



Leveraging Grids: An Application Perspective

Lavanya Ramakrishnan
lavanya@renci.org

**Renaissance Computing Institute
Duke University
North Carolina State University
University of North Carolina - Chapel Hill**

Renaissance Computing Institute


www.renci.org

- **Vision**

- a multidisciplinary joint institute
 - academia, commerce and society
 - Duke, UNC, NCState, ...
- broad in scope and participation
 - from art to zoology

- **Objectives**

- enrich and empower human potential
- enable and sustain economic development
- develop and deploy leading infrastructure
 - computing, communications and data management
 - visualization, collaboration and manufacturing
- regional vision, national visibility
 - NC Bioportal, TeraGrid, LEAD, VGrADS, etc



We are hiring! positions@renci.org

Outline

- **Today's Grid Deployments**
 - Bioportal
 - Ocean Modeling
- **Next Generation Grids**
 - Service Oriented Architecture
 - e.g. LEAD
 - Grid Application Development
 - Virtual Grid Application Development Software
 - Cluster on Demand

Science Communities

- **North Carolina Bioportal**
 - leverage state-wide investment in bioinformatics and grid
 - undergraduate education, graduate education, faculty research
- **The Carolina Center for Exploratory Genetic Analysis**
 - develop collaborative experiences and plans
 - preliminary data to apply for a P50 grant
 - develop a prototype informatics infrastructure
 - data models, methods, tools and portals
 - facilitate use of best practices for existing projects



Bioportal Partners

- **Participating institutions**

- Renaissance Computing Institute (RENCI)
 - strategy, integration and development
- Wake Technical Community College
 - student training and education
- UNC-CH Center for Bioinformatics
 - bioinformatics training
- UNC-CH Information Technology Services (ITS)
 - infrastructure hosting



University of North Carolina at Chapel Hill
Center for Bioinformatics

- **North Carolina Bioportal**

- funded by UNC Office of the President

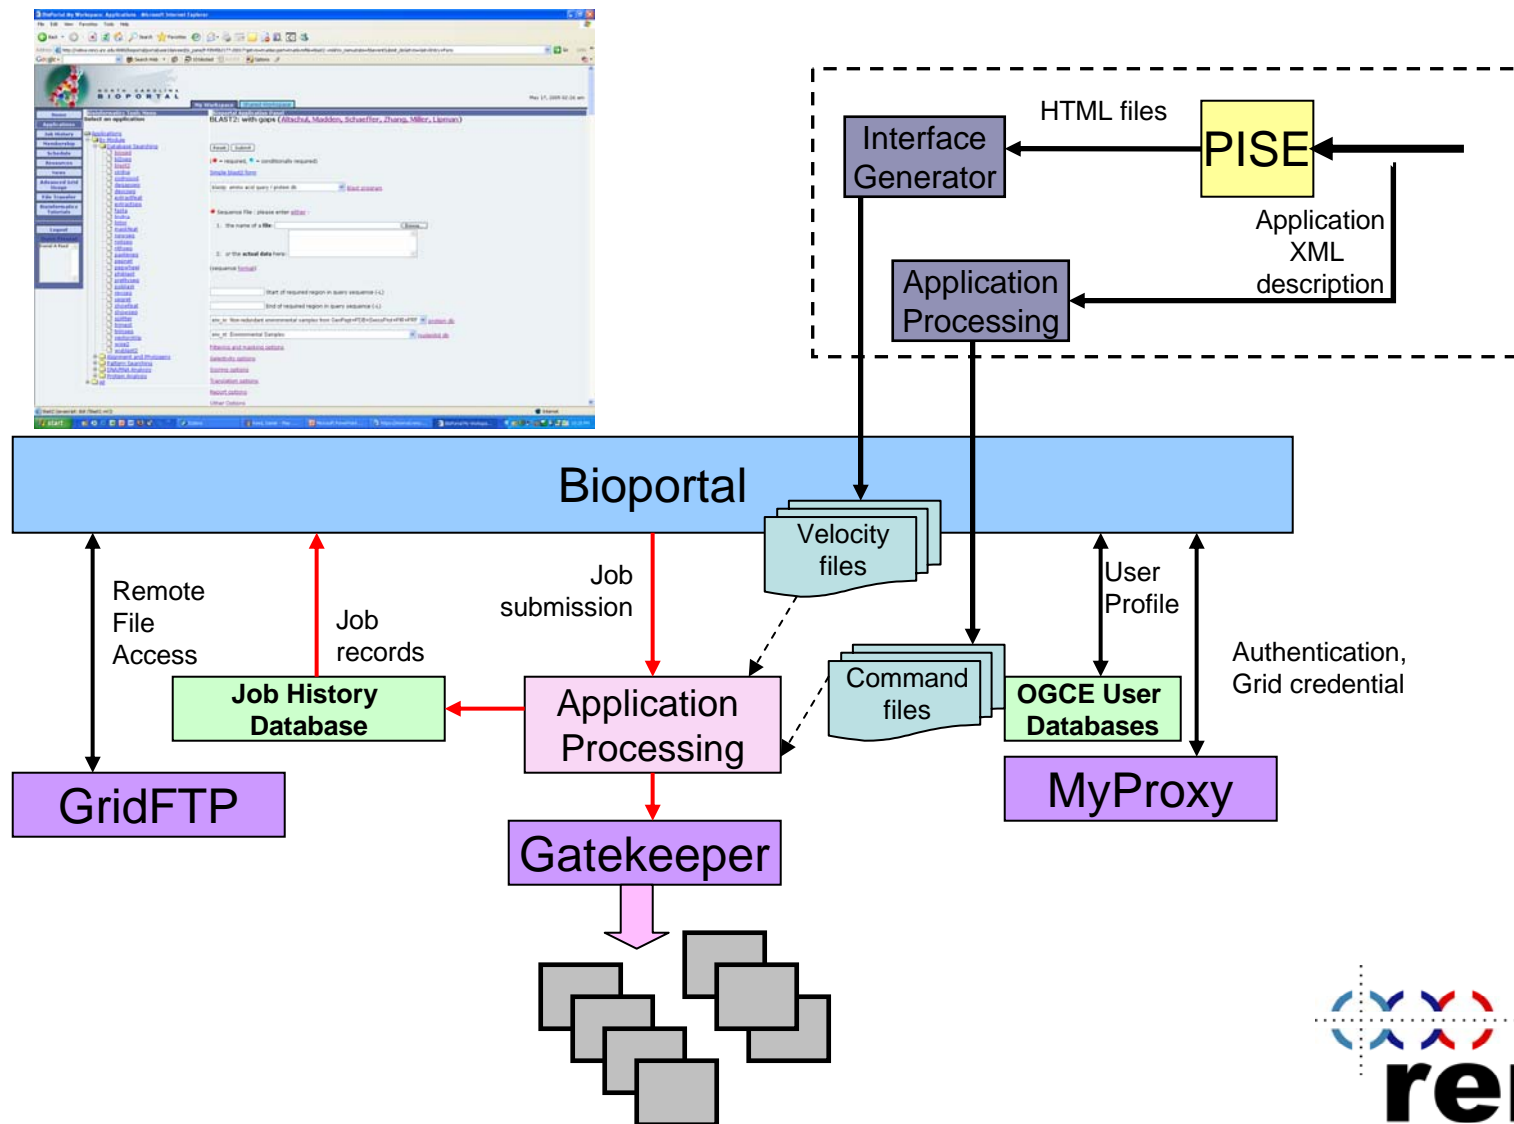


- **National level deployment**

- NSF TeraGrid



Biportal Architecture



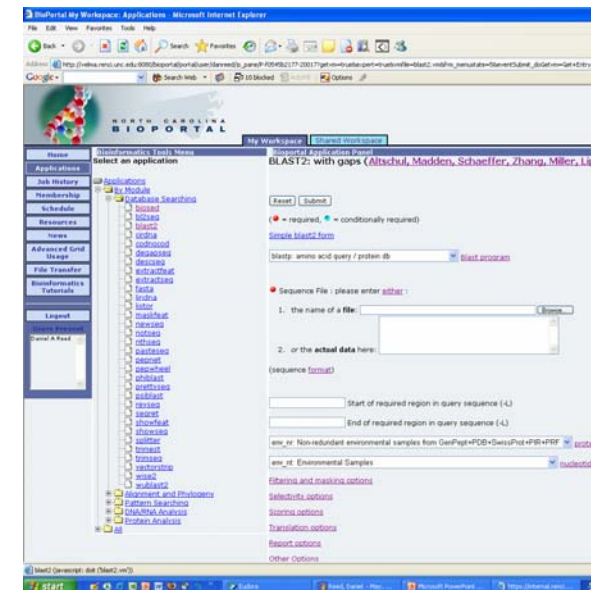
Bioportal

- **Features**

- access to common bioinformatics tools
- extensible toolkit and infrastructure
 - OGCE and National Middleware Initiative (NMI)
 - leverages emerging international standards
- remotely accessible or locally deployable
- packaged and distributed with documentation


- **Education and training**

- hands-on workshops
- clusters, Grids, portals and bioinformatics



North Carolina





- Home
- Applications
- Job History
- Membership
- Schedule
- Resources
- News
- Advanced Grid Usage
- File Transfer
- Bioinformatics Tutorials
- Logout
- Users Present: Lavanya Ramak

- Bioinformatics Tools Menu
Select an application
- Applications
 - By Module
 - Database Searching
 - bioesd
 - bl2seq
 - blast2
 - cirdna
 - degapseq
 - descseq
 - extractfeat
 - extractseq
 - fasta
 - lindna
 - listor
 - maskfeat
 - newseq
 - notseq
 - nthseq
 - pasteseq
 - pepnet
 - pepwheel
 - phiblast
 - prettyseq
 - psiblast
 - revseq
 - secret
 - showfeat
 - showseq
 - splitter
 - trimest
 - trimseq
 - vectorstrip
 - wise2
 - wublast2
 - Alignment and Phylogeny
 - Pattern Searching
 - DNA/RNA Analysis
 - Protein Analysis
 - All

Bioportal Application Panel
EXTRACTFEAT : Extract features from a sequence (EMBOSS)

Reset Submit

(● = required, ● = conditionally required)

[Simple extractfeat form](#)

- [Input section](#)
- [Advanced section](#)
- [Output section](#)

Input section

● sequence -- any [sequences] (-sequence) : please enter [either](#) :

1. the name of a **file**:

2. or the **actual data** here:

(sequence [format](#))

[Return to the main part with your favorite browser's Back function]

Advanced section

Amount of sequence [before](#) feature to extract (-before)

Amount of sequence [after](#) feature to extract (-after)



- Home
- Applications
- Job History
- Membership
- Schedule
- Resources
- News
- Advanced Grid Usage
- File Transfer
- Bioinformatics Tutorials
- Logout
- Users Present
Lavanya Ramak

