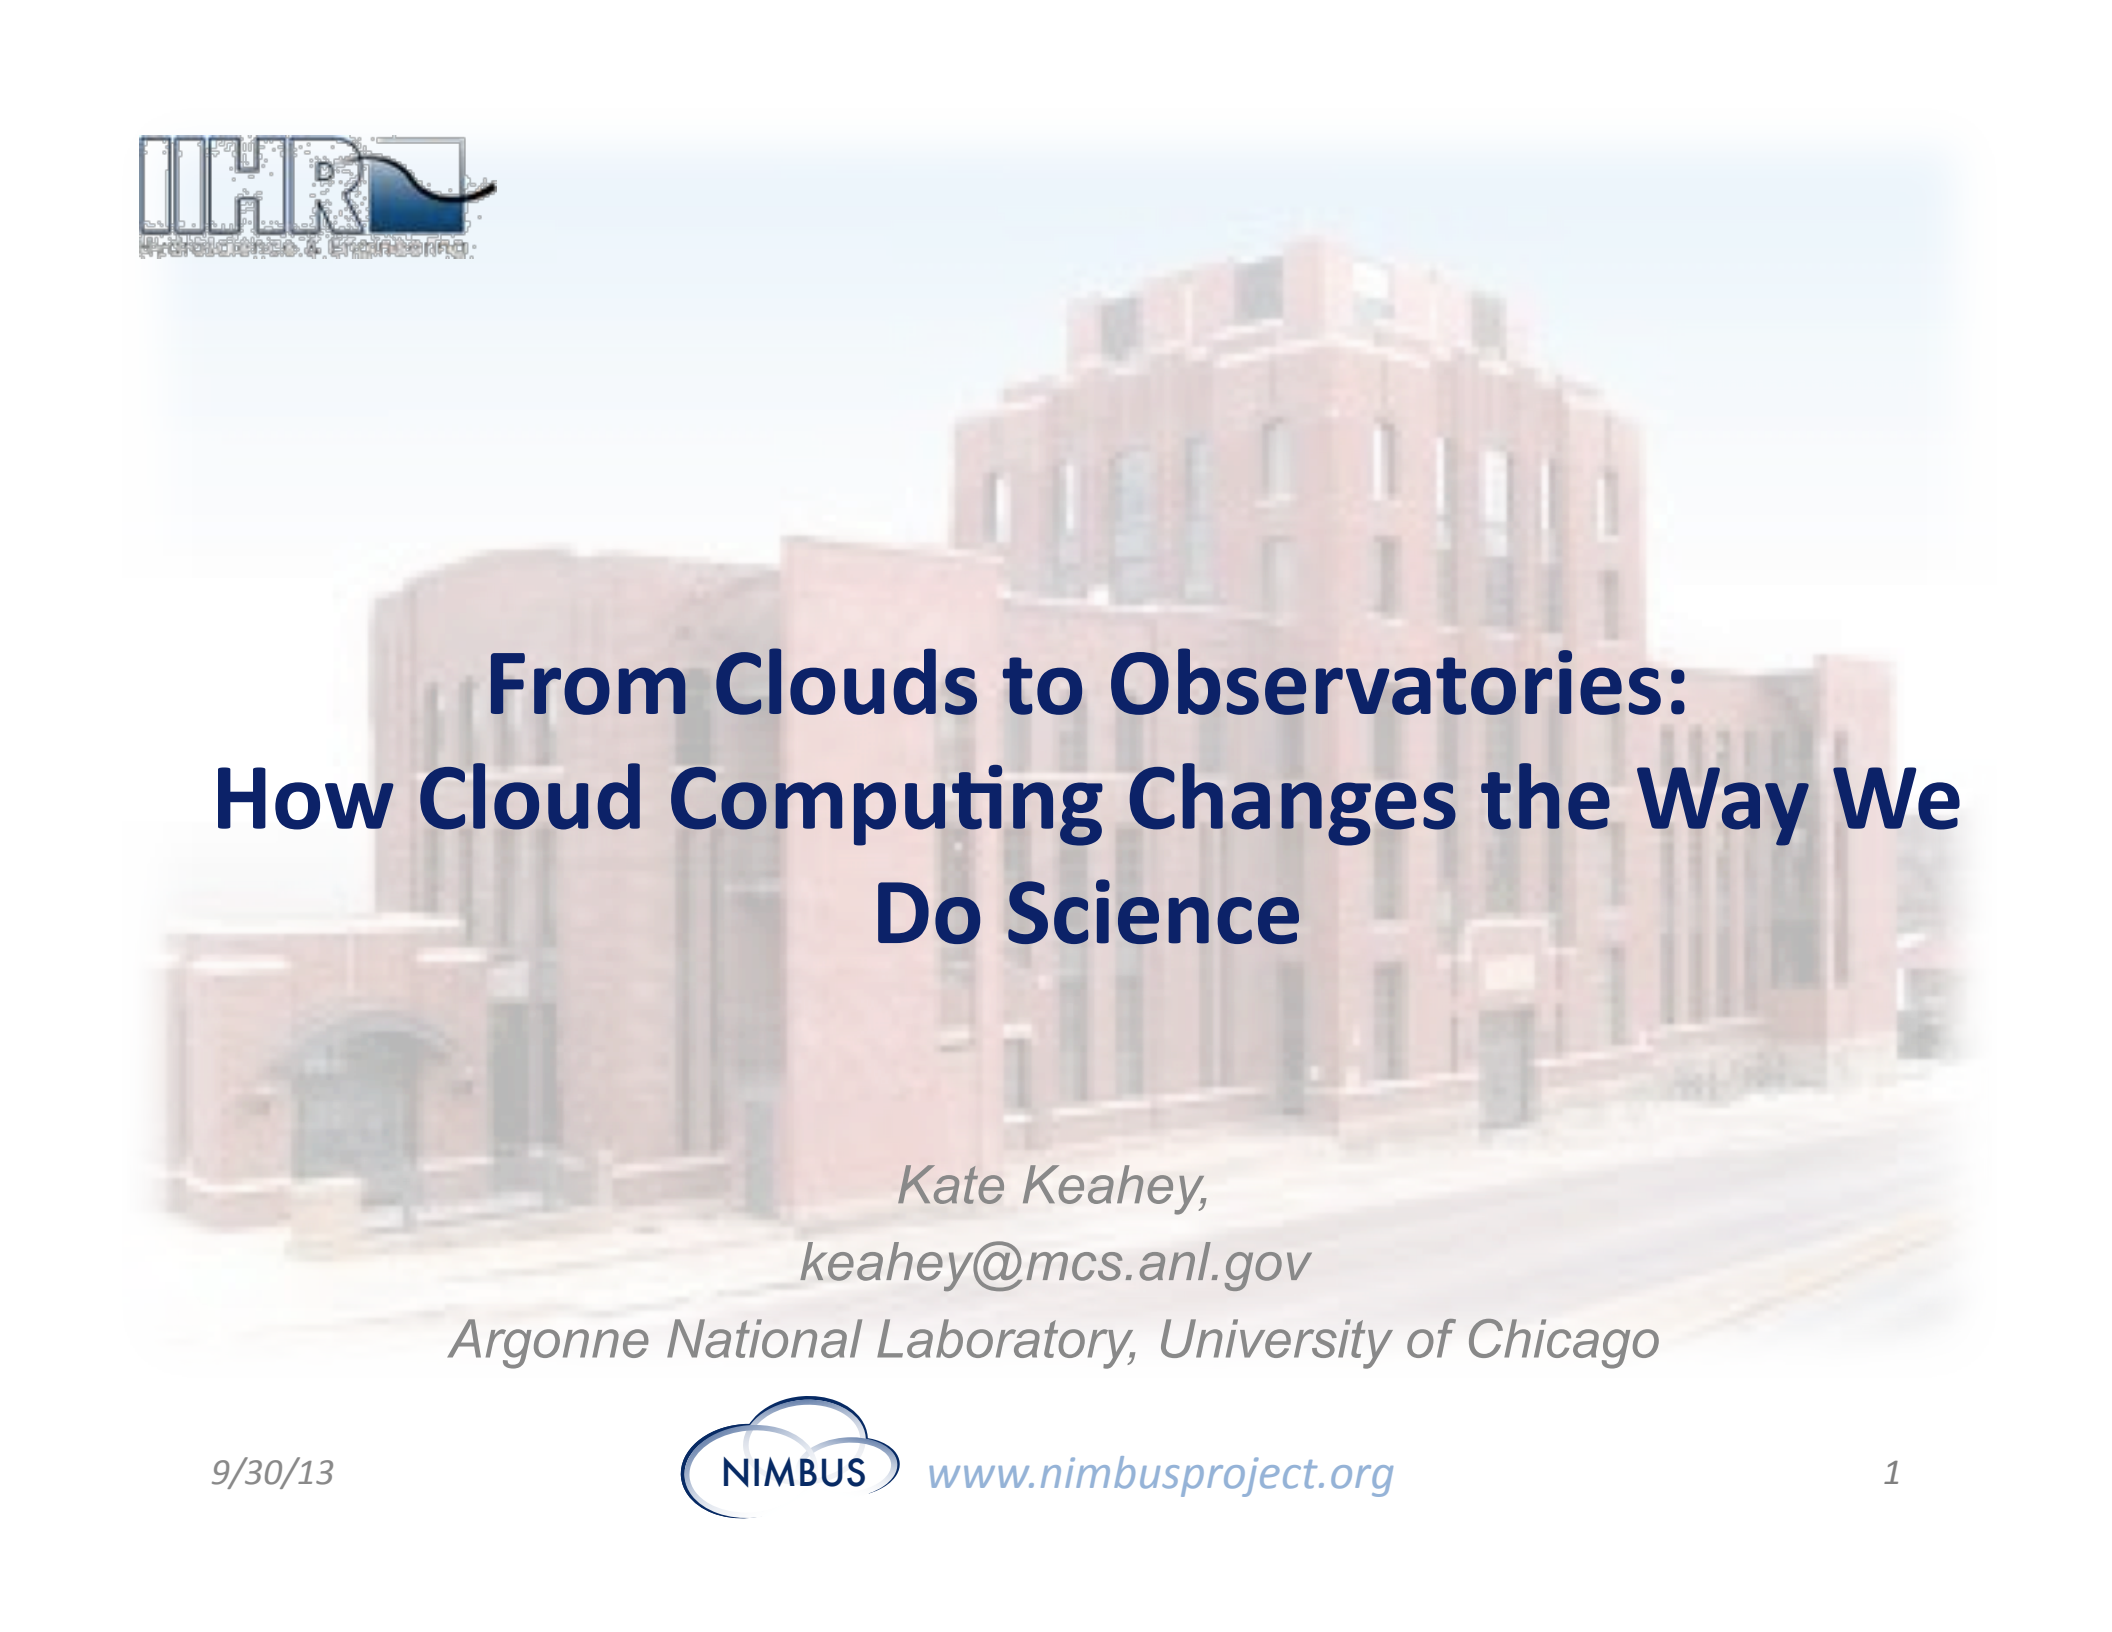

From Clouds to Observatories: How Cloud Computing Changes the Way We Do Science

Kate Keahey,
keahey@mcs.anl.gov
Argonne National Laboratory, University of Chicago

