
LEAD-VGrADS Demo

Status Update

Lavanya Ramakrishnan

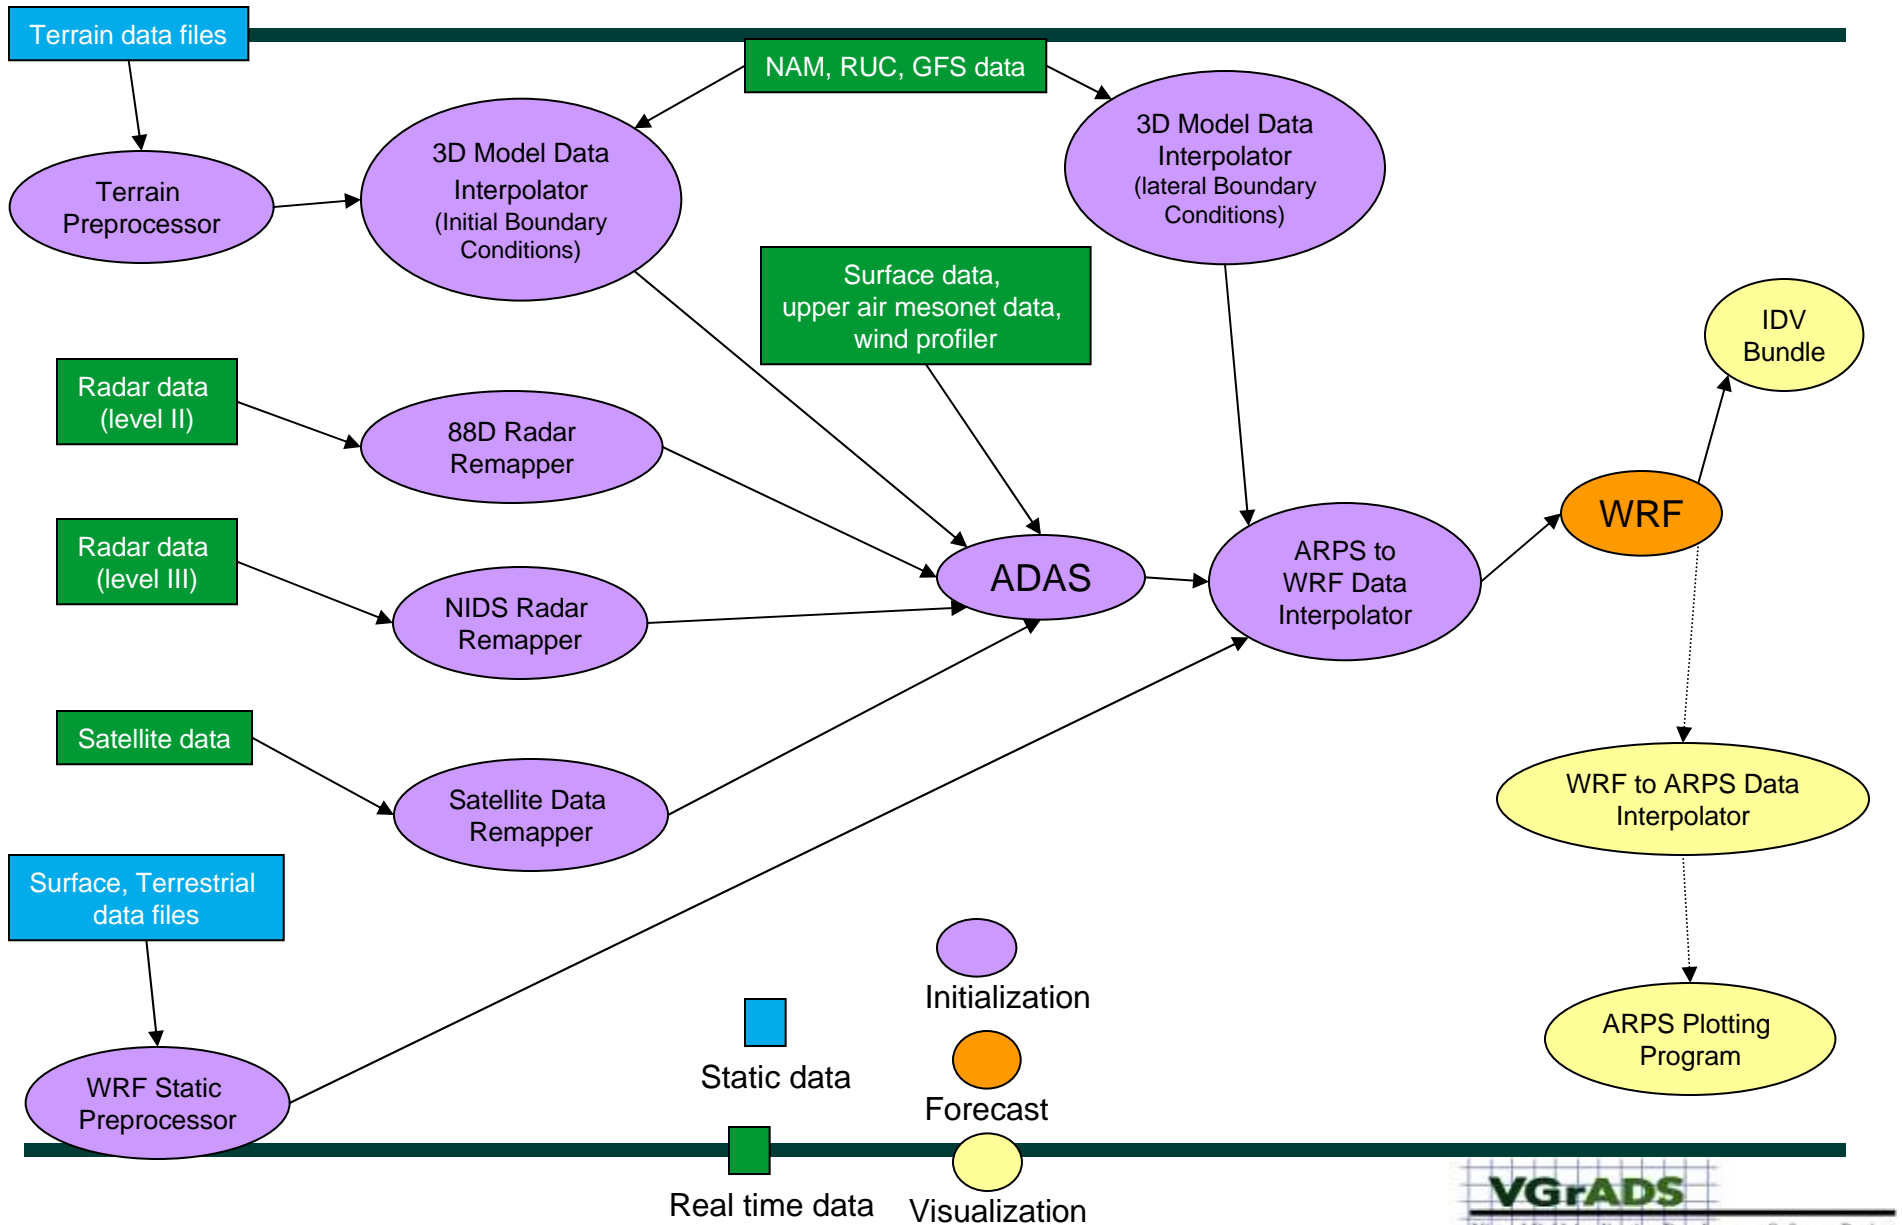
Daniel Nurmi

Yang-suk Kee

Ryan Zhang

September 6, 2006

LEAD Static Workflow





Experiment Builder Portlet

Experiment Wizard

User: lavanya/EMAIL=lavanya@renci.org Project: ThuTest

Specify a name, description, and select workflow

Name:

Description:

Workflow

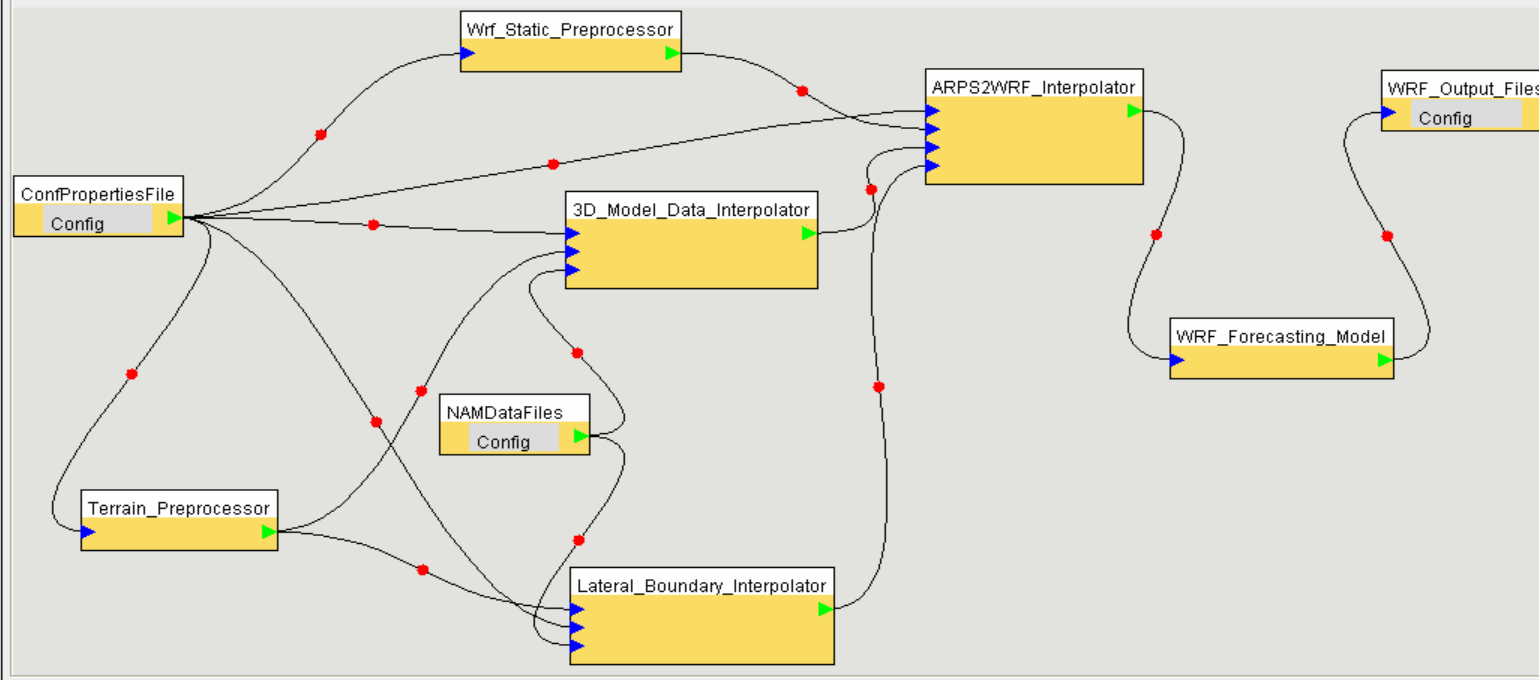