Xportlets : ProxyManager

The following GSI proxy credentials are loaded into your account:

(default proxy) /O=Grid/OU=UNC/OU=RENCI/OU=local/CN=Lavanya Ramakrishnan

Click the button below to add another GSI proxy credential to your account:

Get New Proxy

XPortlets: Gram Job Launcher

Launch a Gram Job

Executable name *

Parameters

Host name *

Executable Directory

Output file *

Launch Gram Job

Fields marked with * are required

Grid Job Submission Portlet

Monitor submitted jobs

Job name

Host name

Executable

Arguments

Directory:

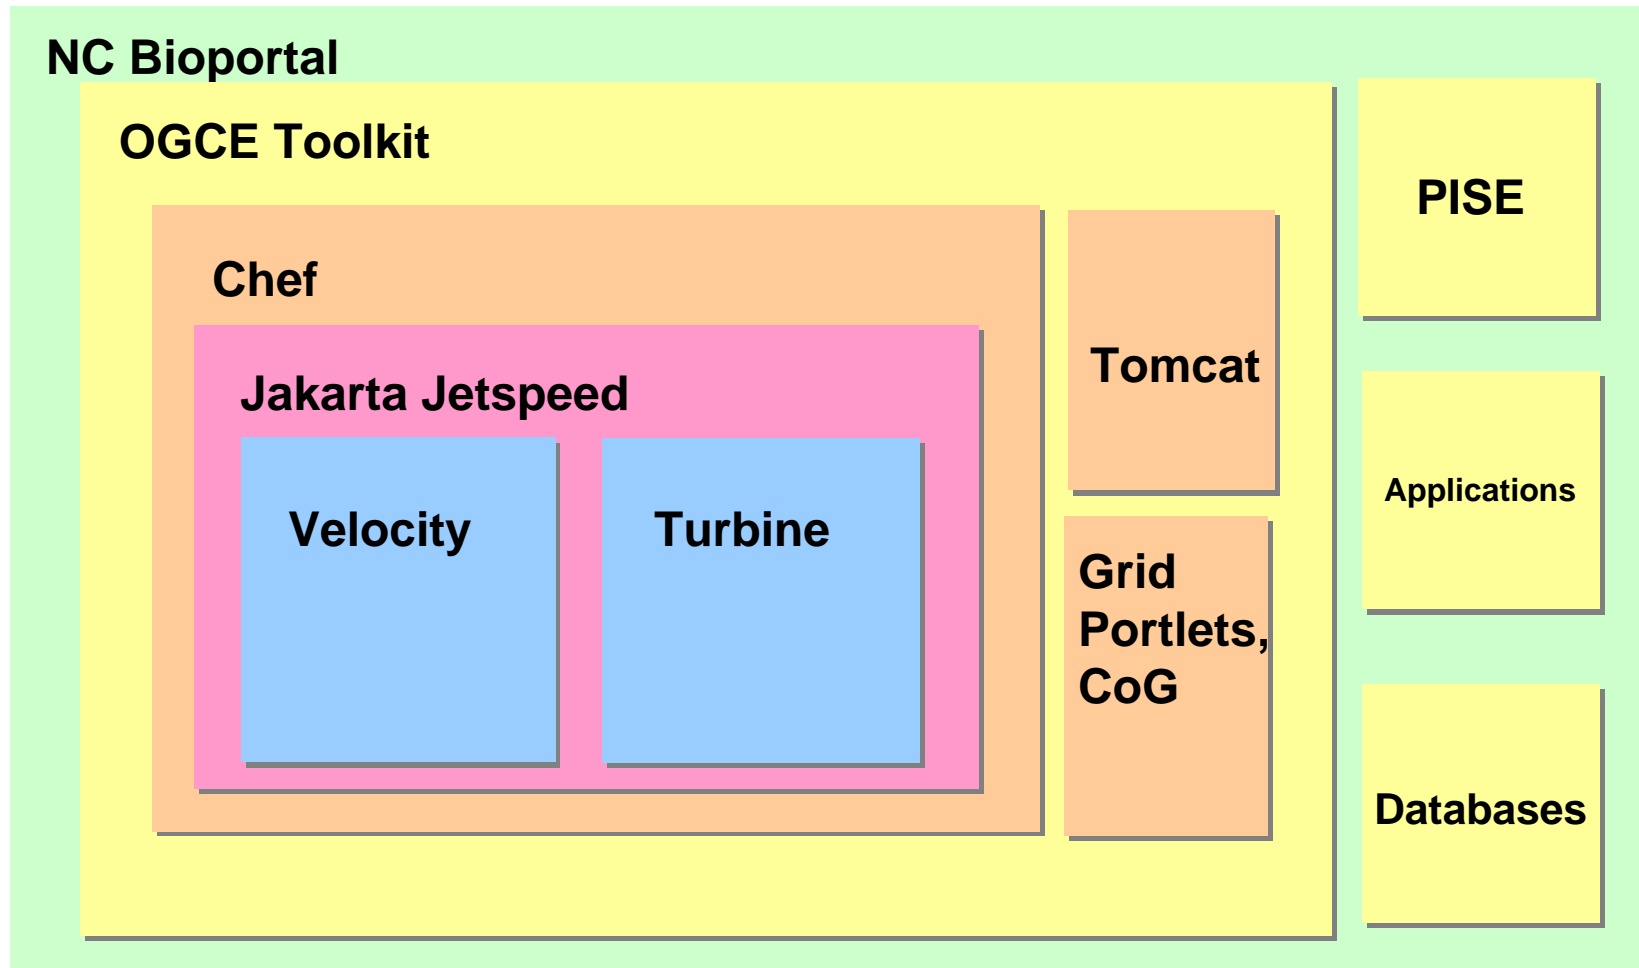
Standard Output File

Standard Error File

Run Job


Transparent to most users

Putting the Technologies Together



Biportal Technologies

- **Globus**
 - Job Submission through Gatekeeper
 - GridFTP for file transfer
 - MyProxy credential repository
- **Open Grid Computing Environment (OGCE)**
 - access to other Globus functionality through Java CoG
 - many application uses
 - LEAD, NEES, TeraGrid, NC Bioportal
- **Michigan Chef/Sakai**
 - collaborative course tool
 - later used for distributed communities (e.g., NEES)
- **Tomcat, Jetspeed, Velocity, Turbine**
 - Apache Java servlet container
 - enterprise information portal
 - Java-based template engine
 - web application framework



The **Apache Jakarta Project**

<http://jakarta.apache.org/>

renci
renaissance computing institute

PISE

- **Pasteur Institute Software Environment (PISE)**
 - generates web interfaces for molecular biology tools
 - XML specification for command line interface
- **Rationale and objectives**
 - simplify specification of program interfaces
 - homogeneous specification mechanisms
 - reuse of existing software interfaces
 - independent development and integration
 - extension for integration with graphical interfaces
 - complexity hiding and commonality
- **Bioportal program described in PISE**
 - semi-automated GUI synthesis from XML via Perl
- **Output is a generated command line, for example**
 - `blastall -p blastp -d env_nr -i query.dat.blast2.1116248106513 -a 2`

