XSEDE BOF: Science Clouds
or
It Takes a Village to Give a Talk on Cloud Computing

Kate Keahey, David Lifka, Manish Parashar, Warren Smith, Carol Song, Shaowen Wang

Clouds Available to Science

FutureGrid Overview

- A distributed infrastructure to support computer and computational science experiments
  - Performance analysis, software testing and evaluation, interoperability
- A rich education and training platform
  - University courses, multi-day training sessions, informal learning and exploration
- Includes Cloud, Grid, and High-Performance Computing environments
  - Allows users to configure environments to meet their needs
  - Typically available interactively
- Provides tools to support rigorous experimentation
  - Documenting configurations, recording experiments
- Allocatable via XSEDE

FutureGrid Deployment Architecture

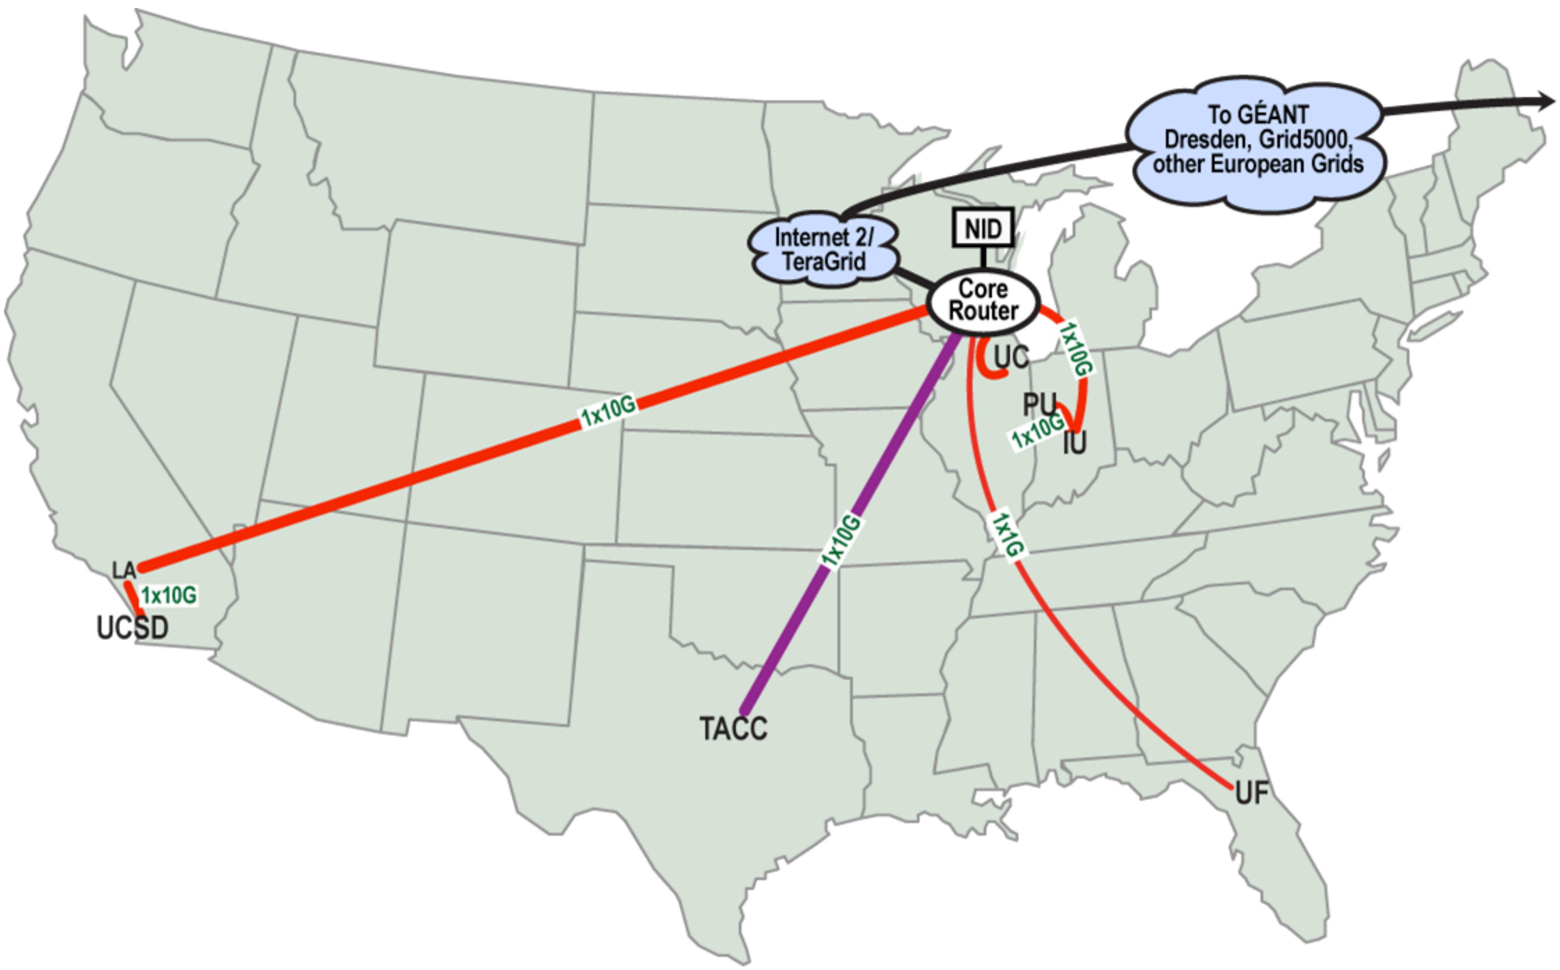

All sites are connected via a core router that is attached to a configurable network impairment device. Dedicated network links are shown in red.

IaaS Partitions

Site	# Cores*	TFLOPS*	Total RAM* (GB)	Secondary Storage* (TB)	Platforms
IU	1308	11	2048+	335	Eucalyptus, OpenStack
TACC	768	8	1152	30	Nimbus, OpenStack
UC	672	7	1344	120	Nimbus
SDSC	672	7	2688	72	Nimbus, OpenStack
UF	256	3	768		Nimbus
Total	3676	42	9344	557	

*All partitions on cluster

- Partitions mostly independent (Nimbus ones share authentication credentials)
- Accessed via implementation-specific interfaces and Amazon interfaces
- IU system includes large memory/disk nodes

Virtual Machine Images and Appliances

- Provide a set of generic Linux images
- Provide appliance images for specific tasks
  - Virtual cluster
  - Condor
  - MPI
  - Hadoop
- Users create and share their own images

red cloud

On-Demand Research Computing

– Infrastructure as a Service –

– Software as a Service –

– Cloud Storage Solutions –

www.cac.cornell.edu/redcloud

Commercial Clouds AWS

- Instances
  - From $0.02 per hour to Top500 membership
  - Xen hypervisor
- Storage
  - Transient: 160 GB-48 TB on VM
  - Persistent: Elastic Block Storage (EBS) 1BG-1TB
- Networking
  - Internal: up to 10 Gb Ethernet for Cluster Compute
  - AWS Direct Connect (for large volume users)
  - Sneakernet
- Pricing
  - Reserved instances (up to 71% savings), 1-3 year commitment
  - Spot instances