My Workflows

Sample Workflows

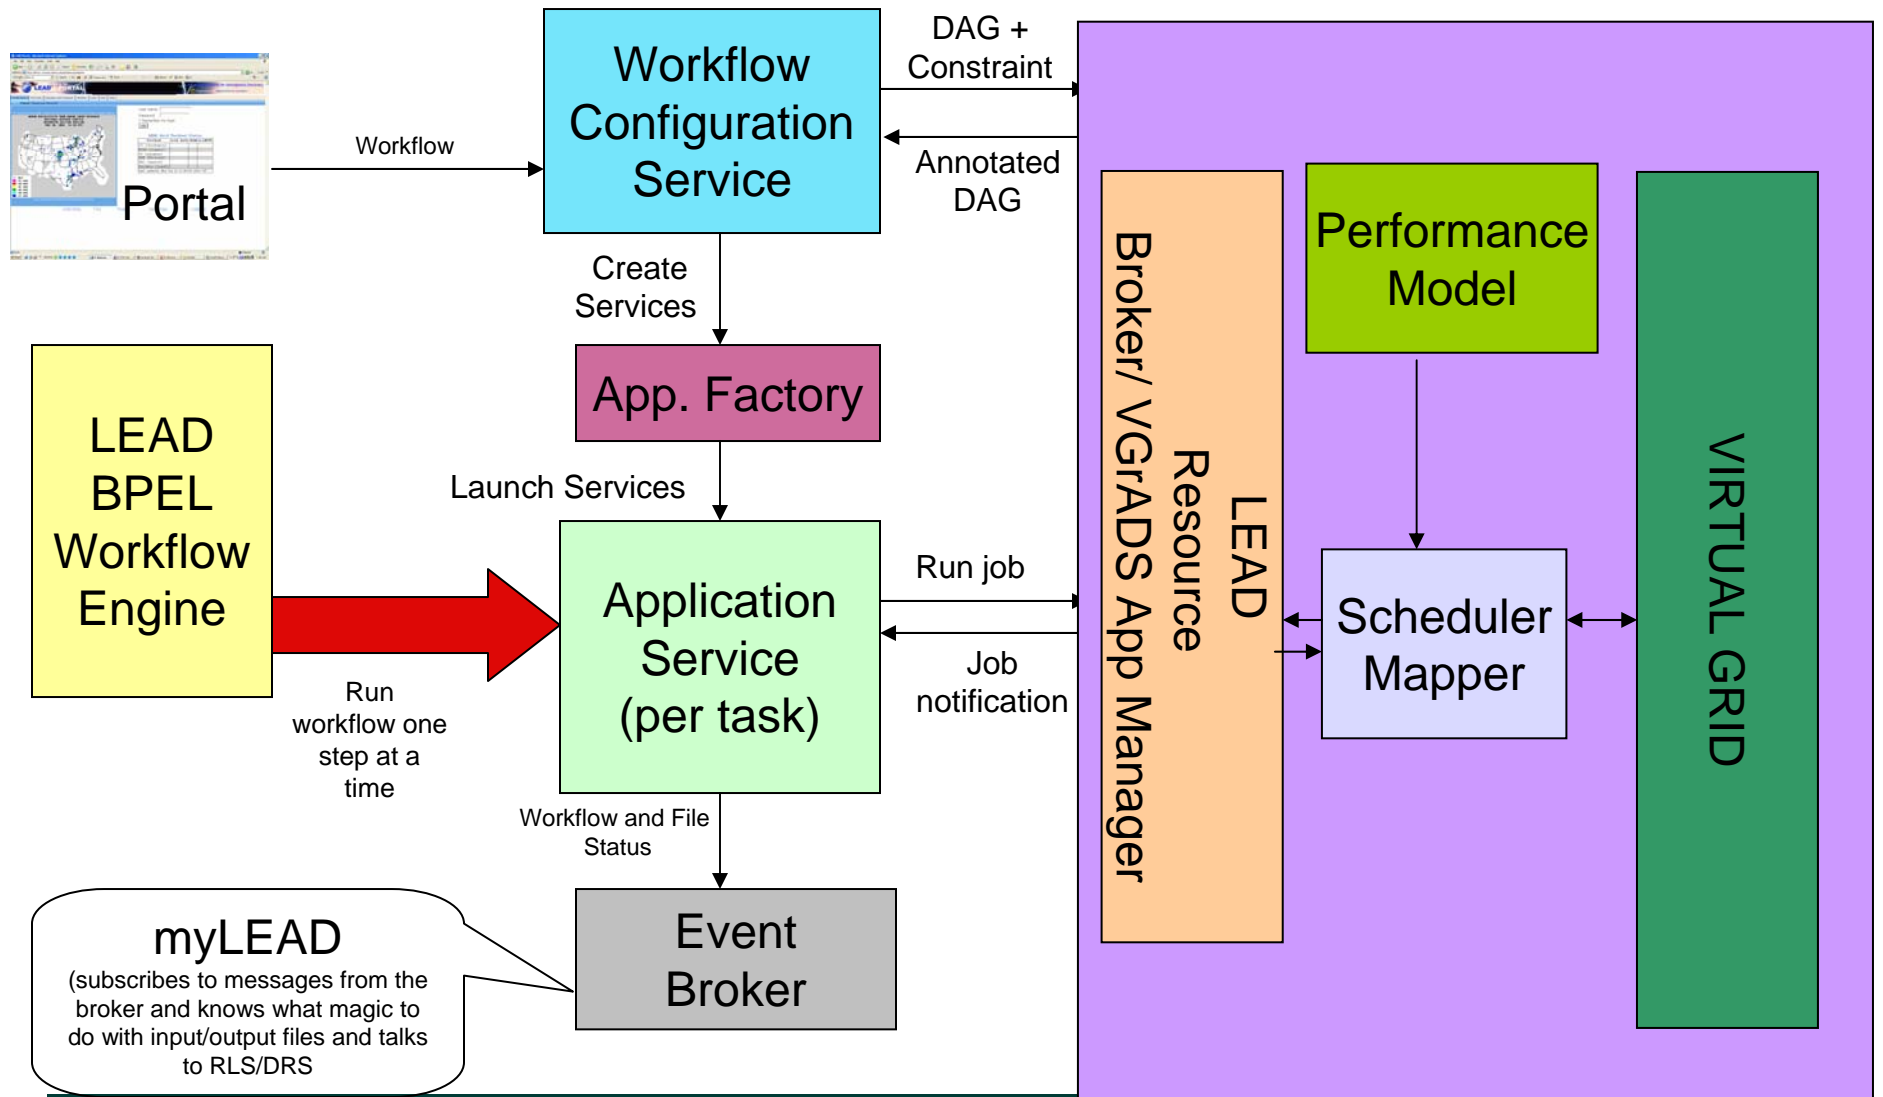