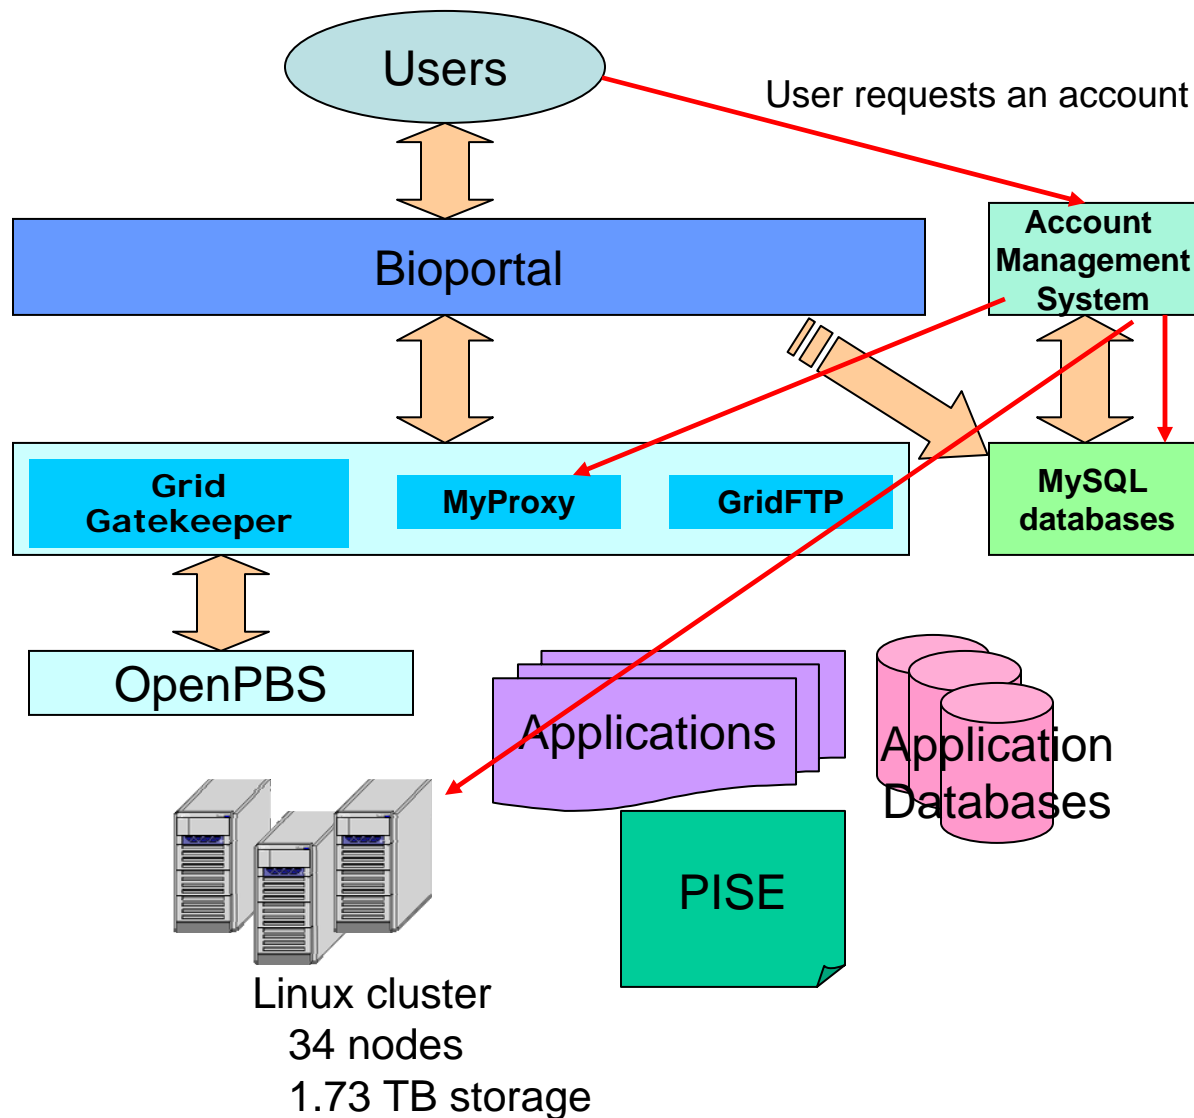


An Example PISE XML

```
<parameter type="Integer">
  <name>start_region</name>
  <attributes>
    <prompt>
      Start of required region in query sequence (-L)
    </prompt>
  </attributes>
</parameter>
```

```
<parameter ismandatory="1" iscommand="1"
  issimple="1" type="Excl">
  <name>blast2</name>
  <attributes>
    <prompt>Blast program</prompt>
    <format>
      <language>perl</language>
      <code>"blastall -p $value"</code>
    </format>
    <vdef><value>blastp</value></vdef>
    <group>1</group>
    <vlist>
      <value>blastn</value>
      <label>blastn: nucleotide query / nucleotide
      db</label>
      <value>blastp</value>
      <label>blastp: amino acid query / protein
      db</label>
      <value>blastx</value>
      <label>blastx: nucleotide query translated / protein
      db</label>
      <value>tblastn</value>
      <label>tblastn: protein query / translated
      nucleotide db</label>
      <value>tblastx</value>
      <label>tblastx: nucleotide query transl. / transl.
      nucleotide db</label>
      <value>psitblastn</value>
      <label>psitblastn: protein query / transl. nucleotide
      db</label>
    </vlist>
  </parameter>
```

Integrated Account Management



Behind the scenes

1. Create Unix account
2. Create a certificate request
3. Sign the certificate request
4. Update MyProxy
5. Add entry to gridmap file
6. Create a portal account

Community Accounts

- Single account for all jobs,
- Per user on portal
- Maintain audit tracks
- Admin portlet access

Current Biportal Applications

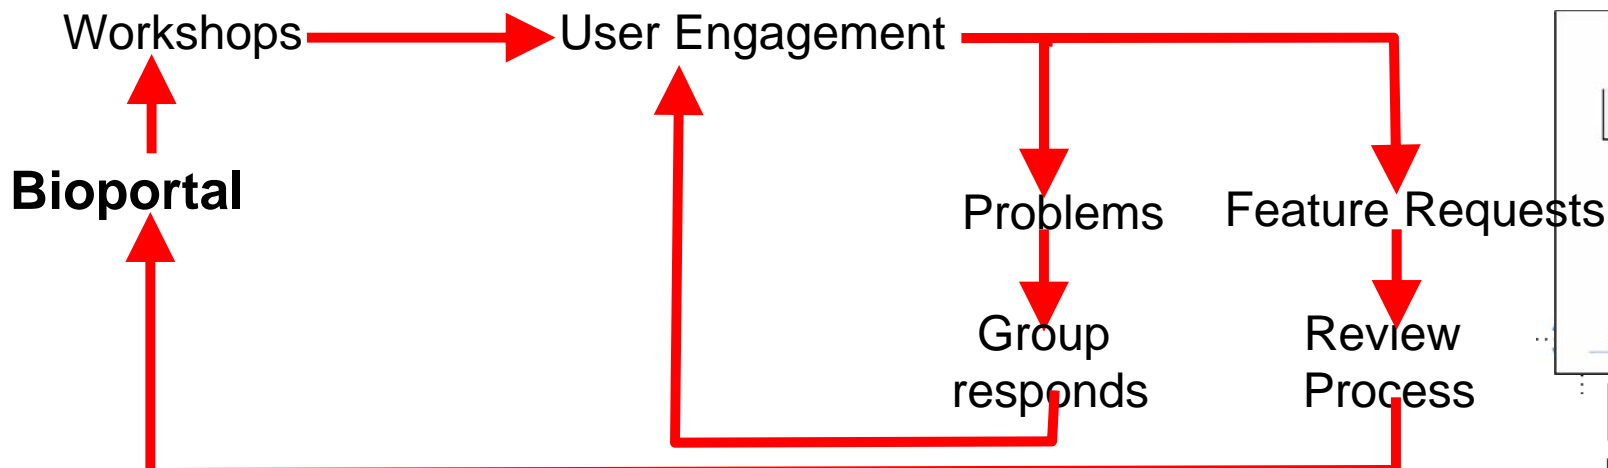
- **Applications**
 - ~140 distinct codes
- **Application Suites**
 - **EMBOSS**
 - European Molecular Biology Open Software Suite
 - **GLIMMER**
 - gene identification in microbial DNA
 - **HMMER**
 - Hidden Markov Model program for profile-based sequence analysis
 - **NCBI**
 - diverse set of tools
 - **PHYLIP**
 - PHYLogeny Inference Package for inferring phylogenies
- **Others (incomplete list)**
 - ClustalW, FASTA
- **Standard bioinformatics databases**
 - NCBI Aggregate (95 GB)
 - three formats: native, BLAST and WUBLAST
 - GenBank (206 GB)
 - GenPept (3 GB)
 - PDB (6.3 GB)
 - Prints (72 MB)
 - RepBase (8.6 MB)
 - UniProt (12 GB)
 - PFam (8.7 GB)
 - ProSite (16 MB)
 - TransFac (36 MB)
- **Database update mechanism**
 - follows the schedule of the distribution source



Community Engagement

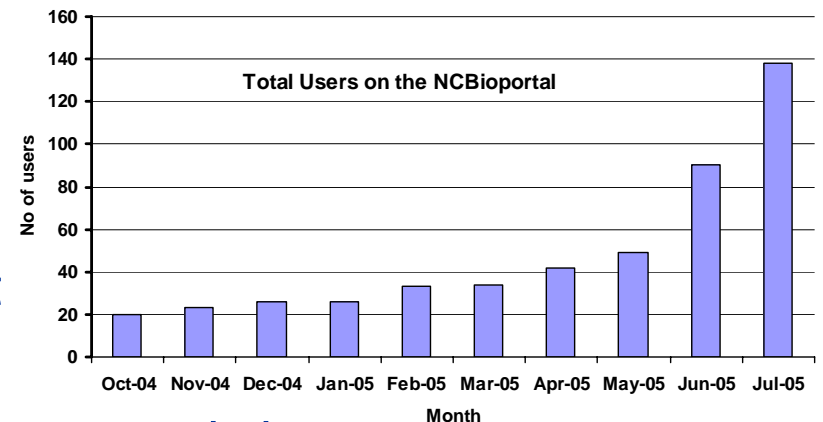


- **Workshops**
 - End users
 - Use the portal
 - System administrators, Developers
 - Local deployments, add your own application
- **Packaging**
 - User Guide, Administrator Guide, Developer Guide
 - PISE, Bio apps I, Bio apps II
 - Portal
 - Database and other helper scripts
 - Torque, Maui, MyProxy, Globus



Building the Biportal Experiences

- **Usability and support**
 - account management
- **Security**
 - account creation and management
- **Policy**
 - 14 day job history policy & 90 day password change
 - size and policy of database updates
- **Community Building**
 - increased usage driven by workshops
 - feature requests e.g. file format, workflow support
- **Deployment Challenges**
 - technology upgrades
 - testing



Biportal: What's Next

- **Community Engagement**
 - workshops, experiences and deployments
 - Software and documentation
- **Infrastructure**
 - dynamic job scheduling across multiple sites
 - load driven based on community use
 - fully automated database updates, possible distributed replication
 - driven by user needs and available disk space
- **Portal tool suite**
 - expand application and databases based on **user feedback**
 - phylogeny, morphology, microarray analysis, ...
- **Multiple deployment**
 - TeraGrid, etc



TeraGrid[™]



Outline

- **Today's Grid Deployments**
 - Bioportal
 - Ocean Modeling
- **Next Generation Grids**
 - Service Oriented Architecture
 - e.g. LEAD
 - Grid Application Development
 - Virtual Grid Application Development Software
 - Cluster on Demand

SURA Coastal Ocean Observing and Prediction (SCOOP) Program

- **National Integrated Ocean Observing System (IOOS)**
 - Rapidly assess, predict, and mitigate the impact of large storms
 - Make information available for emergency response
 - “Plug and play model” to enable next generation operation research
- **Service-oriented Architecture**
 - Compatible with other disciplines
 - Grid, distributed sensors

SCOOP Partnering Institutions

- **SURA Funded Participants**

- U. Florida, U. Alabama, LSU, VIMS, Texas A& M, U. Maryland, U. of Miami, GoMOOS, UNC

- **North Carolina Participants**

- **UNC Marine Sciences**

- Brian Blanton, Rick Luettich, Larry Mason (ITS)

- **MCNC/GCNS**

- Michael Garvin, Steve Thorpe, Chuck Kesler

- **REnaissance Computing Institute (RENCI)**

- Lavanya Ramakrishnan, Brad Viviano, Dan Reed



Science Scenarios

- **Daily operational forecasts**
 - 24/7/365 continuous,
- **Real-time ensemble model prediction**
 - time of hurricane, increased accuracy
- **Retrospective analysis**
 - evaluate model results, innovate new mechanisms
- **Interdisciplinary problems**
 - inundation affected by storm surge, terrestrial hydrology, precipitation

Cross-cutting Components

Directories

Security
(GSI, etc)

...

Monitoring

User Interface Layer

portal (resource access, workflow interfaces, interactive search services, etc.), visualization tools, software libraries

Application and Tools Layer

Models
& Analysis
Tools

Workflow
Tools

...

Data
Visualization

Data
Translation

Management Layer

Archive
Management

Data
Management

Resource
Management

...

