

HFIP Architecture and Real-Time Workflows

14 April 2011



Overview



HFIP System at ESRL

Current Configuration

Upgrade Options

HSM Tape Discussion

Construction Real-Time Workflows

Current HFIP System

tJet

844 Nodes, 10128 cores

Accessed via the **hfip** parallel environment

nJet

59 nodes, 472 cores

Accessed via the **nhfip** parallel environment

All jobs run in **hfipserial** run on nJet

Lfs1

430 TB is for HFIP

10 GB/s

Storage

Existing lfs1 infrastructure allows upgrade of **400 TB**

Low cost, as only disks are needed

Determining if additional performance is required

Compute

With projected budget, considering other costs associated with system

416 nodes – 4992 cores

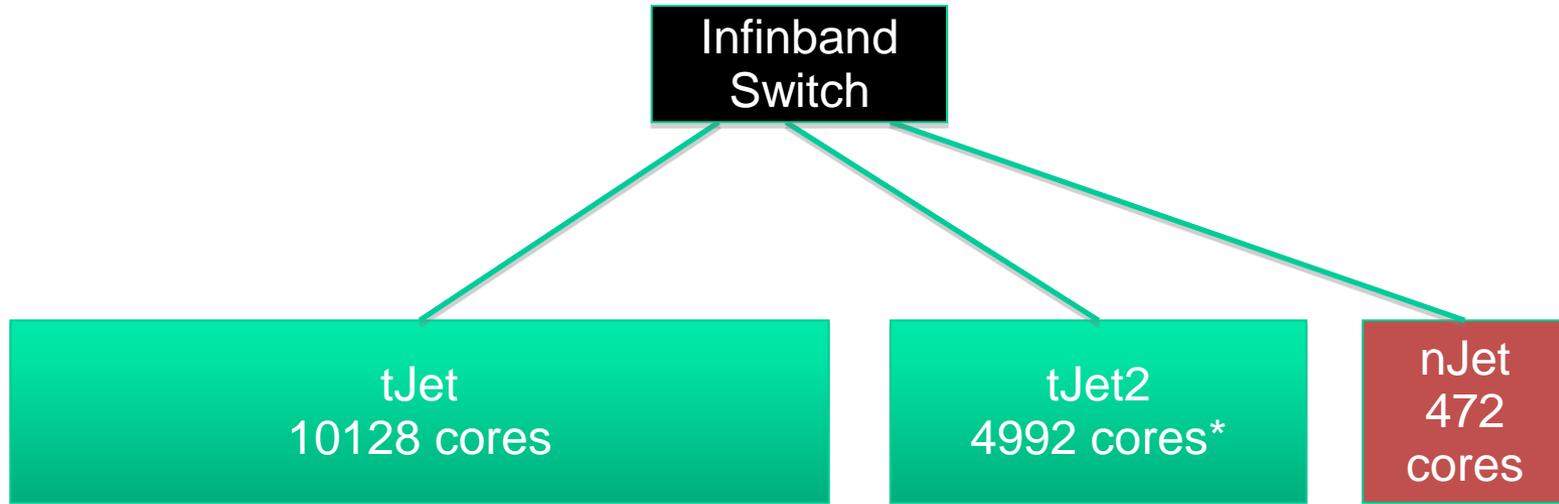
Same tJet nodes, no new technology is available

Why isn't upgrade as large as last time?

- Last time, most of nJet was traded away to increase tJet size
- Other labor costs have come up between HFIP program and RDHPCS program

Can I run a 15,000 core job?

- Not really
- New cores will be apart of a separate tree
- Scheduling across all cores will be transparent users will do nothing different
- To build a single tree is much more expensive



- Users will submit jobs to HFIP PE as normal
- System will place jobs on either tJet or tJet2 where space is available
- We can run scaling tests across both partitions, but only in test mode
 - Largest core count job we have seen is ~2000 cores, so separate trees is not an issue

*Note: Core counts for tJet2 are projected and not guaranteed at this time



Schedule



Procurement will be done in two pieces.
Timeline depends on when funding is approved

Compute – 14-16 weeks after SON

Storage – 12-14 weeks after SON

Yes, we understand how this will affect runs for hurricane season. We will try to compress the schedule.

We are running out of tapes

Usage has doubled over the last few months

Due to budget problems and transition to West Virginia, no money now

Short term solution – Turn off double copying

A Real-time experiment is a workflow that has a time-critical component to it. Specifically, the time-critical component is a deadline after which value of the data or results is lost.

