Building Gateways to Grid Capabilities

Dennis Gannon
(with collaborator Beth Plale)
Department of Computer Science
School of Informatics
Indiana University

Theme of This Presentation

- Our Grid middleware has matured incredibly well over the last 10 years.
  - Globus, g-Lite, OMII all provide an impressive array of basic service.
  - Large Grid deployments number in the dozens (if not hundreds).
- However, the middleware is not the whole story.
- We can now provide “Gateway” portals that can allow users access to Grid capabilities without exposing them to Grid technology.
- The key components:
  - Data virtualization
  - Computational virtualization and workflow tools.

Outline

- Some Early Grid Projects & Lessons
  - The I-Way & the NASA IPG
  - The evolution of the grid portal Science Gateway.
  - The emergence of service oriented architectures.
- LEAD: A Grid data-driven Gateway SOA
  - Changing the Paradigm for Meteorology
- Some Lessons Learned
- Some Challenges for the Future
  - How can we make systems more adaptive in response to data challenges?
  - How do we create real knowledge discovery tools that combine data query and workflow?

The First Stages of Transformation

- From Clusters to Distributed Systems
- The SC97 I-Way
  - An experiment in wide-area HPC resource sharing.
  - A dozen or so projects doing applications distributed across the network.
- Colliding Galaxies
  - Couple together two large simulations running on two supercomputers to build a single application spanning a wide area network.

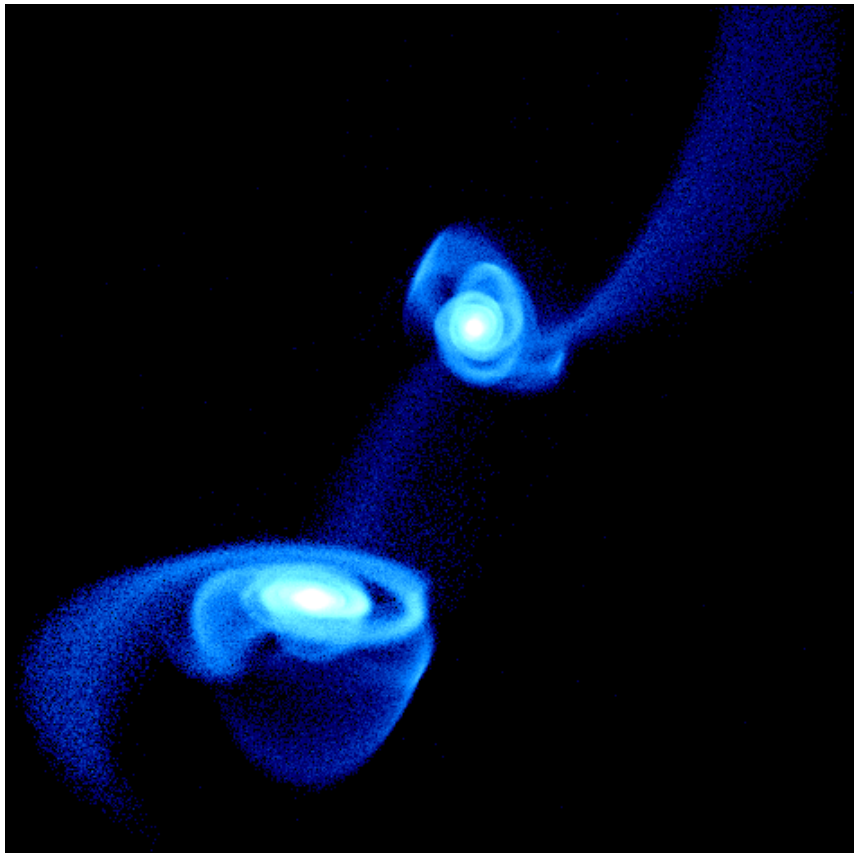

Lesson from the I-Way

- Tightly coupled apps are paced by your lowest bandwidth link.
- Demos are easy. Persistent frameworks are hard to support.
  - Grids need to be persistent.
- Tools to “compose” Grid applications from remote components are essential
  - But don’t expect that scientists will change production codes!
- Debugging Distributed Applications is very hard.

NASA's IPG - 1998

- The goal:
  - To unite all the major NASA computing platforms (Langley, Ames, Glen, Marshal …) into a "Grid".
  - The *Information Power Grid*.
    - The first major grid deployment.
- The plan:
  - Use some new software: Legion or Globus.
    - Flipped a coin. Globus was winner.
- The architecture:
  - A layer of services on top of core gridware.

IPG Lessons

- Respect Local Culture. It is Important
  - Managers of computing centers do not want to cede local authority to “the Grid”.
- Grid security MUST be easy for people to use.
  - Solutions that require certificate management by users are not going to work.
- Grids work only when there is a purpose to building them.
  - The science or business application must always be the driver.

IPG Innovations

Grid Forum

- Became the Global Grid Forum

Grid Portals – Gateways to access Grids

– (along with ncsa and sdsc!)

– To give a community of researchers easy access to the tools, data and computational power needed to solve today's scientific and engineering problems.

– To do this in a discipline specific language that is common to the target community.

– To hide any underlying Grid technology.

Let’s look at a few real example Grid Science Gateways

(about a dozen … many more exist!)

NEESGrid

Realtime access to earthquake
Shake table experiments at remote sites.

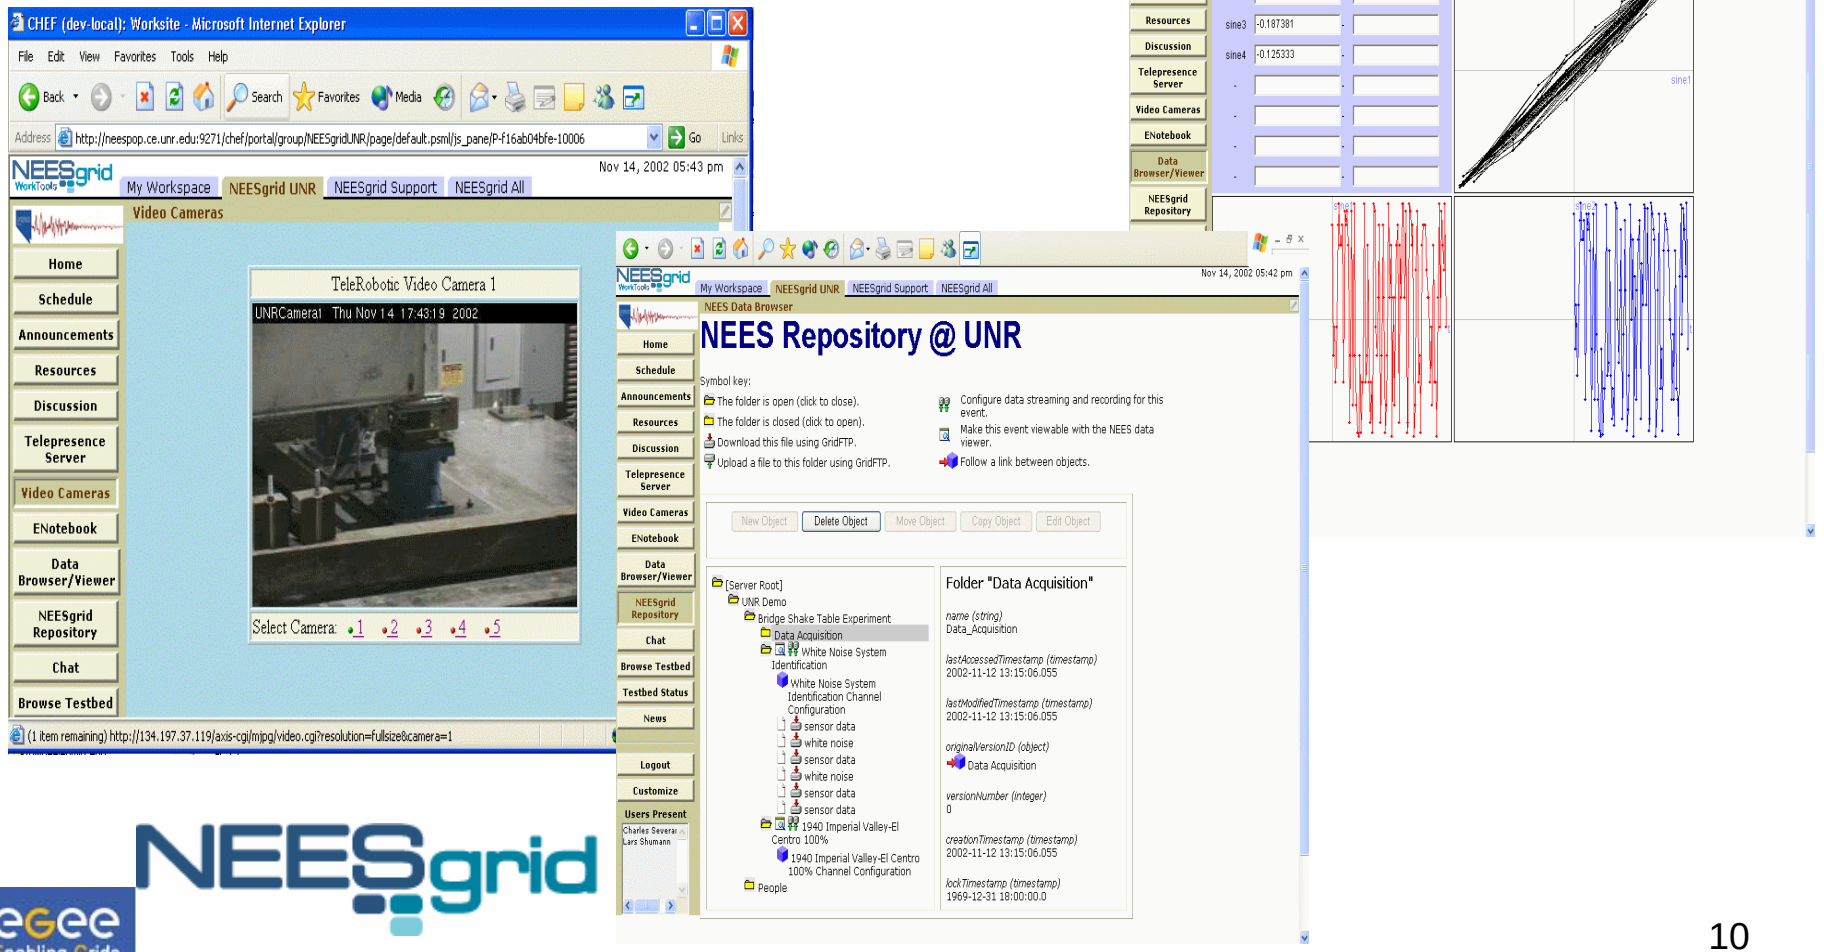

BIRN – Biomedical Information

Geological Information Grid Portal

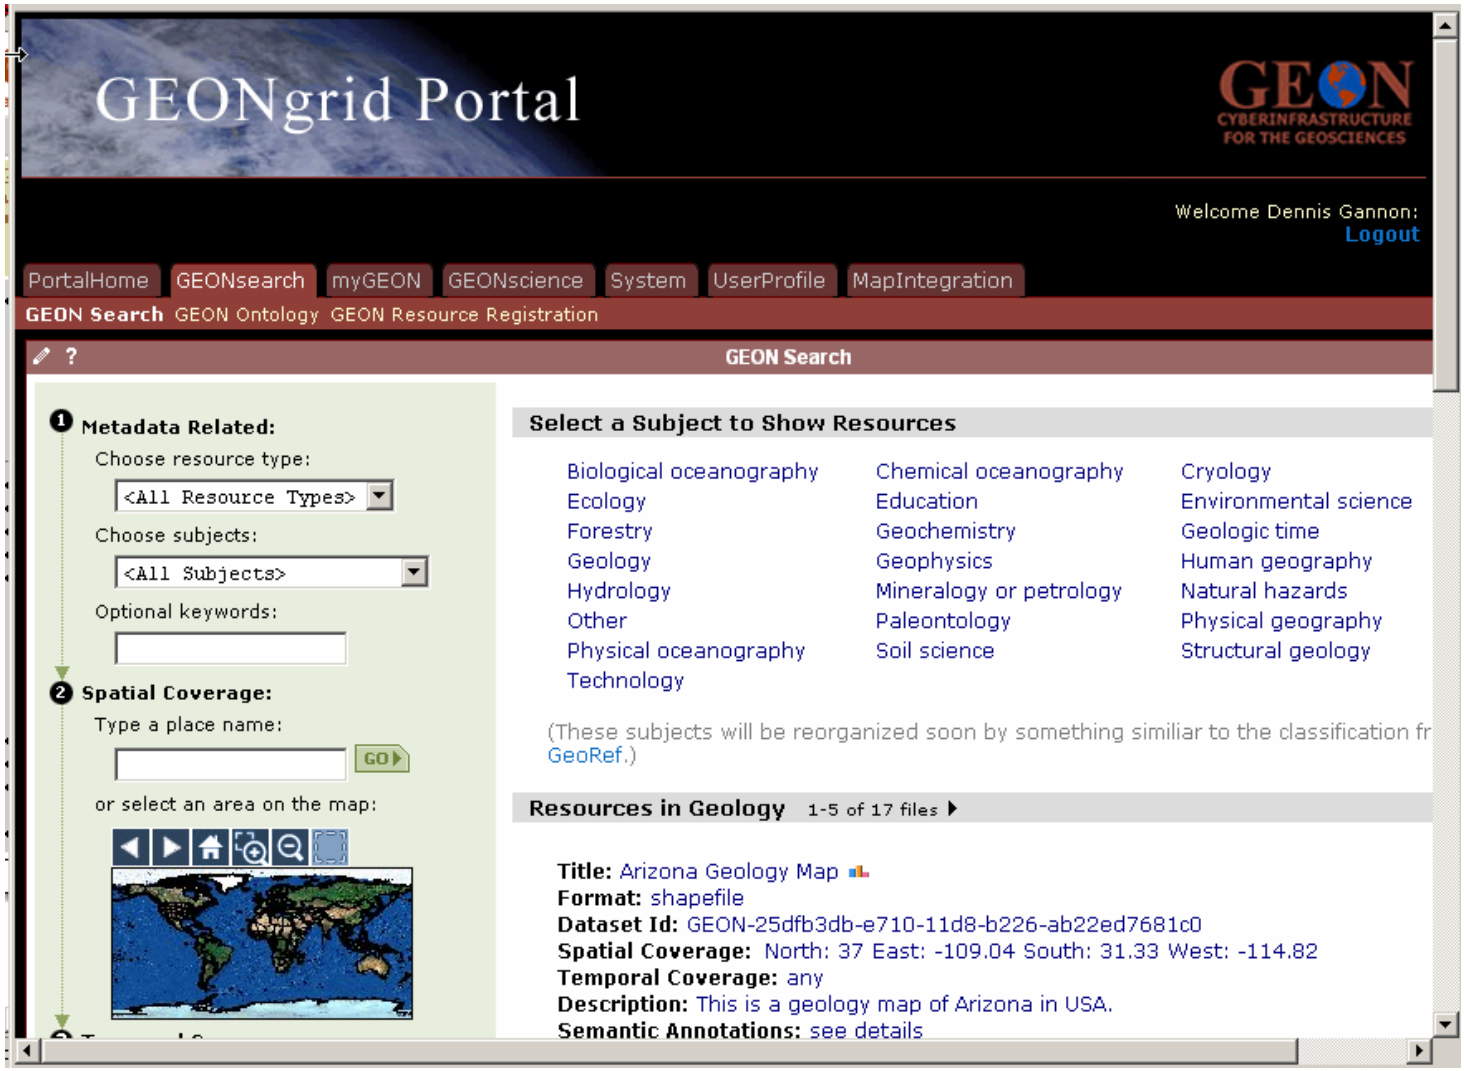

Mesoscale Meteorology

NSF LEAD project - making the tools that are needed to make accurate predictions of tornados and hurricanes.