Application
Management

Resource Access Layer

Data transport

LDM, GridFTP, scp, etc

Web service protocols

SOAP, XML, WSRF, etc

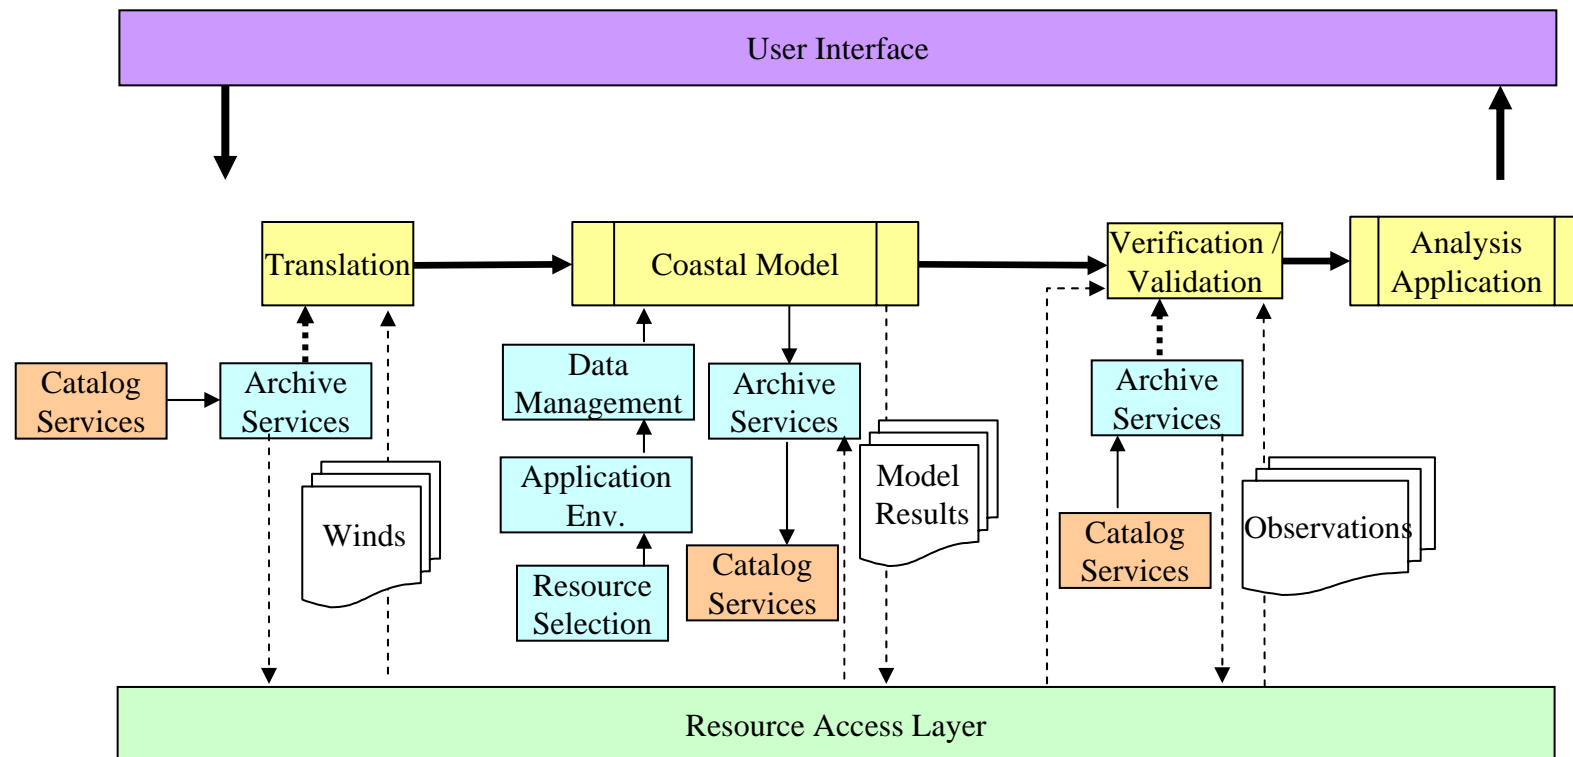
...

Virtualization

⋮

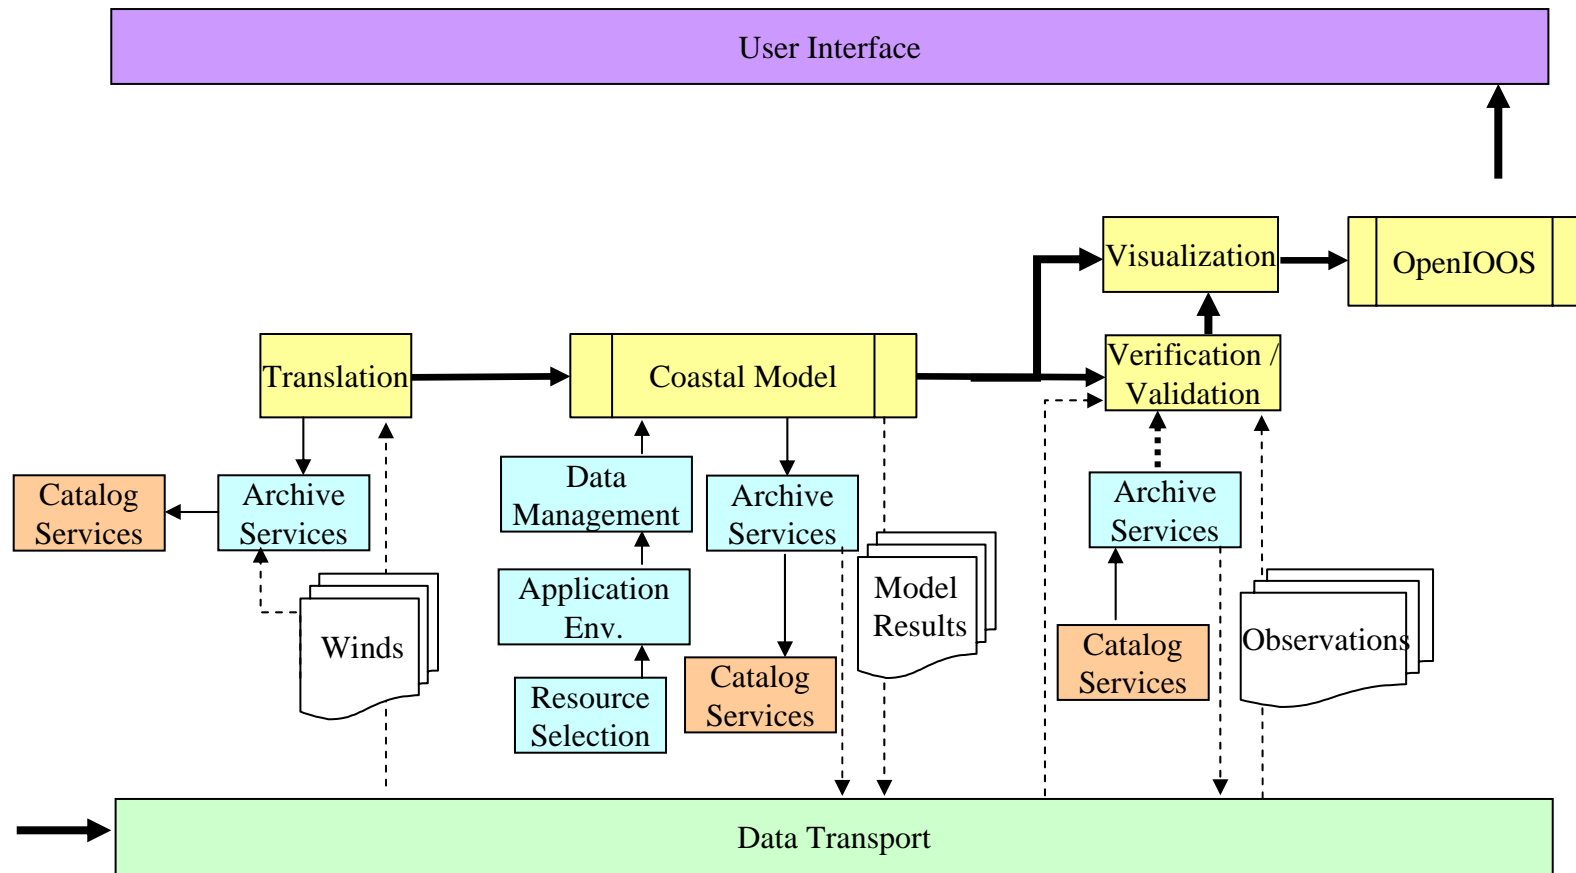
Resources (compute, storage, network)

Arch. Scenario: Retrospective Analysis



Source: SCOOP Team

Arch. Scenario: Forecast Analysis



Source: SCOOP Team

SCOOP Portal: ADCIRC Hindcast

SCOOP Portal Feb 22, 2005 09:02 am

My Workspace
Hindcast Model Runs

Home | Membership | Schedule | Resources | News | Chat | File Transfer | OPeNDAP Server | Proxy Manager | ADCIRC | Logout | Customize | Users Present (Brian Blanton)

Enter range of dates for running the model

Year: 2004 | Month: September | Date: 12


Year: 2004 | Month: September | Date: 26

Model: ADCIRC

Parallel

Submit

Set Run Dates
Hurricane Ivan



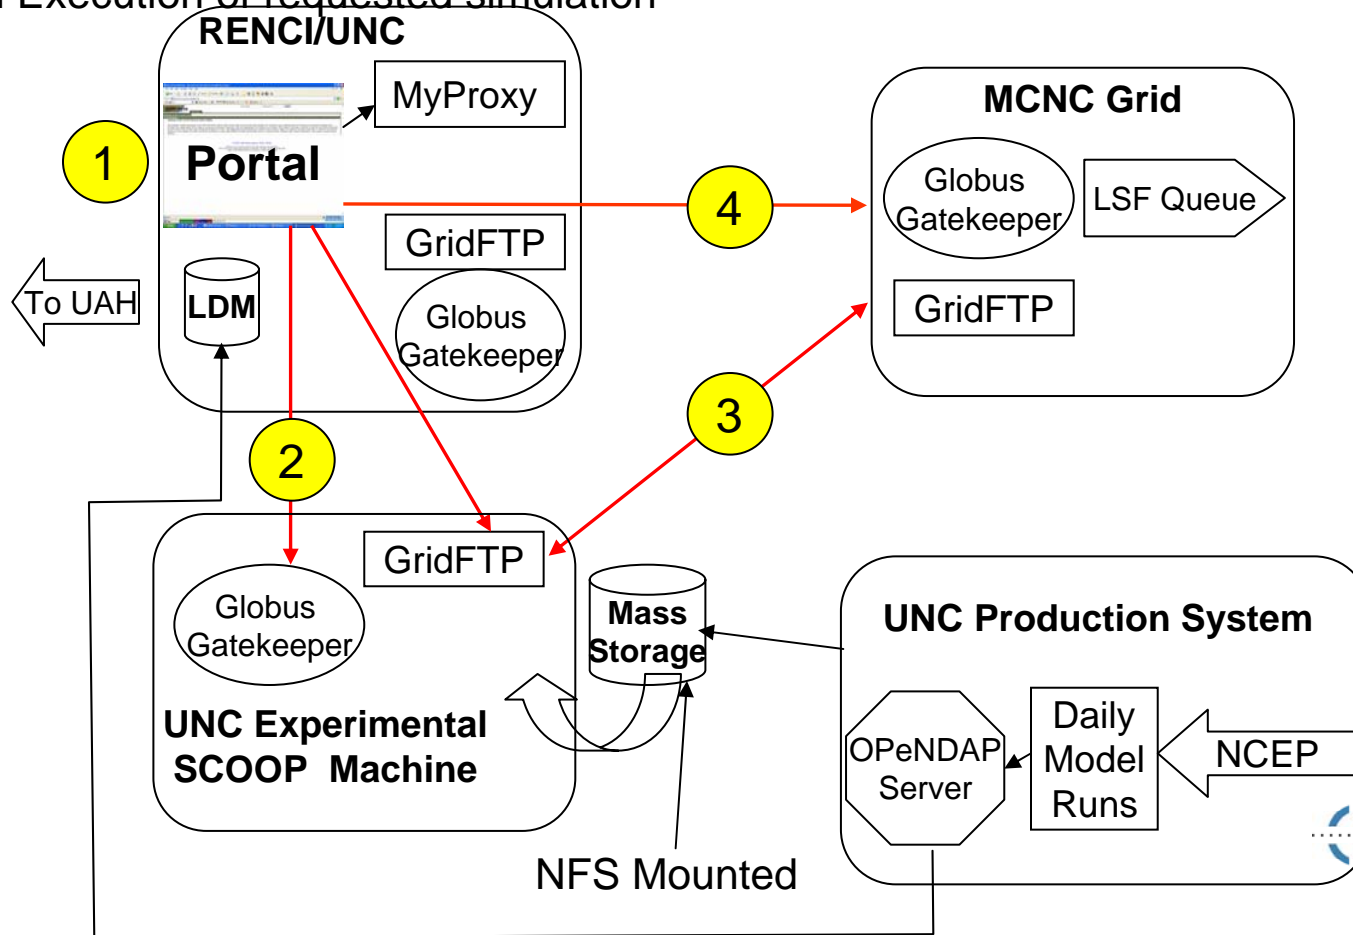
| [SCOOP](#) | [UNC Marine Science](#) | [RENCI](#) | [MCN](#)
© Renaissance Computing Institute. All rights reserved. Copyright © The Open Grid Computing Environments Collaboratory. All rights reserved. CHEF v01.00.07[build #303141], Jetspeed v1.4b2[cvs08oct].

