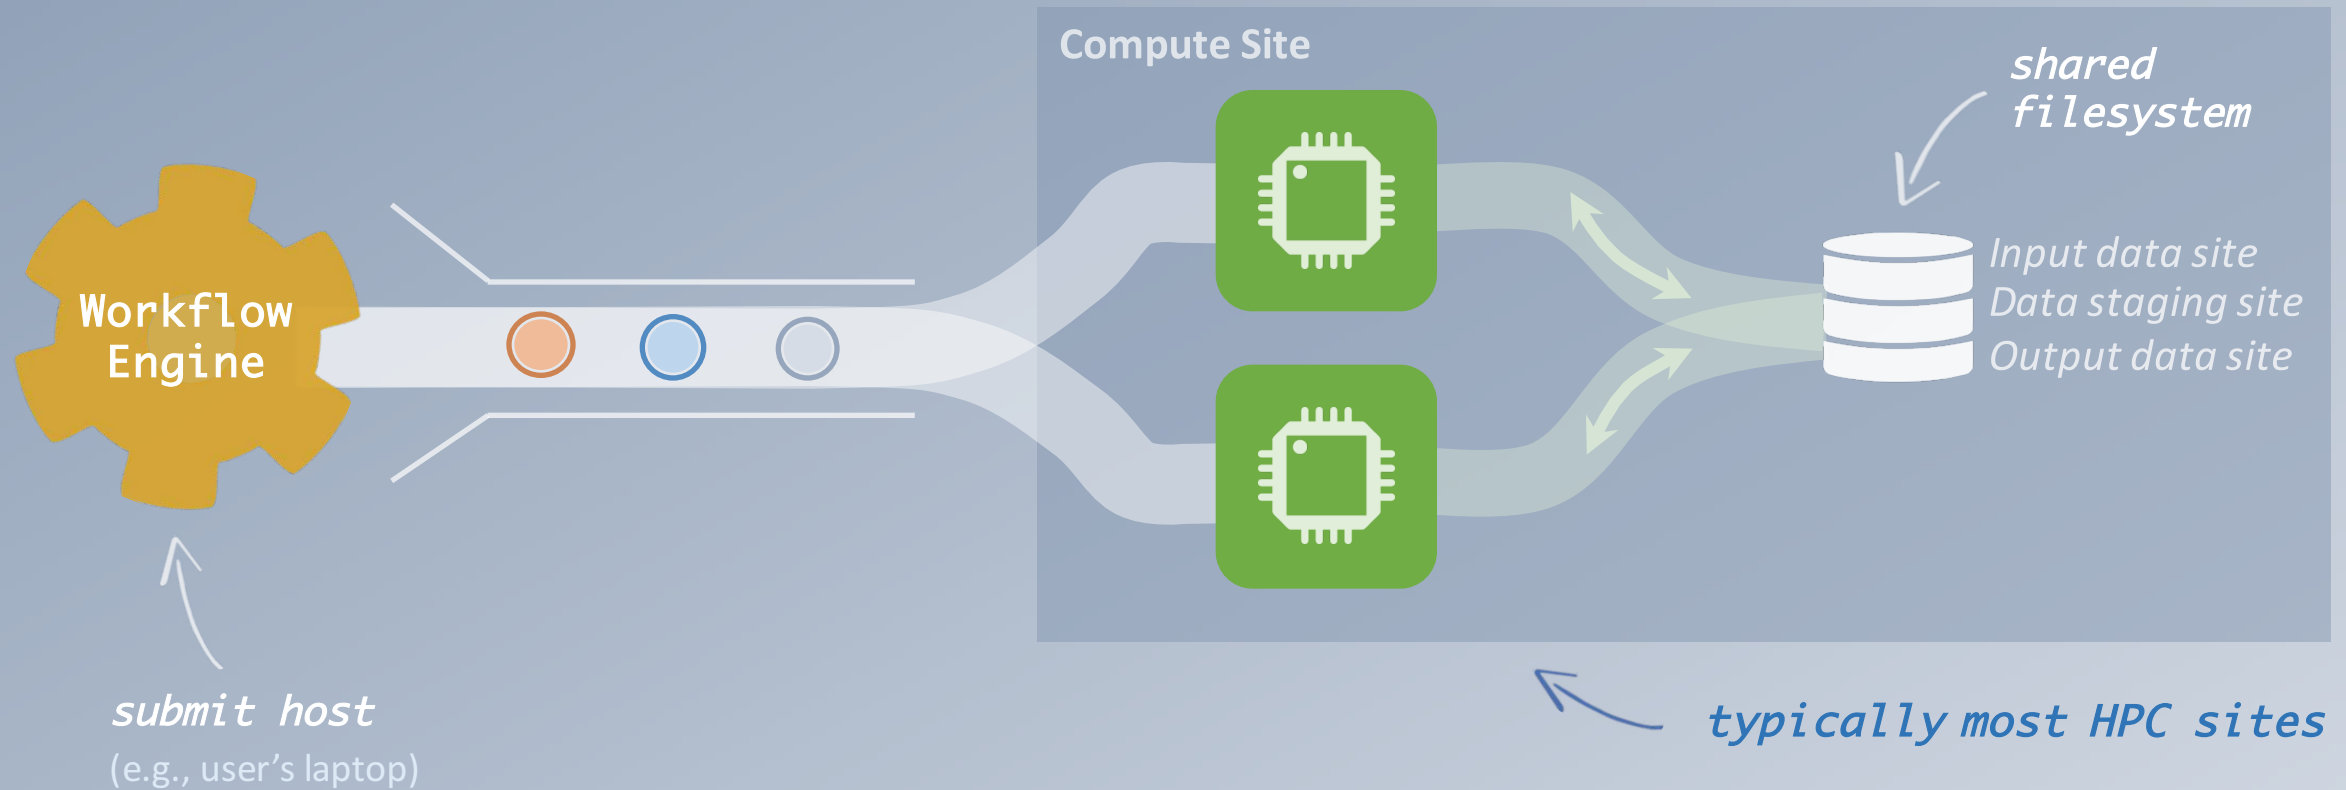
ASTERISM




Where to run scientific workflows?

High Performance Computing

There are several possible configurations...