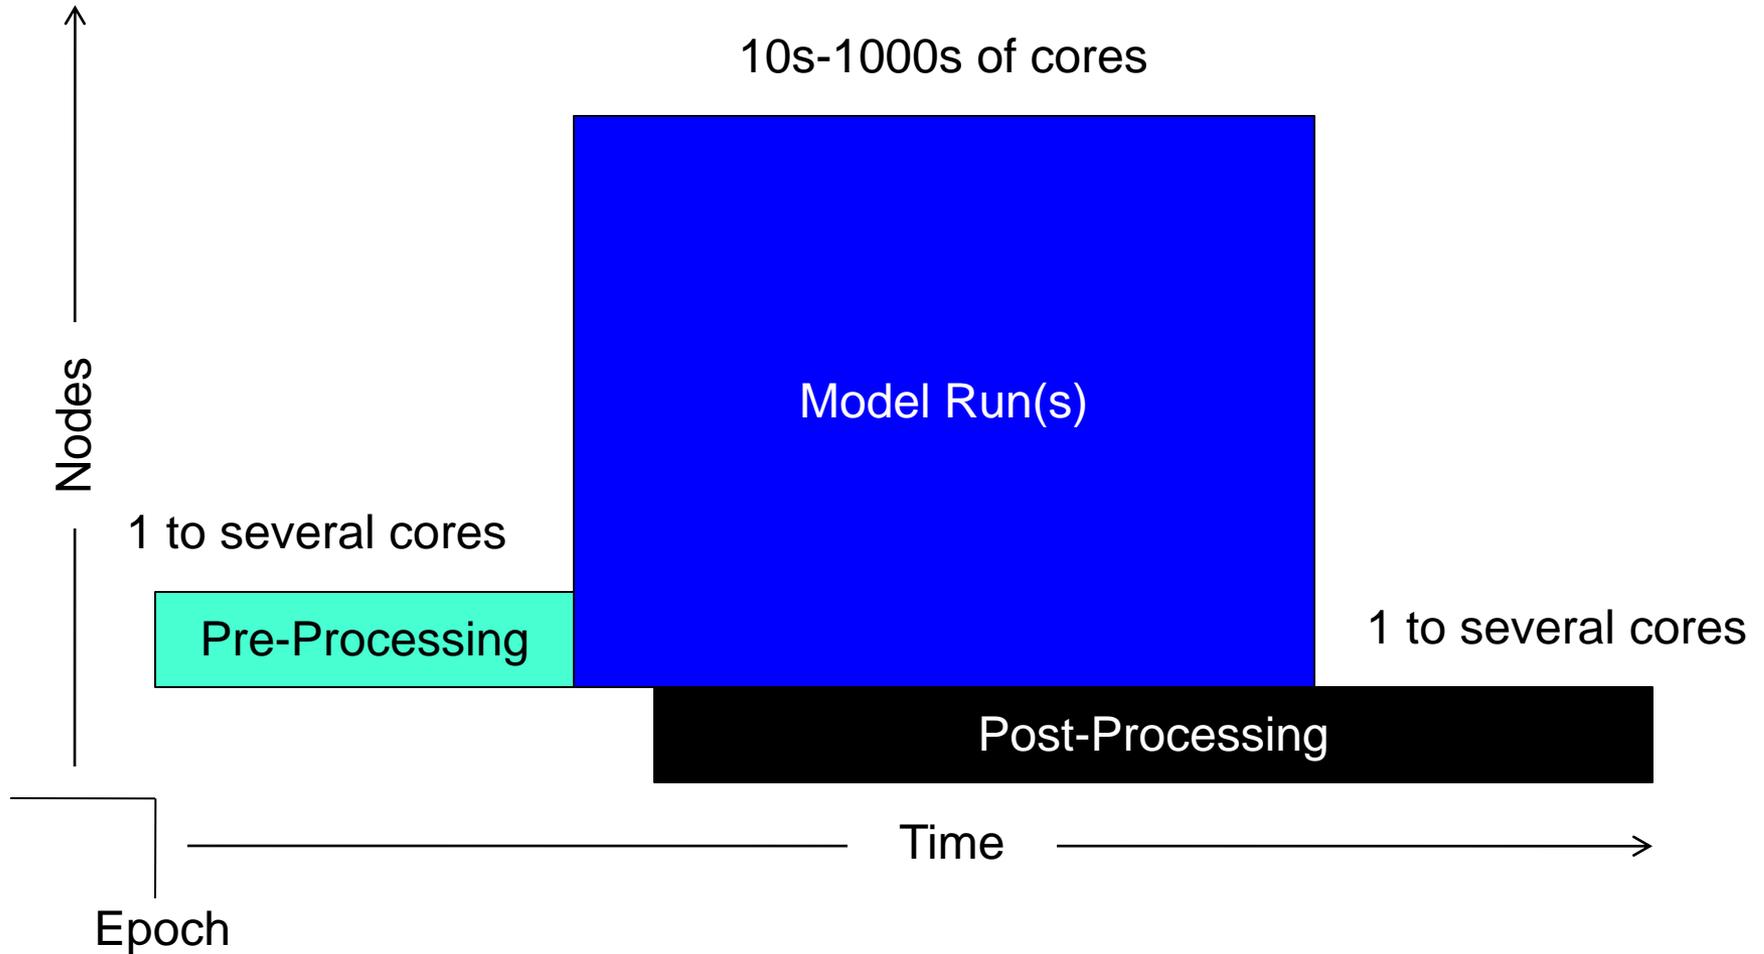
The difference between a real-time experiment and operation job is whether life or property could be affected by the lack of results.

A workflow is a set of interconnected jobs where each job may have a single or multiple dependencies on time, data, or other jobs.