Submit Job
To Grid

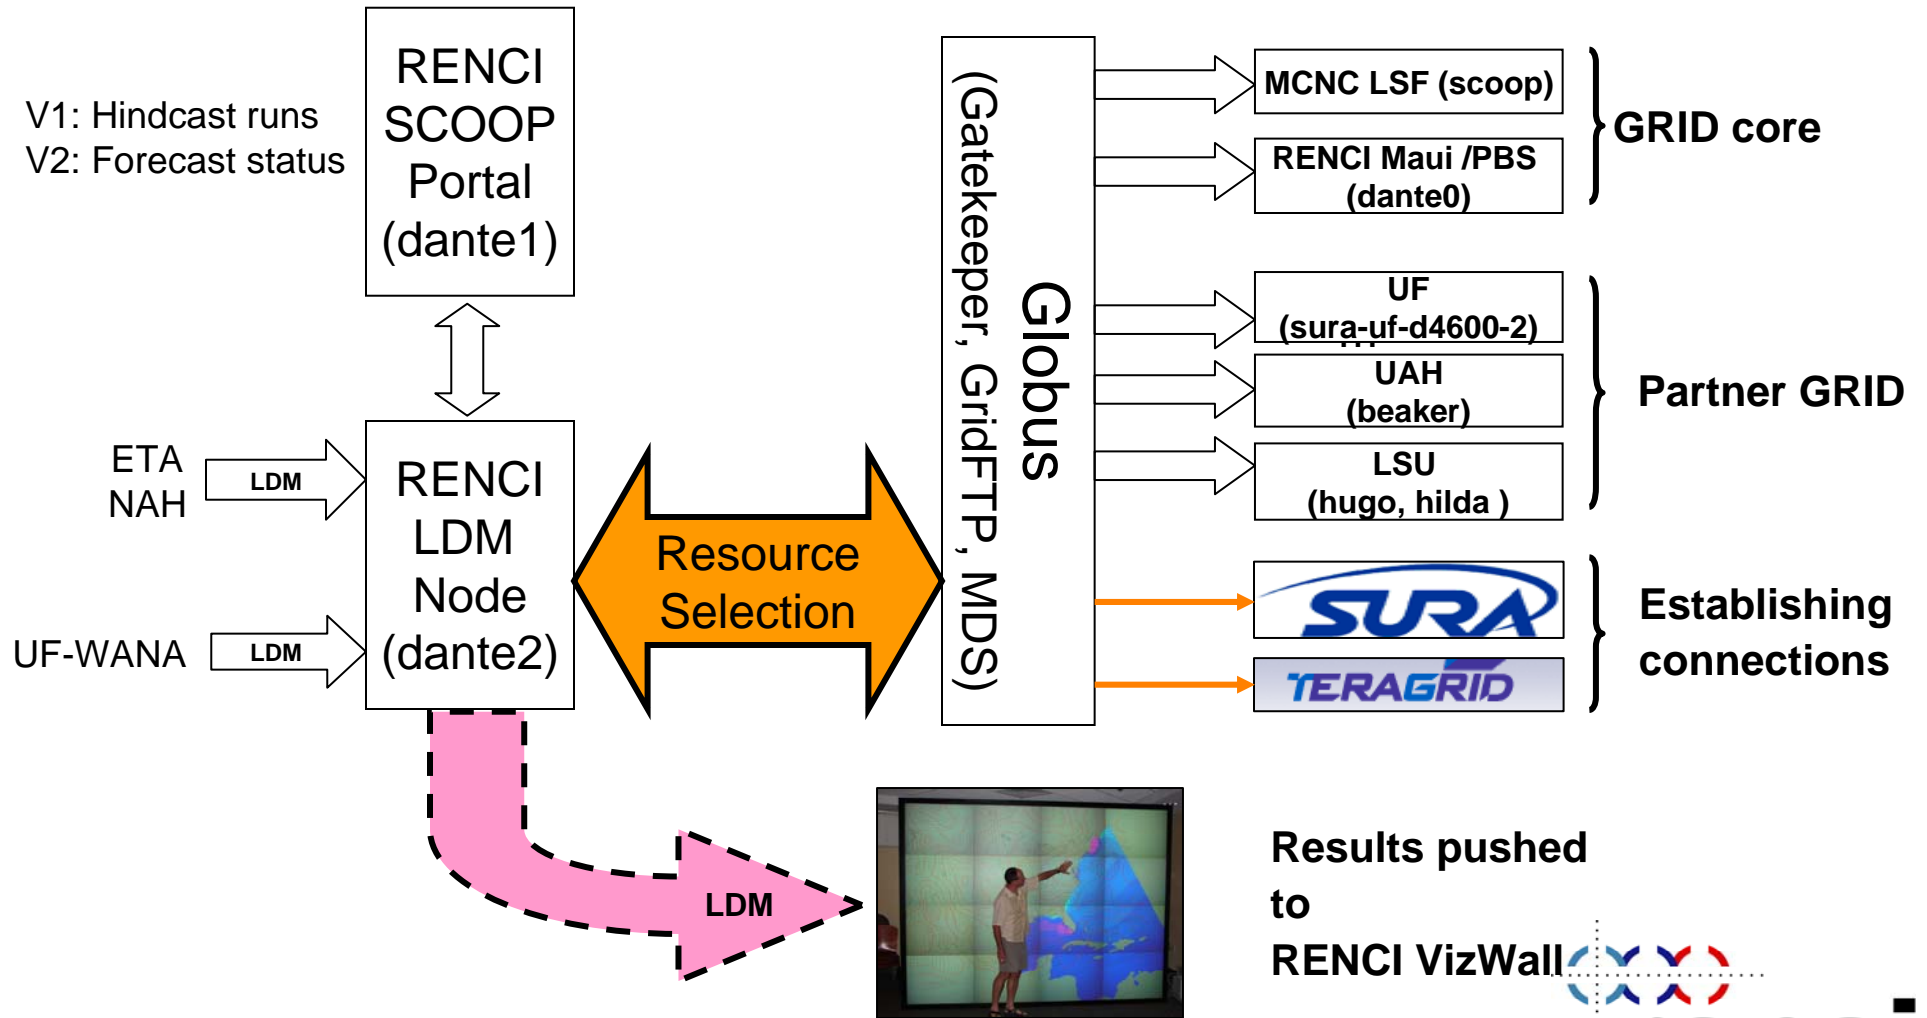
Current ADCIRC grid
16 CPU Decomposition

Hindcast Analysis on the Grid

- 1 Specify model run parameters
- 2 Make tarball of needed Archived Files
- 3 Third-party transfer between Portal host and Compute host
- 4 Portal Execution of requested simulation