Cloud Computing High-scalable object storages

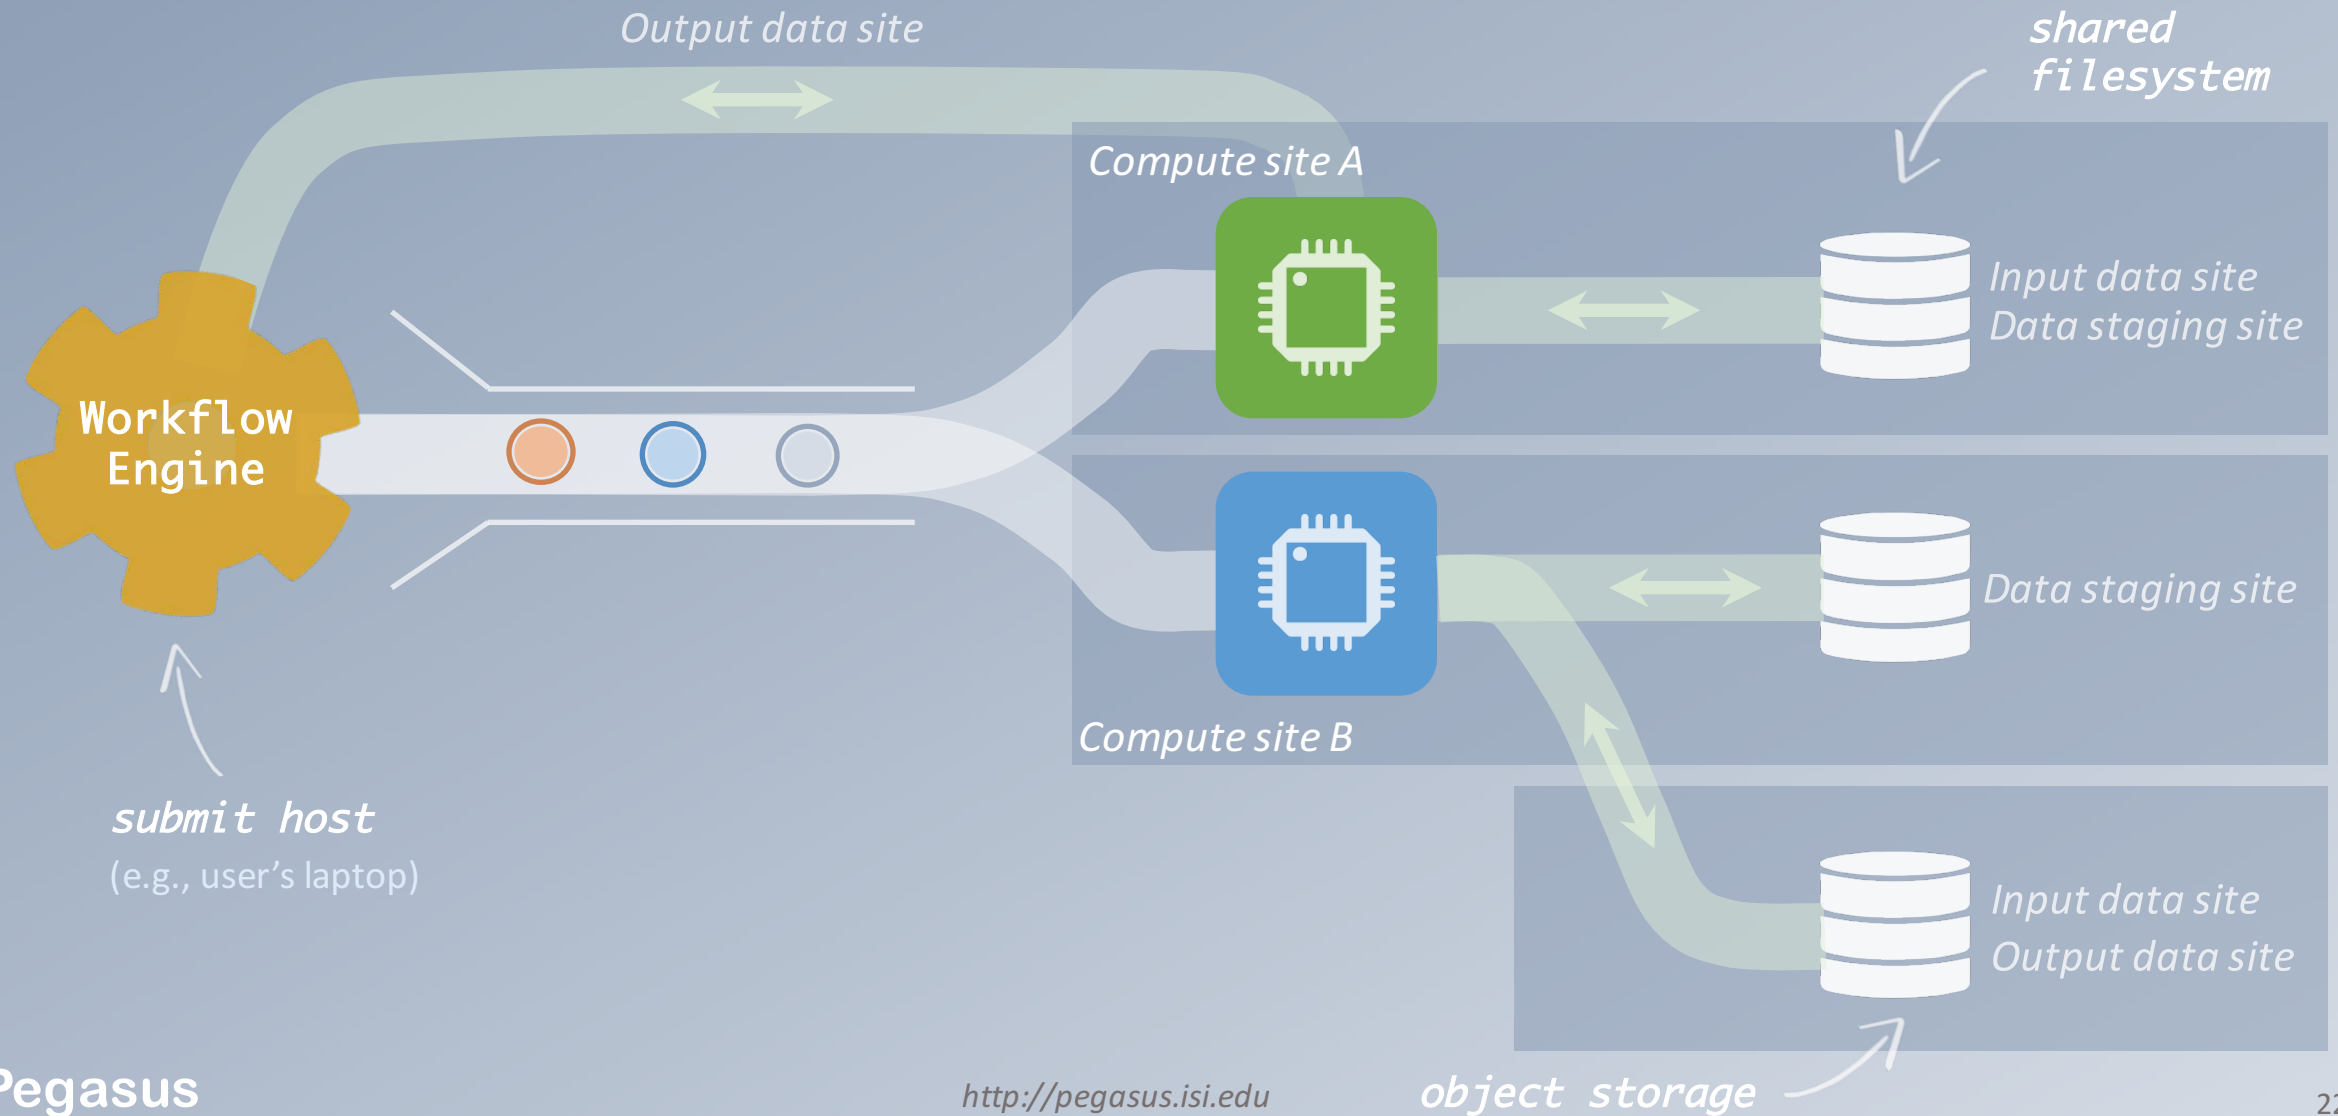


Grid Computing

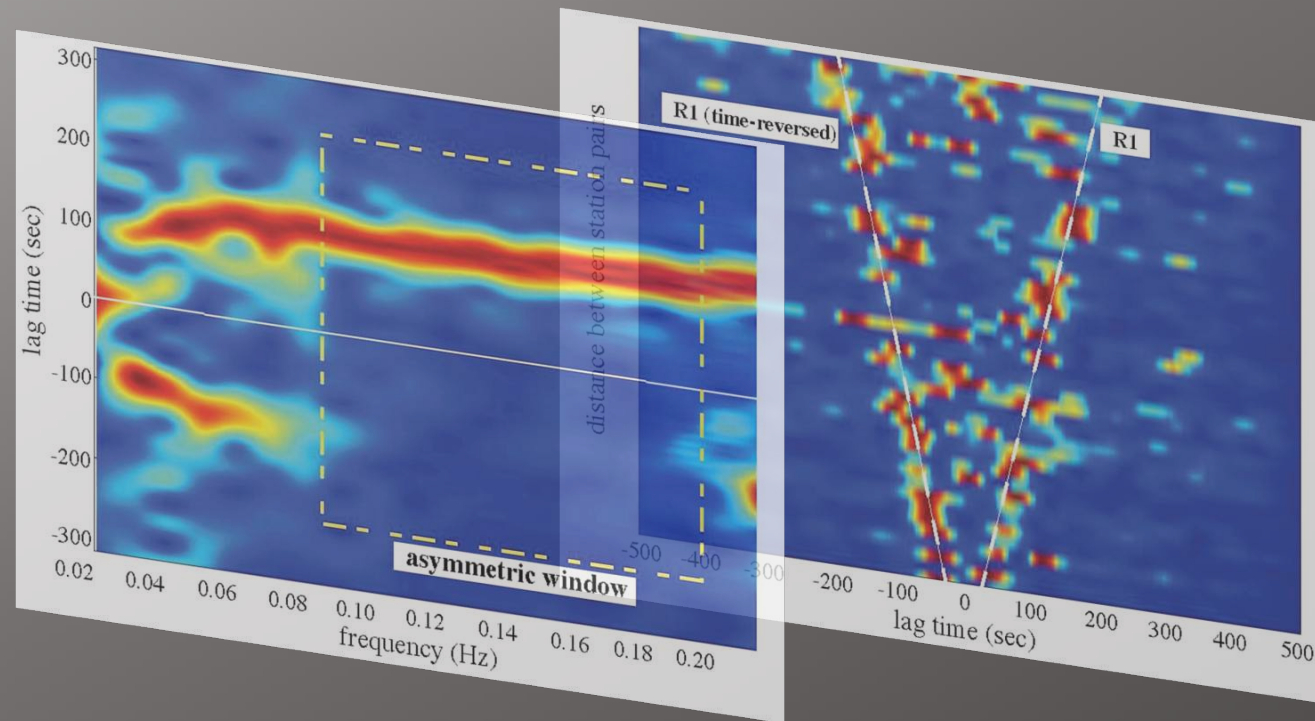
local data management



And yes... you can mix everything!



How do we use Asterism to automate seismic analysis?

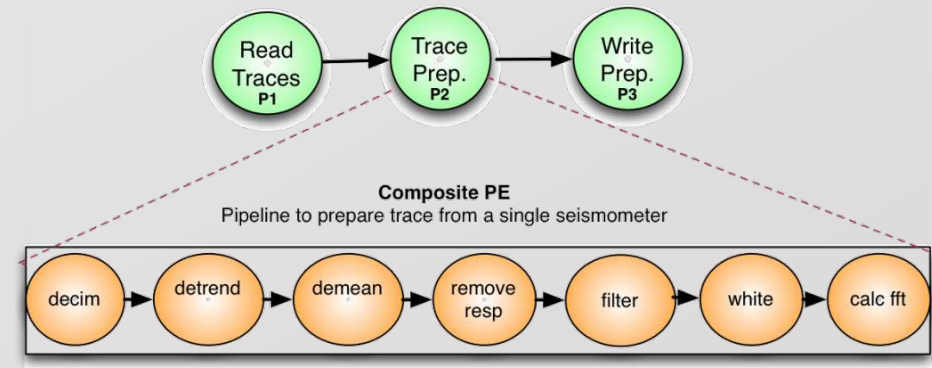


Seismic Ambient Noise Cross-Correlation


Preprocesses and cross-correlates traces (sequences of measurements of acceleration in three dimensions) from multiple seismic stations (IRIS database)

Phase 1: data preparation using statistics for extracting information from the noise

Phase 2: compute correlation, identifying the time for signals to travel between stations. Infers properties of the intervening rock




Phase 1 (pre-process)



Phase 2 (cross-correlation)