Other Major Commercial Offerings

- Google Compute Cloud (since 06/12)
  - Target high performance market, emphasizes consistent performance
  - Uses KVM
  - Pricing by the minute
- Windows Azure (since 2010)
  - Started out as a platform offering
  - Uses Windows Azure Hypervisor (can do Linux)
- Rackspace
  - Based on OpenStack (since 2012)
  - Uses Xen hypervisor

More resources at scienceclouds.org

Science Cloud Solutions: from Private and Community Clouds to Commercial Clouds from Proof-of-Concept to Solution

The Appliances

- Where do VM images come from?
  - BoxGrinder, VMBuilder, rBuilder, veewee, Oz, etc.
- Challenge: interoperability and consistency
- Nimbus Image Portal
  - VM Image creation
  - Generates images for different hypervisors and clouds
  - Based on Packer

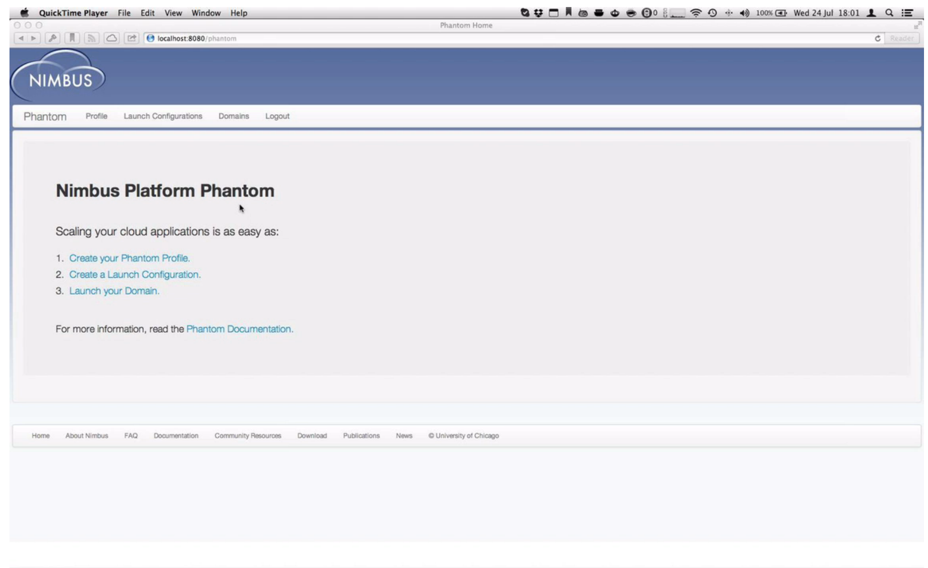

The Contextualization

- Mainstream ctx tools: Chef, Puppet
- Providing abstractions, scalability, repeatability and control
  - StarCluster, Nimbus Context Broker, cloudinit.d