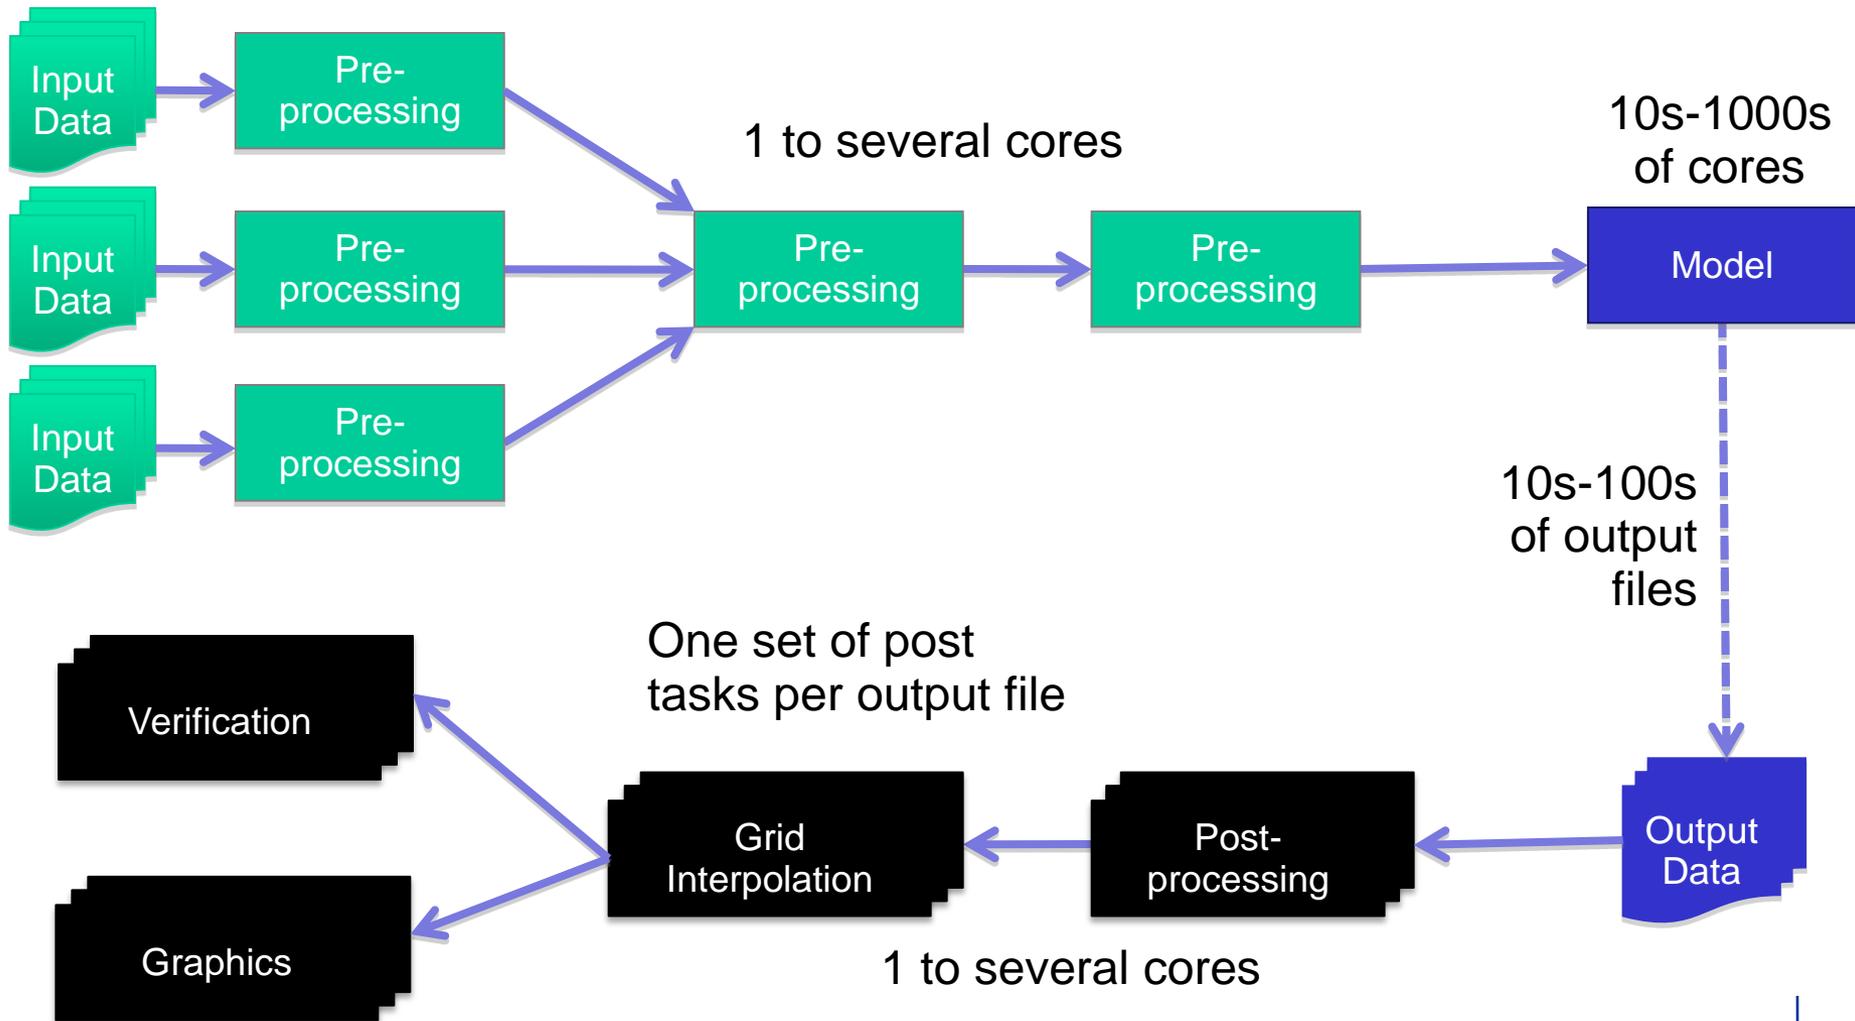
Real-time experiments are supported on Jet with the use of standing reservations