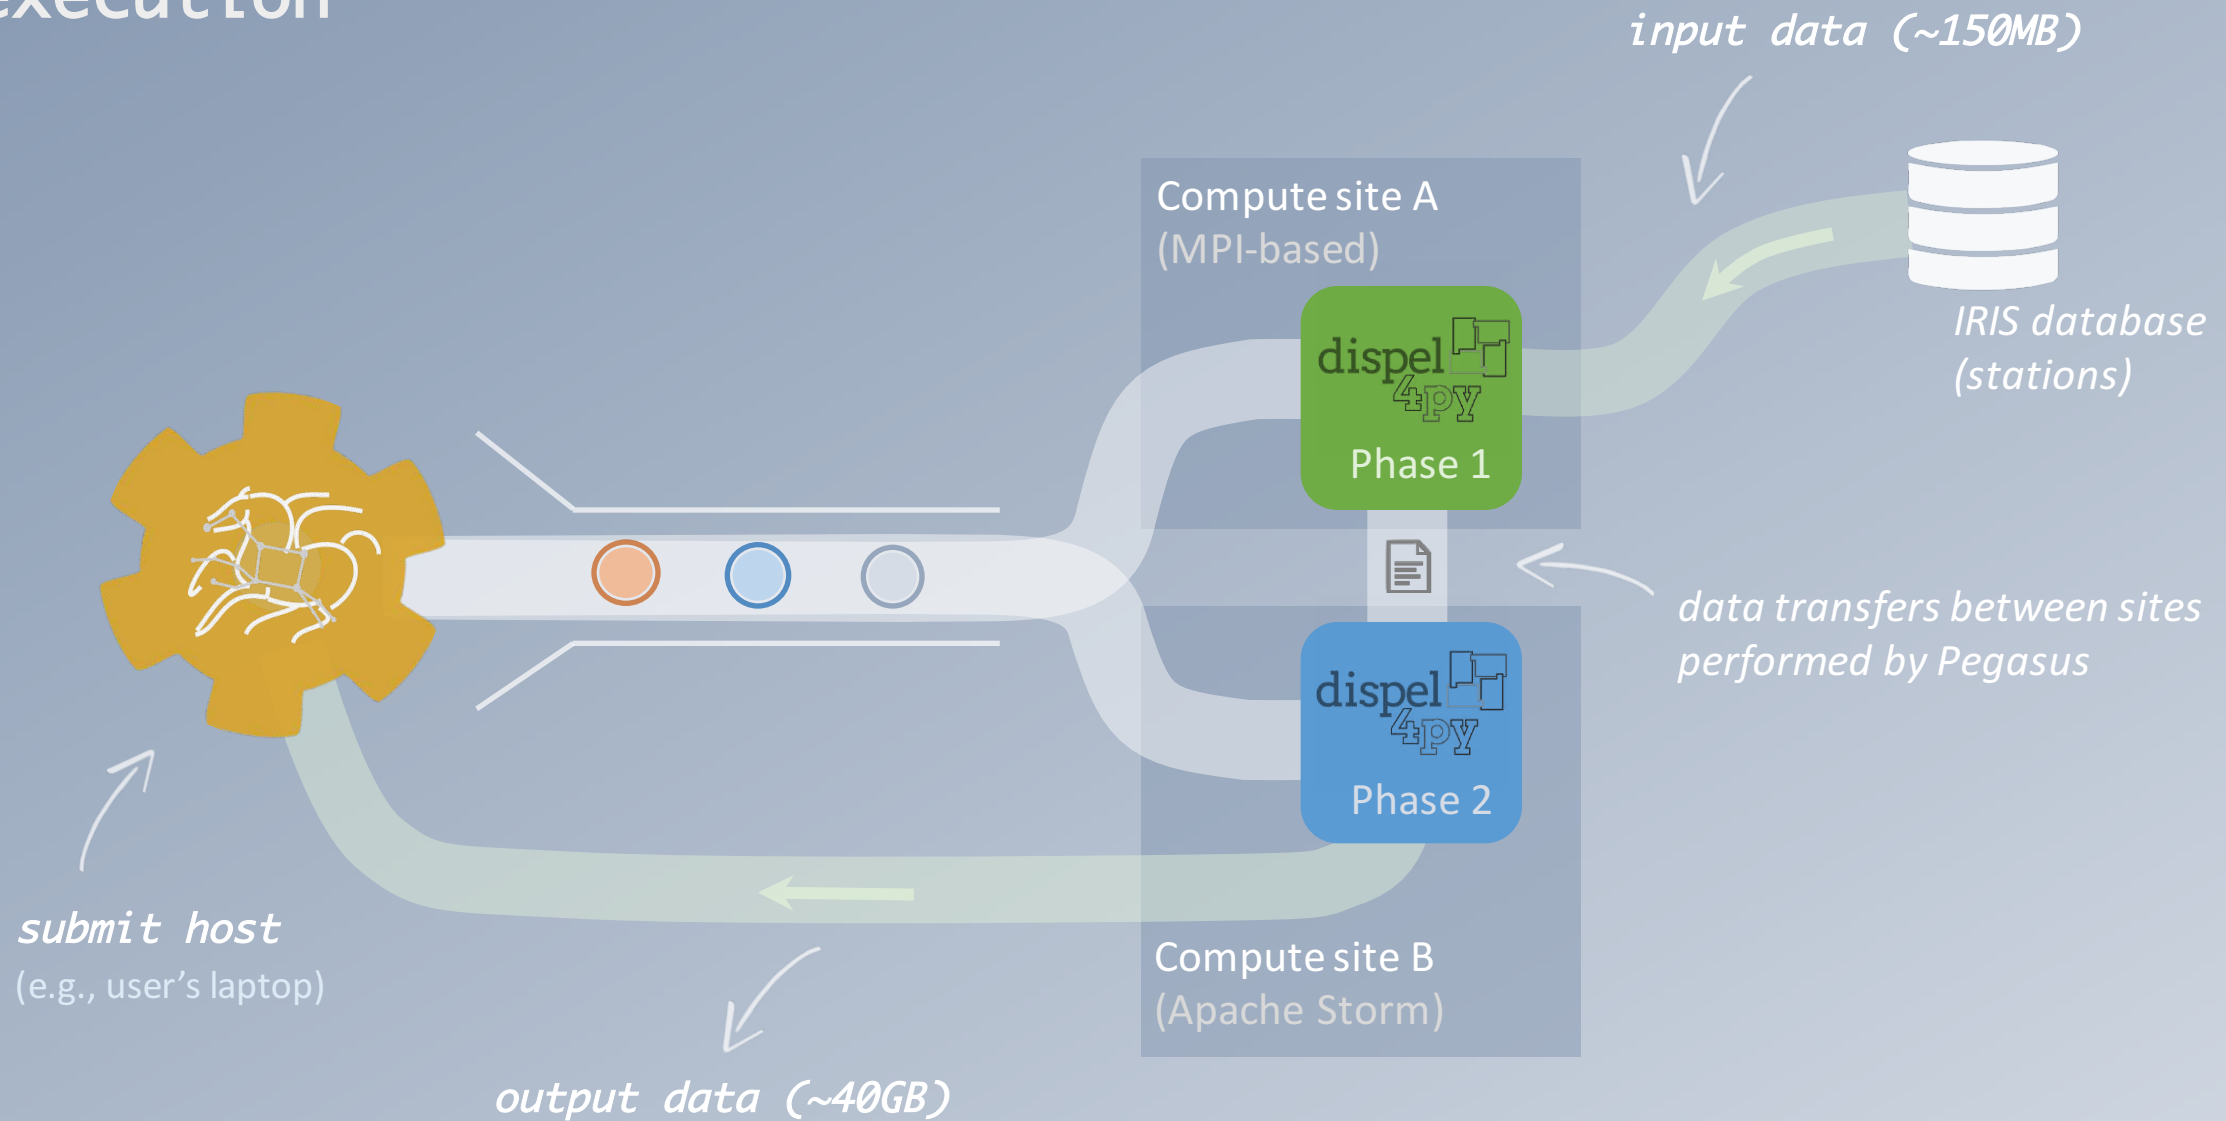
Seismic Ambient Noise Cross-Correlation

WORKFLOW



- ✓ Distributed computation framework for event stream processing
- ✓ Designed for massive scalability, supports fault-tolerance with a “fail fast, auto restart” approach to processes