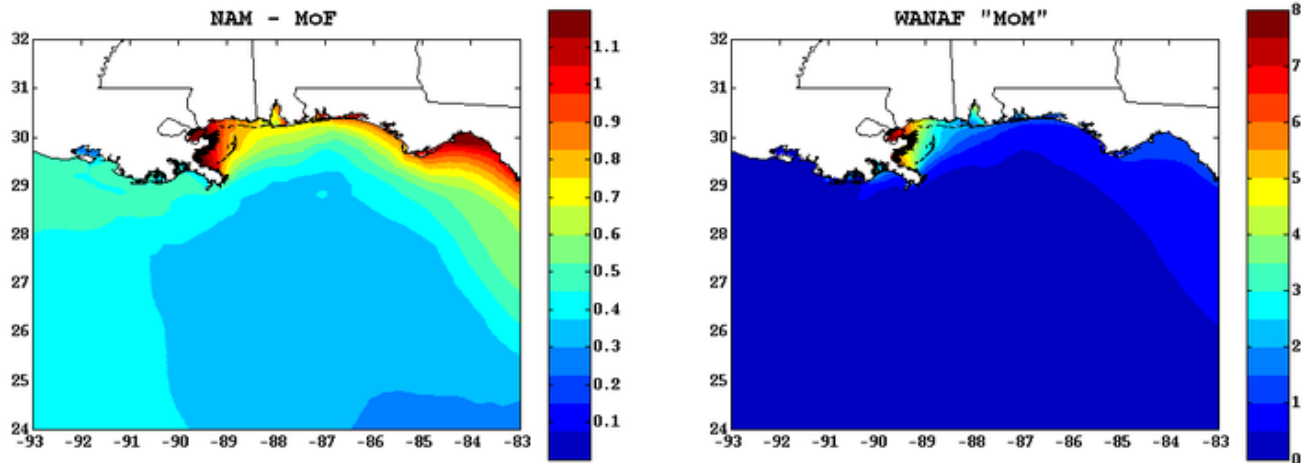


Ensemble Modeling



Results pushed to
RENCI VizWall

Hurricane Katrina



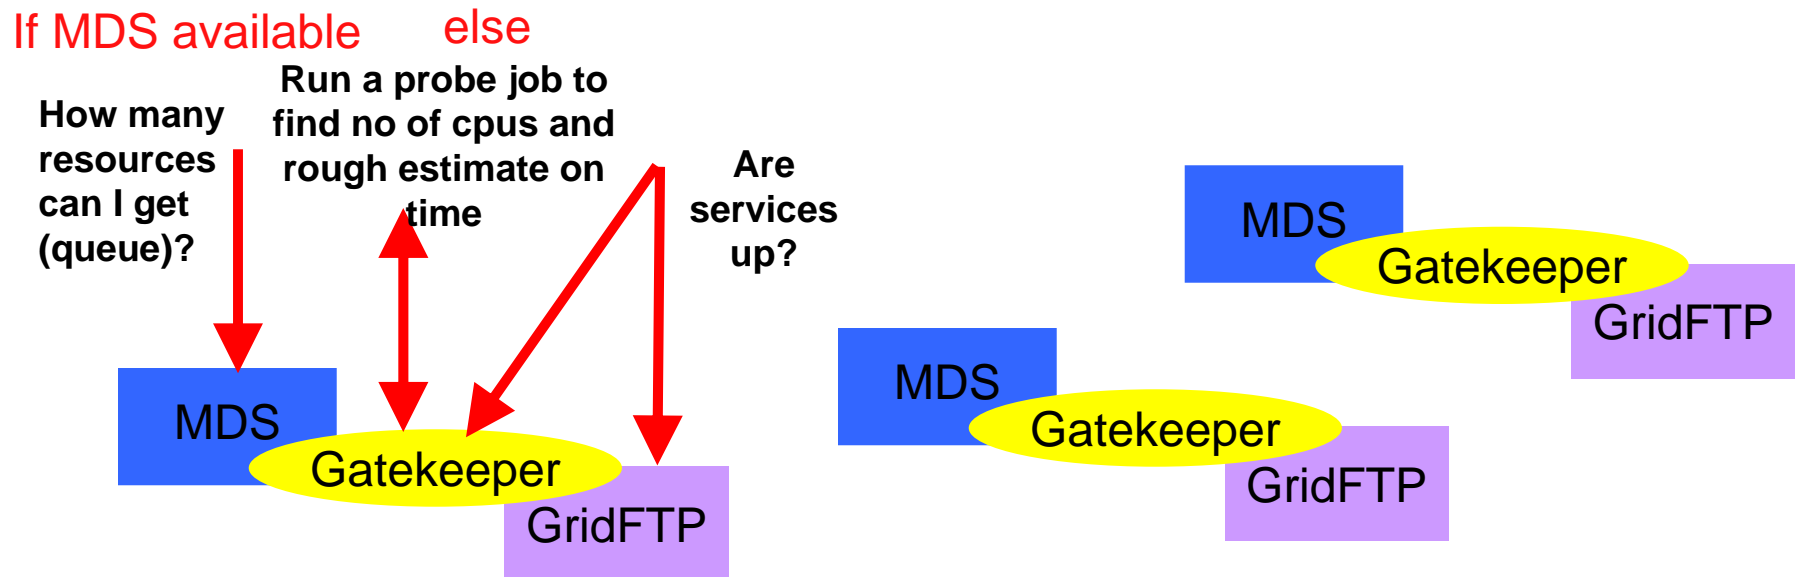
Left: ADCIRC max water level for 72 hr forecast starting 29 Aug 2005, driven by the "usual, always-available" ETA winds.

Right: ADCIRC max water level over ALL of UFL ensemble wind fields for 72 hr forecast starting 29 Aug 2005, driven by "UFL always-available" ETA winds.

Images credit: Brian O. Blanton, Dept of Marine Sciences, UNC Chapel Hill

Real-time Resource Selection

- Given a set of resources, which is the best resource I should run on?



Building the SCOOP Grid Testbed

- **Components at every site**
 - Globus gatekeeper, GridFTP, PBS/LSF (optional), MDS (optional)
- **Globus setup at compute sites**
 - hours of testing, firewall problems
 - MDS not setup
- **Resource availability and setup**
 - e.g: uudecode was not installed
- **Forecast mode**
 - debugging is hard
- **Criteria for resource selection**
 - performance, reliability



SCOOP Portal: Forecast Status

SCOOP Portal My Workspace: Previous ADCIRC Runs - Mozilla Firefox

http://dante1.renci.org:8080/scoop/portal/user/lavanya/js_pane/P-f0545bae05-1022s

Nov 20, 2005 11:31 pm

SCOOP Portal My Workspace

Scoop Status Portlet

Scoop Adcirc Run Status

Color Codes: [Preparing Run](#) [Running](#) [Completed Normally](#) [Failed](#) [Receiving LDM](#) [Sending To Resource](#) [Retrieving From Resource](#)

Data Arrival	Forecast Date	Cycle	Wind Field	Grid Field	Ensemble	Start Date	Finish Date	Resource	CPU Count
N.A.	20050808	00	eta	N.A.	N.A.	Mon Aug 08 03:19:50 EDT 2005	Mon Aug 08 03:27:07 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	32
N.A.	20050810	00	eta	N.A.	N.A.	Wed Aug 10 03:20:25 EDT 2005	Wed Aug 10 03:27:08 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	38
N.A.	20050811	00	eta	N.A.	N.A.	Thu Aug 11 05:25:19 EDT 2005	Thu Aug 11 05:30:24 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	26
N.A.	20050812	00	eta	N.A.	N.A.	Fri Aug 12 03:20:32 EDT 2005	Fri Aug 12 03:25:44 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	20
N.A.	20050814	00	eta	N.A.	N.A.	Sun Aug 14 03:20:13 EDT 2005	Sun Aug 14 03:25:43 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	44
N.A.	20050815	00	eta	N.A.	N.A.	Mon Aug 15 08:20:19 EDT 2005	Mon Aug 15 08:26:23 EDT 2005	dante0.renci.org/jobmanager-pbs	18
N.A.	20050818	00	eta	N.A.	N.A.	Thu Aug 18 03:20:18 EDT 2005	Thu Aug 18 09:59:26 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	36
N.A.	20050819	00	eta	N.A.	N.A.	Fri Aug 19 03:20:55 EDT 2005	Fri Aug 19 09:59:59 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	20
N.A.	20050820	00	eta	N.A.	N.A.	Sat Aug 20 03:20:32 EDT 2005	Sat Aug 20 03:26:22 EDT 2005	scoop.ncgrid.org/jobmanager-lsf	48

Done

start VGrADS: Extending G... SCOOP Portal My Wo... Microsoft Office P... lavanya - Microsoft O...

11:31 PM

SCOOP Portal My Workspace: OPeNDAP Server - Mozilla Firefox

File Edit View Go Bookmarks Yahoo! Tools Help

http://dante1.renci.org:8080/scoop/portal/user/lavanya/js_pane/P-f65d80596c-1000z

Customize Links Free Hotmail Windows Marketplace Windows Media Windows RealPlayer - The bes...

SCOOP Portal Nov 20, 2005 11:28 pm

My Workspace

OPeNDAP server

DODS Index of /data/nc-coos/model_data

<u>Name</u>	<u>Last modified</u>	<u>Size</u>	<u>Description</u>
[DIR] Parent Directory	17-Aug-2005 16:...		
[DIR] adcirc/	20-Nov-2005 11:...		
[DIR] met/	01-M-2005 09:...		
[DIR] quoddy/	23-F-2005 01:...		

Apache/1.3.28 Server at nc-coos.org Port 80

Logout

Customize

Users Present

Lavanya Ramakrishi

Done

start VGrADS: Extending G... SCOOP Portal My Wo... Microsoft Office P... lavanya - Microsoft O...

11:28 PM

OpenDAP Access

Outline

- ✓ **Today's Grid Deployments**
 - Bioportal
 - Ocean Modeling
- **Next Generation Grids**
 - Service Oriented Architecture
 - e.g. LEAD
 - Grid Application Development
 - Virtual Grid Application Development Software
 - Cluster on Demand

Next Generation Problems

- **Service Oriented Architecture**
 - modular, composable, event-driven
- **Grid programming**
 - Specifying resource requirements
 - most scientists want to specify higher level requirements
 - Performability of grid environments
- **Managing application environments**
 - specialized set of requirements for each application
 - resource provider cost and accountability


Linked Environments for Atmospheric Discovery

- **Rationale**

- each year, mesoscale weather – floods, tornadoes, hail, strong winds, lightning, hurricanes and winter storms – causes hundreds of deaths, routinely disrupts transportation and commerce, and results in annual economic losses in excess of \$13B.

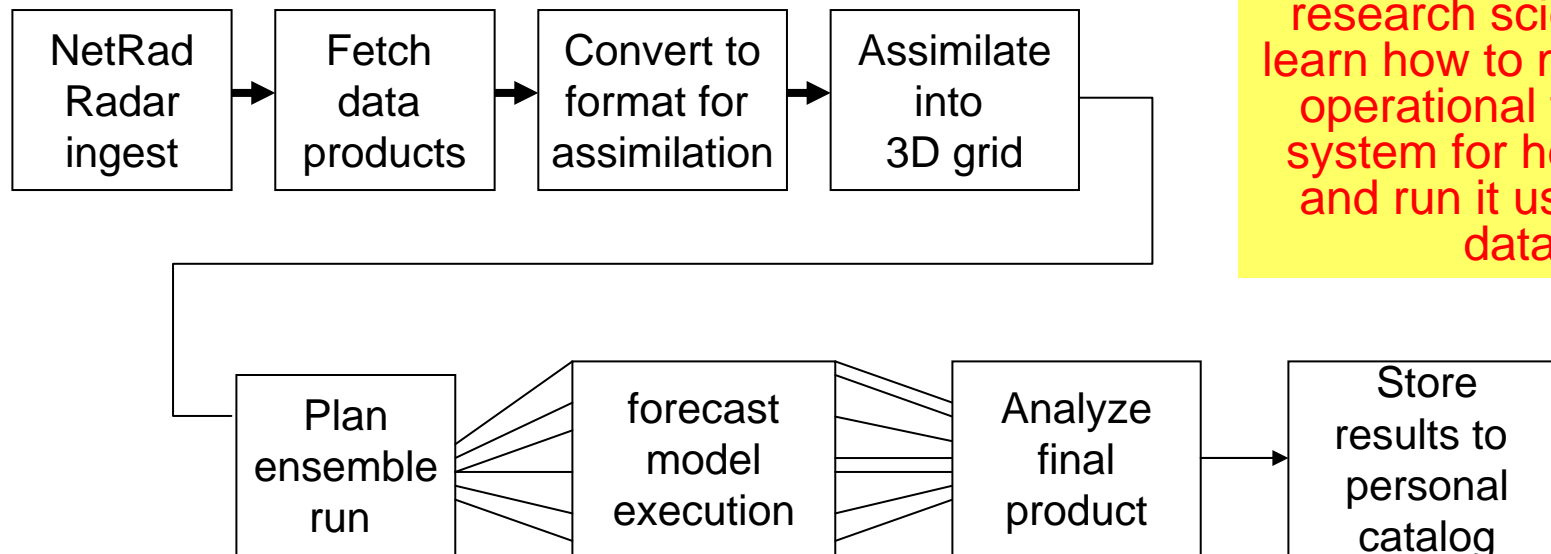
- **From “offline” to “online” forecasting**