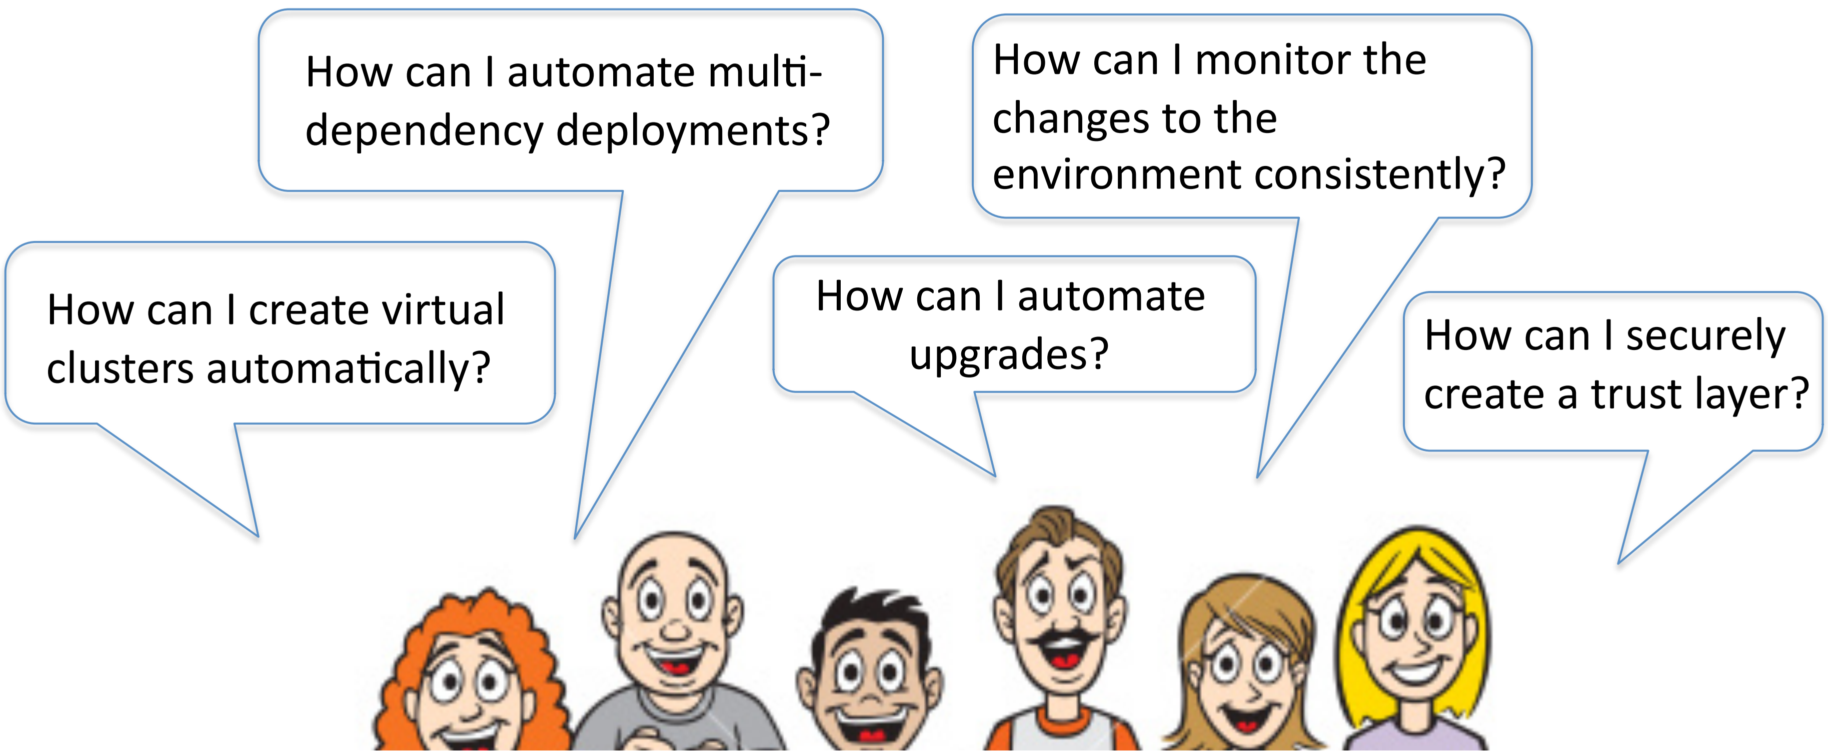

The Player

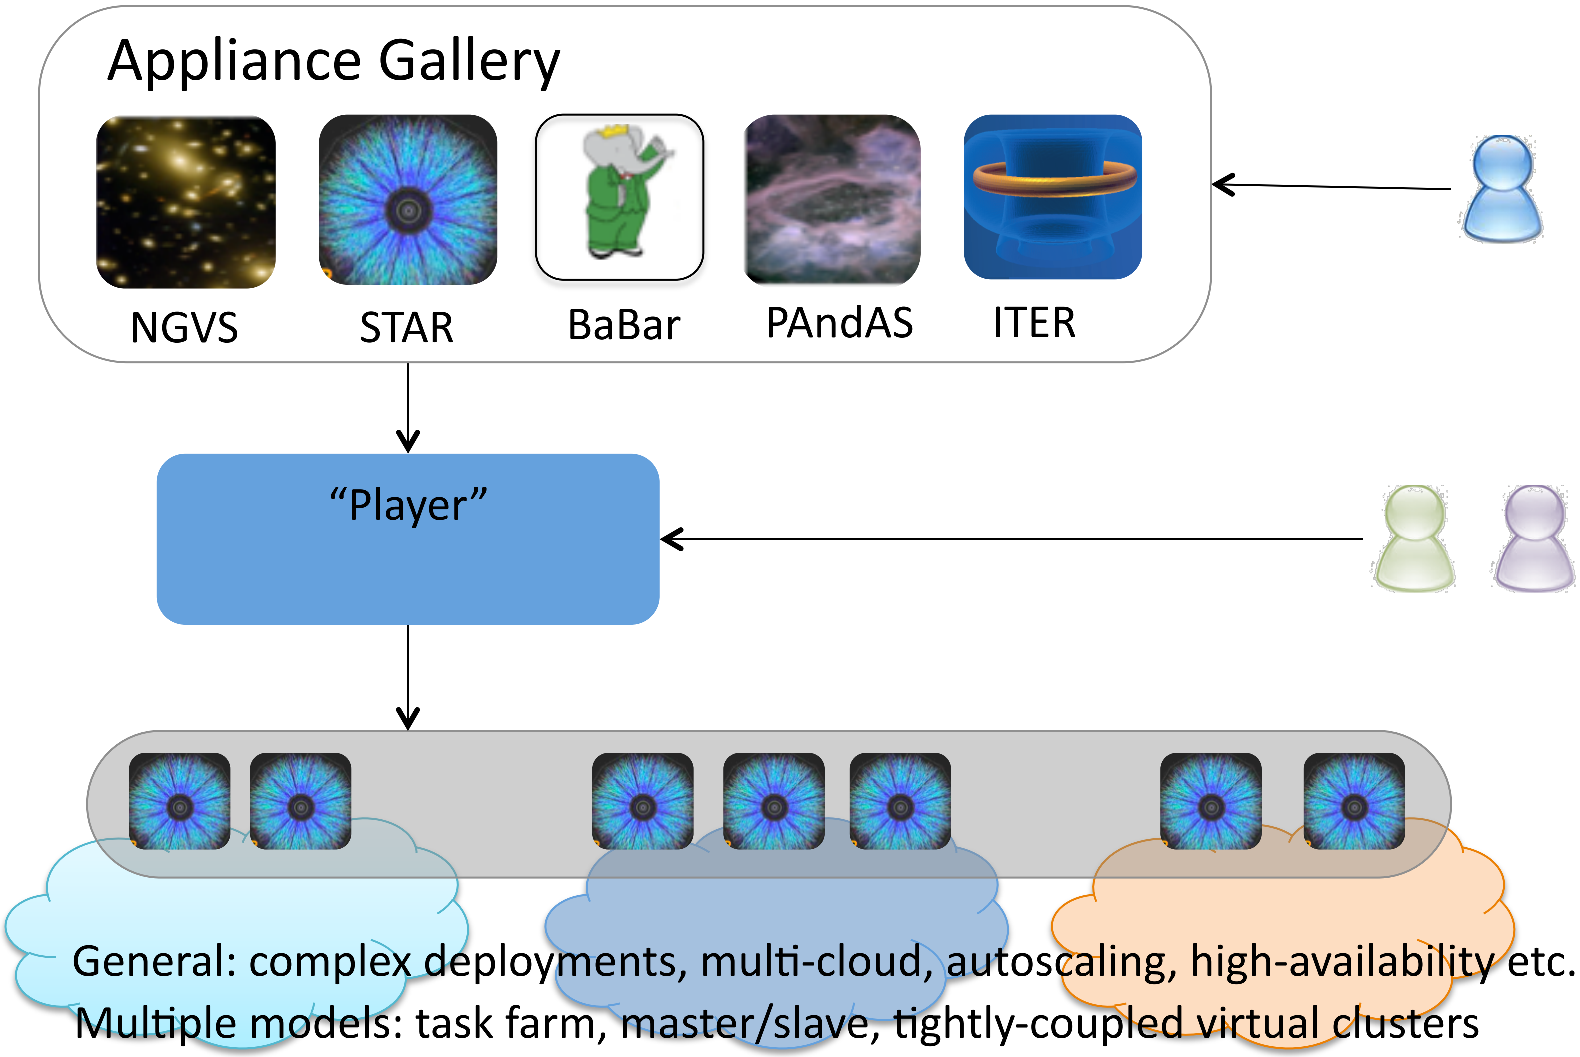

General: complex deployments, multi-cloud, autoscaling, high-availability etc.
Multiple models: task farm, master/slave, tightly-coupled virtual clusters

Multi-Cloud, Availability and Scalability

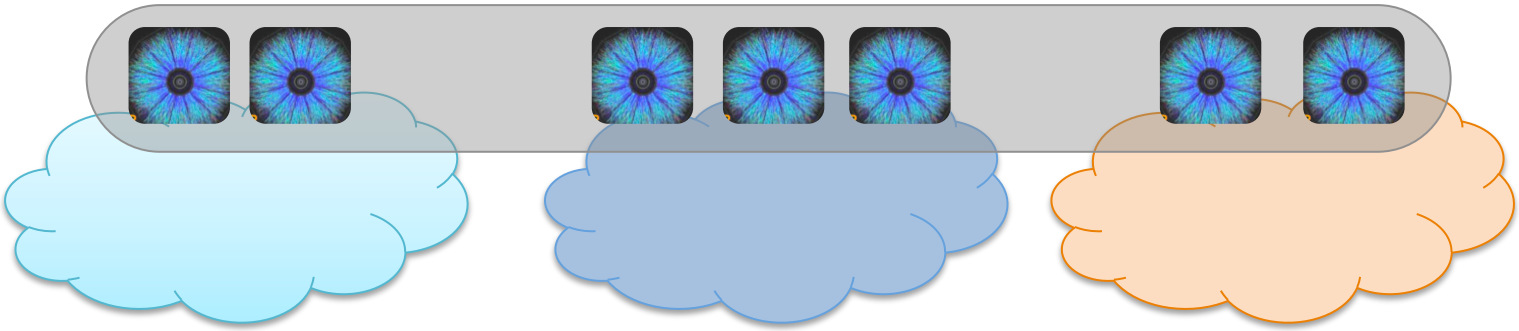

Nimbus Phantom

- Scalability and availability: regulates domain properties (compute, storage) using system and application metrics
- Multi-cloud: works with multiple providers
- Finding the best resource fit
- Extensible monitoring: VM-based (OpenTSDB, traffic sentinel), provider-based (CloudWatch) and custom
- Policy-driven: from pre-defined policies to python programs