- ✓ Rich array of available spouts specialized for receiving data from all types of sources
- ✓ Hadoop of real-time processing, very scalable

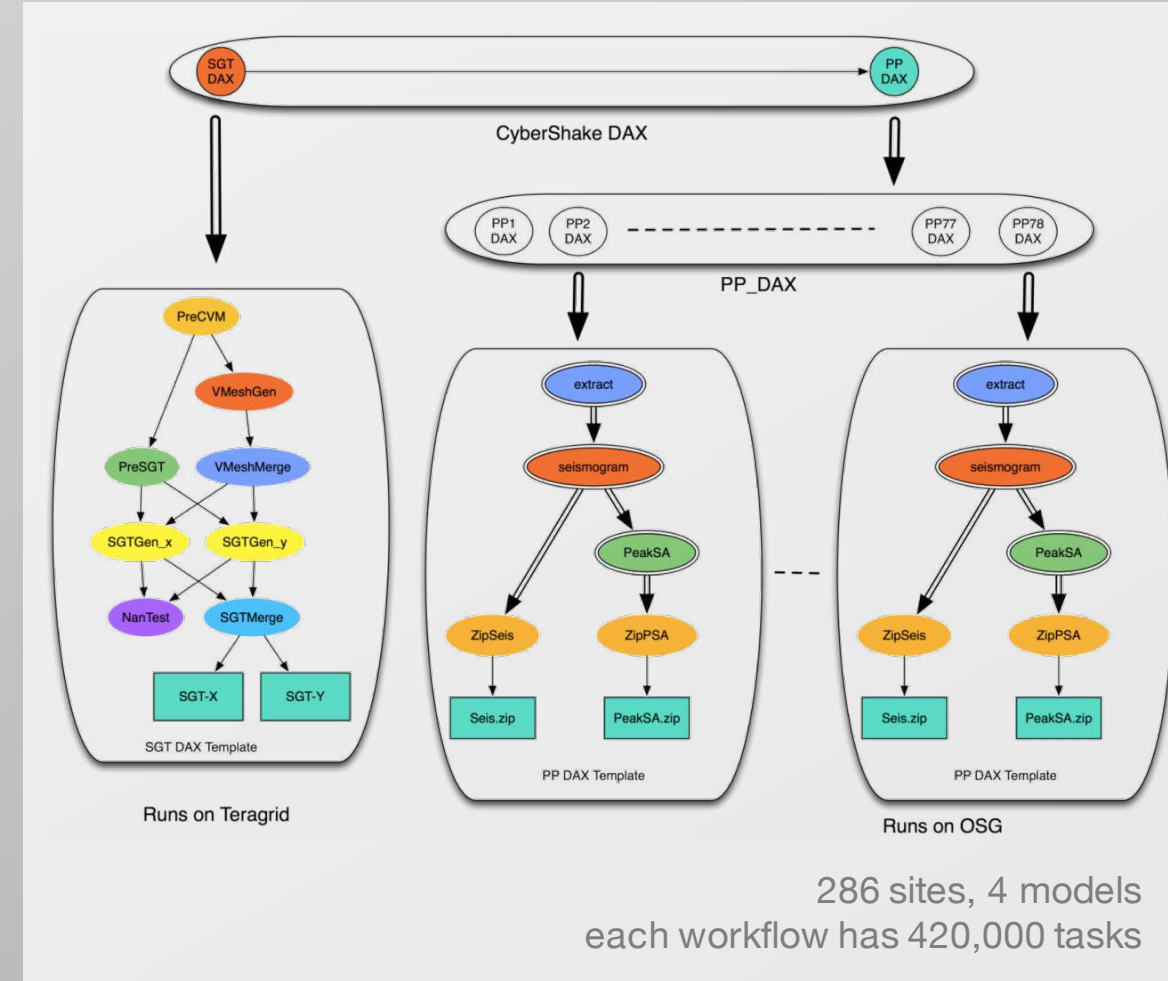
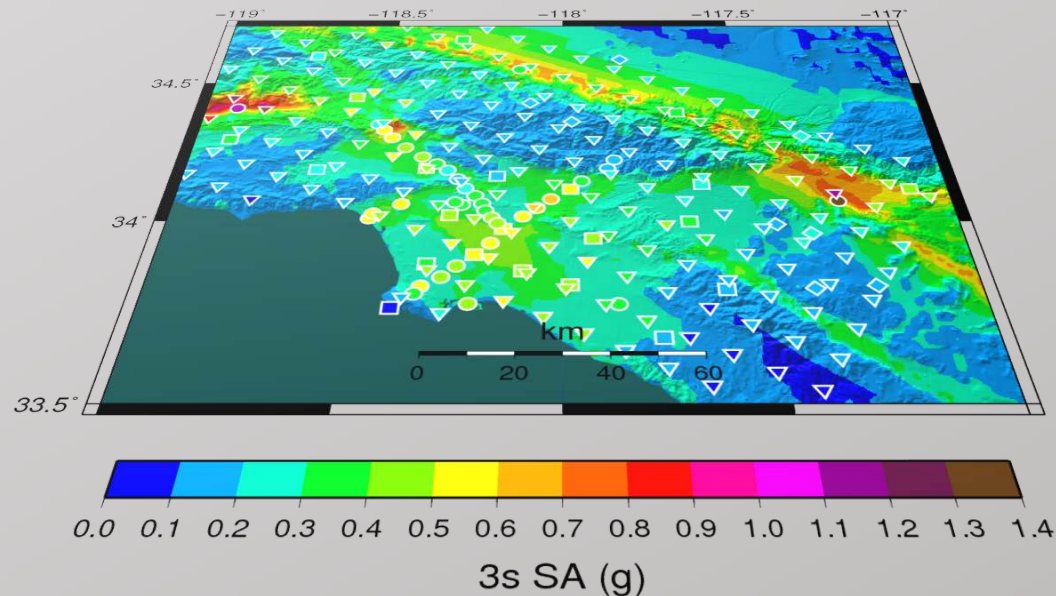
Seismic workflow execution