- data assimilation and adaptive evaluation



Static Forecasting

- **Large-scale tightly coupled components**
 - not adaptable to weather or resour
 - modifications are not trivial



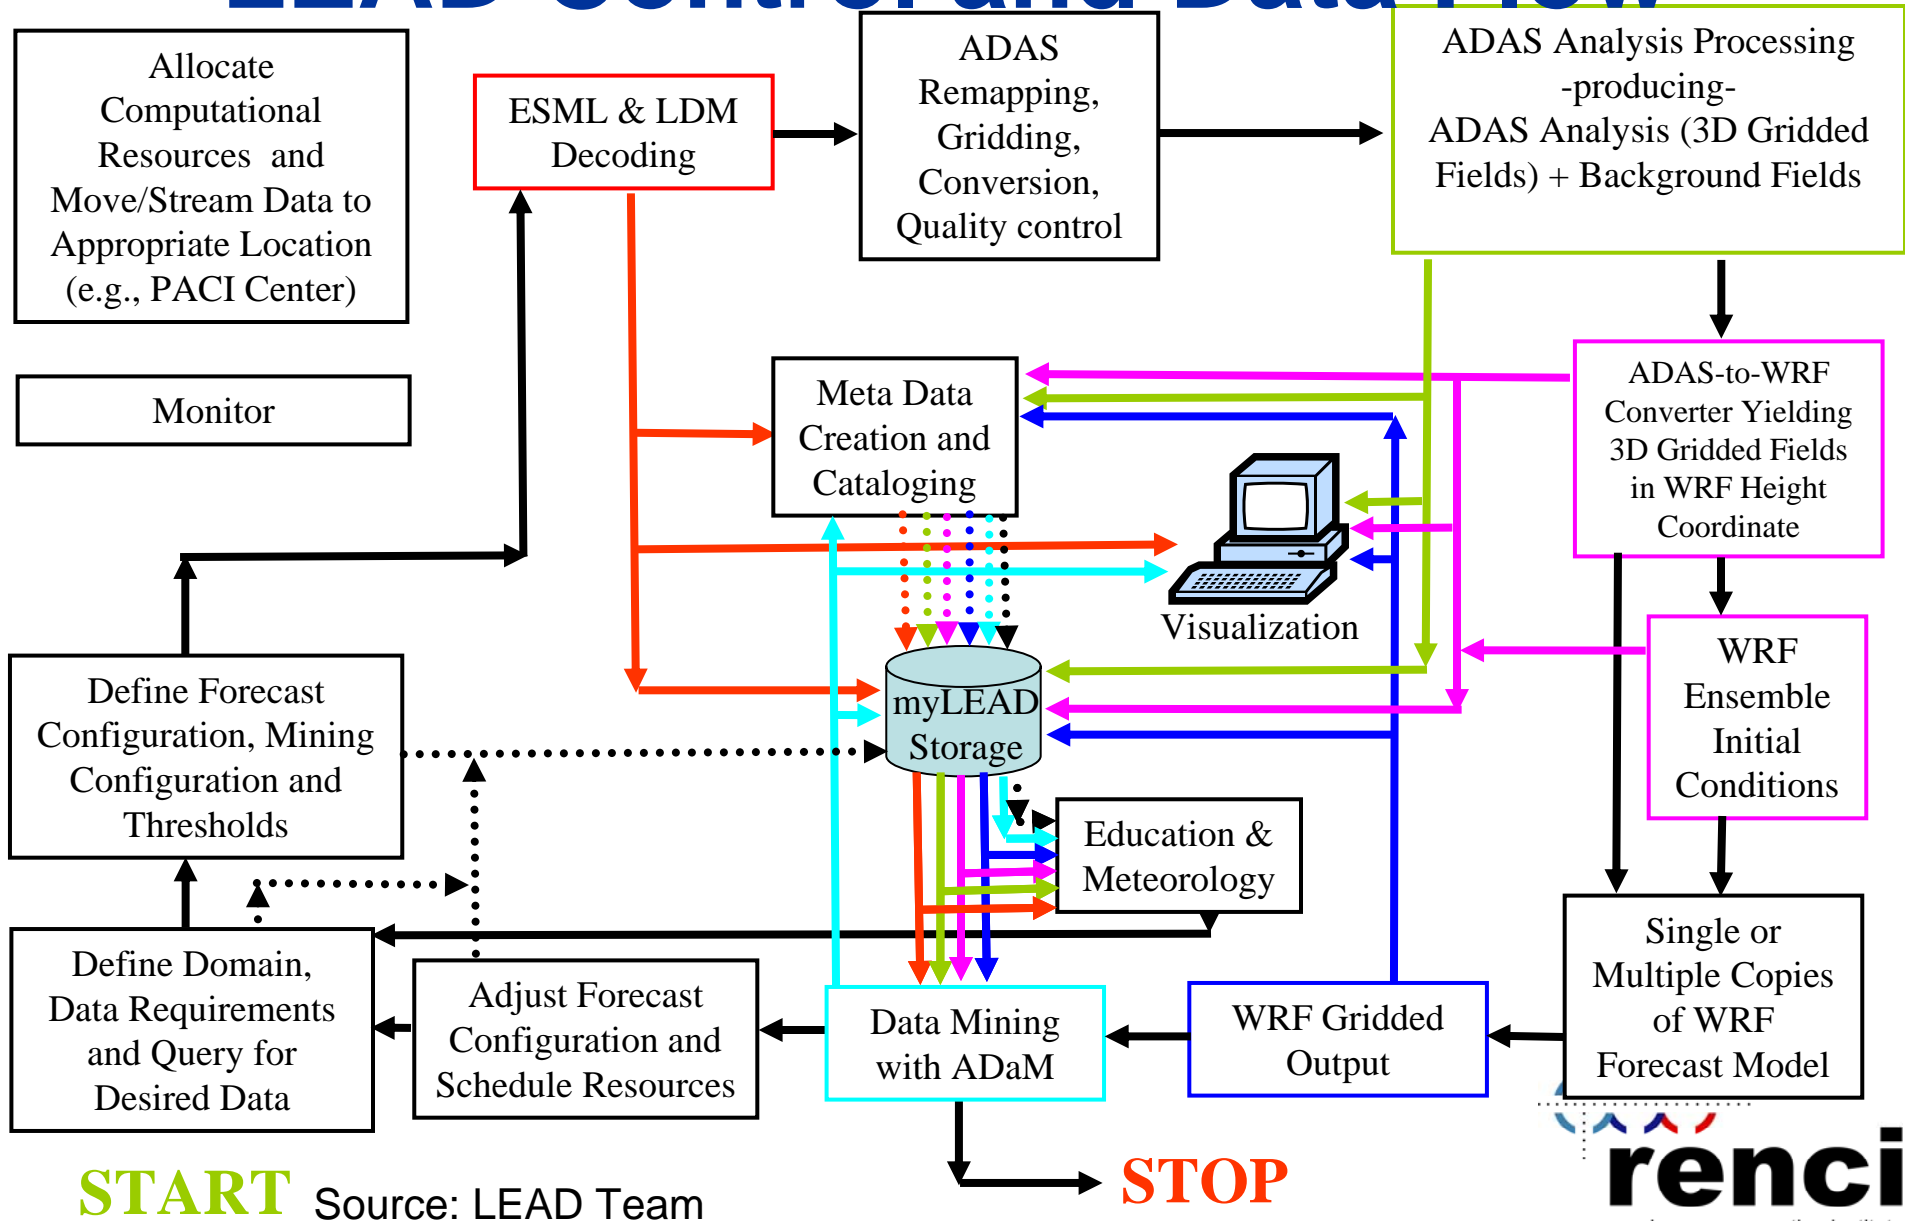
Center for Analysis and Prediction of Storms (CAPS, OU):
It took an MS student 6 months, working with a research scientist, to learn how to modify the operational forecast system for her needs and run it using real data

Dynamic Adaptive LEAD System

- **Meteorology goal**
 - to provide timely and accurate forecasts using dynamic adaptation
- **Computer Science goal**
 - map application requirements to resource capabilities
 - redundant runs, scheduling policies
 - adapt to weather as well as resource behavior