NAM-Initialized-WRF-Forecast

Edit Refresh New

Description

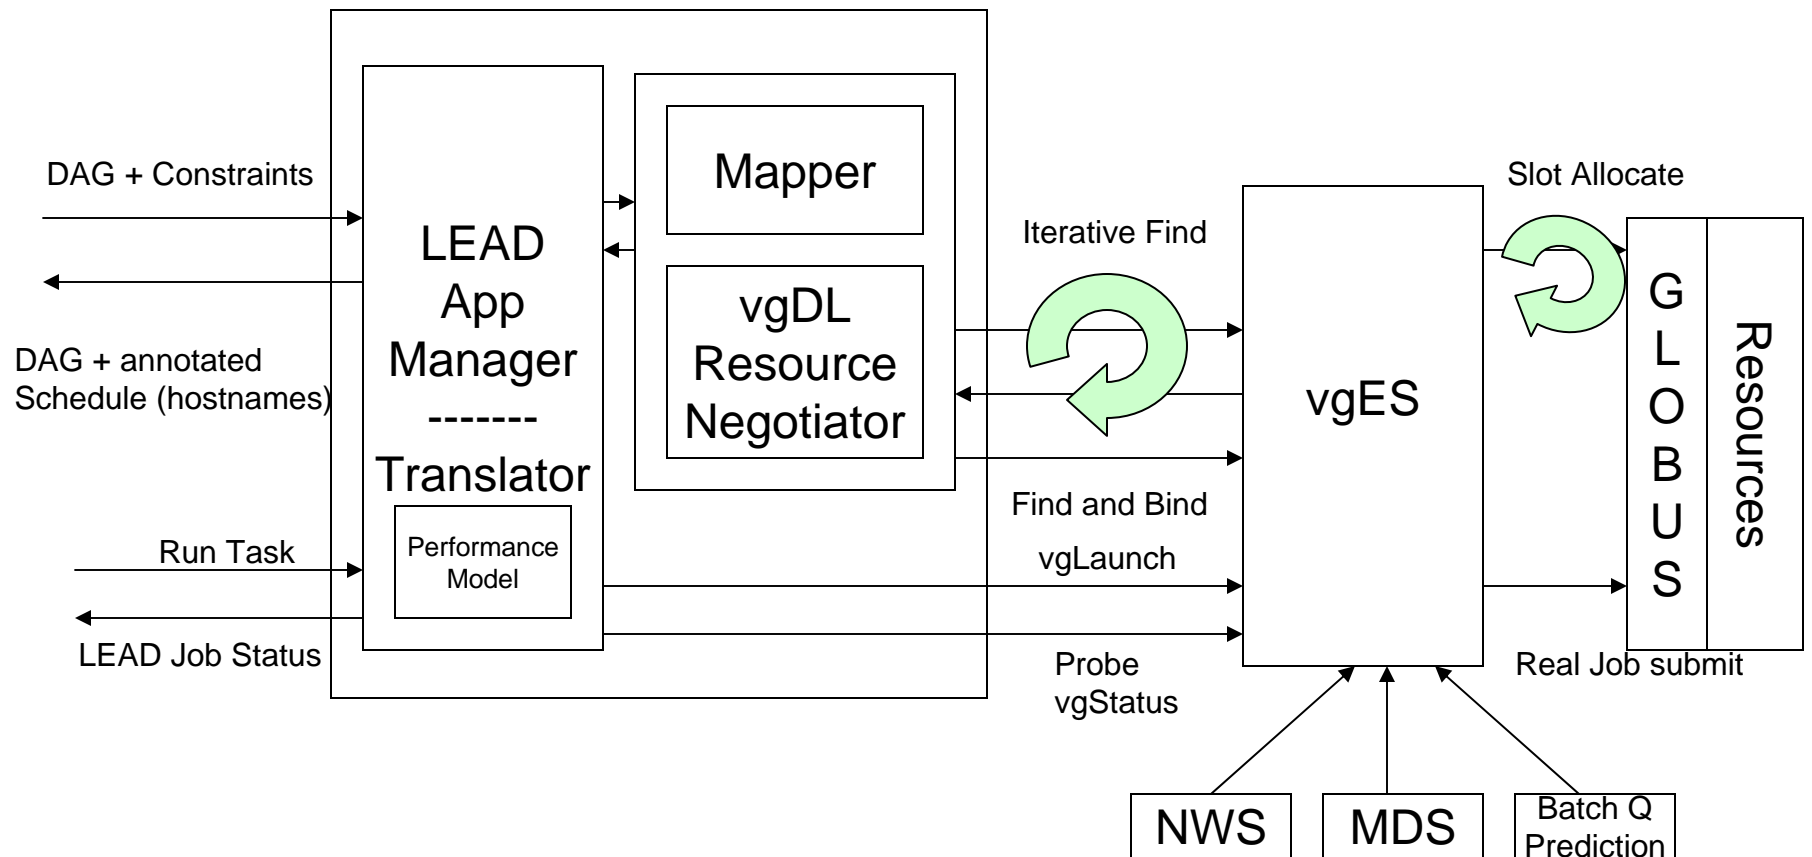


Integrated Architecture (LEAD-VGrADS)