Process Management with Swift

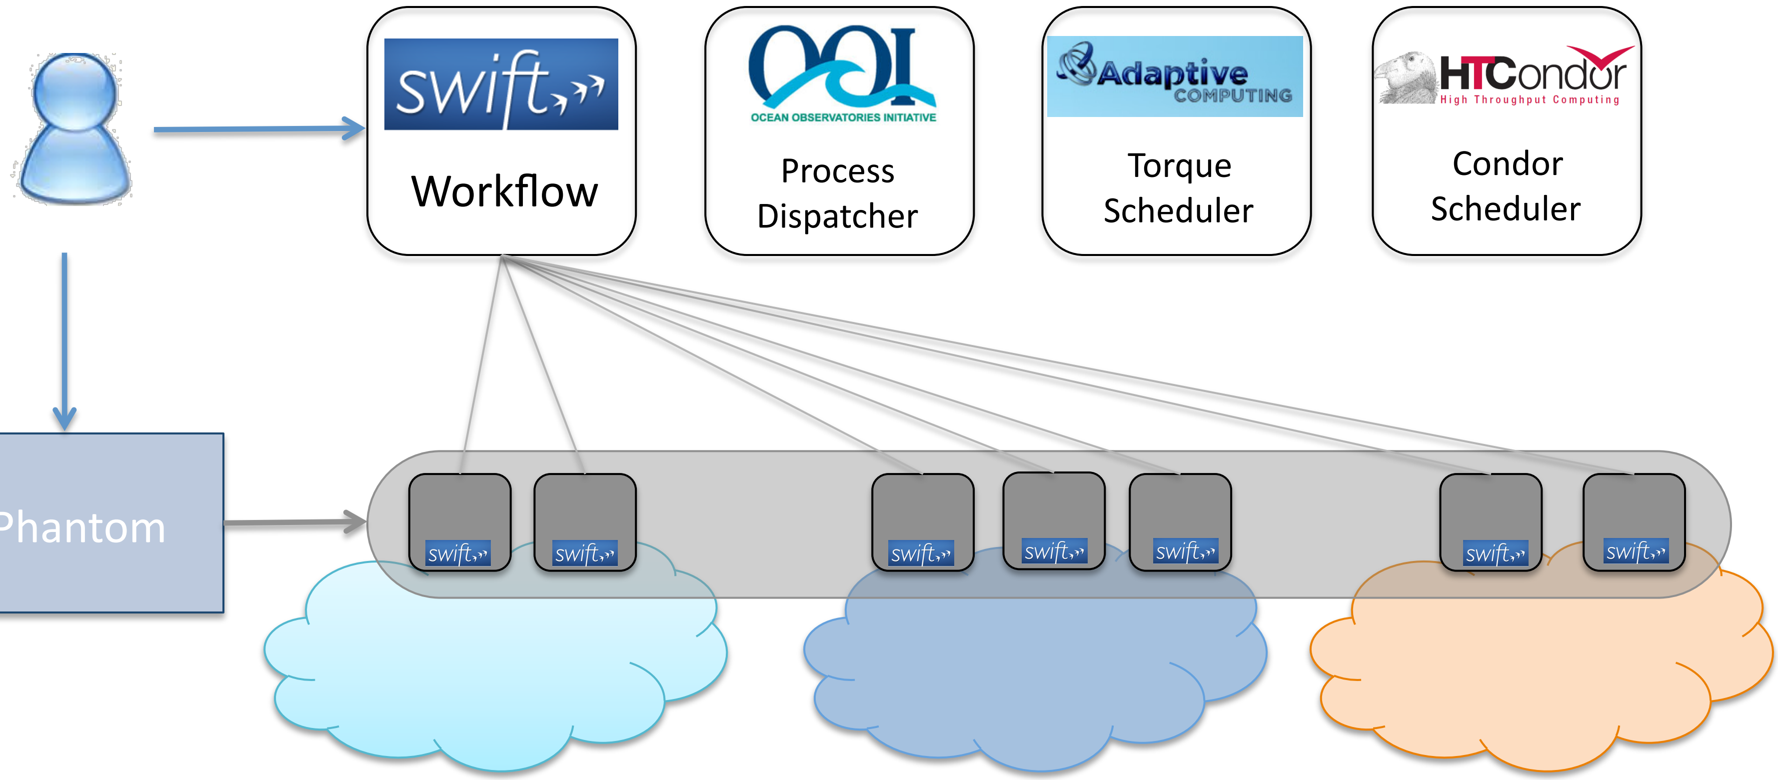

Science Cloud Applications and Application Patterns

Ocean Observatory Initiative

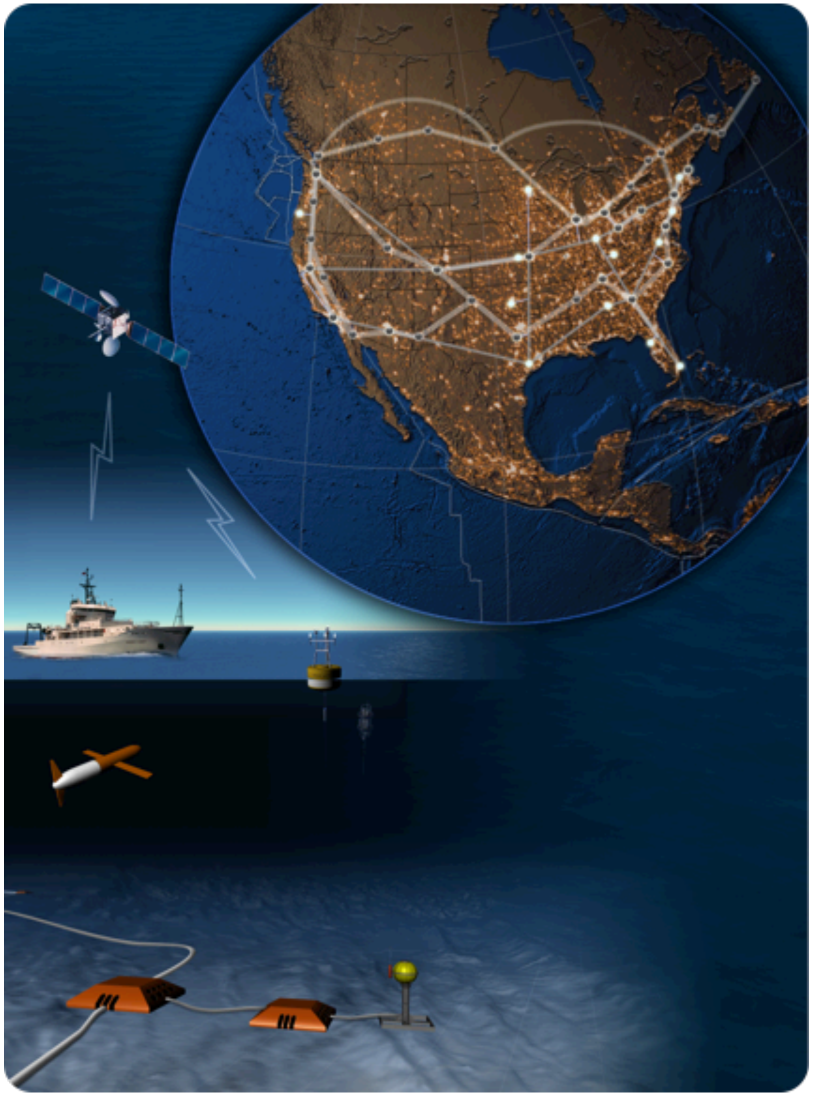

- Towards Observatory Science
- Sensor-driven processing
  - An “always-on” service
  - Real-time event-based data stream processing capabilities
  - Highly volatile need for data distribution and processing
- Nimbus team building platform services for integrated, reactive support for on-demand science
  - High-availability
  - Auto-scaling
- From regional Nimbus clouds to commercial clouds