Southern California Earthquake Center's CyberShake

Builders ask seismologists: What will the peak ground motion be at my new building in the next 50 years?

Seismologists answer this question using Probabilistic Seismic Hazard Analysis (PSHA)

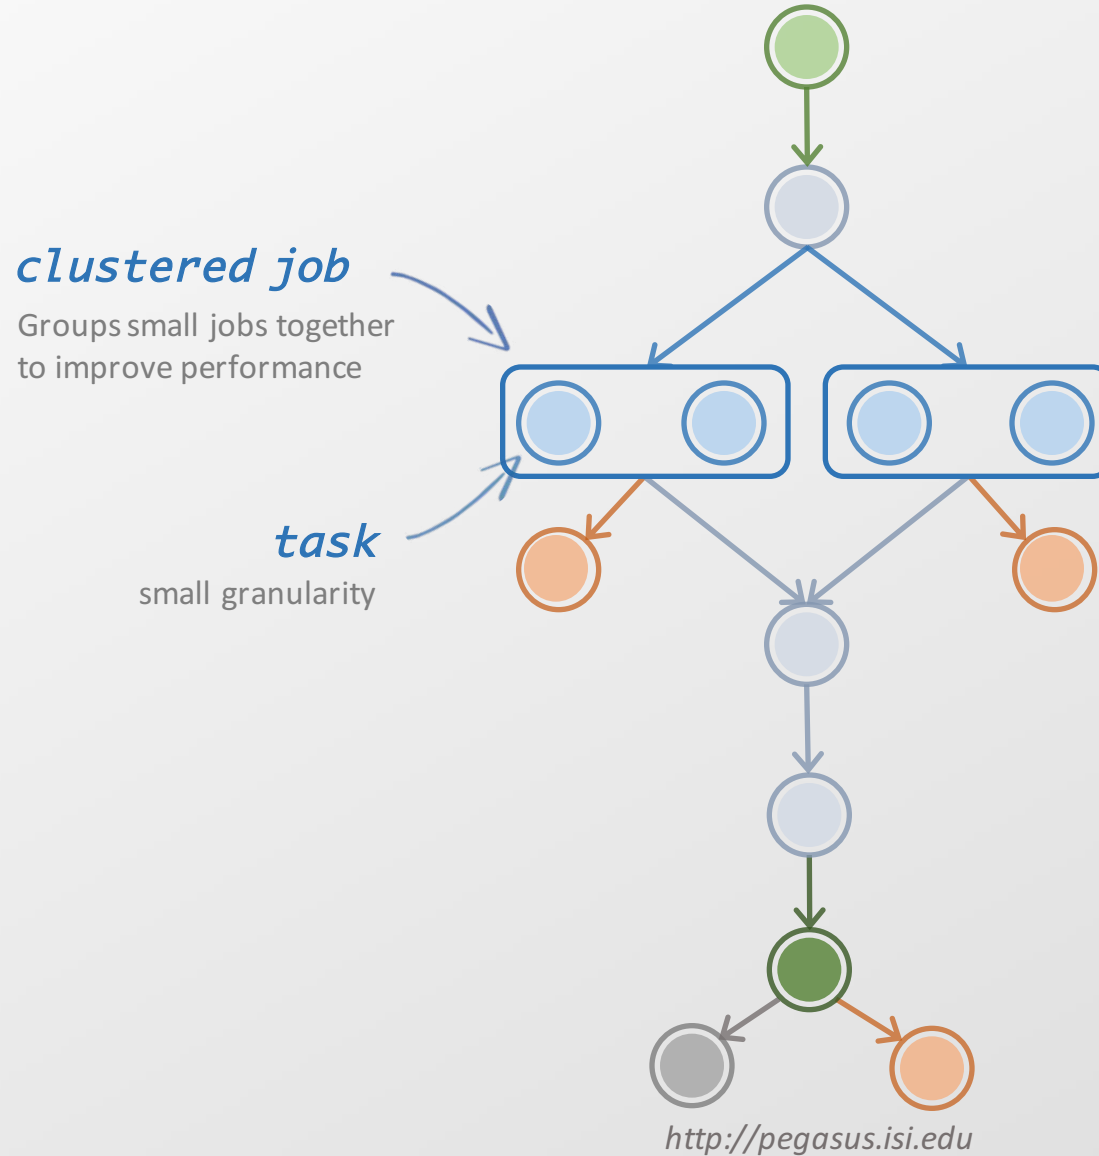


286 sites, 4 models
each workflow has 420,000 tasks

A few more workflow features...