Schedule towards a workflow deadline

VGrADS Demo Components



Milestones

- ☑ Aug 20 - First round exchange of jars, libraries etc
- Aug 29/30 - End to end integrated testing with 3 resource sites (real dag, fake applications)
- Sept 6 - Document issues from testing and resolve issues from testing
- Sept 15 - All demo features implemented and testing begins
- Sept 30 - End to end testing with lead applications complete, test case running every day
- Oct 7 - decide the scenario and artificial constraints
- Oct 15 - Test, test, demo viz, finalize demo details etc
- Oct 30 - keep it working!
- Nov 11 Demo time

Current Status and Open Issues [1/2]

- Minor problems and bugs were detected and fixed
- NCSA TeraGrid machine inaccessible
 - queues were being flushed
 - debugging is hard with queue wait times
- Missing information for Scheduler
 - population of vgES database
 - Missing MDS information, will populate manually
 - need for NWS data for data size estimates
 - Now this is setup

Current Status and Open Issues [2/2]

- **GRAM/PBS issues**
 - PBS Glide in through GRAM doesn't get environment properly on some machines
 - Launch of second-level GRAM
 - vgES slot manager detects status of pbs setup and stands up the GRAM
- **Slightly modified immediate goal**
 - Get software running on one machine for simple tasks
 - Simplified shorter DAG for end to end testing

Target Machines Update

- **NCSA TG - Mercury**
 - Should we get reservations for testing
- **UC TG**
 - Will add since lower load
 - Need to collect performance model data
- **IU TG/ Big Red**
 - Smaller cluster
- **SDSC TG**
- **Rice Cray ADA**
- **RENCI dante**
- **RENCI dante**
 - Launching pad for the VGrADS stack
 - Development machine
- **IU Tyr/Other nodes**
 - For LEAD web services, etc
 - Portal

Performance Model and Application Manager

Lavanya Ramakrishnan

Application Manager [1/2]

- **Standard LEAD Resource brokering interface**
 - Header inputs: workflow id, task id, etc
 - Inputs: DAG, workflow constraints (deadline)
 - Output: AnnotatedDAG with host information (hostnames, URL for services, no procs.)
- **Scheduler interface**
 - Input: Performance model, parsed DAG, constraints
 - Output:
 - Start time, End time, VG information (save state)
 - cluster head node, no processors (for LEAD)

Application Manager [2/2]

- **Job submission interface**
 - Header inputs: workflow id, task id, etc
 - Input: contact string, RSL
 - Synchronous Output: status
 - Asynchronous Notification: Status from VG
- **Fault Tolerance in the App manager**
 - Check the start and end times as jobs finish to keep track of if schedule is being met
 - Check if expected host names and mapped hostnames match (incase there was a change in schedule)
 - Collect performance model data

Performance Model Methodology

- **Timing data collection**
 - Define resource classes
 - Look-up table for application and resource with respect to no CPUs
 - Calculate optimal number of nodes
 - Long term: calculate timing information based on user input
- **Data size collection**
 - consider single case
 - Long term: project data sizes based on user input