Cloud Computing: Looking Back

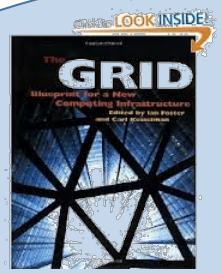

Enabler

Using remote resources

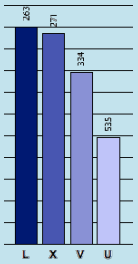

Enabler

Environment ownership
Isolation

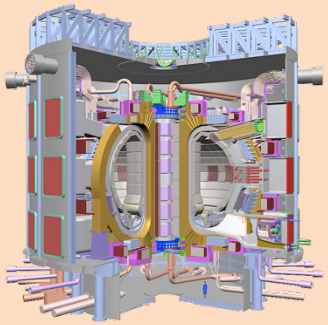

The Need

Environment ownership
On-demand processing

Infrastructure Clouds

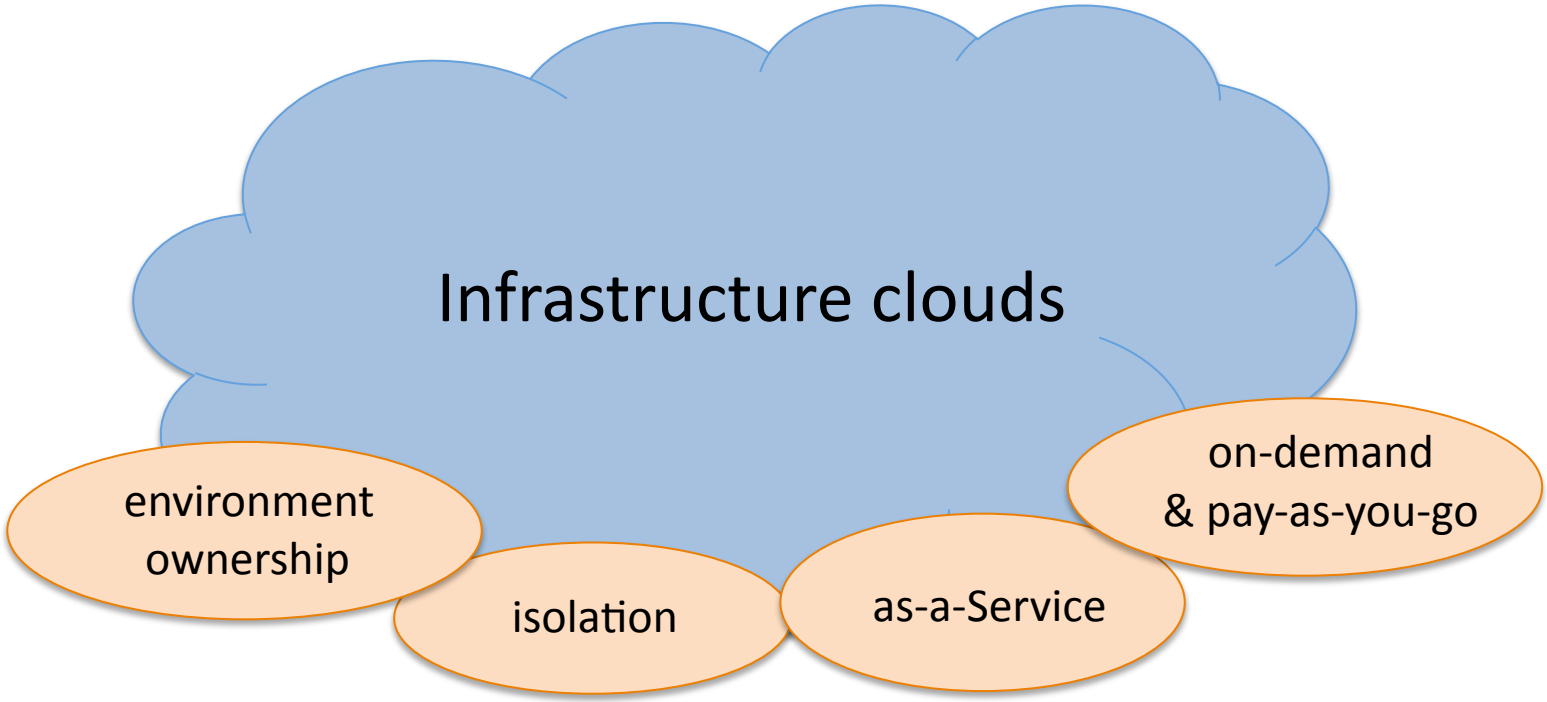

Cloud Computing: Looking Forward

Cloud computing

On-demand, elastic outsourcing

Continuous Data Acquisition

Cheap and reliable "sensing devices"