- Data exploration and Grid workflow

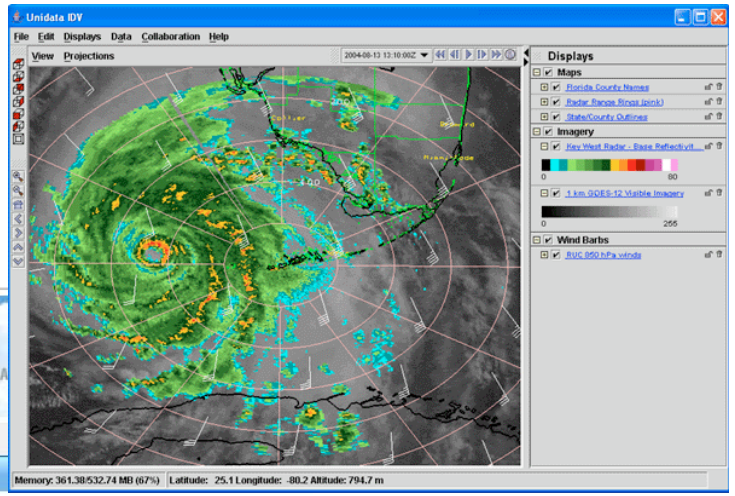

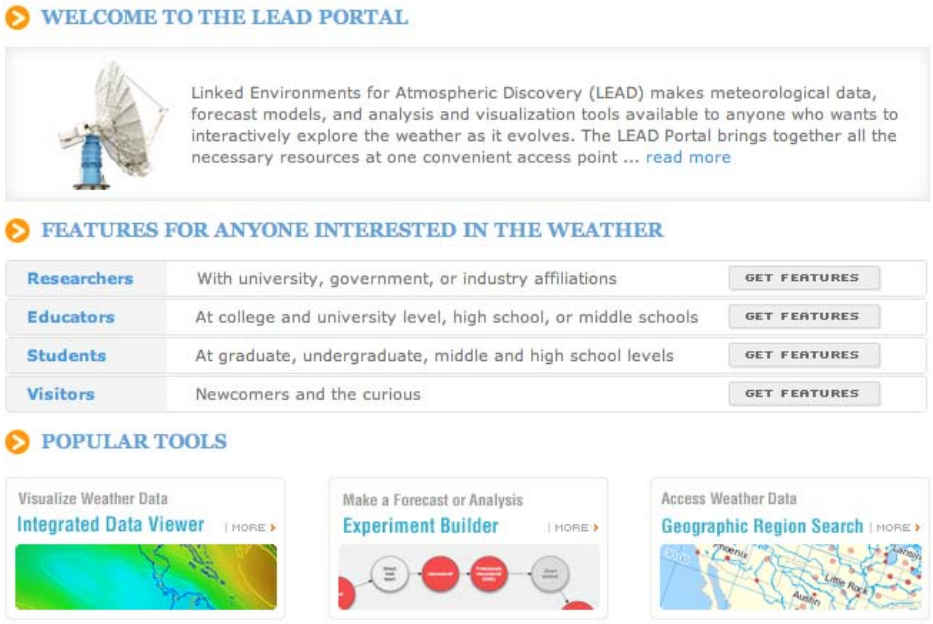

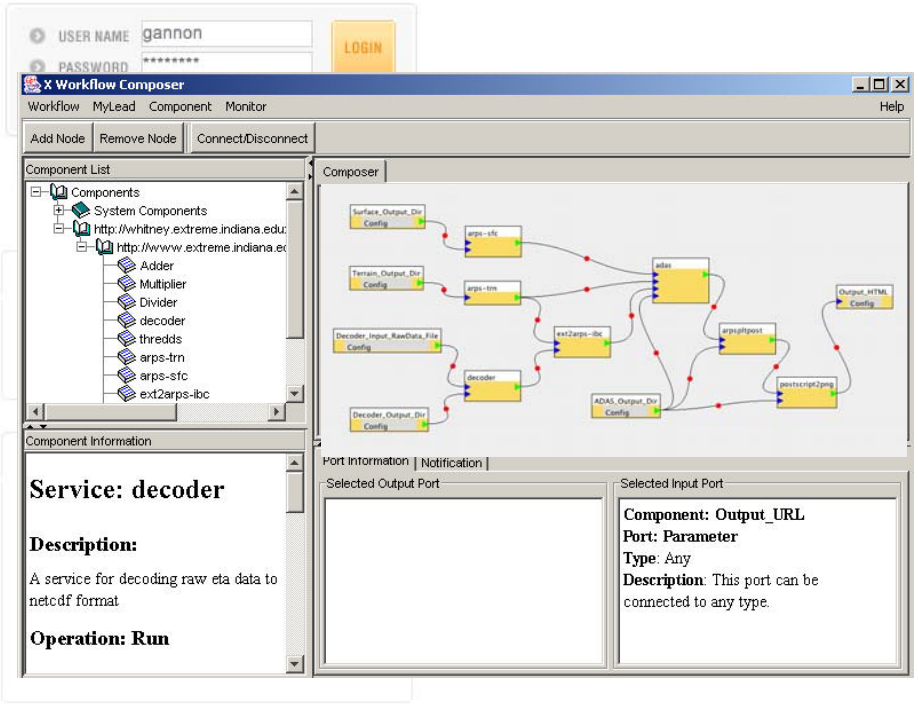

Renci Bio Portal

Providing access to biotechnology tools running on a back-end Grid.

- leverage state-wide investment in bioinformatics
- undergraduate & graduate education, faculty research
- another portal soon: national evolutionary synthesis center