Performance Model Queries

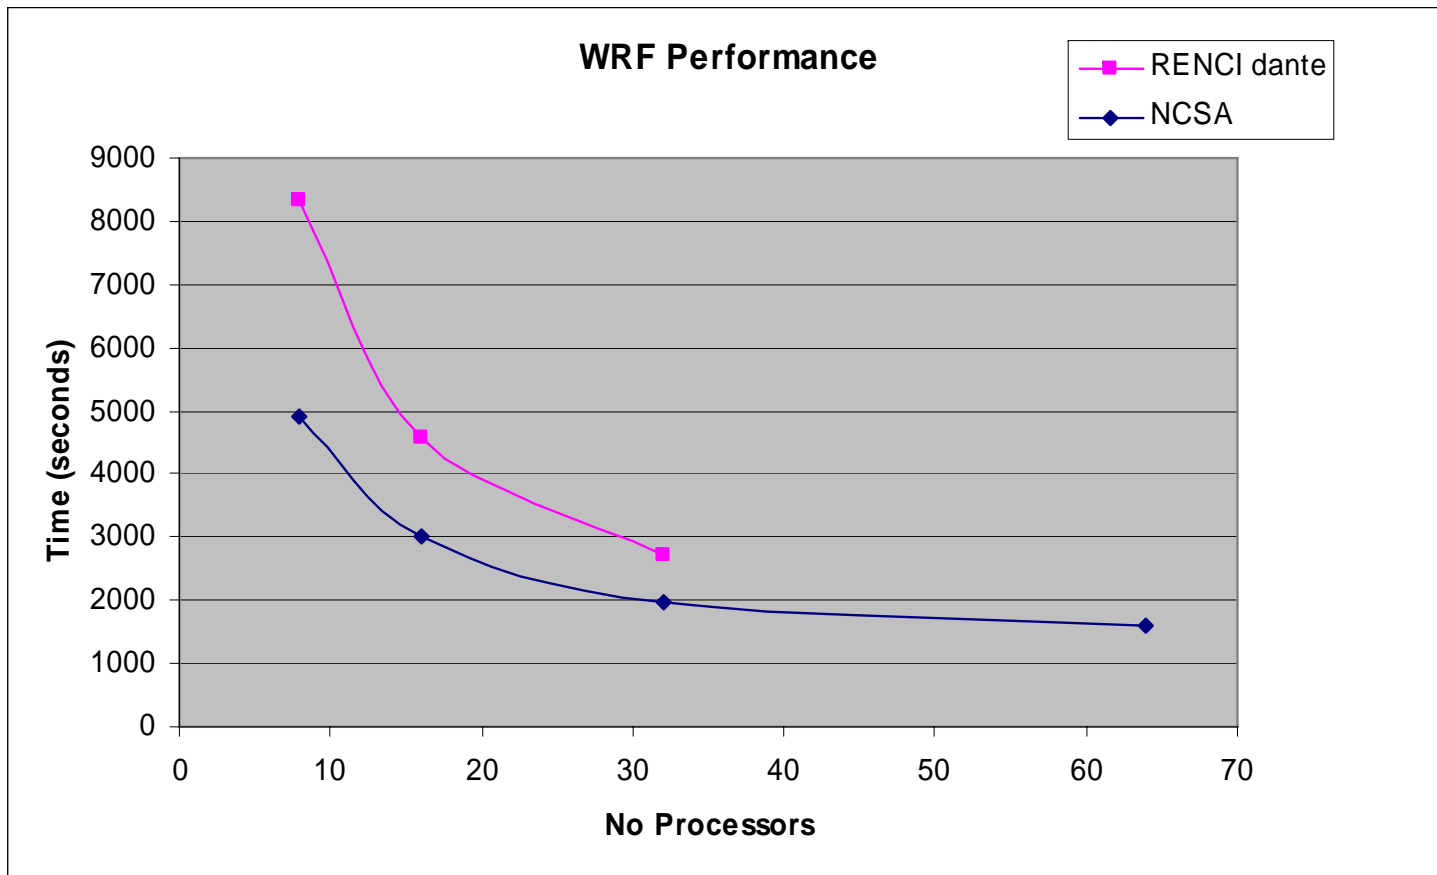
- **Get Performance Model for Task T**
- **Get Performance Model for Task T on Resource Class R**
- **Get Optimal Nodes for Task T on Resource Class R**
- **Get Optimal Task Time for Task T on Resource Class R**
- **Get Resource Classes**

Performance Model (Sample from NCSA)

Application	Execution Time (secs) [procs]	Input data	Output data
arps2wrf	42[8 procs]	668M	198M
arpsintrp	97	1.2G	930M
arpstrn	7	32K	256K
ext2arps	32[8 procs]	15M	184M
lateral	89[8 procs]	59M	1.9G
wrf	4892[8 procs]	198M	32K
wrfstatic	354	32K	19M

WRF Performance

- Intricacies of the science of splitting processors on the X and Y domains is not considered here



Questions?