Opportunity

Virtual Observatories

Ocean Observatory Initiative

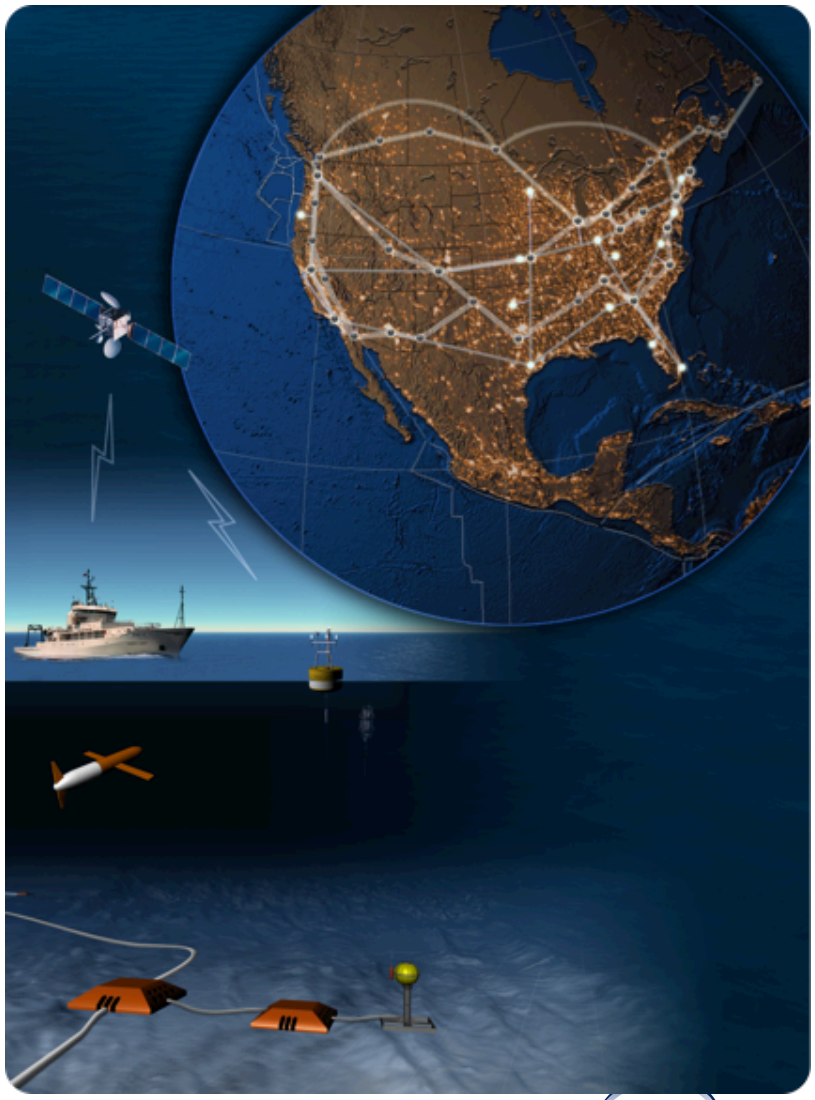

- Towards Observatory Science
- Sensor-driven processing
  - An “always-on” service
  - Real-time event-based data stream processing capabilities
  - Highly volatile need for data distribution and processing
- Nimbus team building platform services for integrated, repeatable support for on-demand science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