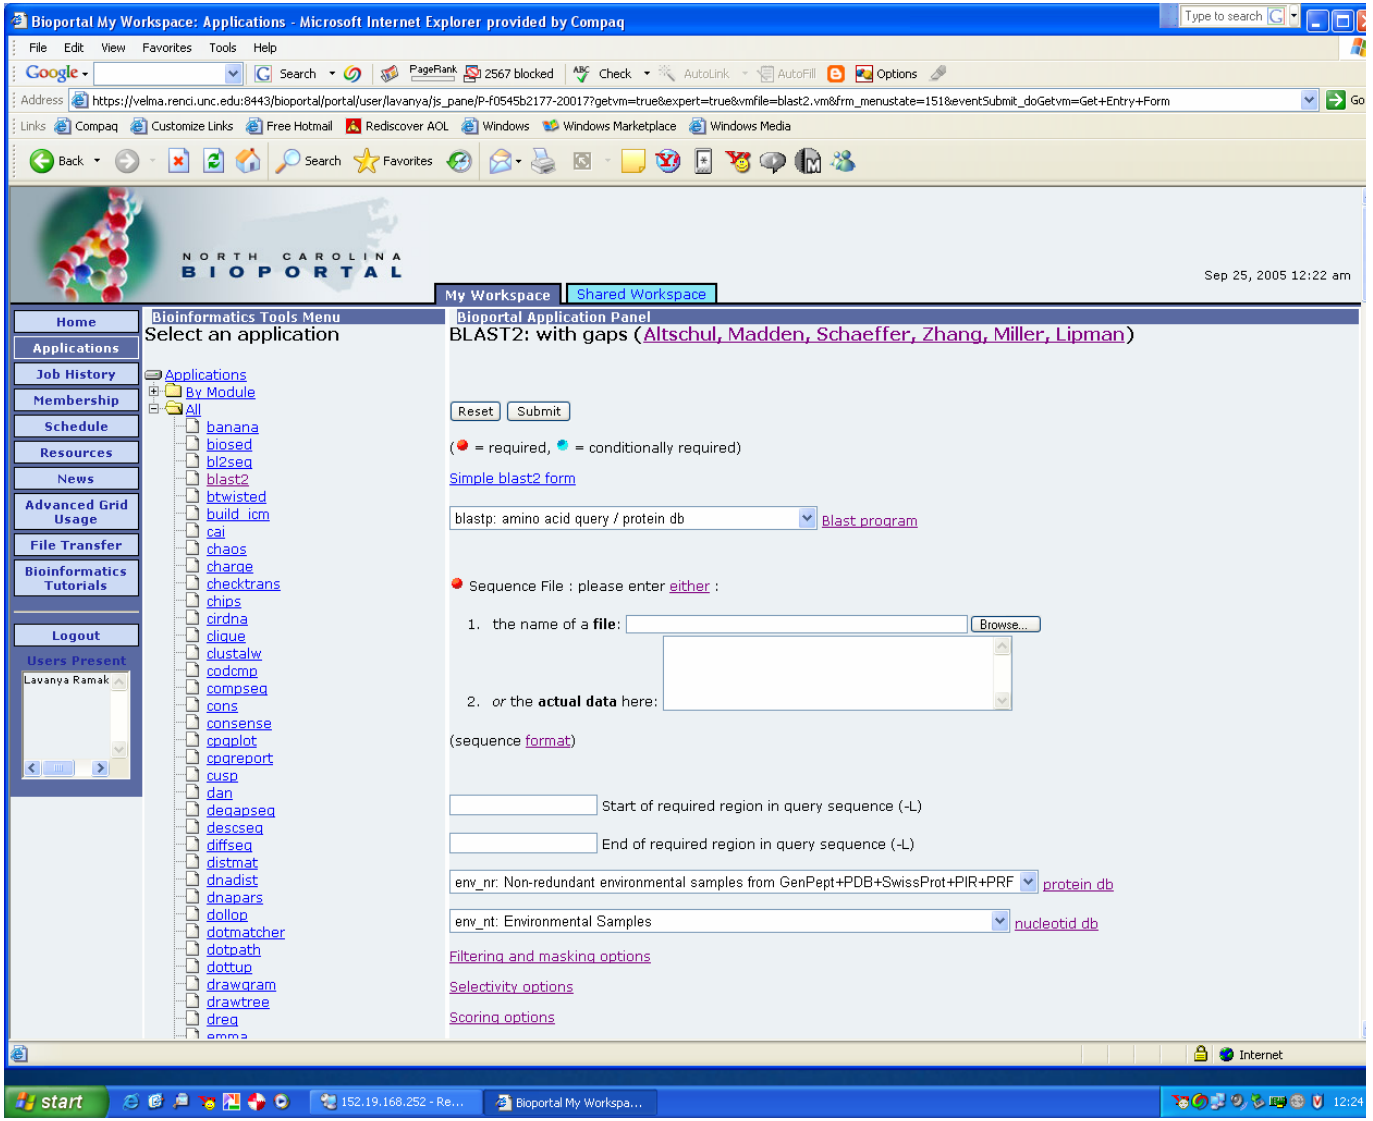

Nanohub - nanotechnology

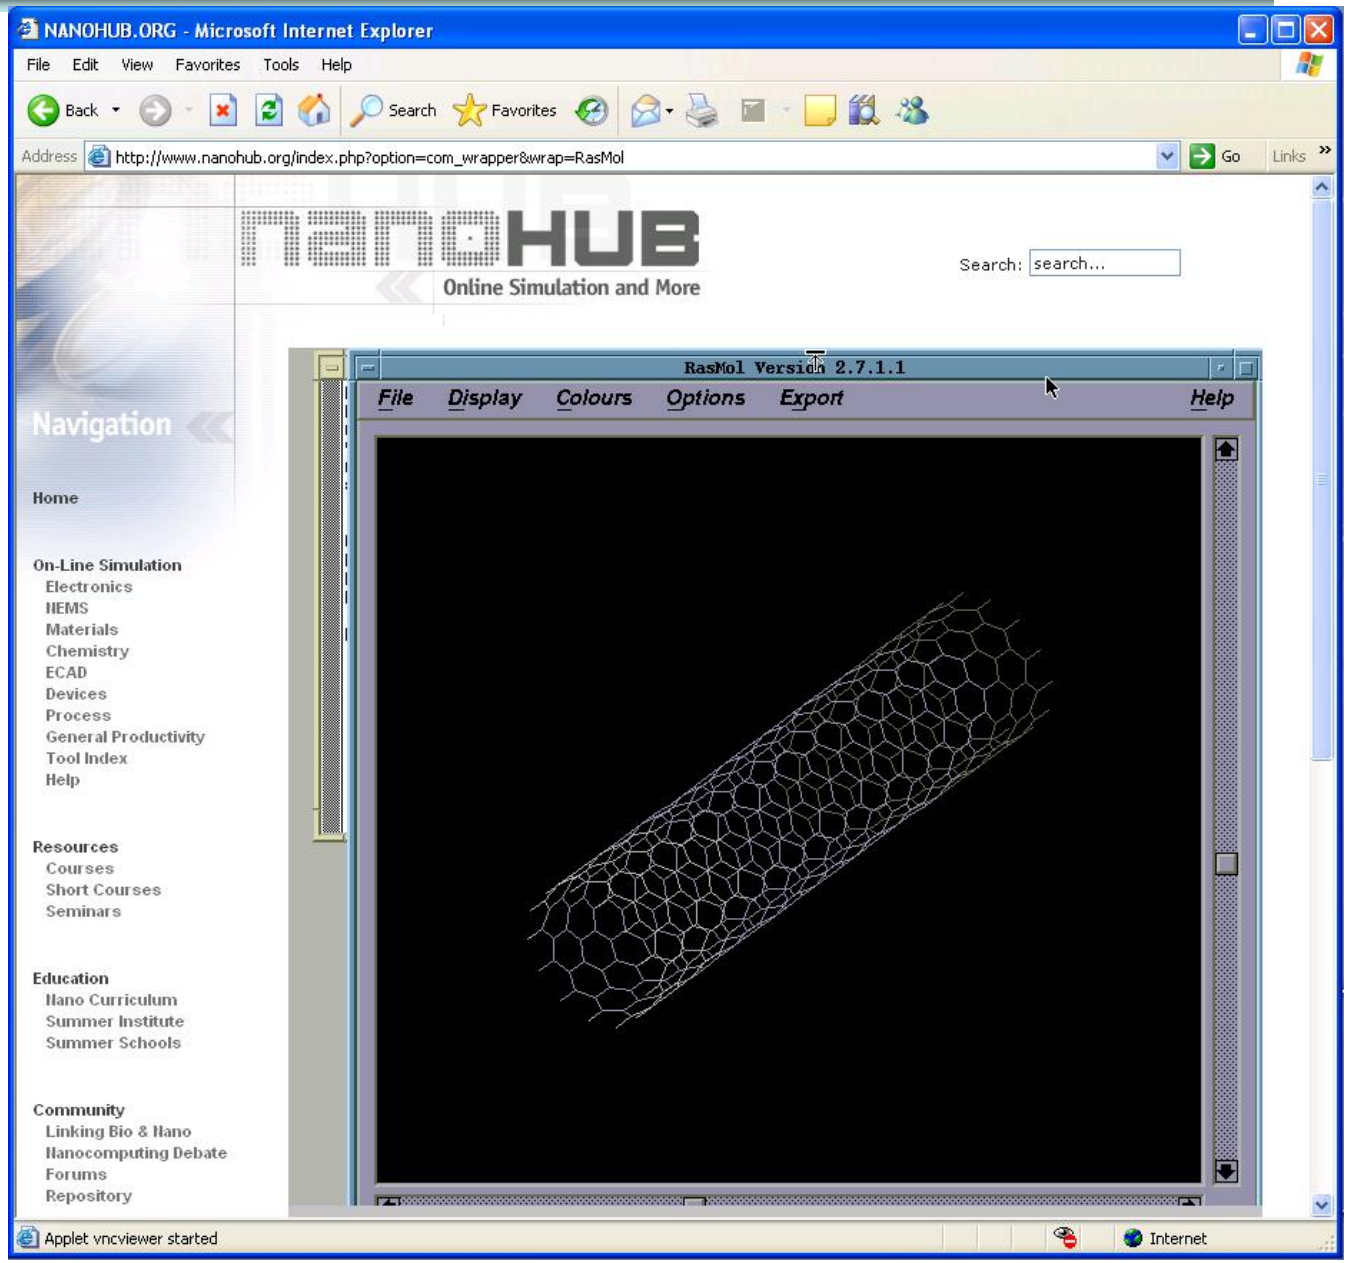

X-Ray Crystallography

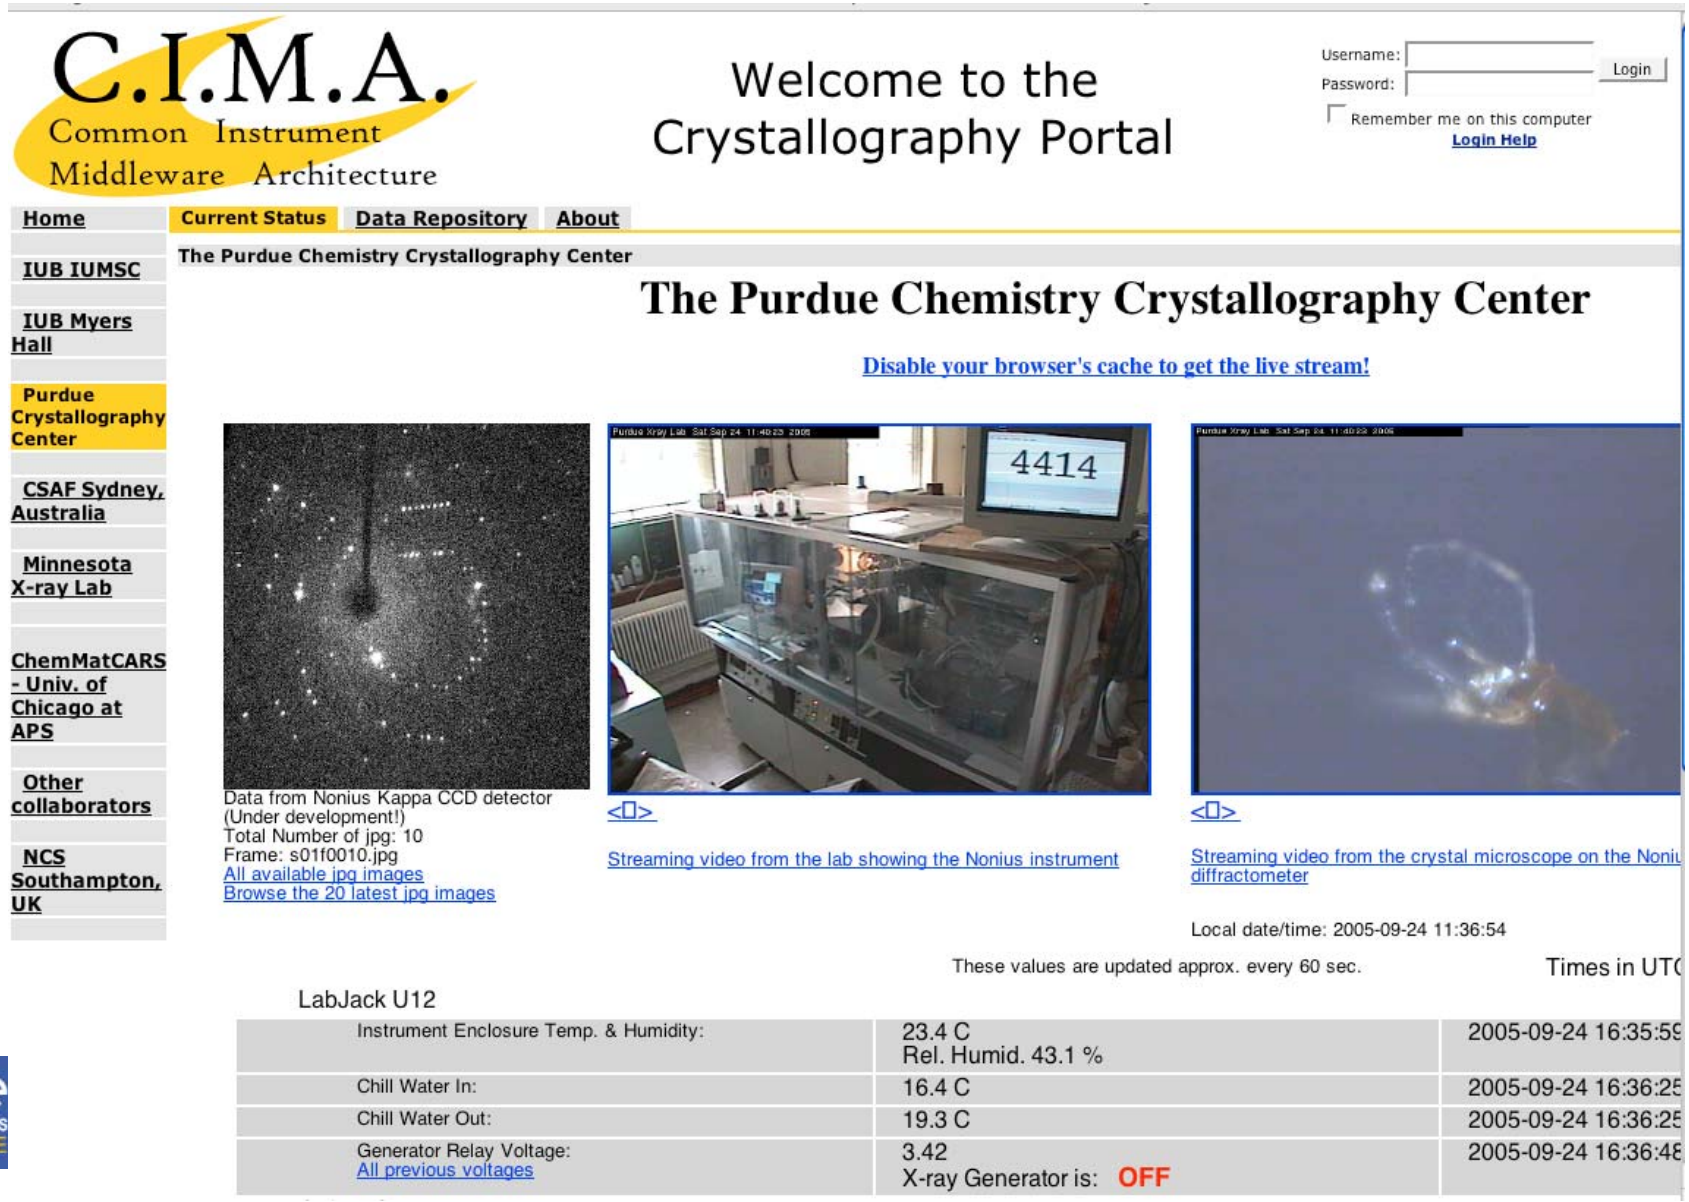

ServoGrid Portal

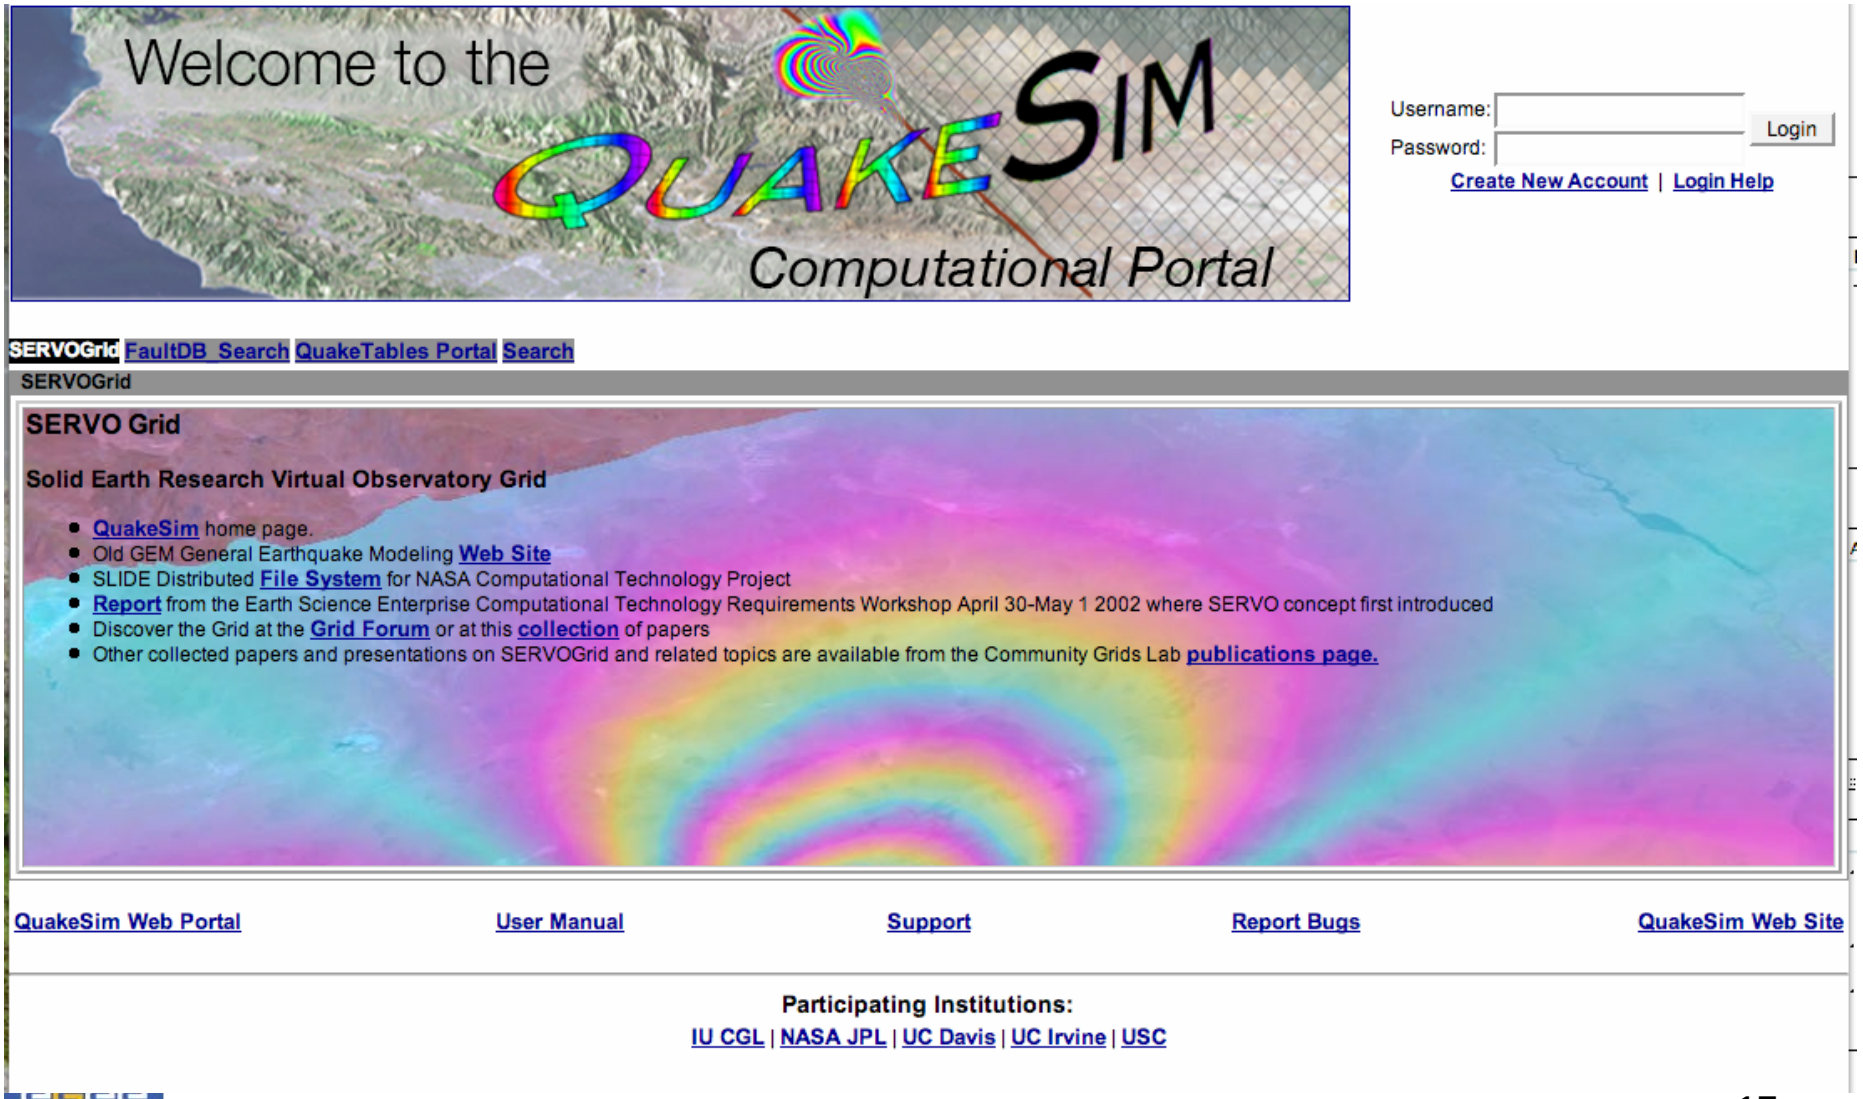

Belfast Gene Grid Portal

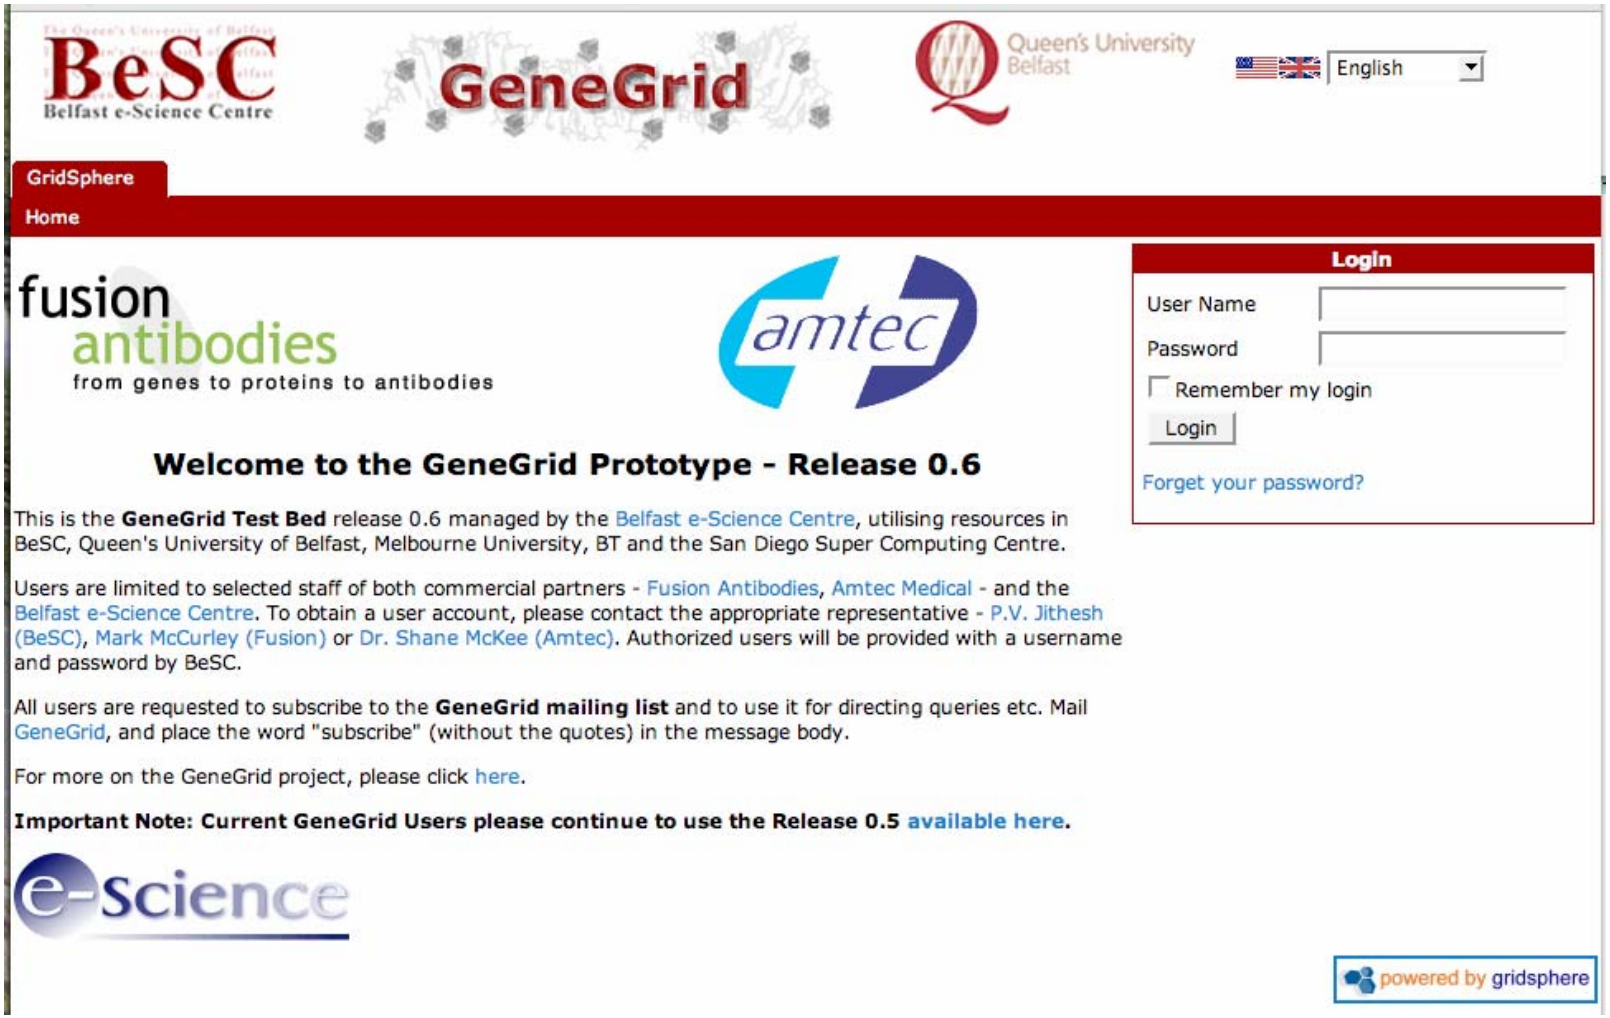

MyGrid - Bioinformatics

myGrid

Navigate

- Home
- About
- Downloads
- Components
  - Component Overview
  - Research Components
- Using myGrid
- Research Using myGrid
- Links
- Publications
- Contact