A standing reservation is where a block of nodes is reserved for a specific window of time and dedicated to a specific user or account.

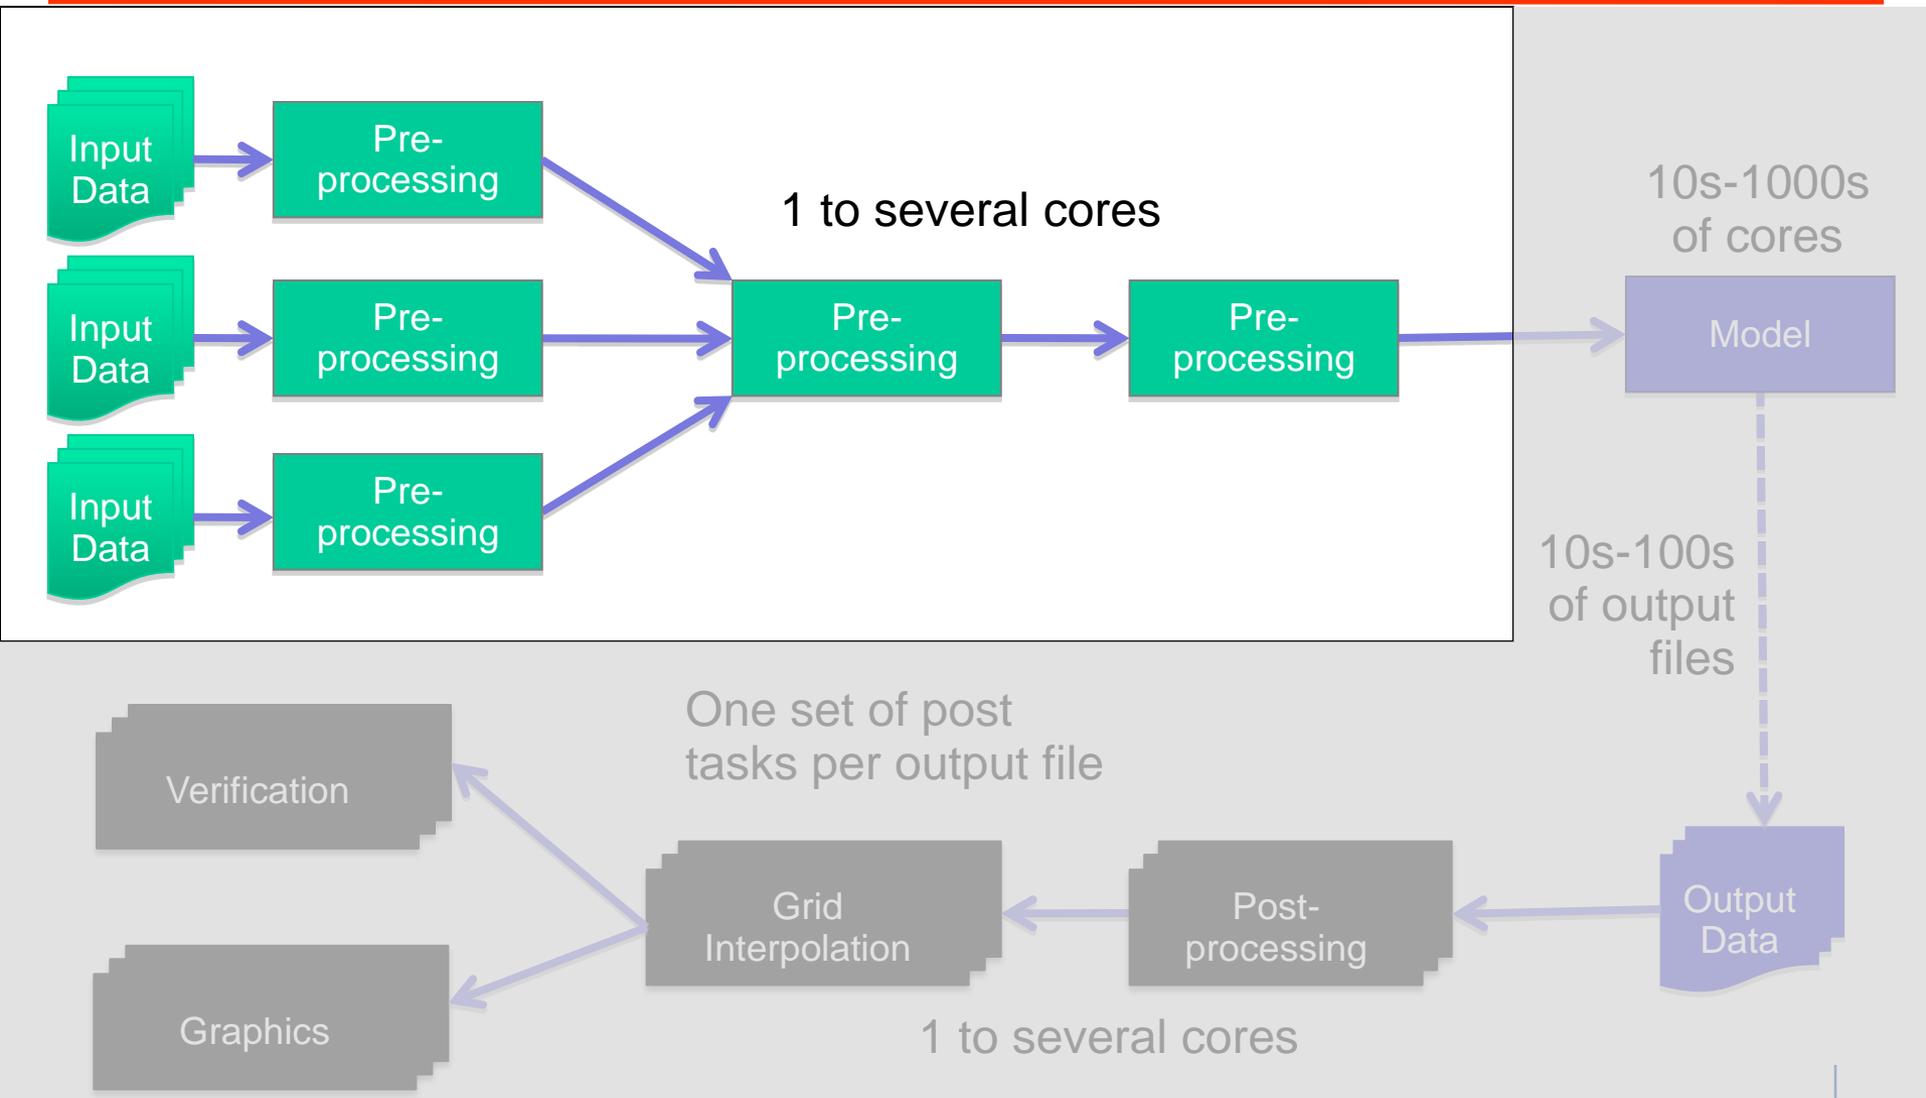
Typical Reservation



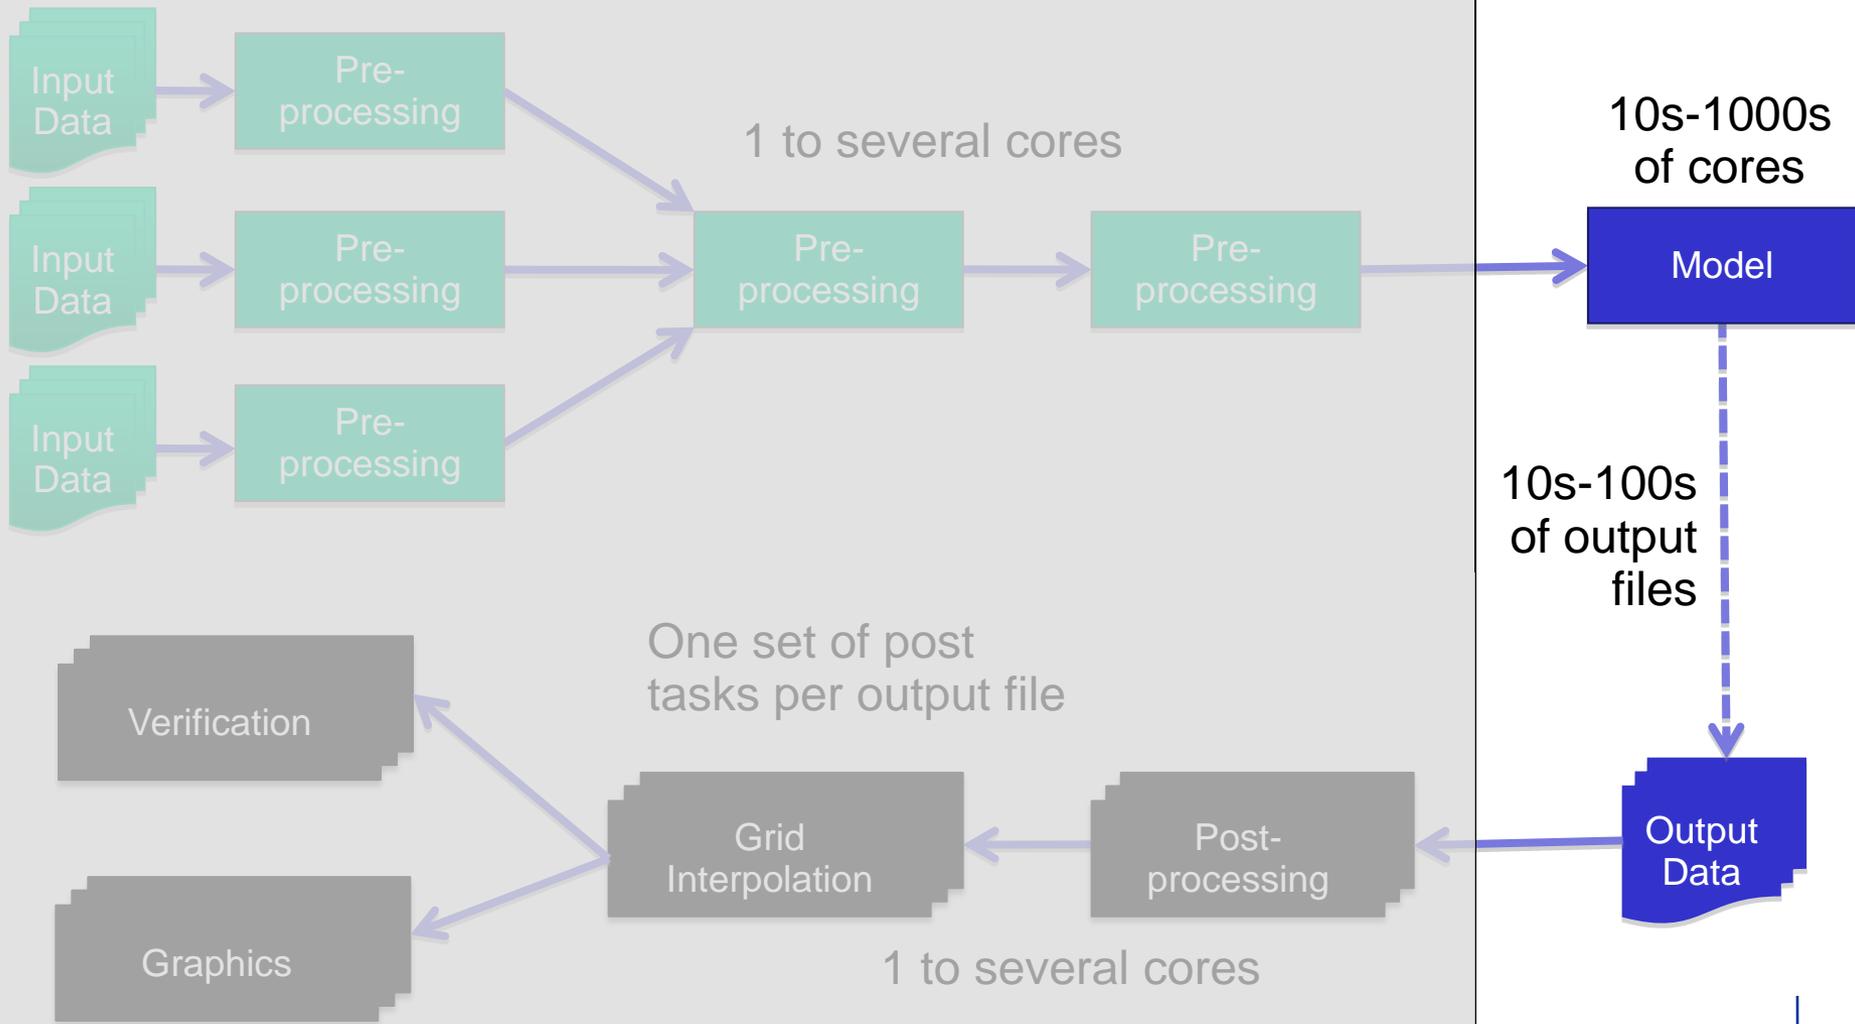
A Typical Workflow



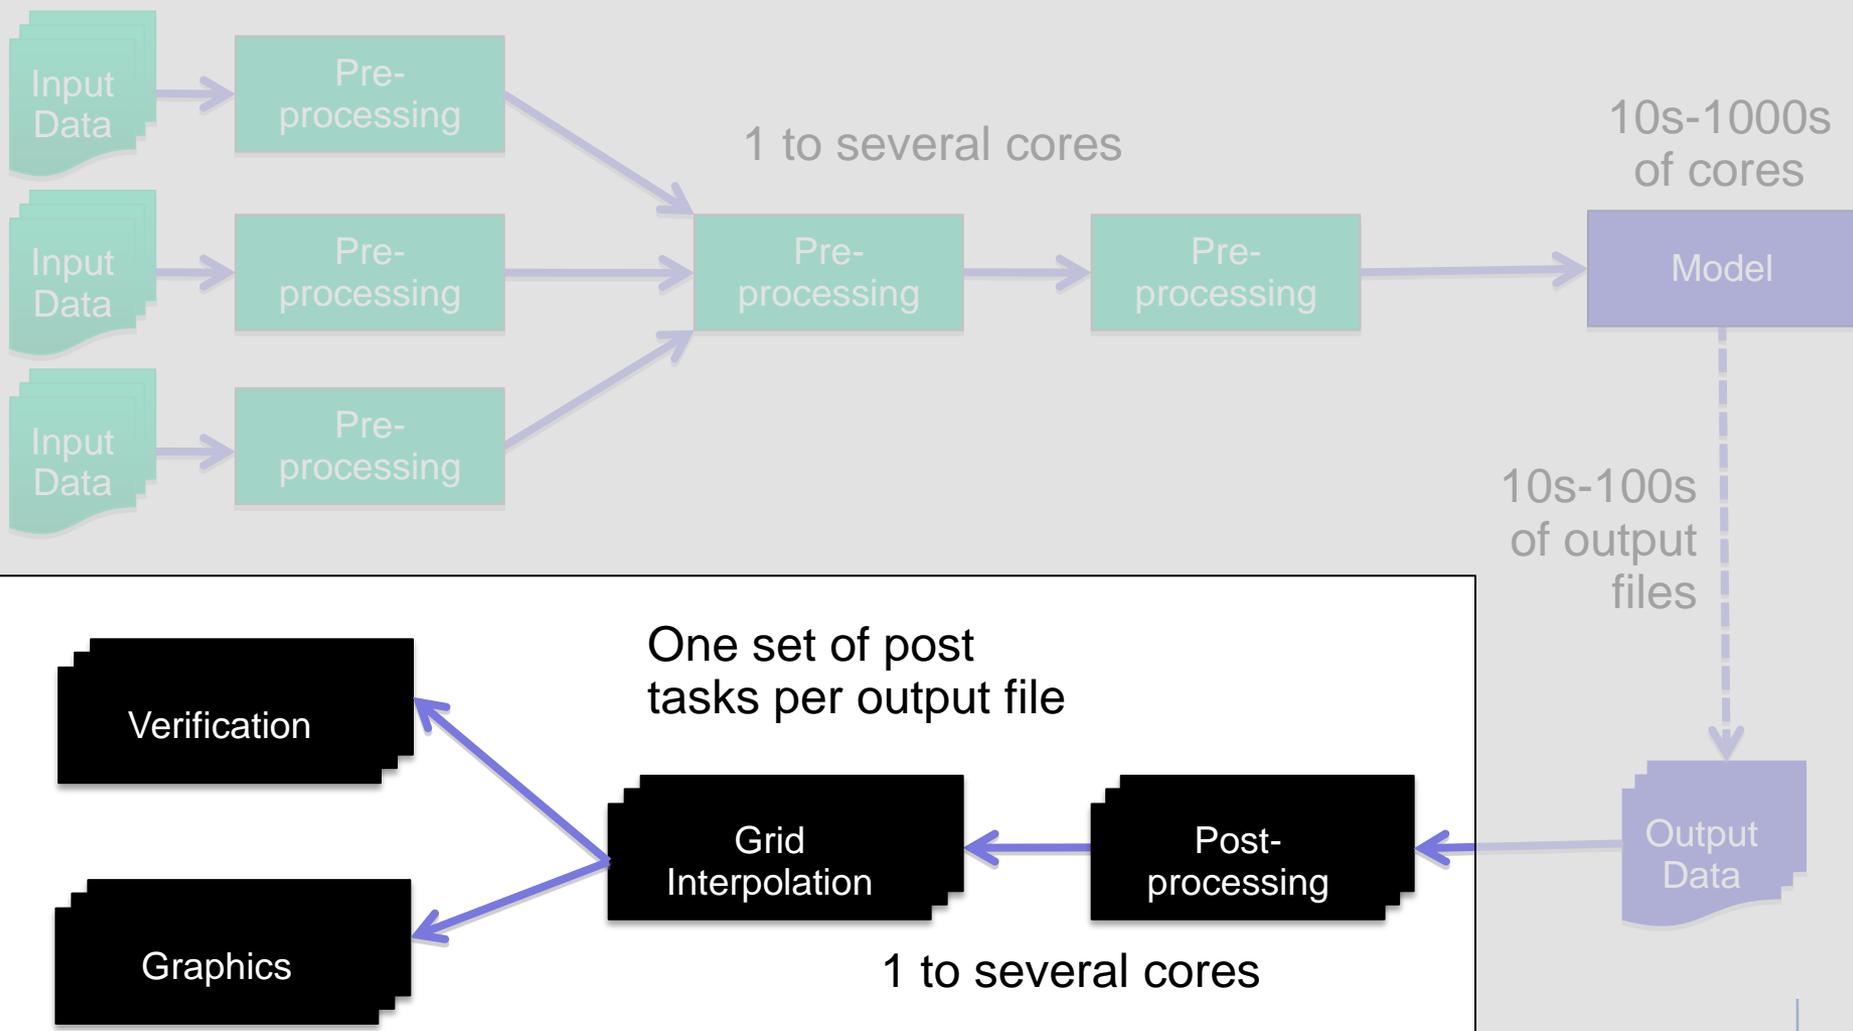