The Forest Project: a Plant Observatory

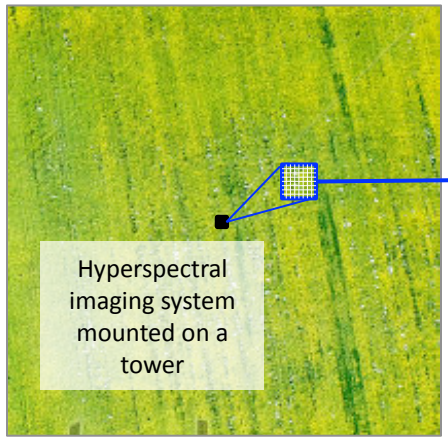

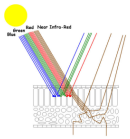

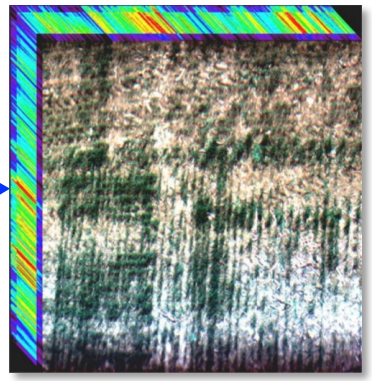

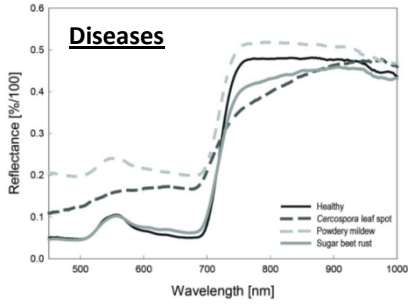

Spectral Indices*
NDVI • SAVI • EVI • NDWI

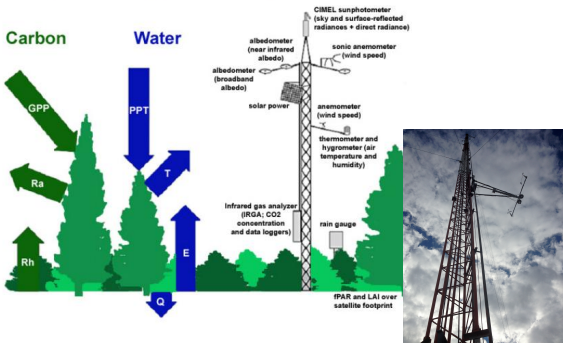

Multi-Temporal Hyperspectral Cubes and Spectral Index Images

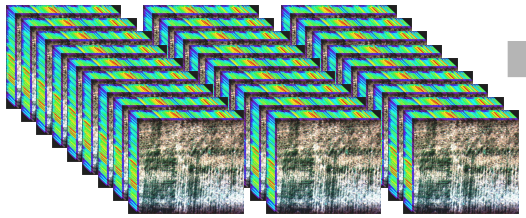

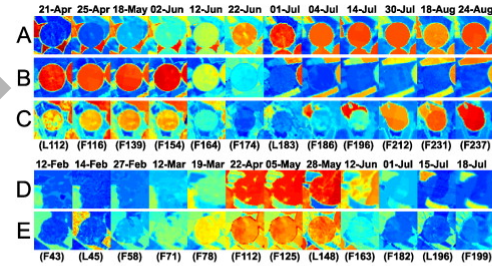

application images courtesy of Yuki Hamada, ANL

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

The Forest Project: Data Stream Aggregation

Aggregation of different data sources
(Temperature)