Log In

Username

Password

Log In

New Account Signup
Forgot Password

myGrid Architecture

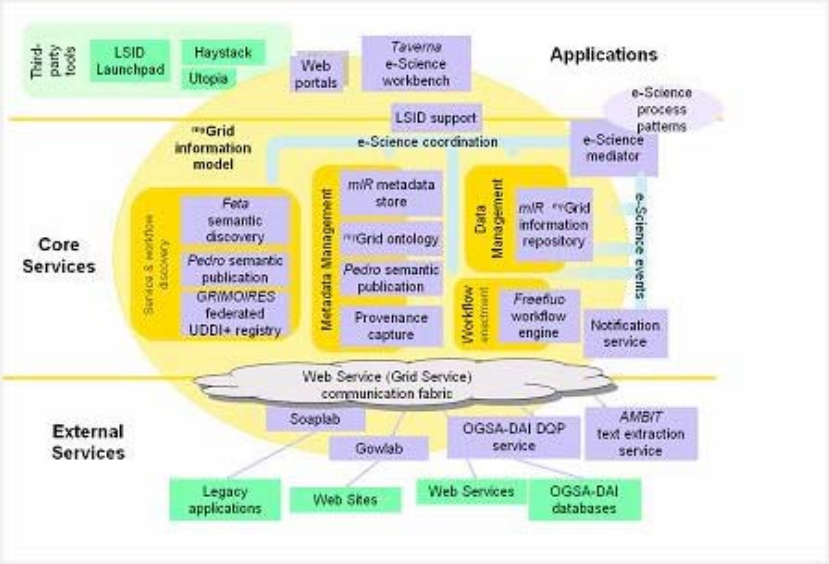

myGrid components - overview

myGrid is a collection of services and components that allows the high level integration of biological applications. The architecture provides the infrastructure necessary, in a web service environment, e-science workbench that actively supports the scientific lifecycle. Each component or service contri system that allows the e-scientist to perform complex in-silico experiments across distributed bioinfo resources.

Lessons from Grid Portals?

- Access to data is the most important feature
  - Allow the user community access to important shared data resources
    - Visualize it, publish it, download it, curate it.
  - Data Discovery
    - Searchable metadata directories.
- Web access to important tools
  - Web-form interfaces to allow users to run important community codes
  - Webstart access to common java-based tools
  - Limited shell access - perhaps to a VM
- Workflow tools
  - Allow users to combine community codes into workflows managed by the portal back-end resources.
- Service Oriented Architecture (SOA)
  - Promotes a pluggable framework to add new features and to virtualize access to resources

You need a Grid Foundation

- A Distributed Infrastructure of services that provides
  - Security
    - Authentication: how is a member of the VO recognized?
    - Authorization: how do you know who is authorized to do what?
    - Privacy: protecting the work of individuals and the VO as a whole.
  - Data Virtualization
    - Services to search for data, catalog and curate it, manage duplicates, publish it
      - All without reference to underlying file systems.
  - Computation Virtualization
    - Users need access to applications and not raw computers. Services must be provided to schedule and monitor and manage compute resources,
      - All the user sees is the ability to interact with the application
- Foundation: Service Oriented Architectures

Data is the core of e-Science

- The amount of data generated by research and instruments dwarfs the stuff you can google.
  - Sometimes known as the deep web.
- How do we discover and catalog binary data?
  - We build searchable metadata catalogs
  - We build methods to query for data that may not even exist yet

From www.icioffshore.com/kyc/page2d.asp

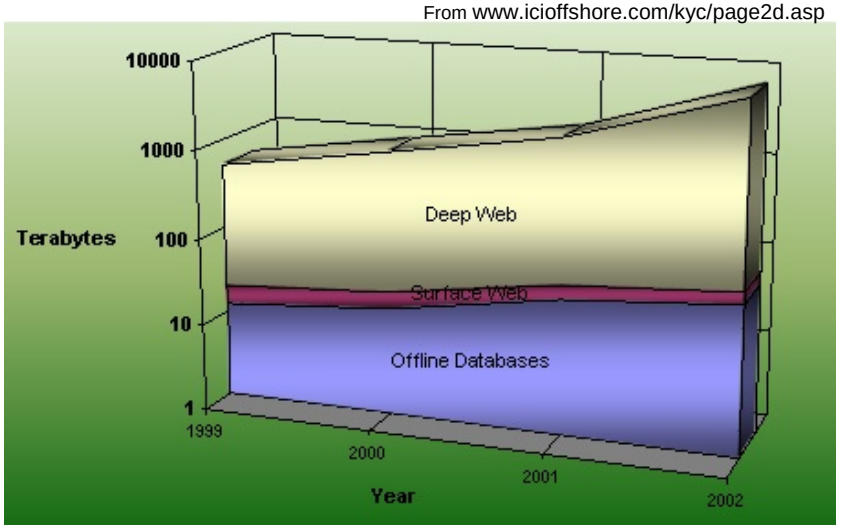

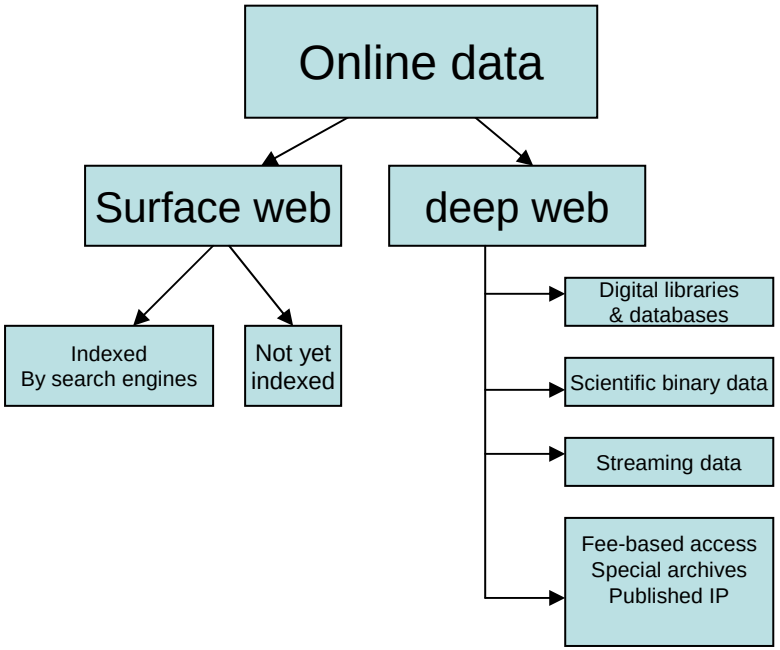

Data Services

- Name resolution - mapping an abstract unique object identifier to a file-object name.
- Replication location - find or generate replicas of a given service
- Resource catalogs - a searchable index of community data products and services. Generated by crawling and direct registration
- Metadata directory - a searchable database of data object metadata

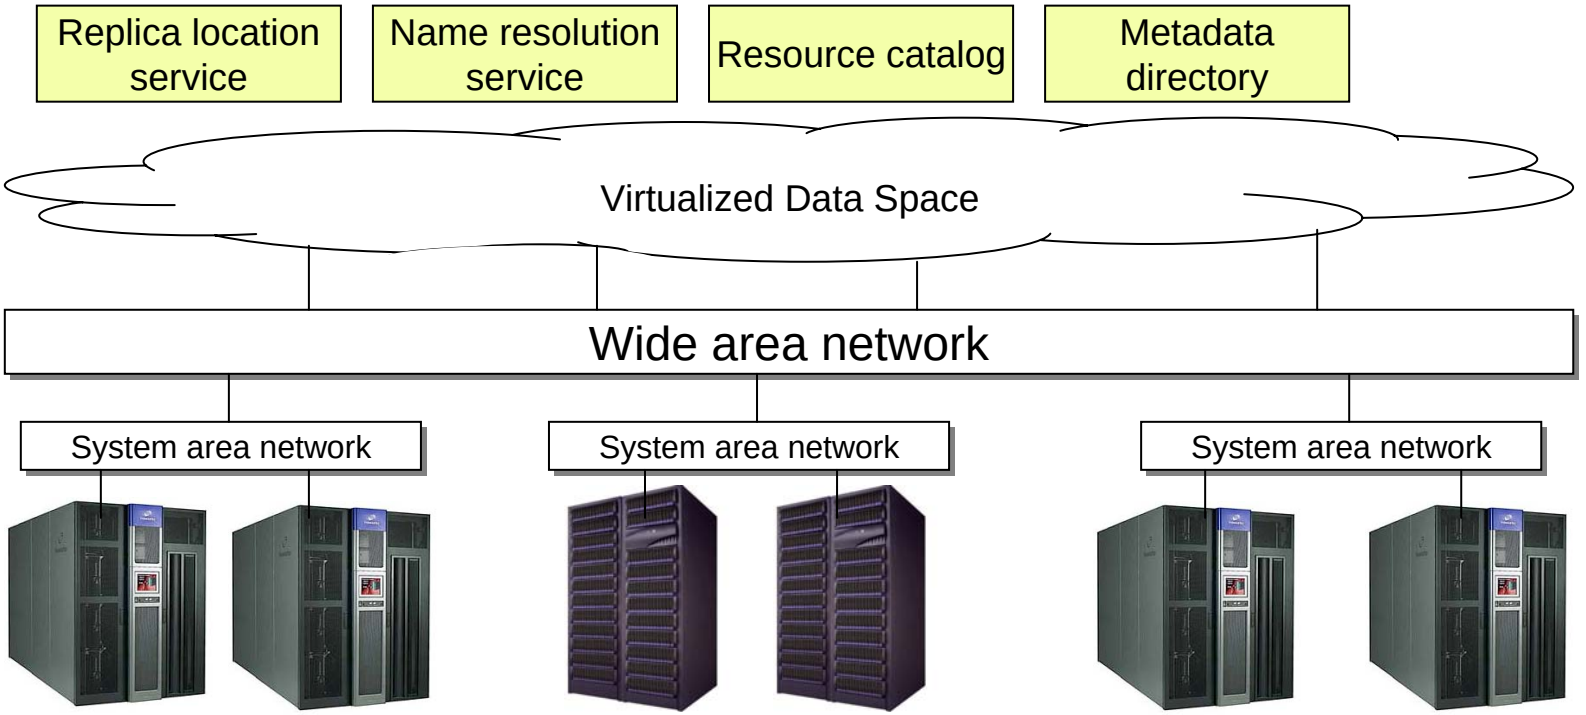

Basic Gateway SOA

Grid
Portal Server

Gateway Services

Proxy Certificate
Server / vault

Application
Deployment

Workflow engine

User Metadata
Catalog MyLEAD

Application Events

Resource Broker

App. Resource
catalogs

Core Grid Services

Security
Services

Information
Services

Self
Management

Resource
Management

Execution
Management

Data
Services

OGSA-like Layer

Physical Resource Layer

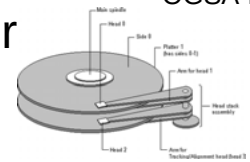

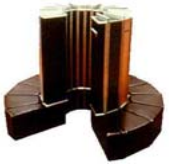

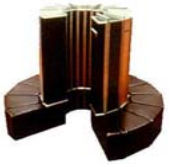

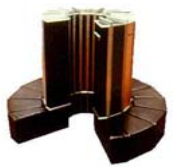

eGee
Enabling Grids
for E-sciencE

One Example in Detail: The LEAD Project

A Grid designed to change the paradigm for mesoscale weather prediction

Building dynamic, adaptive workflows from data streams and better-than-real-time execution constraints.

Tornados and Hurricanes

- The loss of life and property due to tornados and hurricanes each year is large and growing.
- We seem to be entering a period where the situation may becoming worse.
- We are also on a threshold of far more advanced and accurate prediction capabilities.

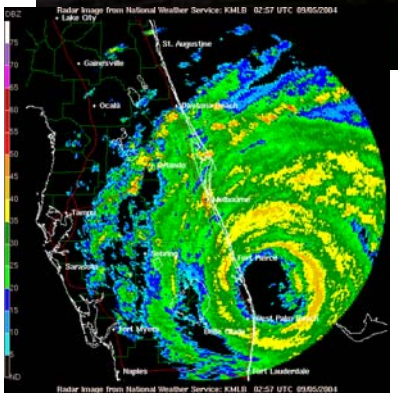

The LEAD Project

LINKED ENVIRONMENTS FOR ATMOSPHERIC DISCOVERY

Traditional Methodology

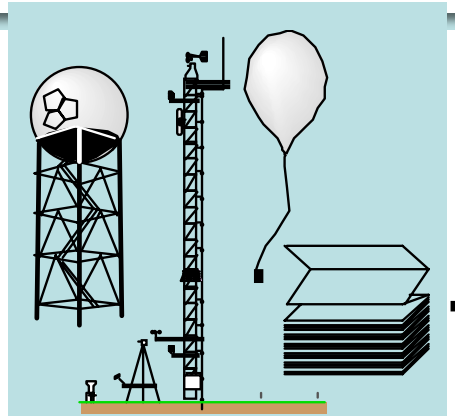

STATIC OBSERVATIONS

Radar Data
Mobile Mesonets
Surface Observations
Upper-Air Balloons
Commercial Aircraft
Geostationary and Polar Orbiting Satellite
Wind Profilers
GPS Satellites

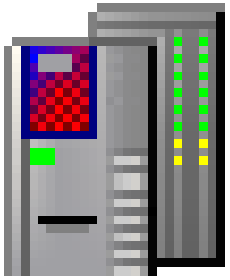

Analysis/Assimilation

Quality Control
Retrieval of Unobserved Quantities
Creation of Gridded Fields

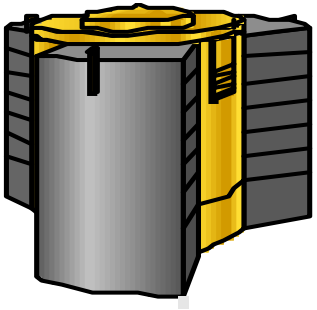

Prediction/Detection

PCs to Teraflop Systems

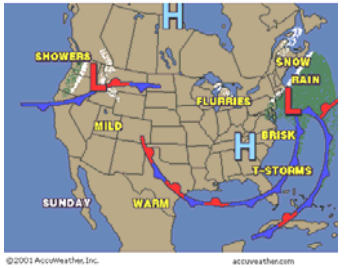

Product Generation, Display, Dissemination

The Process is Entirely Serial and Static (Pre-Scheduled): No Response to the Weather!

End Users

NWS
Private Companies
Students

Am Major Paradigm Shift: CASA NETRAD adaptive Doppler Radars.