A Typical Workflow

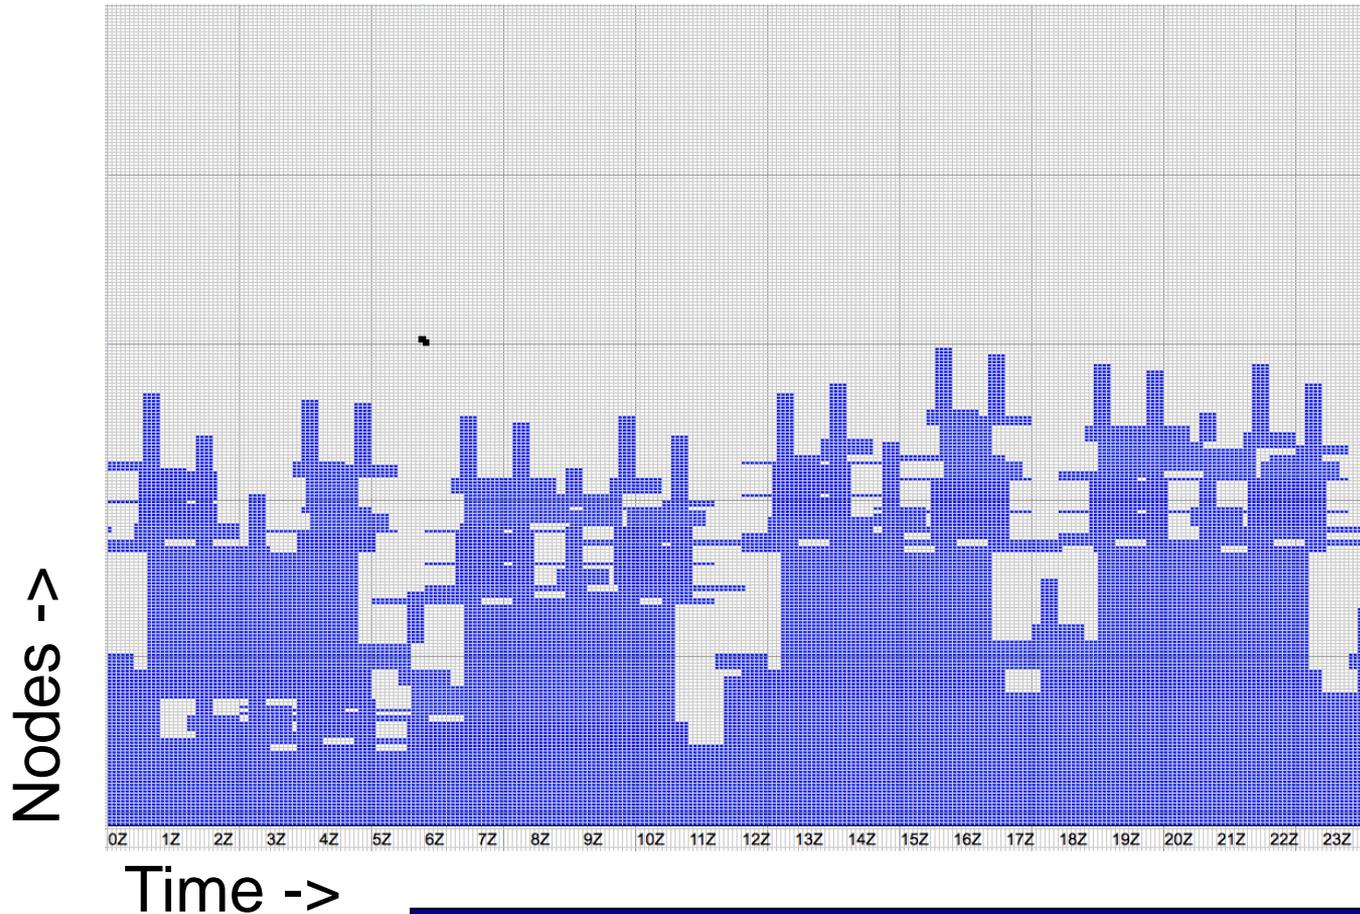


A Typical Workflow



A Typical Workflow





There are 165 separate reservations per day, 38% of total corehours available

- Allows us to gather specific information about workflow
 - Number of jobs, size
 - Deadline
 - Input/Output sizes
 - External dependencies
- Allows to have complete view of all requests, and ensure all can be finished on time
 - We need flexibility to ensure all real-time experiments can be packed on
- Ensure that good software engineering practices are used to construct the workflows
 - IE. Require use of Workflow manager or other tools to build reliable workflows.

https://jetdocs.rdhpcs.noaa.gov/wiki/index.php/Using_the_Workflow_Manager

- Why so many questions?

- Reservations are resource intensive, and block other workloads
- To make sure users have thought about what resources they really need
- Optimize resource usage
- Maximize likelihood of successful completion.
- So we can make sure that requests don't conflict
- You should already know the answers to the questions



Downloading Real-Time Experiment Form



The latest form can be downloaded from:

https://jetdocs.rdhpcs.noaa.gov/wiki/index.php/Real-Time_Experiment_Application_Form

Questions???