OCEAN OBSERVATORIES INITIATIVE

Building a Plant Observatory

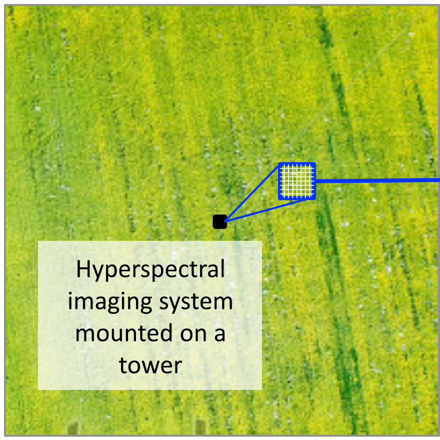

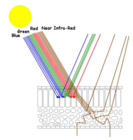

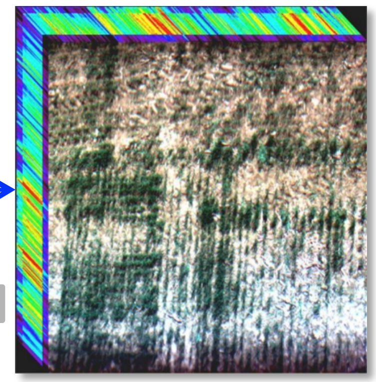

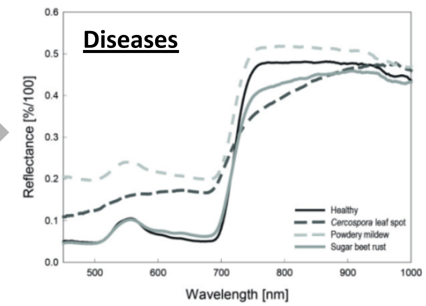

Spectral Indices*
NDVI • SAVI • EVI • NDWI

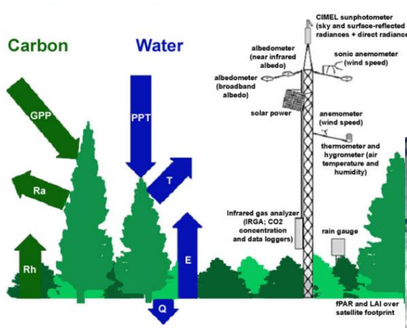

Multi-Temporal Hyperspectral Cubes and Spectral Index Images

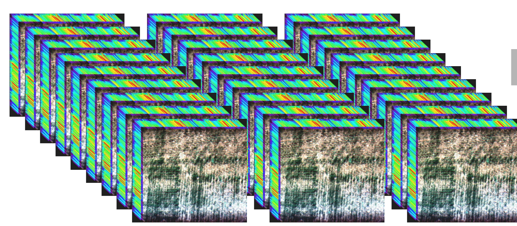

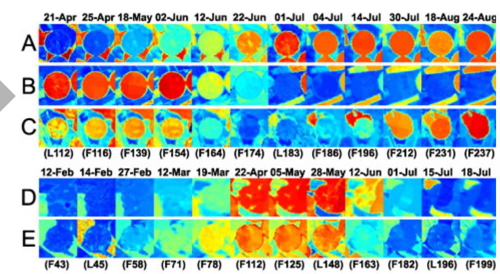

application images courtesy of Yuki Hamada, ANL

Joint project with Pete Beckman, Nicola Ferrier, Yuki Hamada, Rao Kotamarthi, Rajesh Sankaran and others (ANL)

The iPlant Collaborative

- Challenge: to build a lasting, community driven **Cyberinfrastructure** for the ***Grand Challenges*** of **Plant Science**

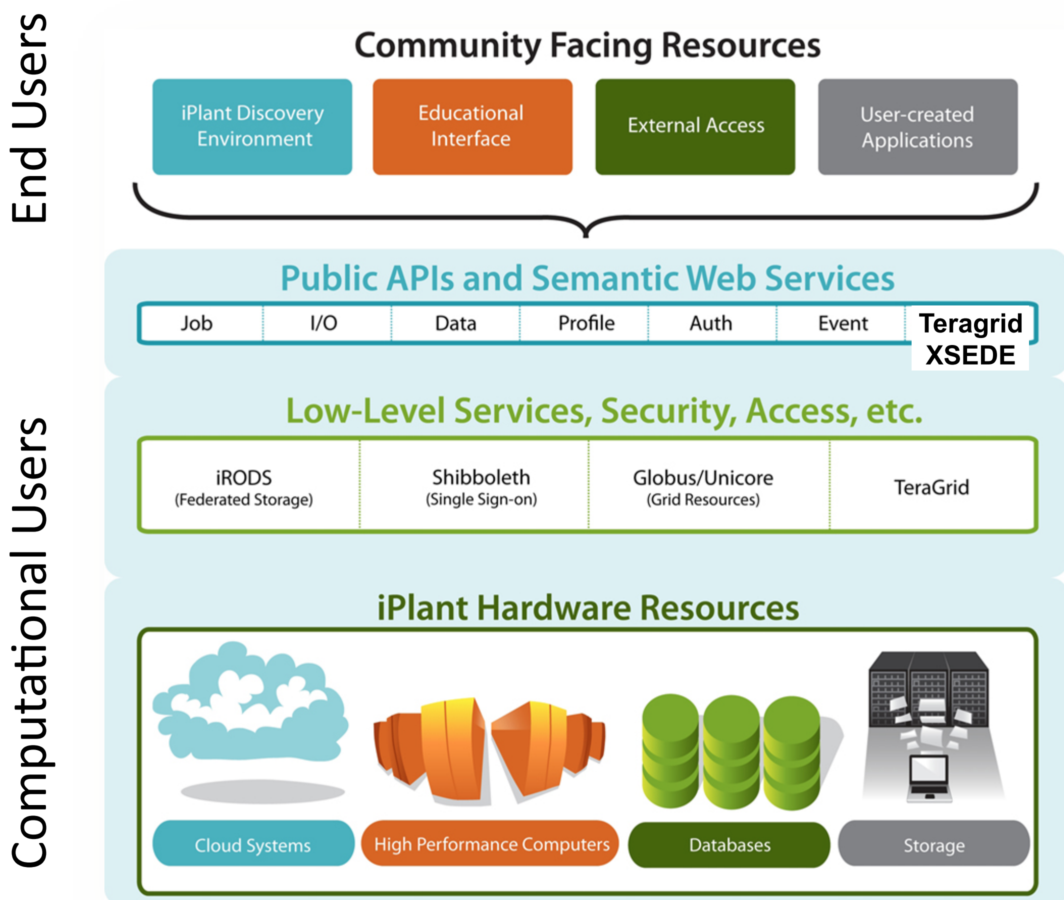

- ✓ Storage
- ✓ Computation
- ✓ Hosting
- ✓ Web Services
- ✓ Scalability

Building a **platform** that can support diverse and constantly evolving needs.