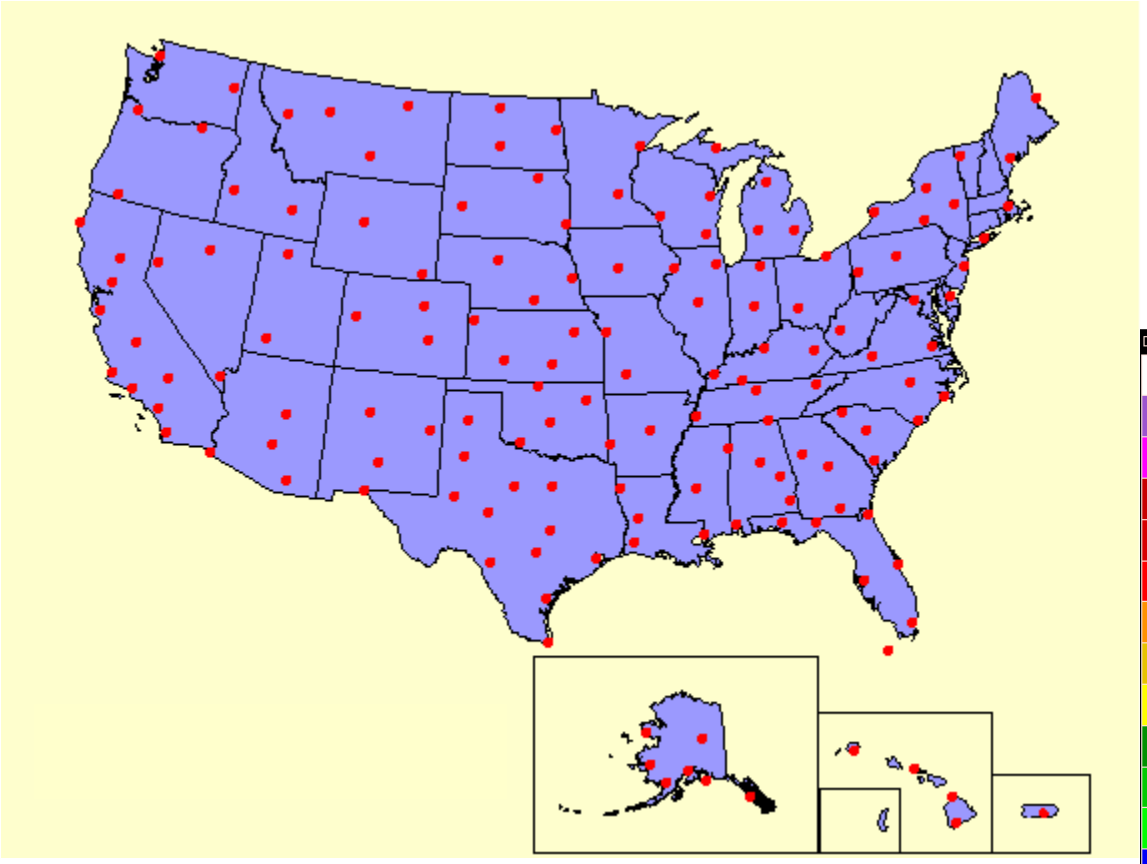

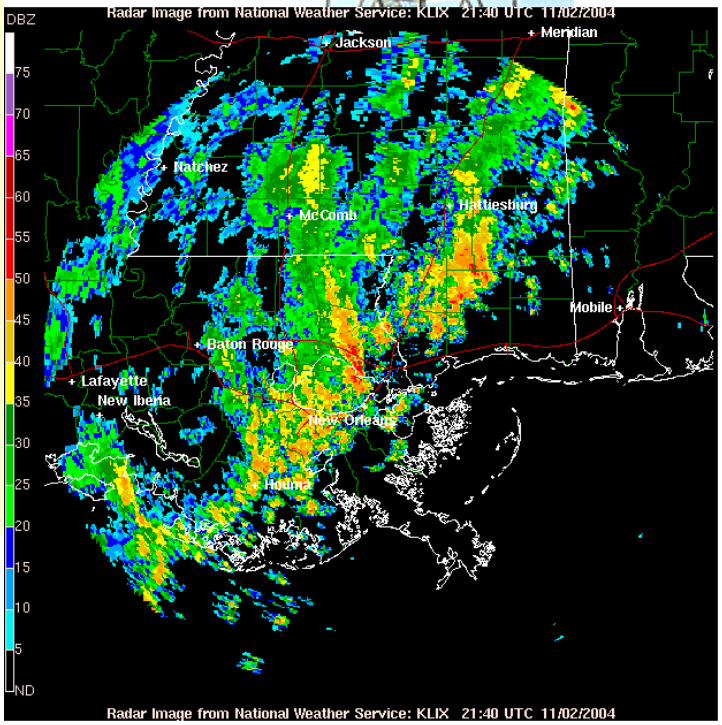

The LEAD Vision: Adaptive Cyberinfrastructure

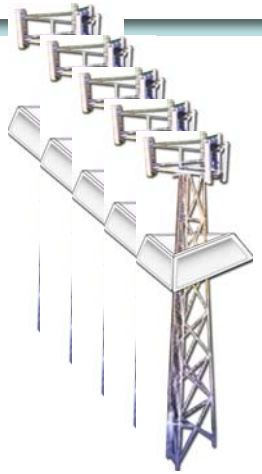

DYNAMIC OBSERVATIONS

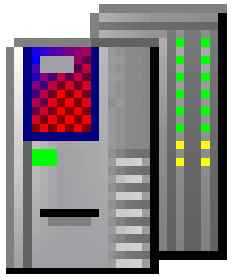

Analysis/Assimilation

Quality Control
Retrieval of Unobserved Quantities
Creation of Gridded Fields

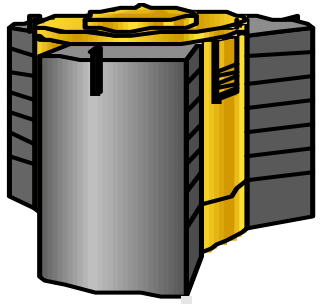

Prediction/Detection

PCs to Teraflop Systems

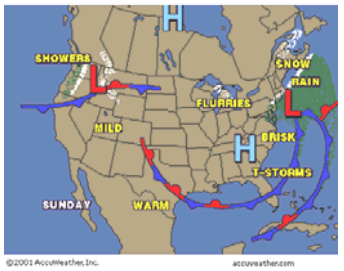

Product Generation, Display, Dissemination

Models and Algorithms Driving Sensors

End Users

NWS
Private Companies
Students

The CS challenge: Build cyberinfrastructure services that provide adaptability, scalability, availability, useability, and real-time response.

Change the Paradigm

- To make fundamental advances we need:
  - Adaptivity in computational model.
- But also Cyberinfrastructure to:
  - Execute complex scenarios in response to weather events
    - Stream processing, triggers
    - Close loop with the instruments.
  - Acquire computational resources on demand.
    - Need supercomputer-scale resources
    - Invoked in response to weather events
  - Deal with data deluge
    - User can no longer manage his/her own experiment products

Reaching the LEAD Goal

- Understanding the role of data is the key
  - from streams, from mining, from simulations to visualizations
- Enabling Discovery
  - The role of the experiment
  - Sharing both process and results
  - Creating an educational “context”
- Each user should be able to see a history of data that is generated for each computational experiment.
- This should be a private data space where the scientist sees “metadata” about each data product produced during the computation.
- Logging into the LEAD portal, you see:

Log in and see your MyLEAD Space

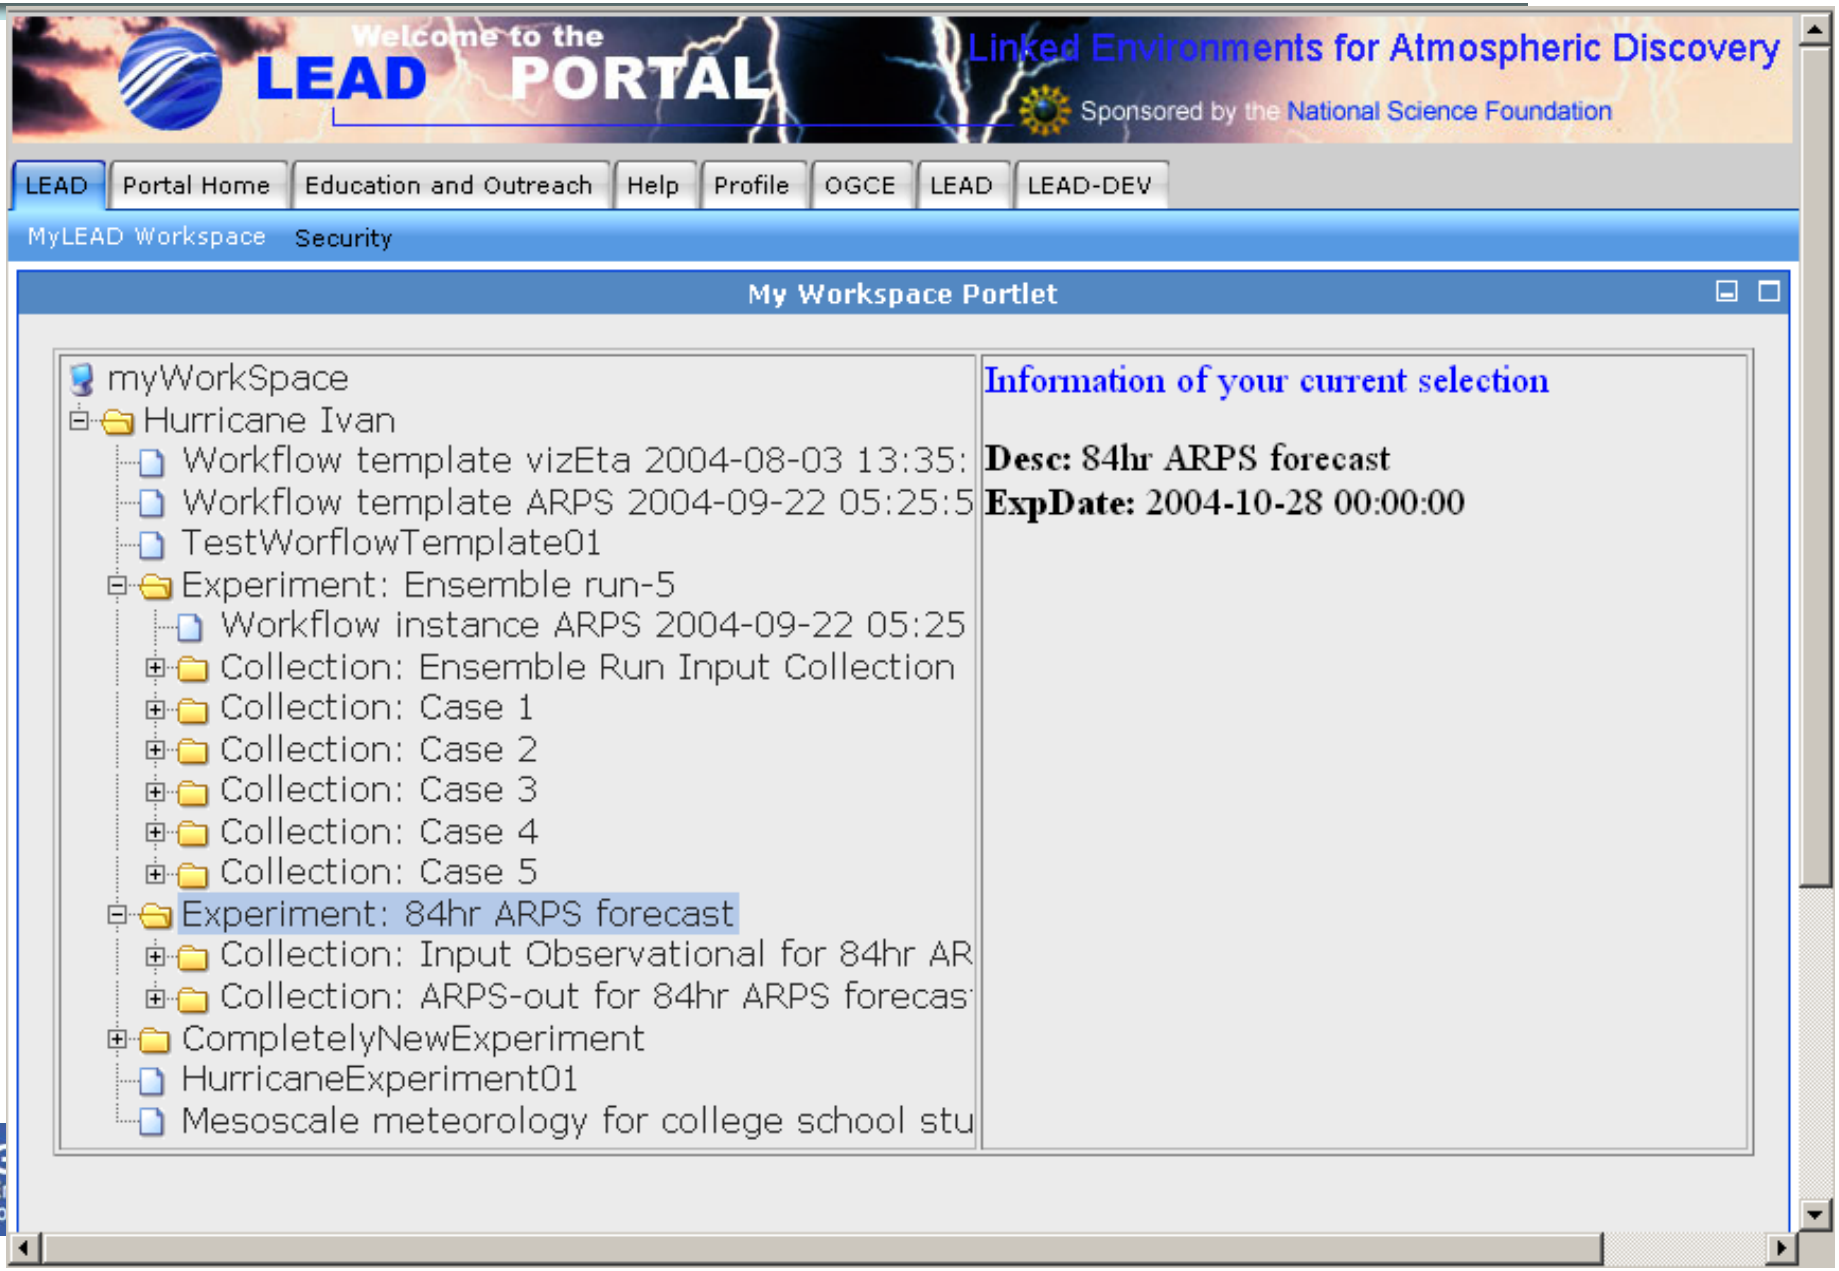

MyLEAD Metadata Catalog

- A catalog of resource metadata.
- Actual data storage is virtualized
  - The data itself my reside anywhere.
  - Replicas may exist. Data may move

- What the user sees should be automatically generated, searchable metadata about their scientific experiments.
  - The metadata describes the actual data object and provides a link to the data services that can resolve an abstract unique identifier for a data product to a real, on-line data “file”.