Performance, why not improve it?

workflow restructuring
workflow reduction
hierarchical workflows
pegasus-mpi-cluster

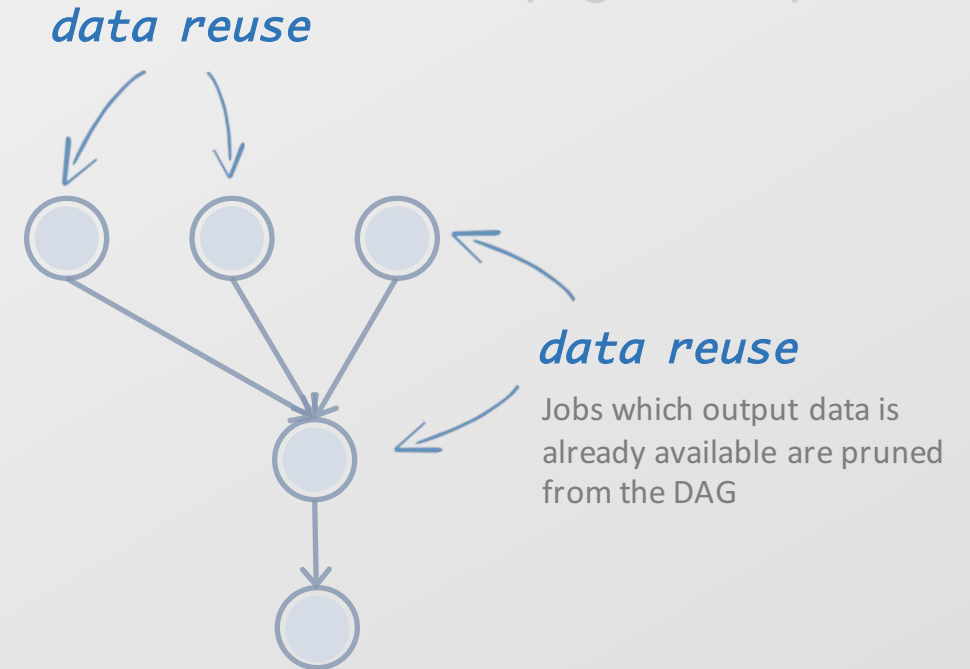


What about data reuse?

workflow restructuring
workflow reduction
hierarchical workflows
pegasus-mpi-cluster