iPlant Cloud Investments

- Atmosphere (pre-configured virtual machines)
  - Private science cloud @ University of Arizona
  - Eucalyptus -> OpenStack
  - Custom UI & API
- Discovery Environment (web-based access to tools and data)
  - Several services running in the cloud
  - DevOps in the cloud
- Foundation API (programmatic access to iPlant)
  - Running on multiple clouds (FutureGrid, HP, Rackspace)
  - Leveraging PaaS and SaaS (HP, Iron.io)
  - DevOps in the cloud
  - Supports execution to arbitrary IaaS providers.

Courtesy of Rion Dooley, TACC

Apache Airavata
Science Gateway Framework

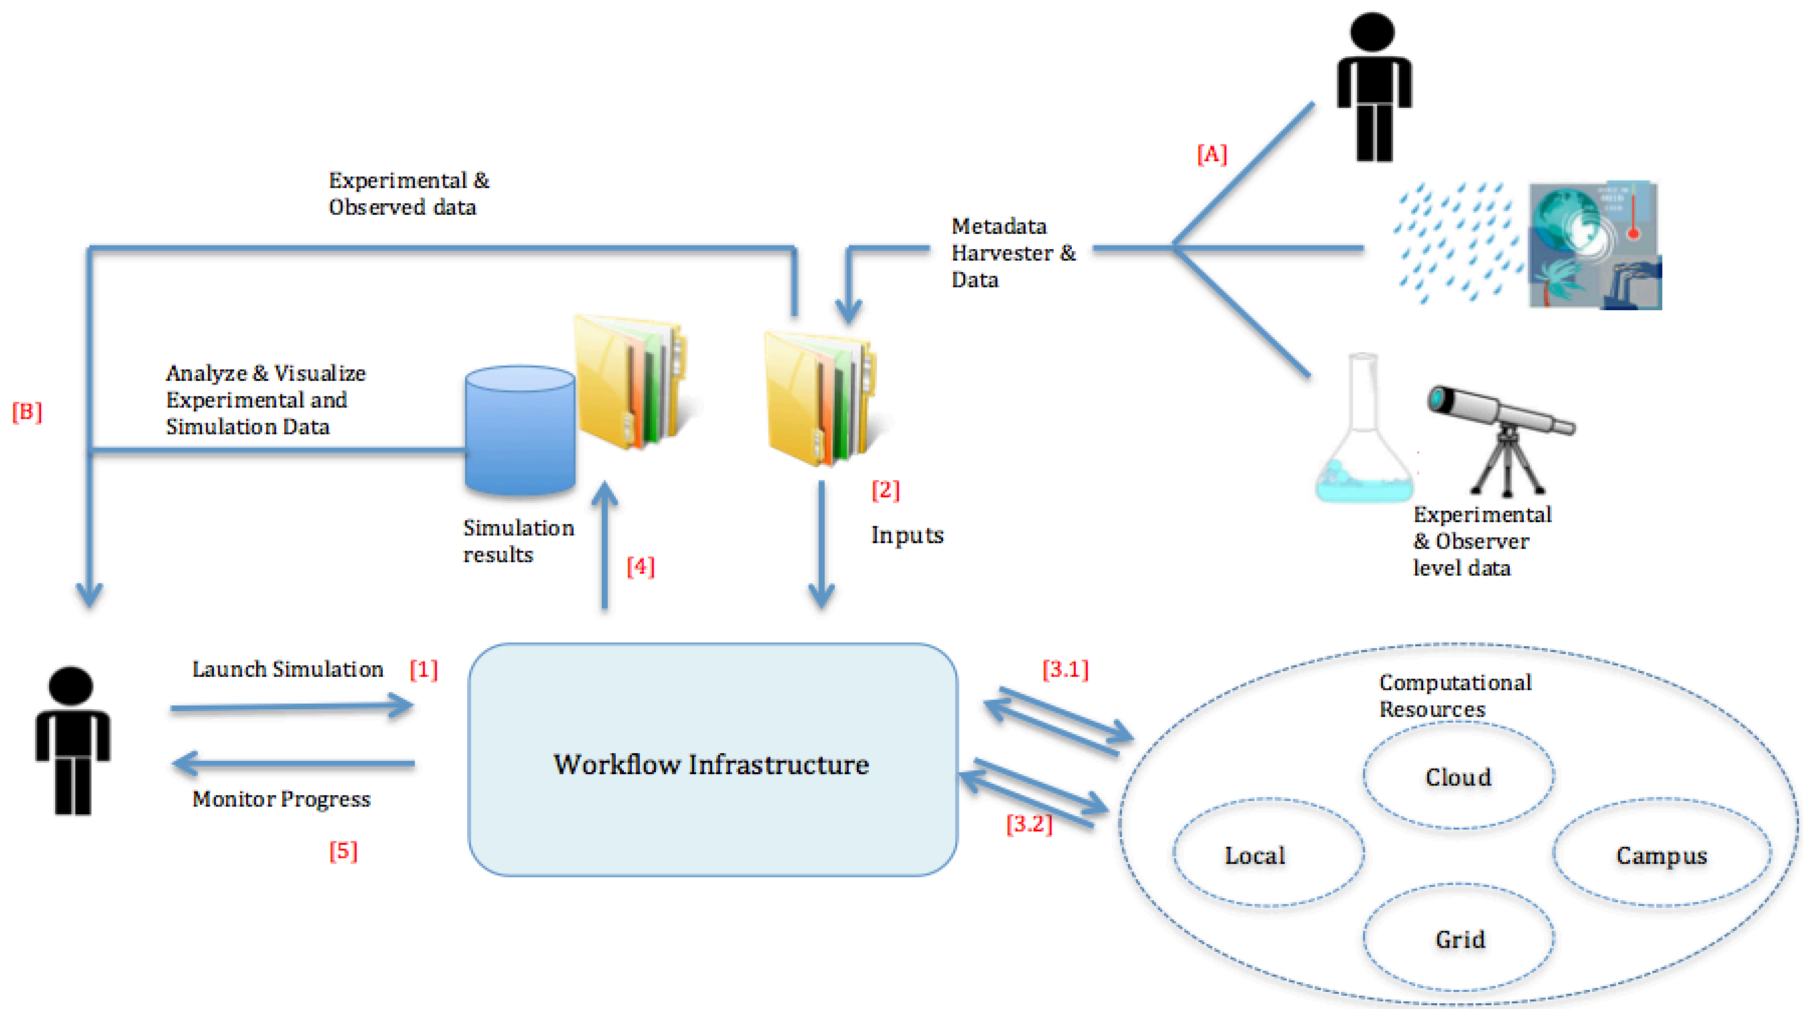

XSEDE 2013, Science Clouds BOF

Courtesy of Suresh Marru, Indiana University

Airavata Cloud Usage

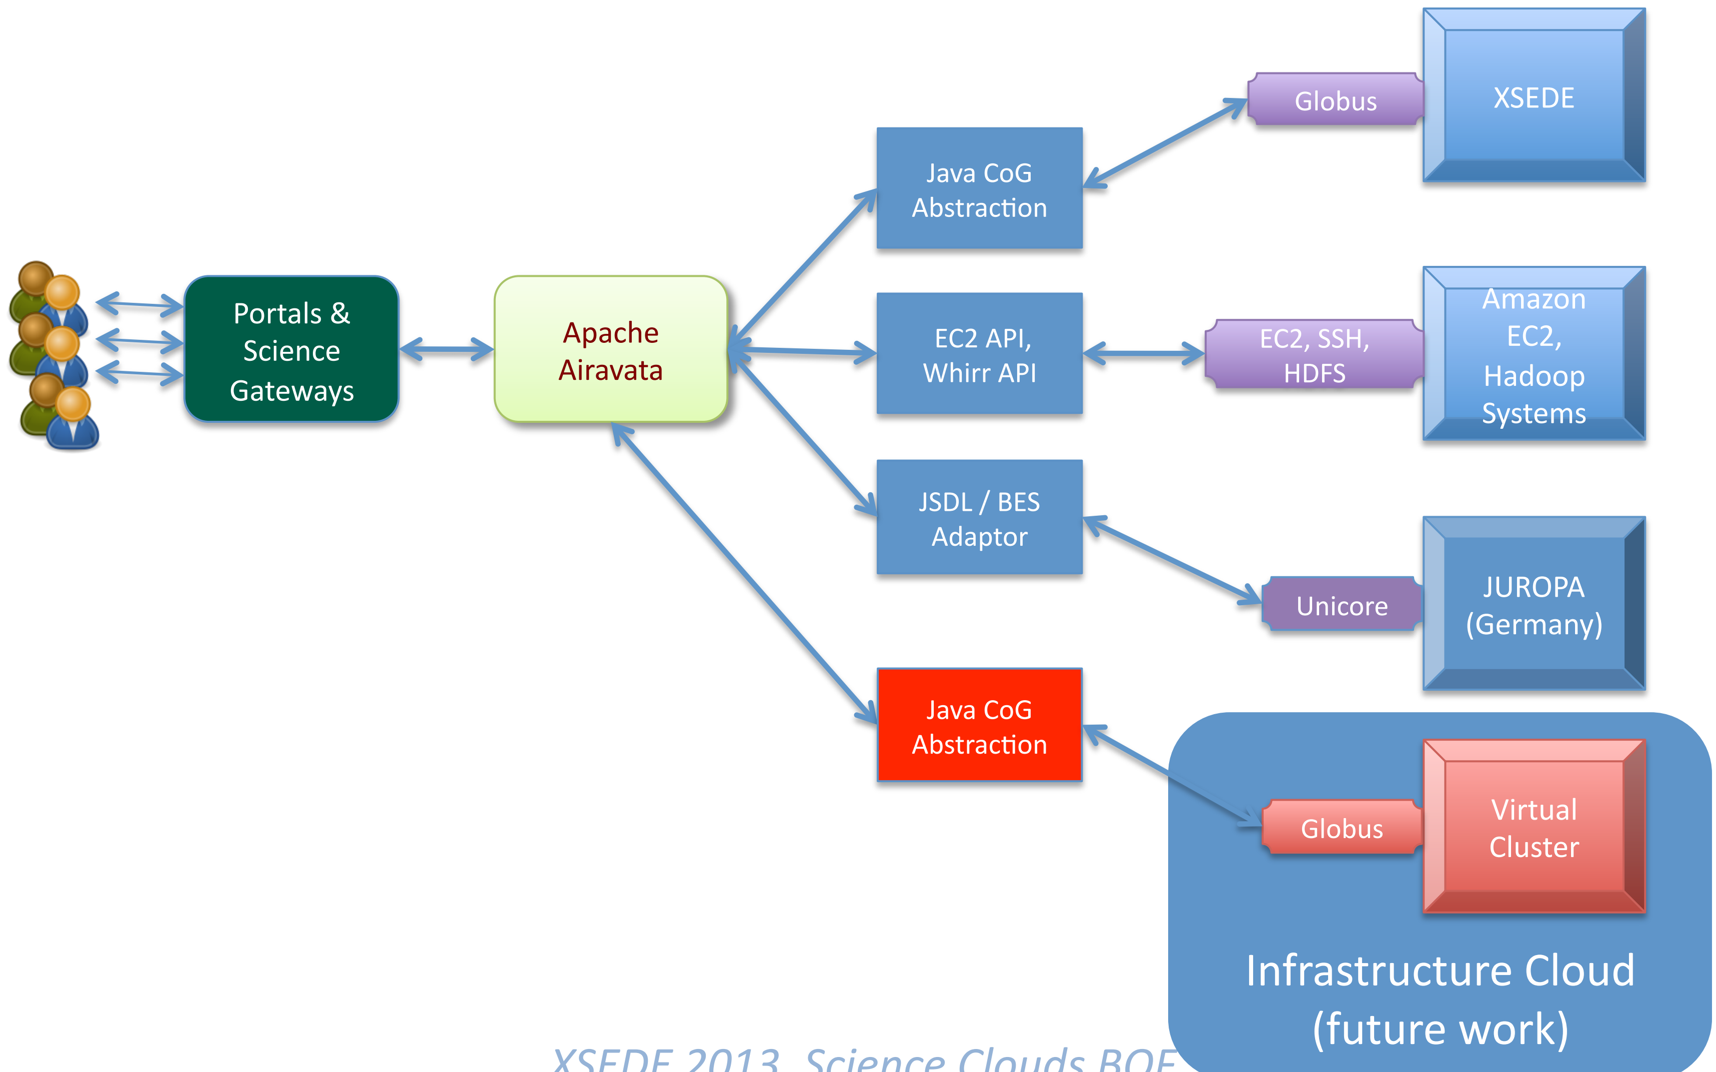

SWATShare –Online SWAT Simulation and Model Sharing

http://water-hub.org

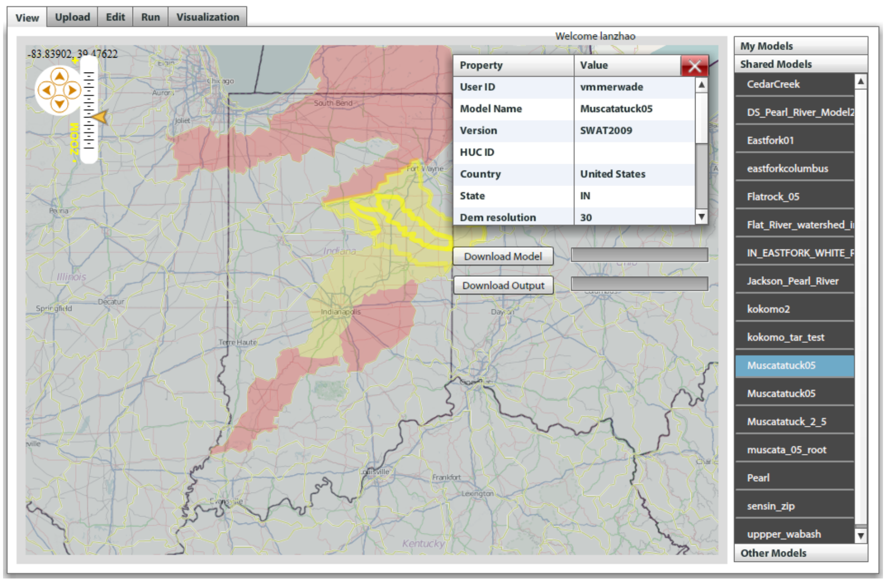

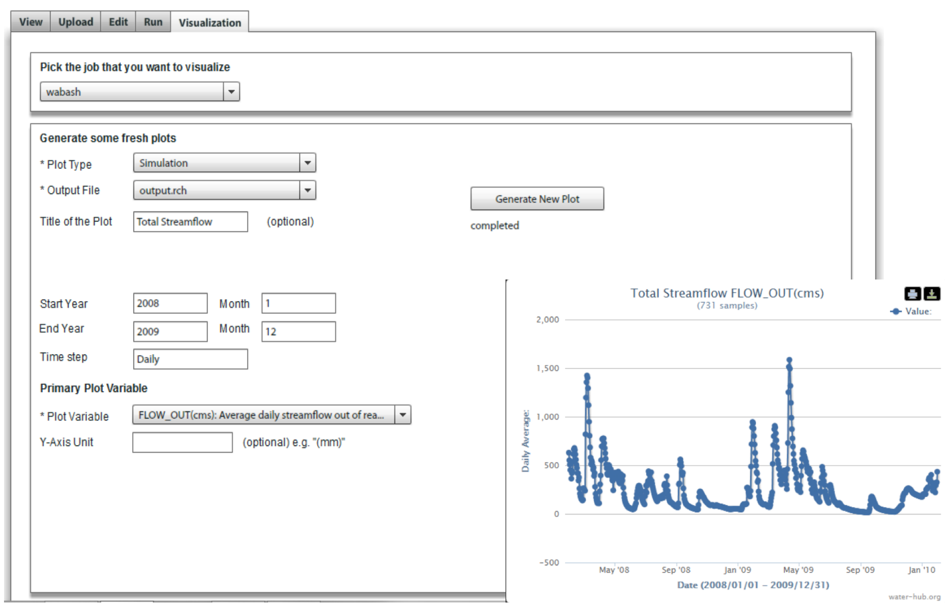

- An XSEDE science gateway
- Available to broad SWAT user community
- Being integrated with Hydroshare project funded by NSF SI2 award
- Used by researchers and classes
  - Simulation of daily streamflow for water resources management in Wabash River Basin
  - Simulation of monthly flows for drought prediction and crop management
  - Simulation of daily, monthly and annual streamflows for nutrient loading
- Technical implementation
- Using different resources based on job nature:
  - Simulation
  - Calibration
  - Sensitivity experiment
- Challenges relevant to cloud computing
  - On-demand scalability
  - Adaptability