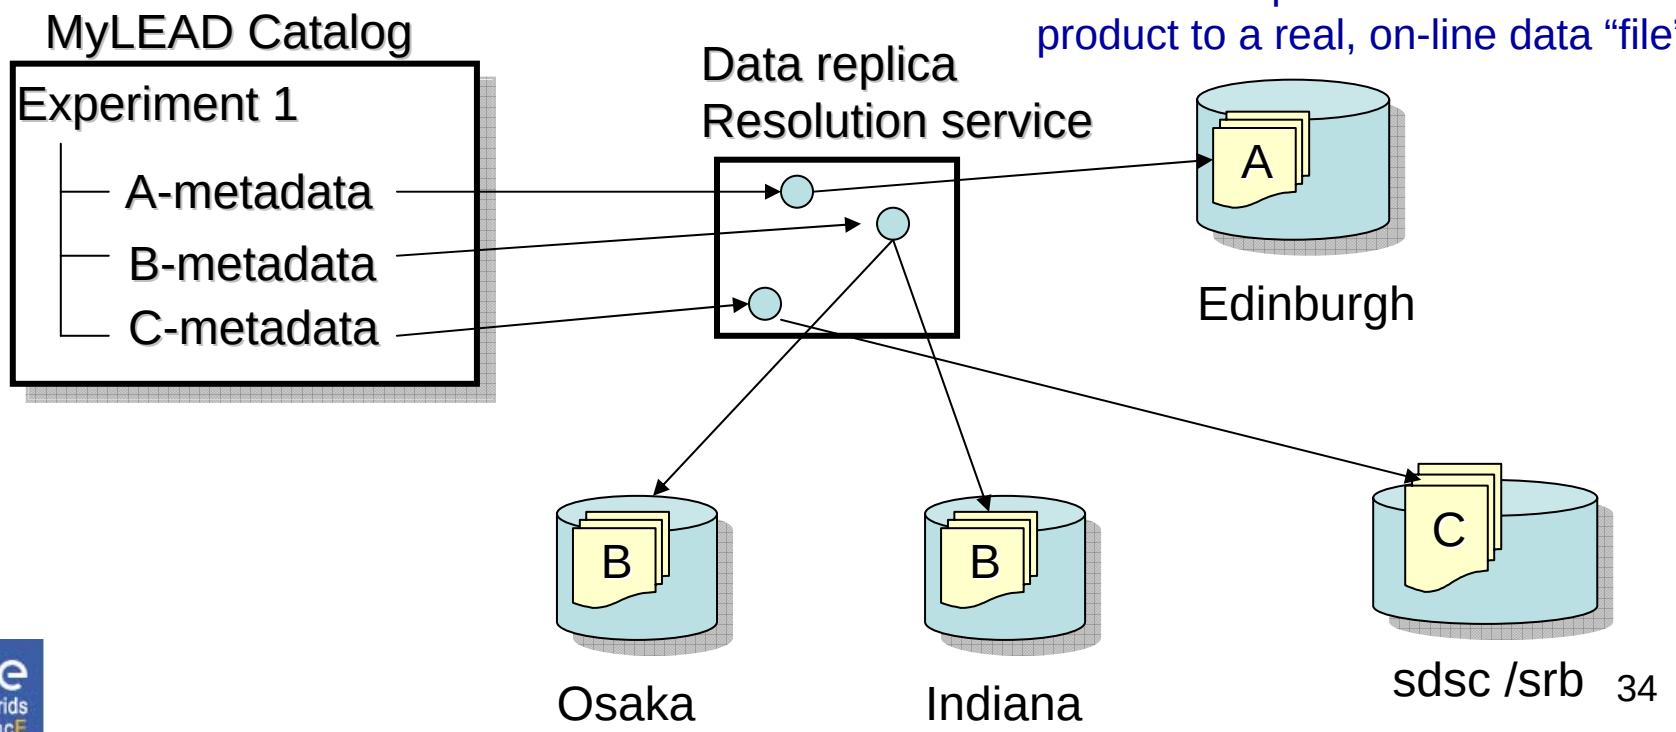

Data search can take many forms

- Select a region and a time range and desired attributes

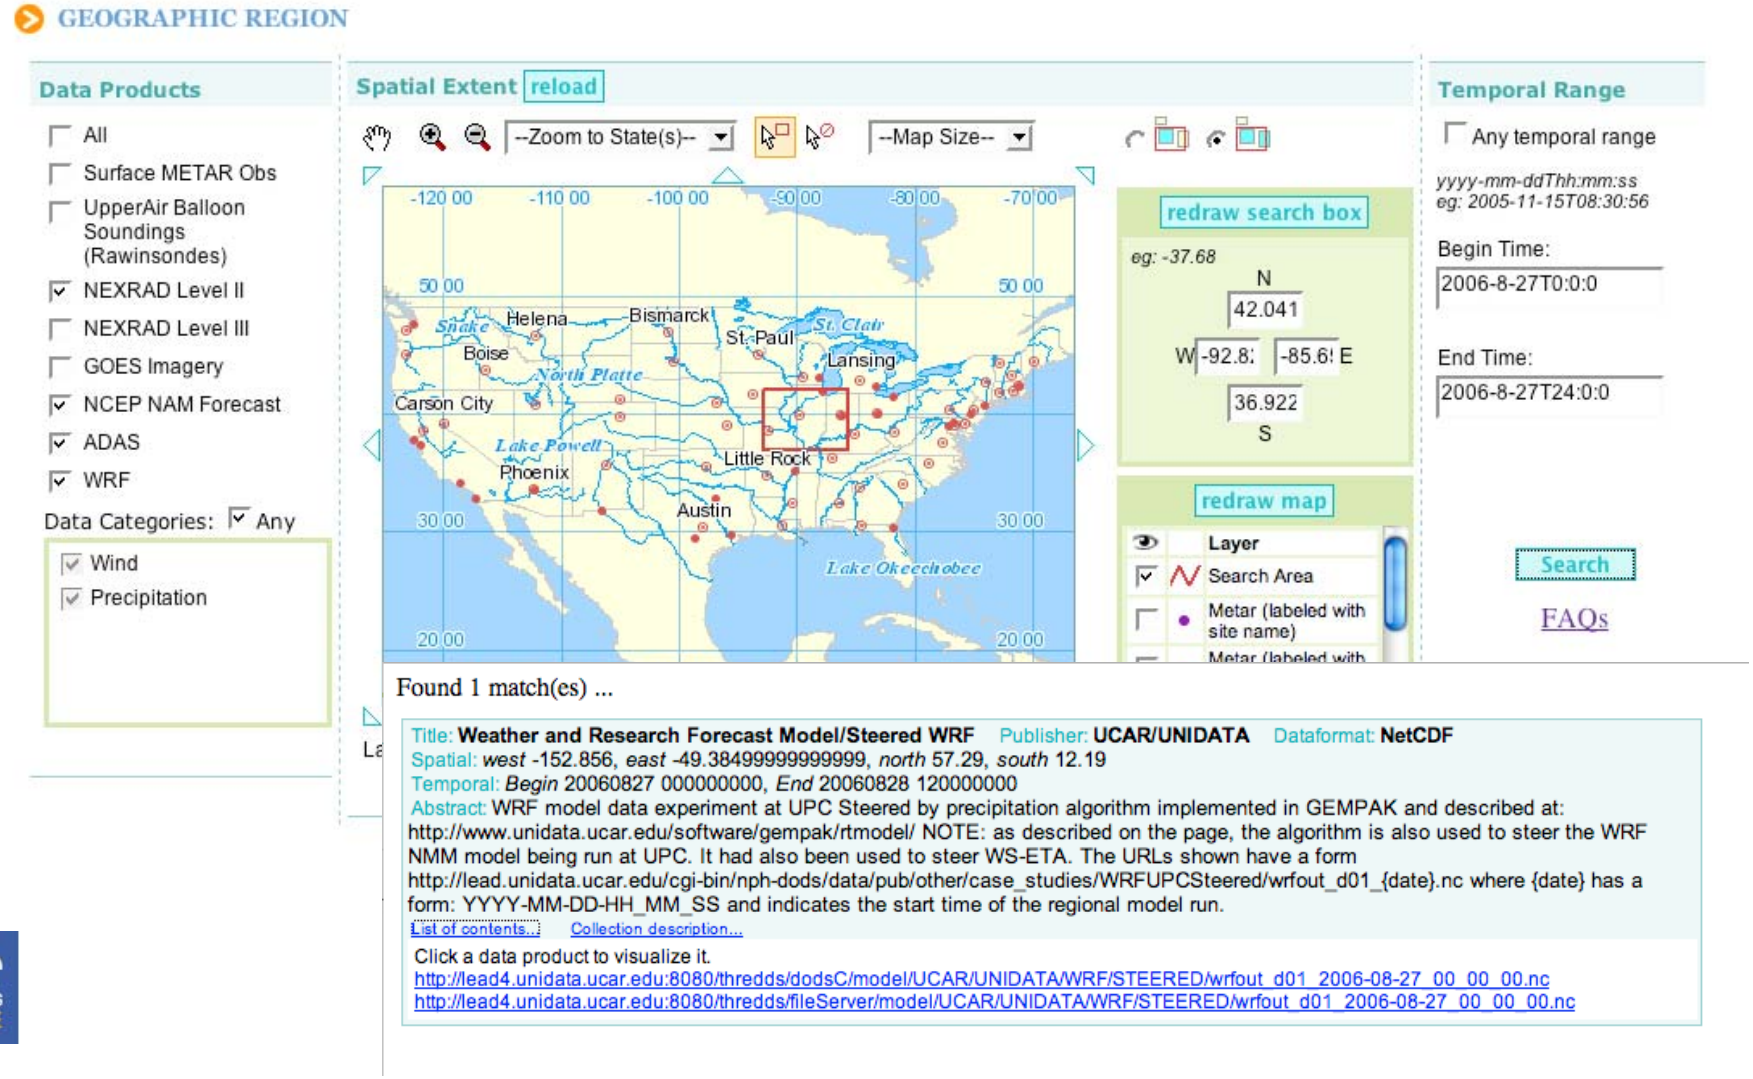

Experiment as Workflow

- Another Paradigm Shift for the users.
- Each activity a user initiates in LEAD is an Experiment which consists of
  - Data discovery and collection.
  - Applied analysis and transformation
    - A graph of activities (workflow)
  - Curated data products and results
- Each workflow activity is logged using an event system and stored as metadata in the users workspace.
  - Provides a complete provenance of work.

Workflow: Composing Computational Tools to build new Tools

- Workflow is a term that describes the process of moving data through a sequence of analysis and transformational steps to achieve a goal.
- In e-Science most data analysis and transformational tools can be abstracted as services.
- A workflow is then a composition of services in which data moves from one end of an abstract process “pipeline” to another.

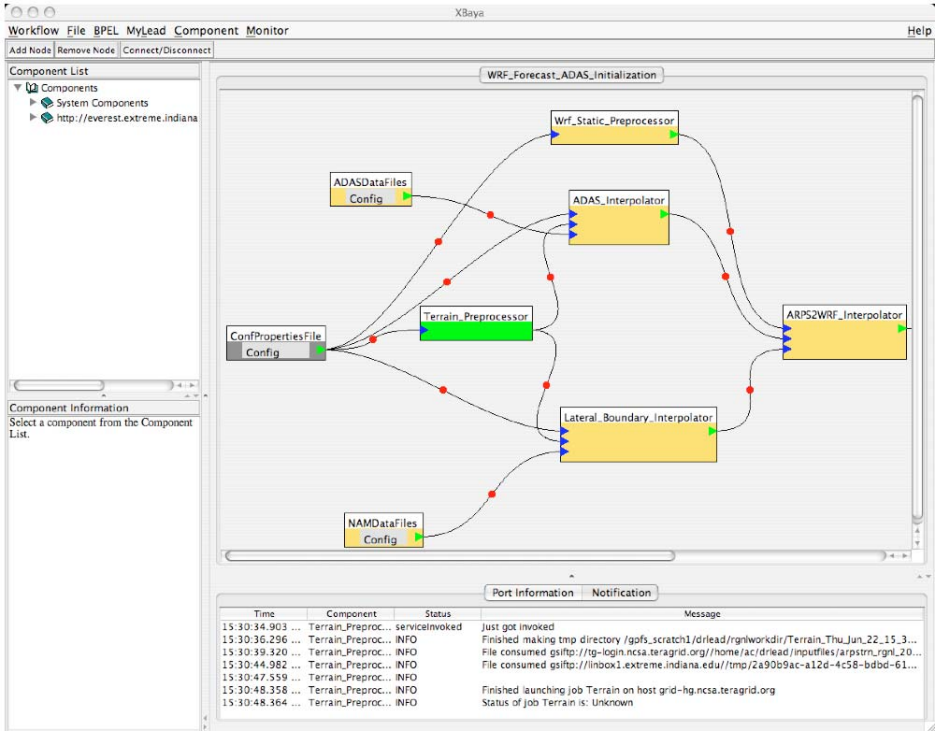

Requirements of Workflow Systems

- Hide the Grid
  - Data access is virtualized
  - Computation is virtualized
    - - applications as services
  - Security is automatic.
- Automatic metadata generation
- Fault Tolerance
  - Workflows may be verrrry long running.
- Dynamic, Adaptive workflows.
  - Respond and adapt to external events.
- Compositional Orthogonality
- A workflow can be a component in another workflow.

Architecture of a LEAD workflow

- A workflow is a (visual) script that
  - Responds to events
  - Orchestrates the interaction of web services that gather data or run applications.

- Notification Broker is the nervous system.
  - Relays “events” between all services
- Workflow Engine Service executes scripts