Need real time monitoring to make adaptation decisions

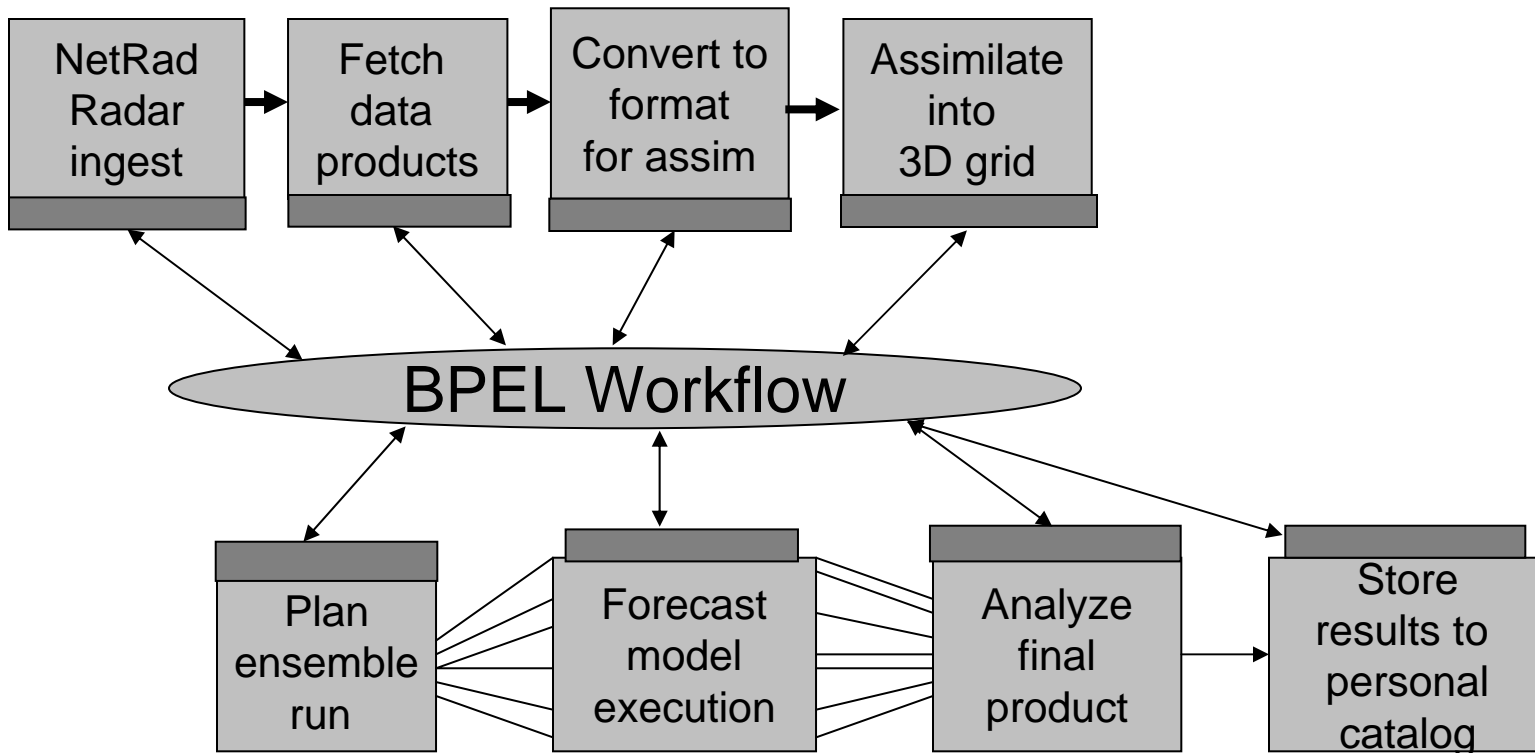
LEAD Control and Data Flow



START Source: LEAD Team

STOP


Service Oriented LEAD Workflows



Source: Beth Plale

Characteristics of LEAD Workflows


- **Coupled analysis and assimilation tools, data repositories**
 - change configuration rapidly and automatically in response to weather
 - Streaming data, steer remote observing technologies
- **Multilevel monitoring and intelligent control**
 - workflow, resource, application, service
 - performance and reliability guarantees of the resources



Adapting Workflow Structure

- **Reactive Adaptation**


- e.g. service failure, resource behavior



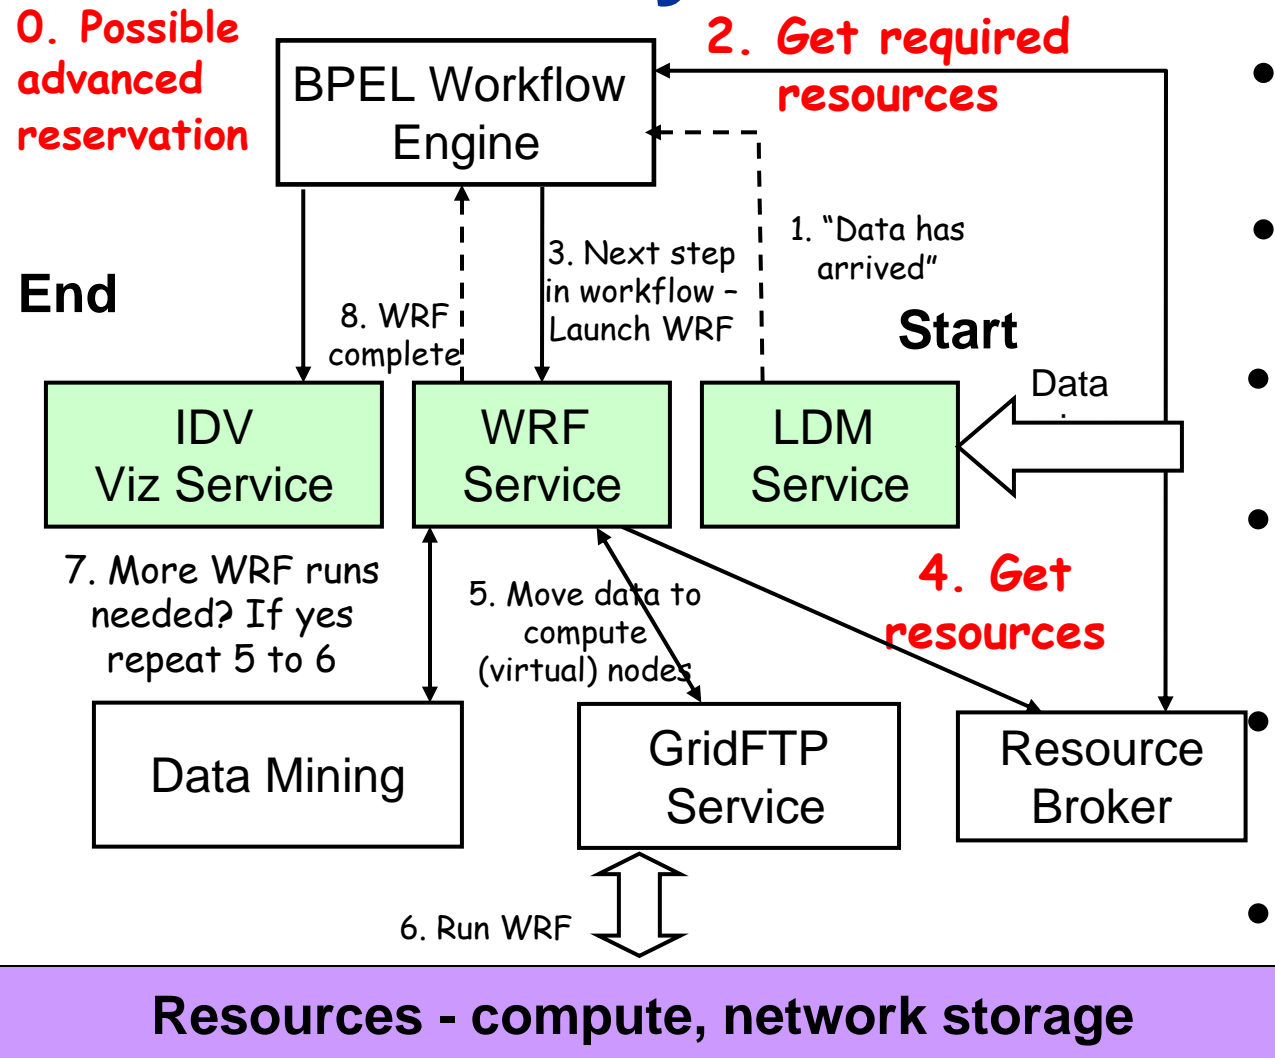
- **Proactive prediction and decisions**

- adjust i, j, k (weather science meets the infrastructure science) to meet individual workflow guarantees

- global optimization of number of workflows being serviced



LEAD Dynamic Workflows




- **Multilevel monitoring**
- **Multiple decision points**
- **Streaming data**
 - Unidata LDM
- **Service monitoring**
 - web service load
- **Application**
 - behavior on resources
- **Resource**
 - performance
 - reliability

Qualitative Resource Selection

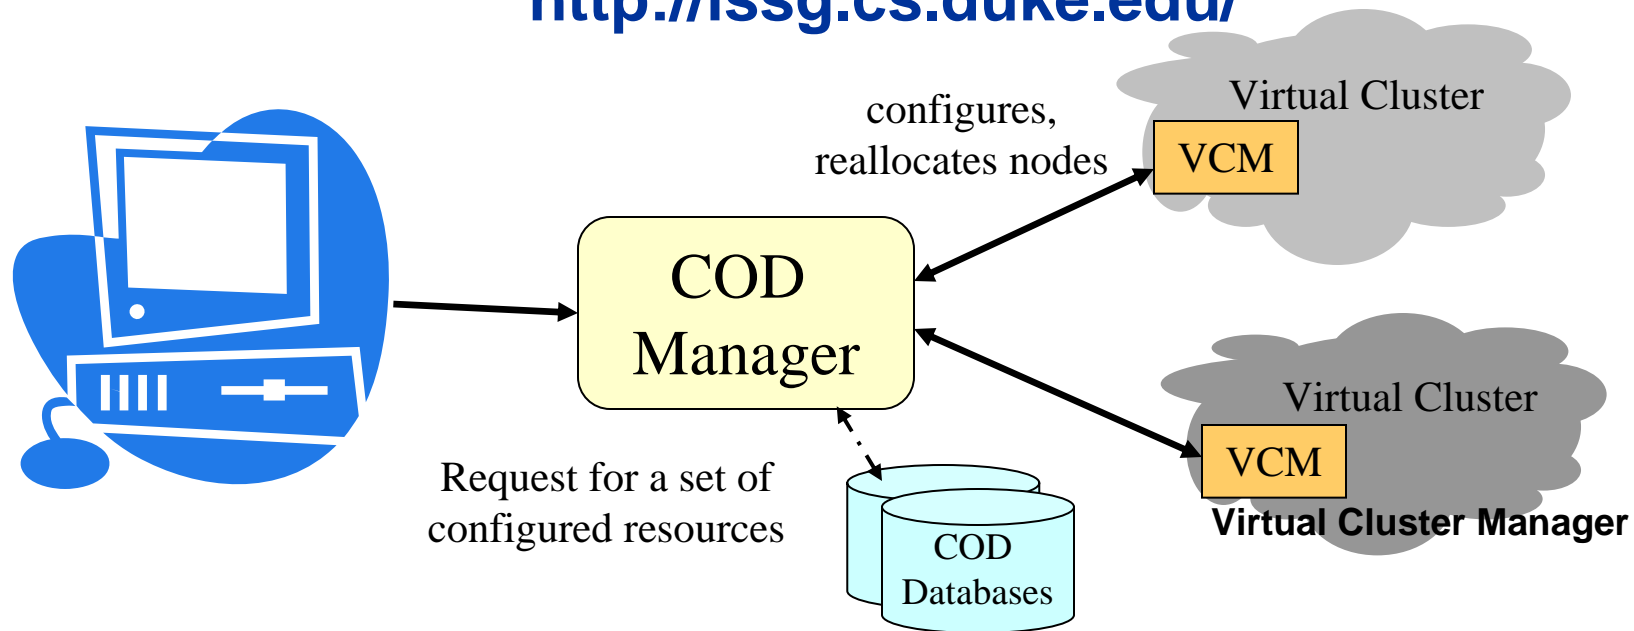
<http://vgrads.rice.edu>

- **Dynamic, scalable resource abstraction framework**
 - scheduler, resource broker
- **Integrated monitoring and notification of resource behavior**
 - NWS, NWS-HAPI



Managing Application Environments

<http://issg.cs.duke.edu/>



- **Dynamic configuration of software environments**
- **Automation of cluster management**
- **Dynamic re-allocation of resources**



Conclusions

- **Grids today**
 - academic and industry balances
 - technology upgrades and how to handle them?
 - software integrations problems
 - common infrastructure, customized interfaces
 - need to work closely with application domain
- **Grids tomorrow**
 - instruments, streaming data, persistent and transient services
 - co-scheduling of data and compute
 - performance and reliability guarantees
 - resource provider incentives



Questions?