    - E.g., migrate from one resource to next

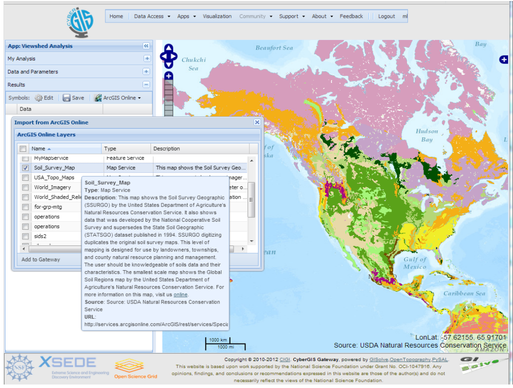

www.cybergis.org

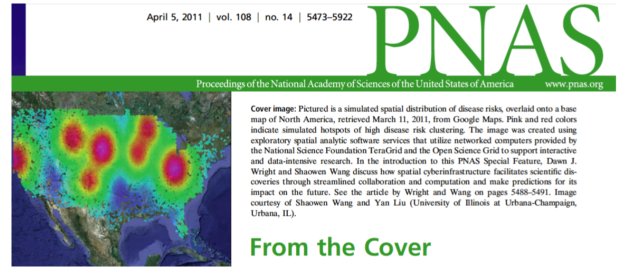
April 5, 2011 | vol. 108 | no. 14 | 5473–5922

PNAS

Proceedings of the National Academy of Sciences of the United States of America www.pnas.org

Cover image: Pictured is a simulated spatial distribution of disease risks, overlaid onto a base map of North America, retrieved March 11, 2011, from Google Maps. Pink and red colors indicate simulated hotspots of high disease risk clustering. The image was created using exploratory spatial analytic software services that utilize networked computers provided by the National Science Foundation TeraGrid and the Open Science Grid to support interactive and data-intensive research. In the introduction to this PNAS Special Feature, Dawn J. Wright and Shaowen Wang discuss how spatial cyberinfrastructure facilitates scientific discoveries through streamlined collaboration and computation and make predictions for its impact on the future. See the article by Wright and Wang on pages 5488–5491. Image courtesy of Shaowen Wang and Yan Liu (University of Illinois at Urbana-Champaign, Urbana, IL).

From the Cover

Spatial

Thinking

Data-Intensive Applications and Sciences

Computational Geospatial Sciences

CyberGIS

Advanced Cyberinfrastructure

Integration and Synthesis

Emergency Management, Geosciences, Health, Sustainability, etc.

Spatial Computational Theories / Methods

GISolve

Clouds, Extreme-Scale Computing, NSF XSEDE, Open Science Grid

Thinking

Computational

Discussion:
How can clouds complement the existing XSEDE resources?

XSEDE Cloud Investigation

- Investigation Team
  - Ian Foster, Steve Tuecke, *ANL/University of Chicago*
  - David Lifka, Susan Mehringer, Paul Redfern, *Cornell University CAC*
  - Craig Stewart, *Indiana University*
  - Manish Parashar, *Rutgers University*
- Cloud Survey Motivation
  - The goal of XSEDE is to enhance research productivity
  - XSEDE must embrace cloud
  - XSEDE must have a clear understanding of how researchers using the cloud today and why
  - Based on this information XSEDE plans to integrate cloud services into its portfolio to support use cases that are not well served by its current resource offerings
- Survey Status *(www.xsede.org/cloudsurvey - closed end of March; report by July 1)*
  - ***80 cloud projects from around globe***, broad participation (21 disciplines + HASS), extensive technical data (19 categories), user perspectives, e.g., cloud benefits/ challenges
  - Focused exclusively on use of cloud for ***research and education***

User Identified Benefits

1. Pay as You Go
2. Lower Costs
3. Compute Elasticity
4. Data Elasticity
5. Software as a Service
6. Education as a Service
7. Broader Use
8. Scientific Workflows
9. Rapid Prototyping
10. Data Analysis

User Identified Challenges

1. Learning Curve
2. Virtual Machine
3. Bandwidth
4. Memory Limits
5. Databases
6. Interoperability
7. Security
8. Data Movement
9. Storage
10. Cost/Funding

Questions

- What other benefits/challenges do you see?
- Ideally, what services would you like to see available?
- What can XSEDE do to help cloud-challenged users?