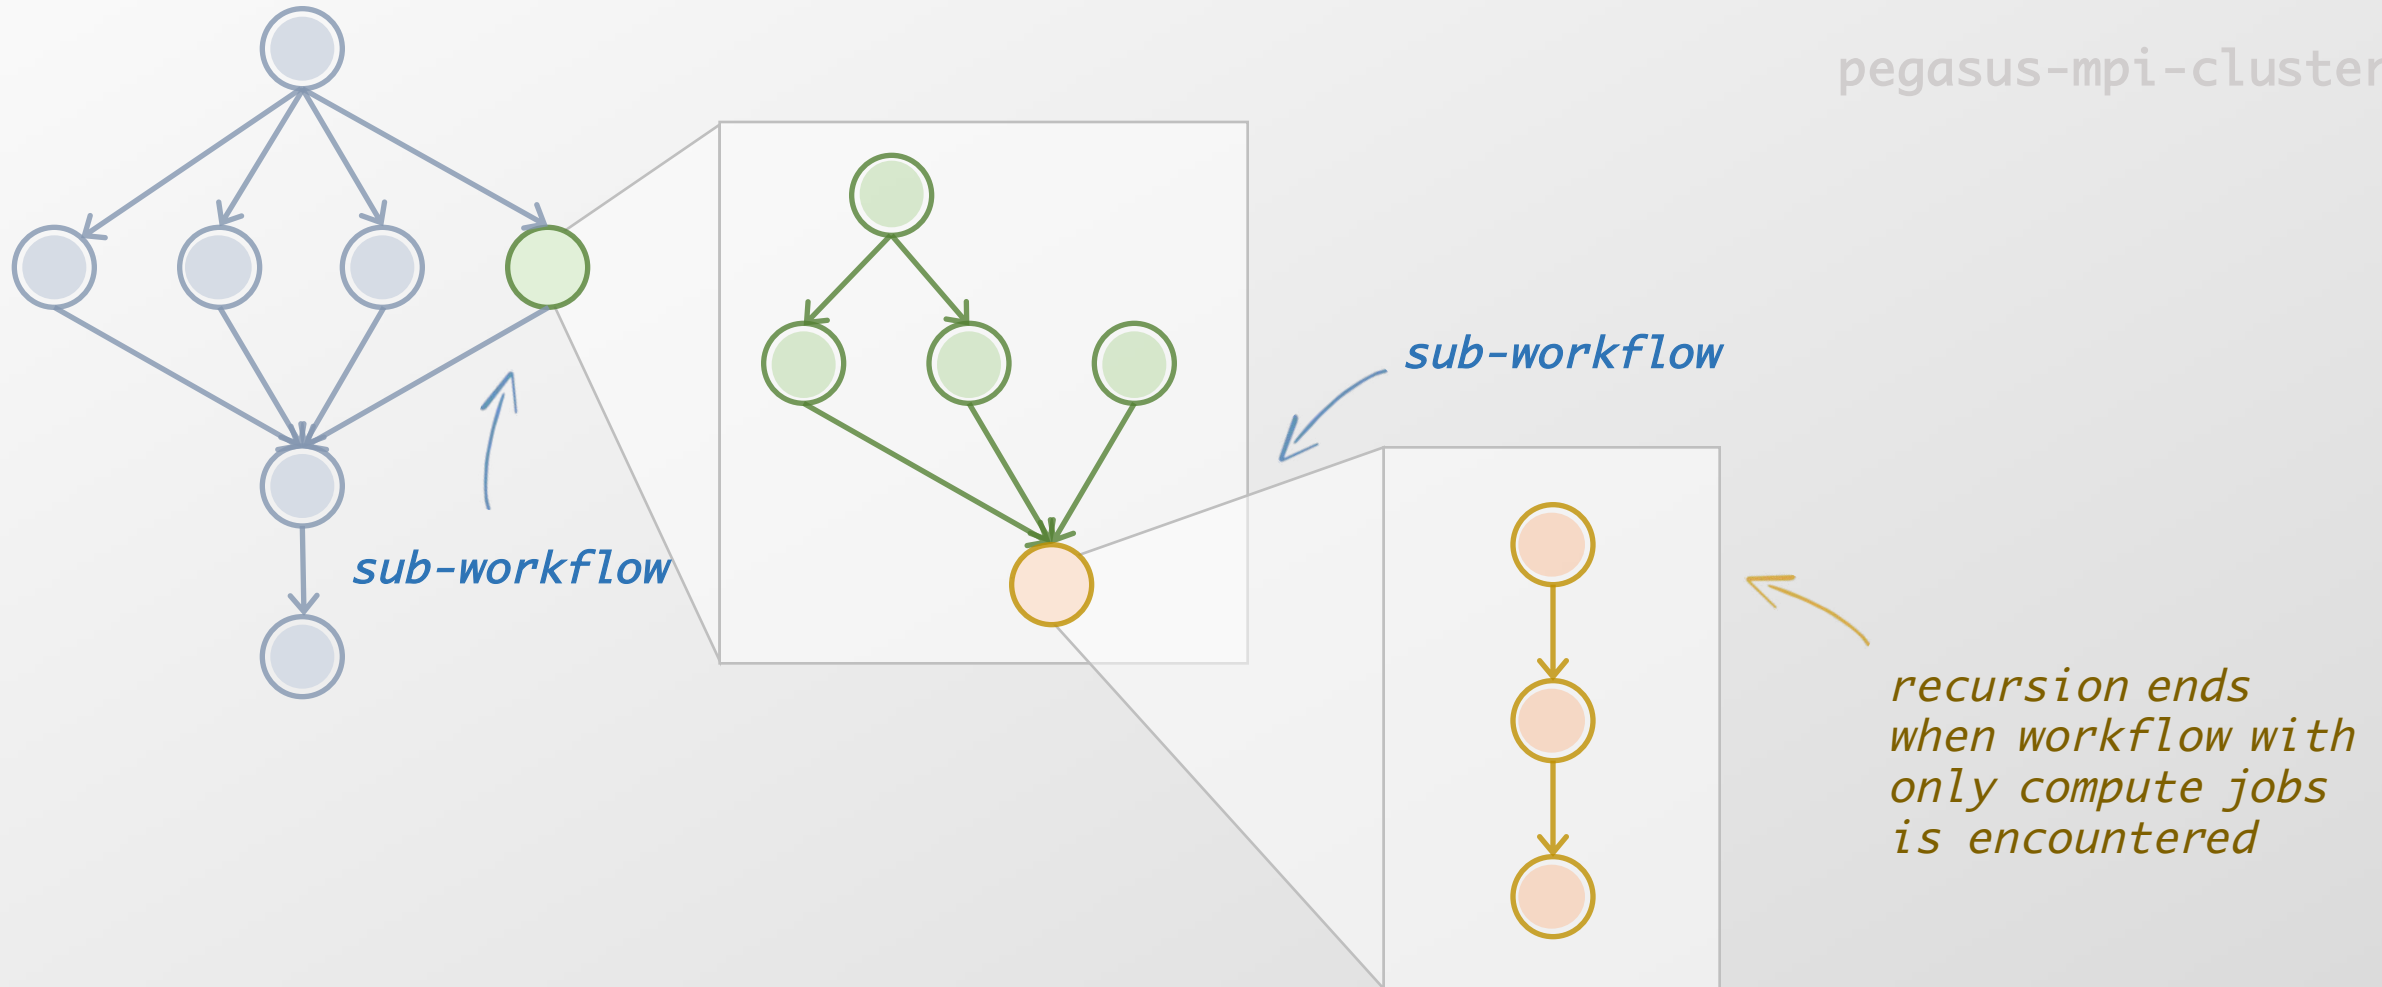


*work flow
reduction*



Handling large-scale workflows


workflow restructuring
workflow reduction
hierarchical workflows
pegasus-mpi-cluster



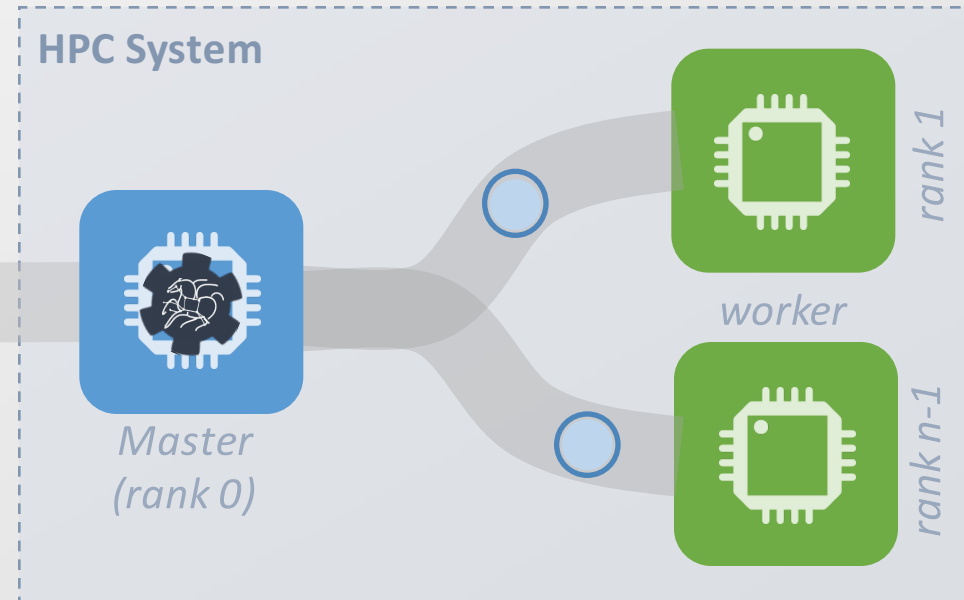
Running fine-grained workflows on HPC systems...

workflow restructuring
workflow reduction
hierarchical workflows
pegasus-mpi-cluster

submit host
(e.g., user's laptop)



workflow wrapped as an MPI job
Allows sub-graphs of a Pegasus workflow to be submitted as monolithic jobs to remote resources





Pegasus

est. 2001

Automate, recover, and debug scientific computations.

Get Started

Pegasus Website

<http://pegasus.isi.edu>

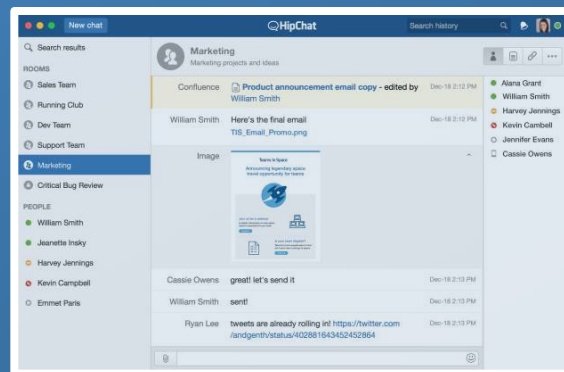
Users Mailing List

pegasus-users@isi.edu

Support

pegasus-support@isi.edu

HipChat





Automating Real-time Seismic Analysis

Through Streaming and High Throughput Workflows

Thank You

Questions?

Rafael Ferreira da Silva, Ph.D.
rafsilva@isi.edu

USC Viterbi
School of Engineering
Information Sciences Institute

Meet our team



Ewa Deelman



Rafael Ferreira da Silva



Karan Vahi



Mats Rynge



Rajiv Mayani