Workflow composer

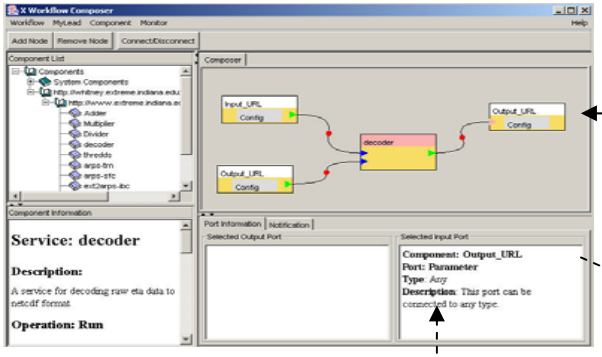

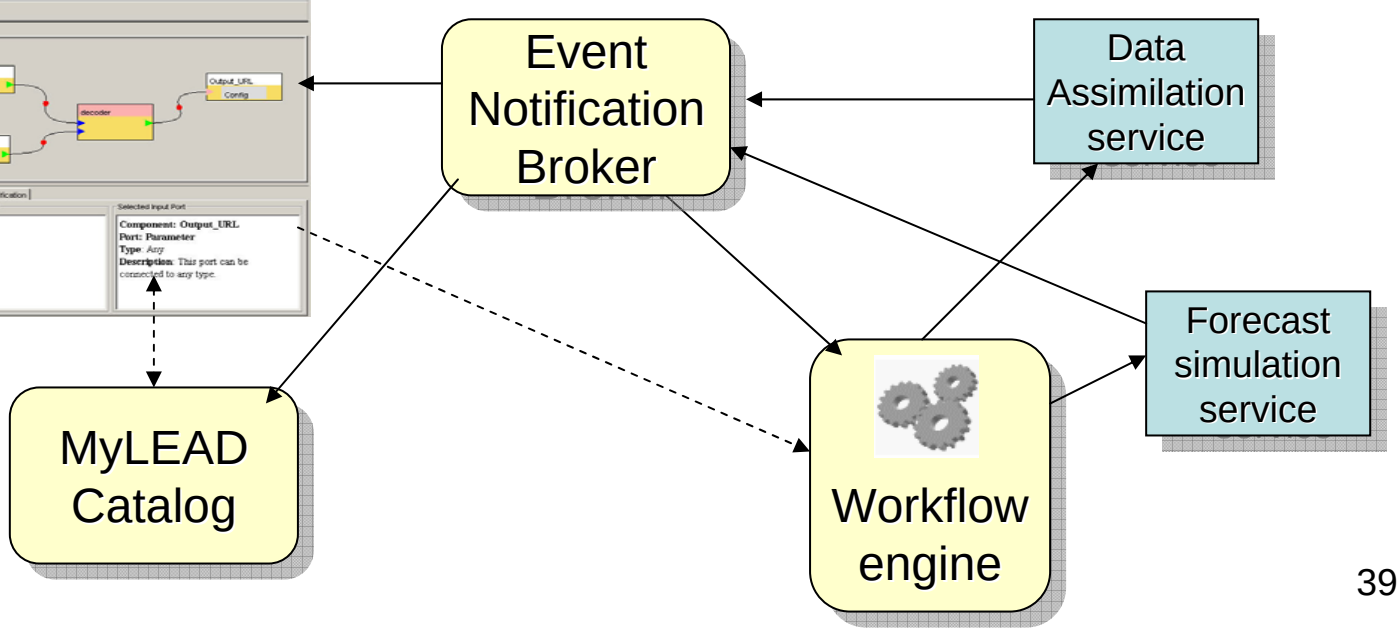

Monitoring the Event Stream

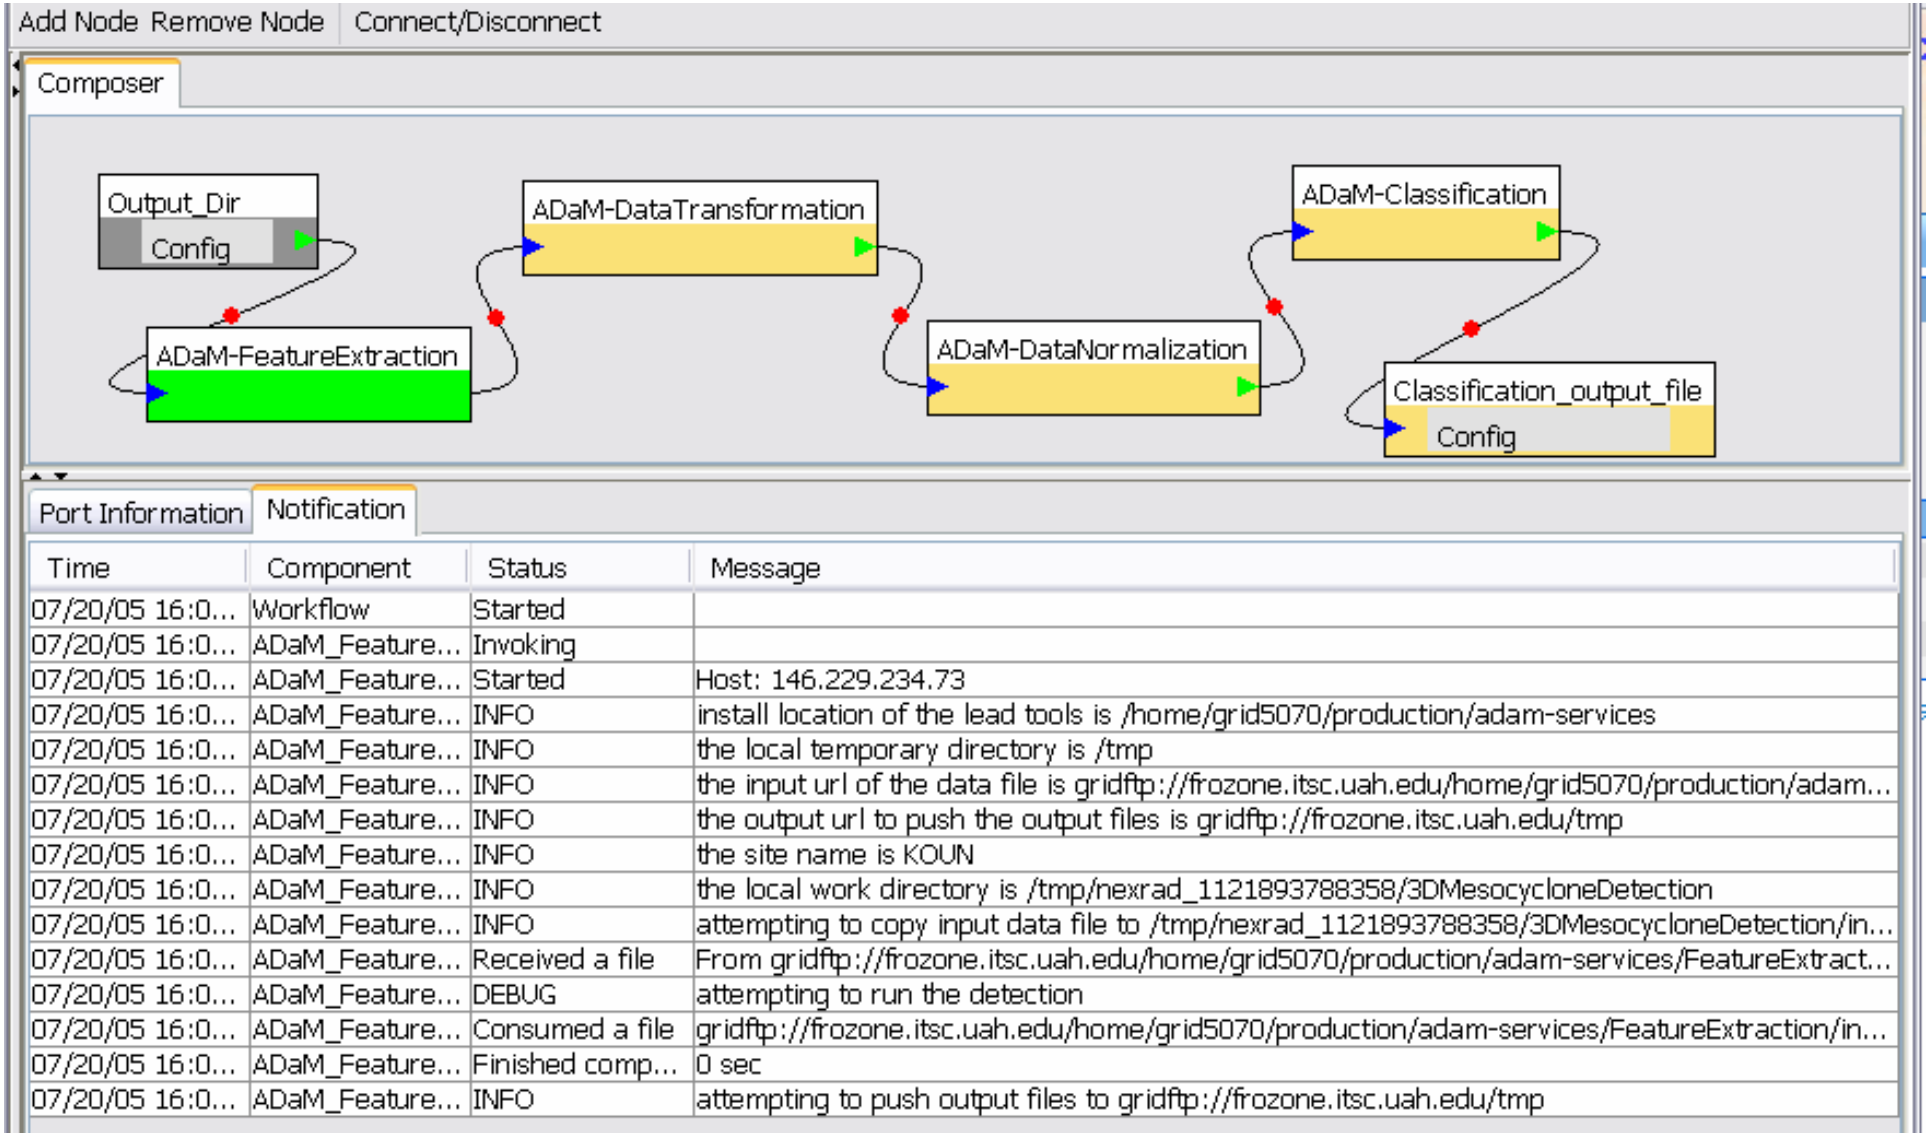

Time	Component	Status	Message
07/20/05 16:0...	Workflow	Started	
07/20/05 16:0...	ADaM_Feature...	Invoking	
07/20/05 16:0...	ADaM_Feature...	Started	Host: 146.229.234.73
07/20/05 16:0...	ADaM_Feature...	INFO	install location of the lead tools is /home/grid5070/production/adam-services
07/20/05 16:0...	ADaM_Feature...	INFO	the local temporary directory is /tmp
07/20/05 16:0...	ADaM_Feature...	INFO	the input url of the data file is gridftp://frozone.itsc.uah.edu/home/grid5070/production/adam...
07/20/05 16:0...	ADaM_Feature...	INFO	the output url to push the output files is gridftp://frozone.itsc.uah.edu/tmp
07/20/05 16:0...	ADaM_Feature...	INFO	the site name is KOUN
07/20/05 16:0...	ADaM_Feature...	INFO	the local work directory is /tmp/nexrad_1121893788358/3DMesocycloneDetection
07/20/05 16:0...	ADaM_Feature...	INFO	attempting to copy input data file to /tmp/nexrad_1121893788358/3DMesocycloneDetection/in...
07/20/05 16:0...	ADaM_Feature...	Received a file	From gridftp://frozone.itsc.uah.edu/home/grid5070/production/adam-services/FeatureExtract...
07/20/05 16:0...	ADaM_Feature...	DEBUG	attempting to run the detection
07/20/05 16:0...	ADaM_Feature...	Consumed a file	gridftp://frozone.itsc.uah.edu/home/grid5070/production/adam-services/FeatureExtraction/in...
07/20/05 16:0...	ADaM_Feature...	Finished comp...	0 sec
07/20/05 16:0...	ADaM_Feature...	INFO	attempting to push output files to gridftp://frozone.itsc.uah.edu/tmp

Dynamic Adaptive Workflow

LEAD requires ability to construct workflows that are

- Data Driven
  - Weather data streams define nature of computation
- Persistent and Agile
  - Data mining of data stream, detects “interesting” feature, event triggers workflow scenario that has been waiting for months.
- Adaptive
  - In response to weather: weather changes.
  - Nature of workflow may have to change on-the-fly.
  - Resources
    - More may be needed, sometimes they become unavailable.
    - Need to be self-aware

Knowledge Discovery from Streaming Data

- It is possible to issue a query which will automatically start a distributed network of data mining filters.
- As results are processed matches to the query can trigger deeper analysis or simulations of possible outcomes.
- The e-science architecture becomes responsive and adaptive to the environment.

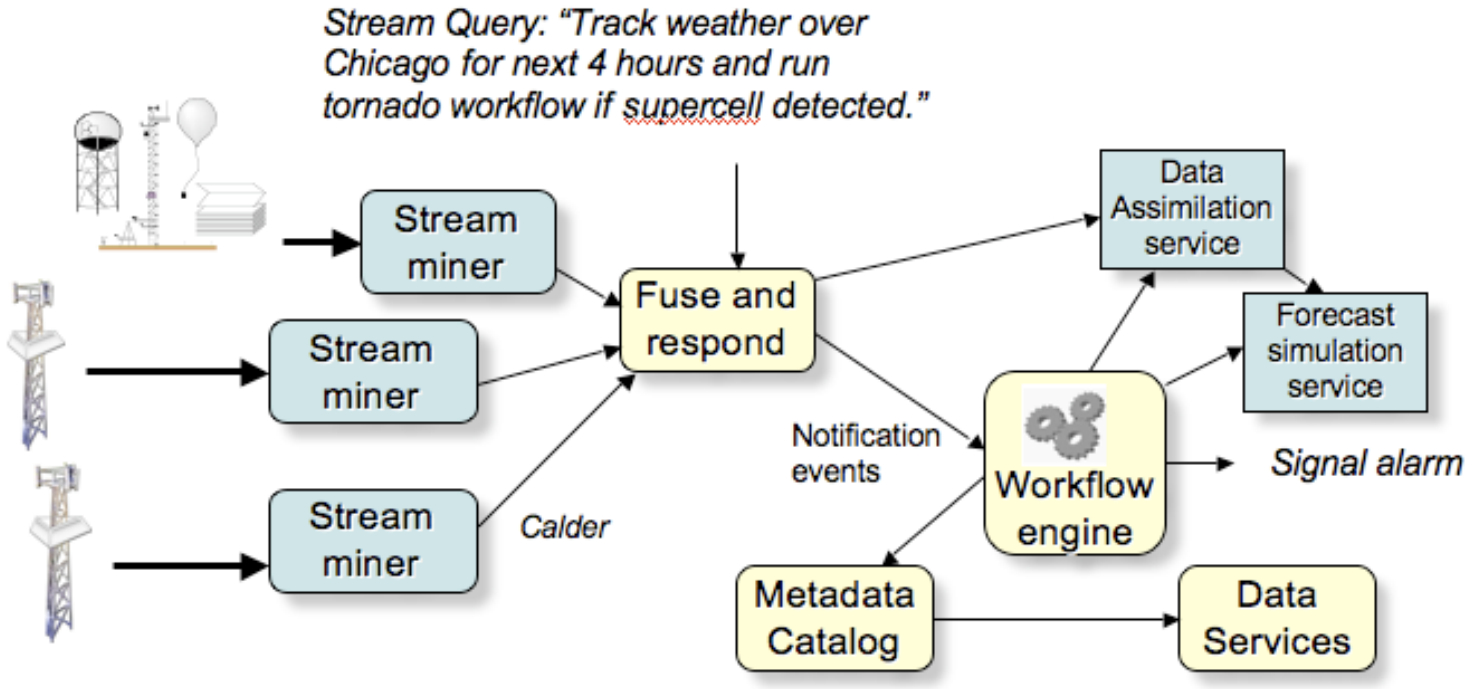

Assimilation-forecast workflow

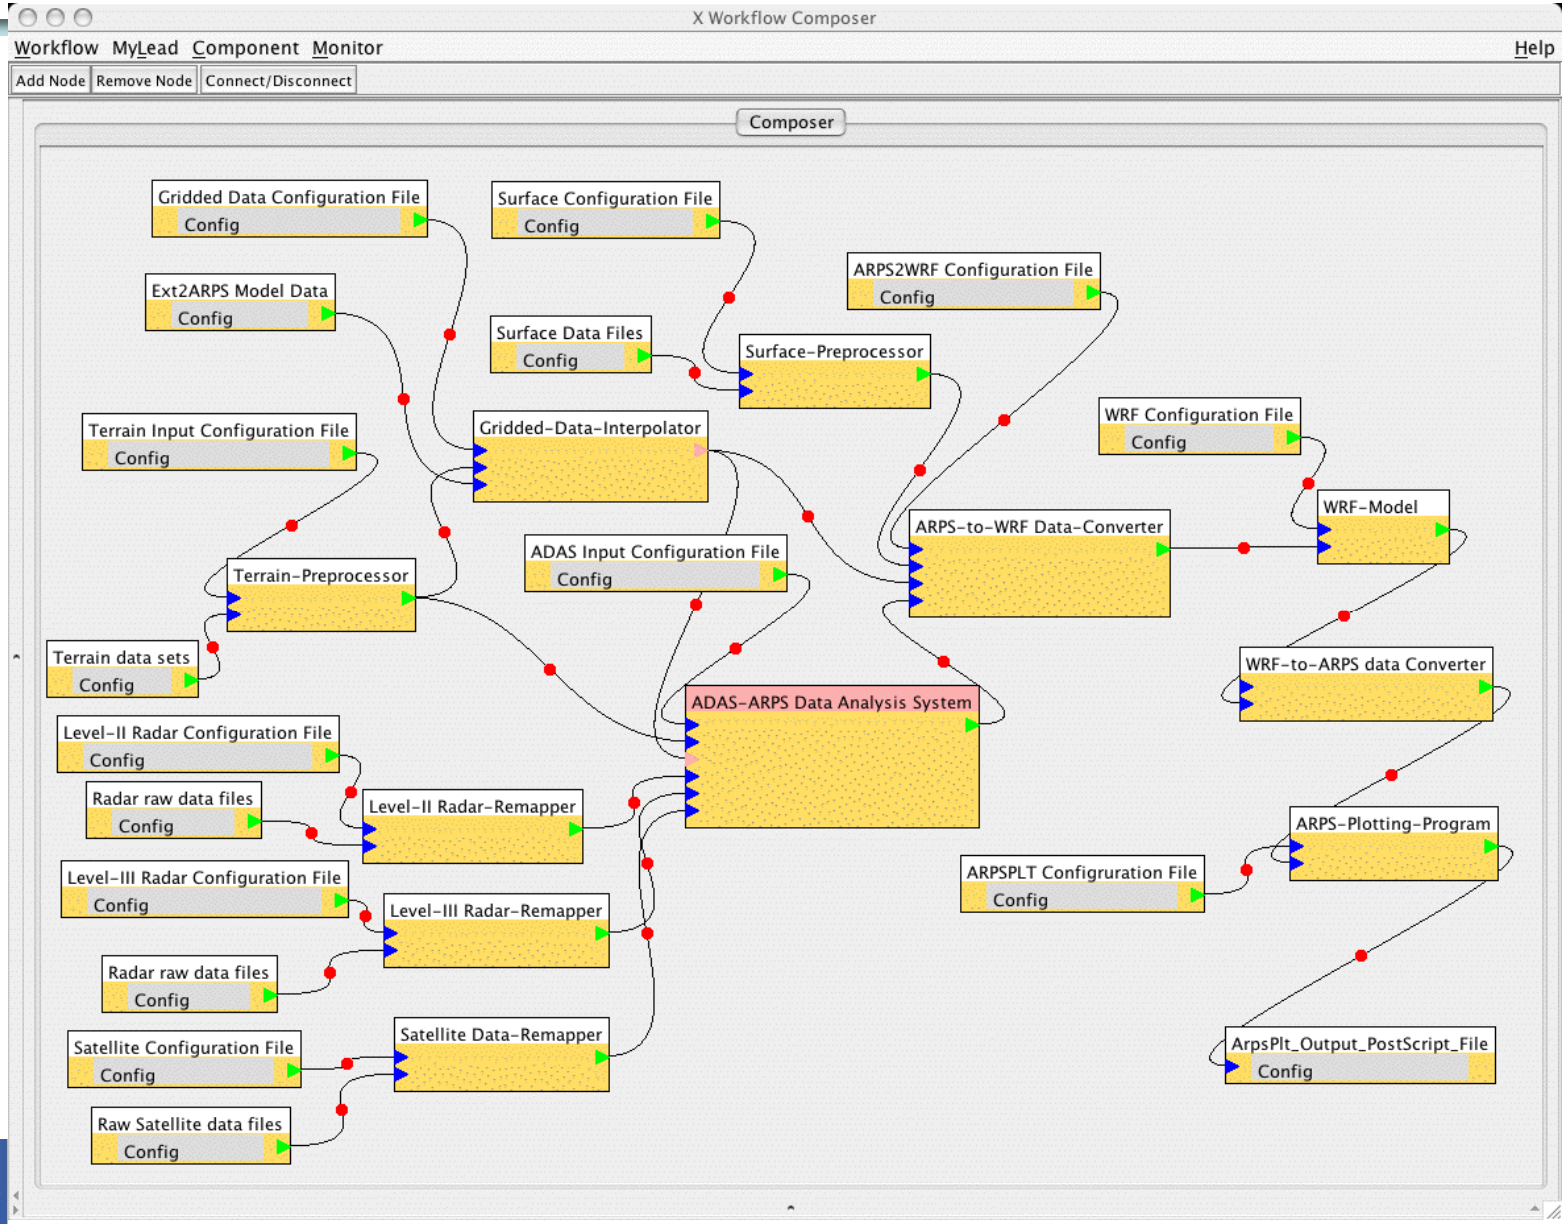

Workflow applied to Katrina

2D image of sea level
generated by ARPS Plotting Service

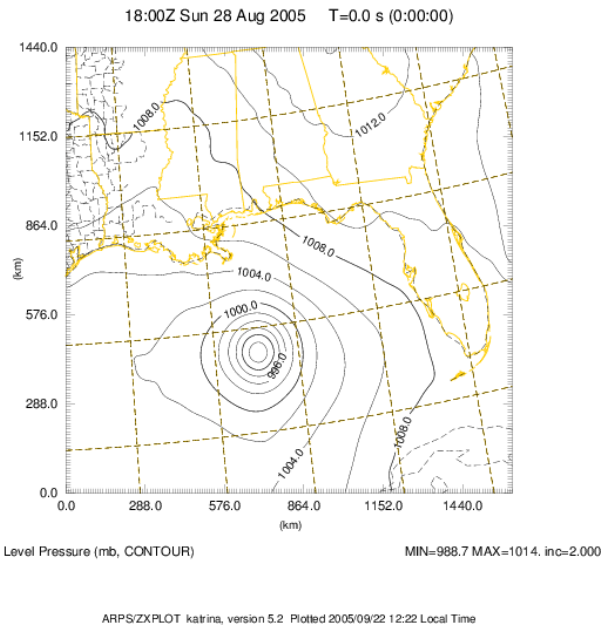

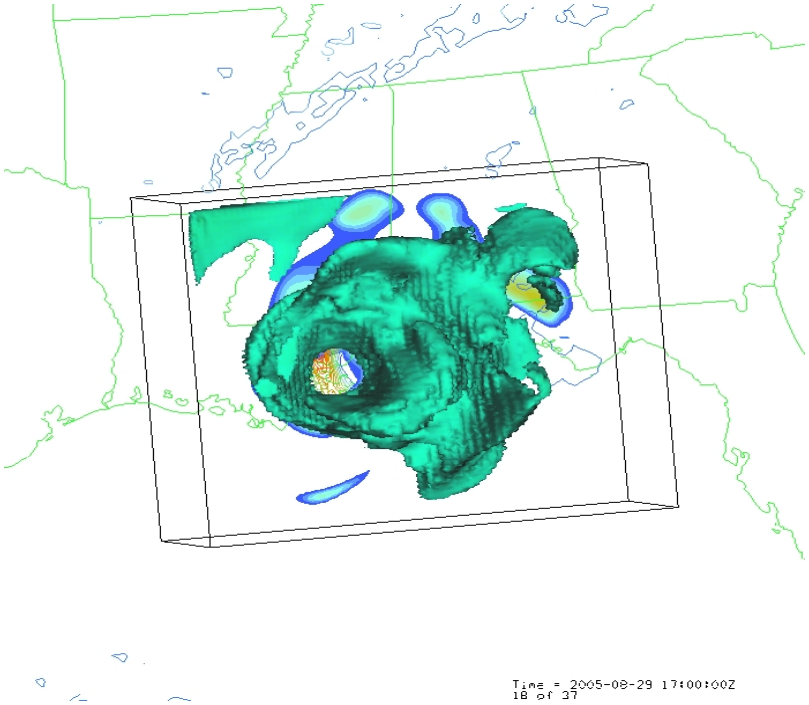

3D Image generated by IDV

Lessons from LEAD and other Grid Gateways

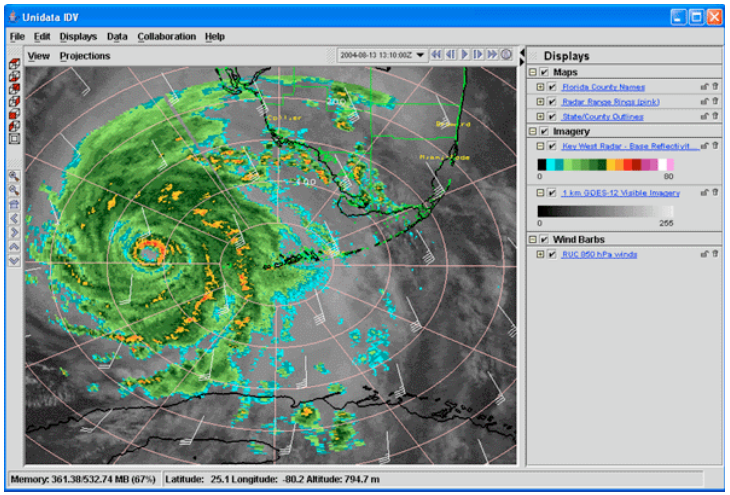

An Evolution of the Grid Concept

- Service Orientation
  - Virtualization of Resources
    - The first step in the Evolution of the Grid concept
  - Virtualization of Data and Applications
    - The Key role of the SOA.
    - Data is managed by services, not humans.
      - Data products never go away and they may live anywhere on the grid.
    - A Metadata catalog contains searchable knowledge about the data product.
      - Contains a name that can be resolved to find a valid replica if the data is needed.

Virtualization of the Applications

- Workflows only ask that a particular application be run.
  - Pass it the location-independent names of data object and parameters
- **An Application service** consults a resource catalog that knows where the application is deployed
- **A resource broker service** is contacted to find the best compute engine to run the app.
  - Depends on requirements
    - Size of request, authorization, data locality.

Some Challenges Ahead

Data Discovery Agents

- Agents for Search
  - Advanced support data discovery.
  - We will soon be able to pose queries for future discovery
    - "I am interested in all new data relating to chemical structures of the following form … When you find them, run the following analysis workflow against it and notify me if the result is interesting".

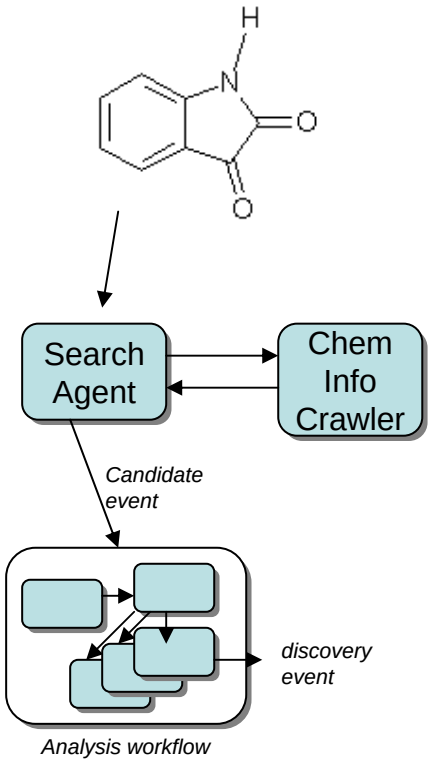

Synthesizing Data Search and Discovery Queries

- Take an abstract query and “compile” it to a data search and analysis workflow.
  - *“track the weather over Chicago for the next 400 hours and run a tornado workflow if a supercell is detected.”*

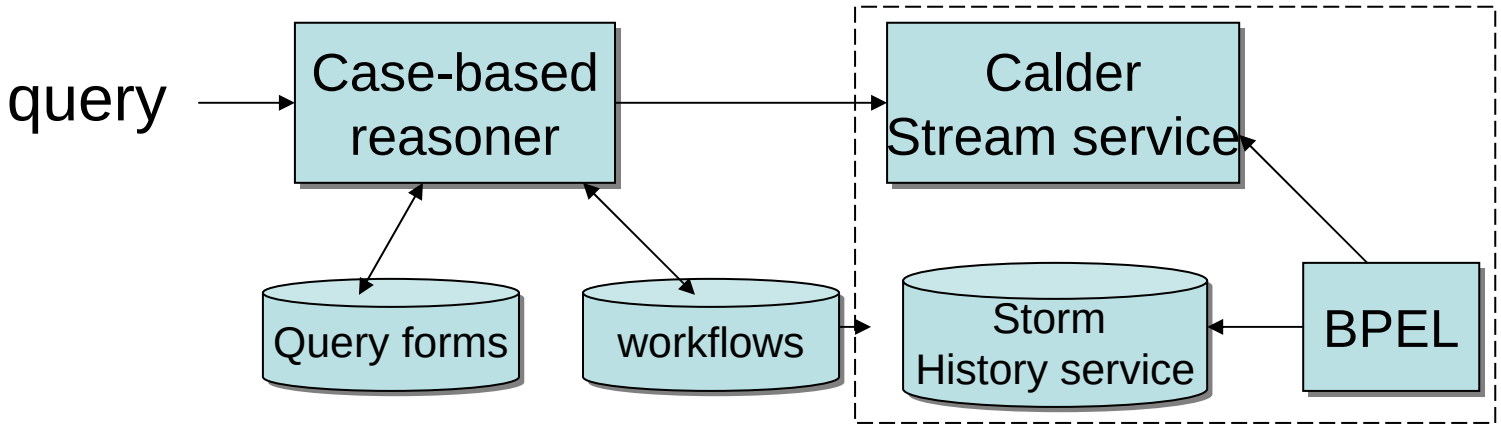

Validating Scientific Discovery

- The portal is becoming part of the process of science by being an active repository of data provenance
- Disks are cheap, so why not record everything?
- The portal records each computational experiment that a user initiates
  - A complete audit trail of the experiment or computation
  - Published results will include link to provenance information for repeatability and transparency.
- Many portals have done this on a smaller scale
  - CIMA, PubChem+NIH cancer screening centers, LEAD, SERVO/Quakesim, ...
  - But this should be standard practice

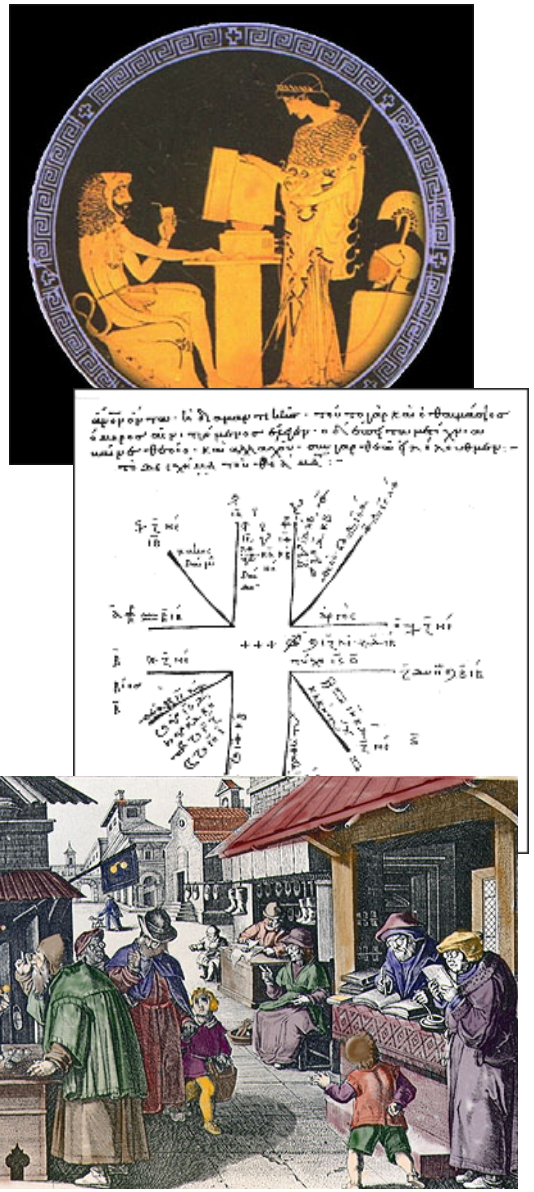

Workflows that share

- What if a workflow can use the provenance service to tell when an intermediate data product is already available.
- Use the virtualized data space to optimize the execution!
- Depends upon the quality of metadata search available.

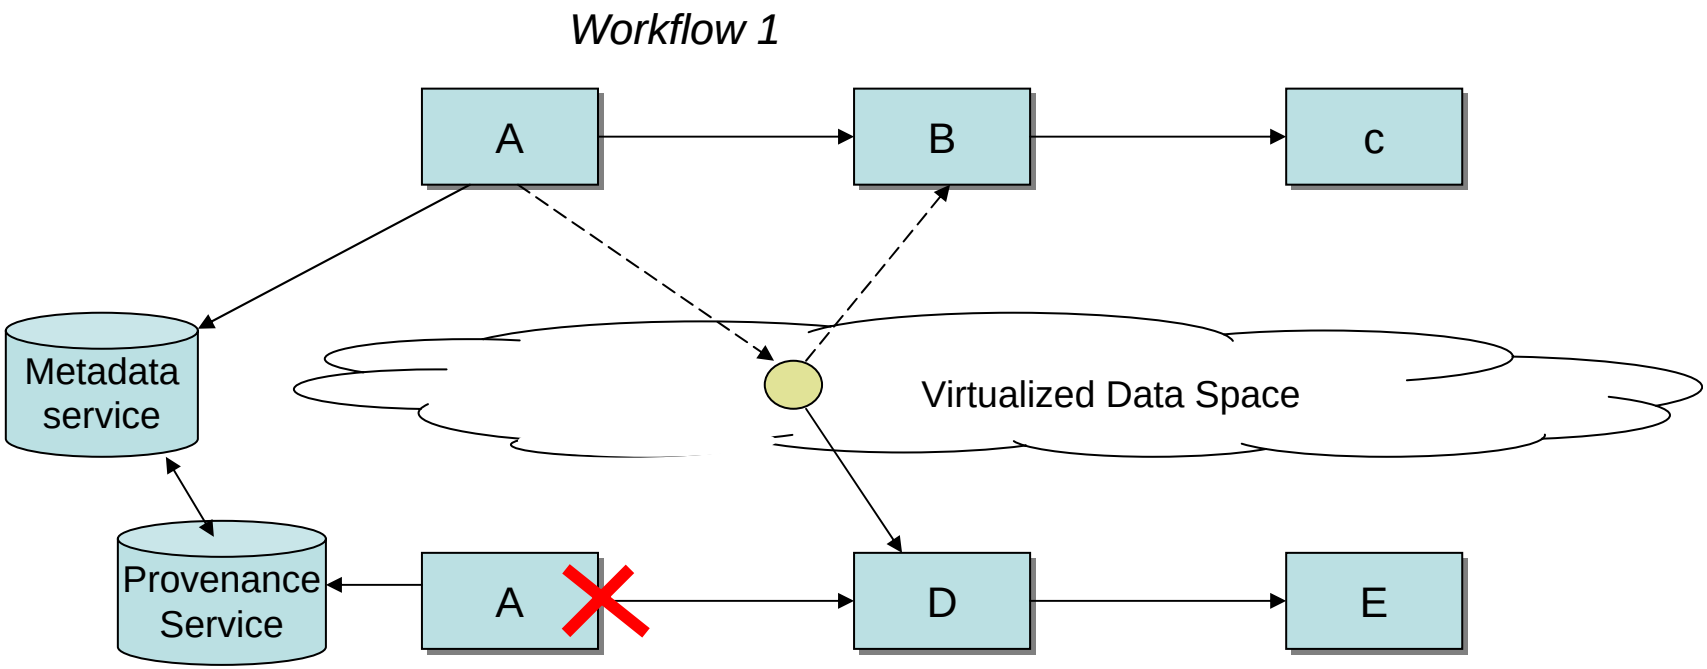

Conclusion

- Science Gateway portals have been enabled by an emerging standard Grid architecture
  - Multi-domain Security
  - Virtualized data management
  - Virtualized computation.
- Grid application construction as Workflow design
  - Over 20 different Workflow tools exist.
  - The need to meet some high standards to be powerful
    - Dynamic, adaptive, fault tolerant, metadata/provenance aware …
- Data query and workflow are intimately connected.
  - *The interesting questions are not just table lookup… they require work and thought.*