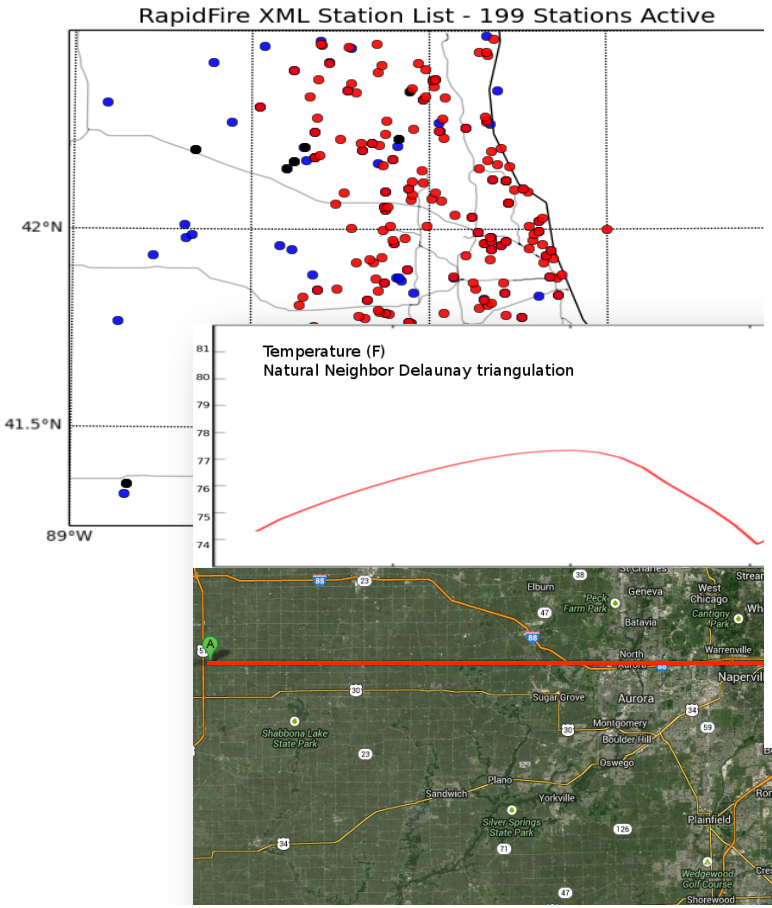

Visualization of time-series data
(Temperature change July 20, 2013 5:00 am-11:55 pm)

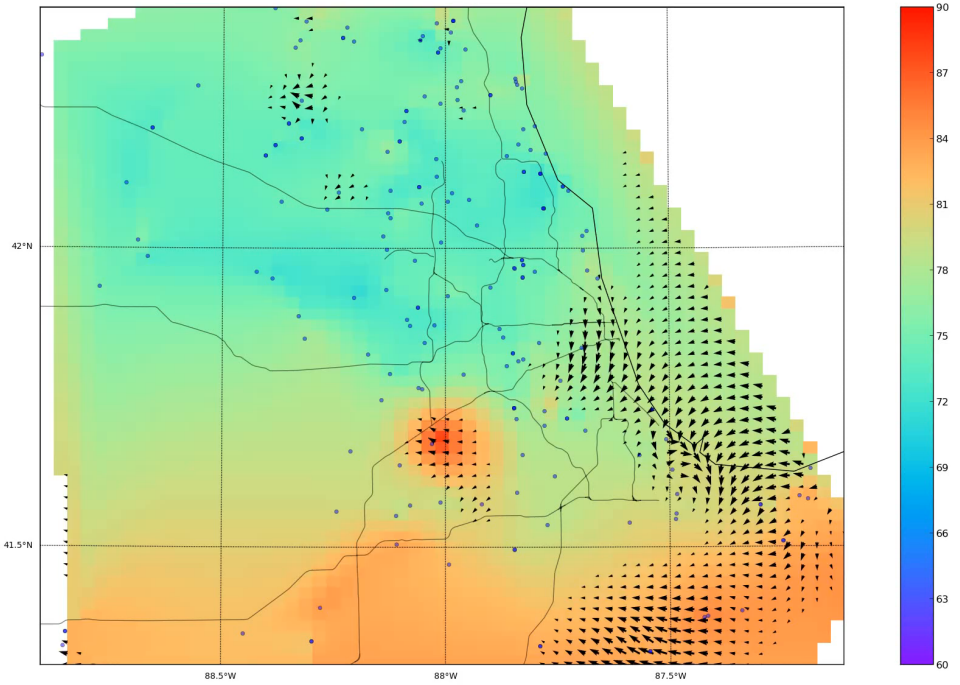

Spatially explicit analysis & semi-continuous representation
(Temperature (F) profile of the transect from A to B)

Social Media Observatory

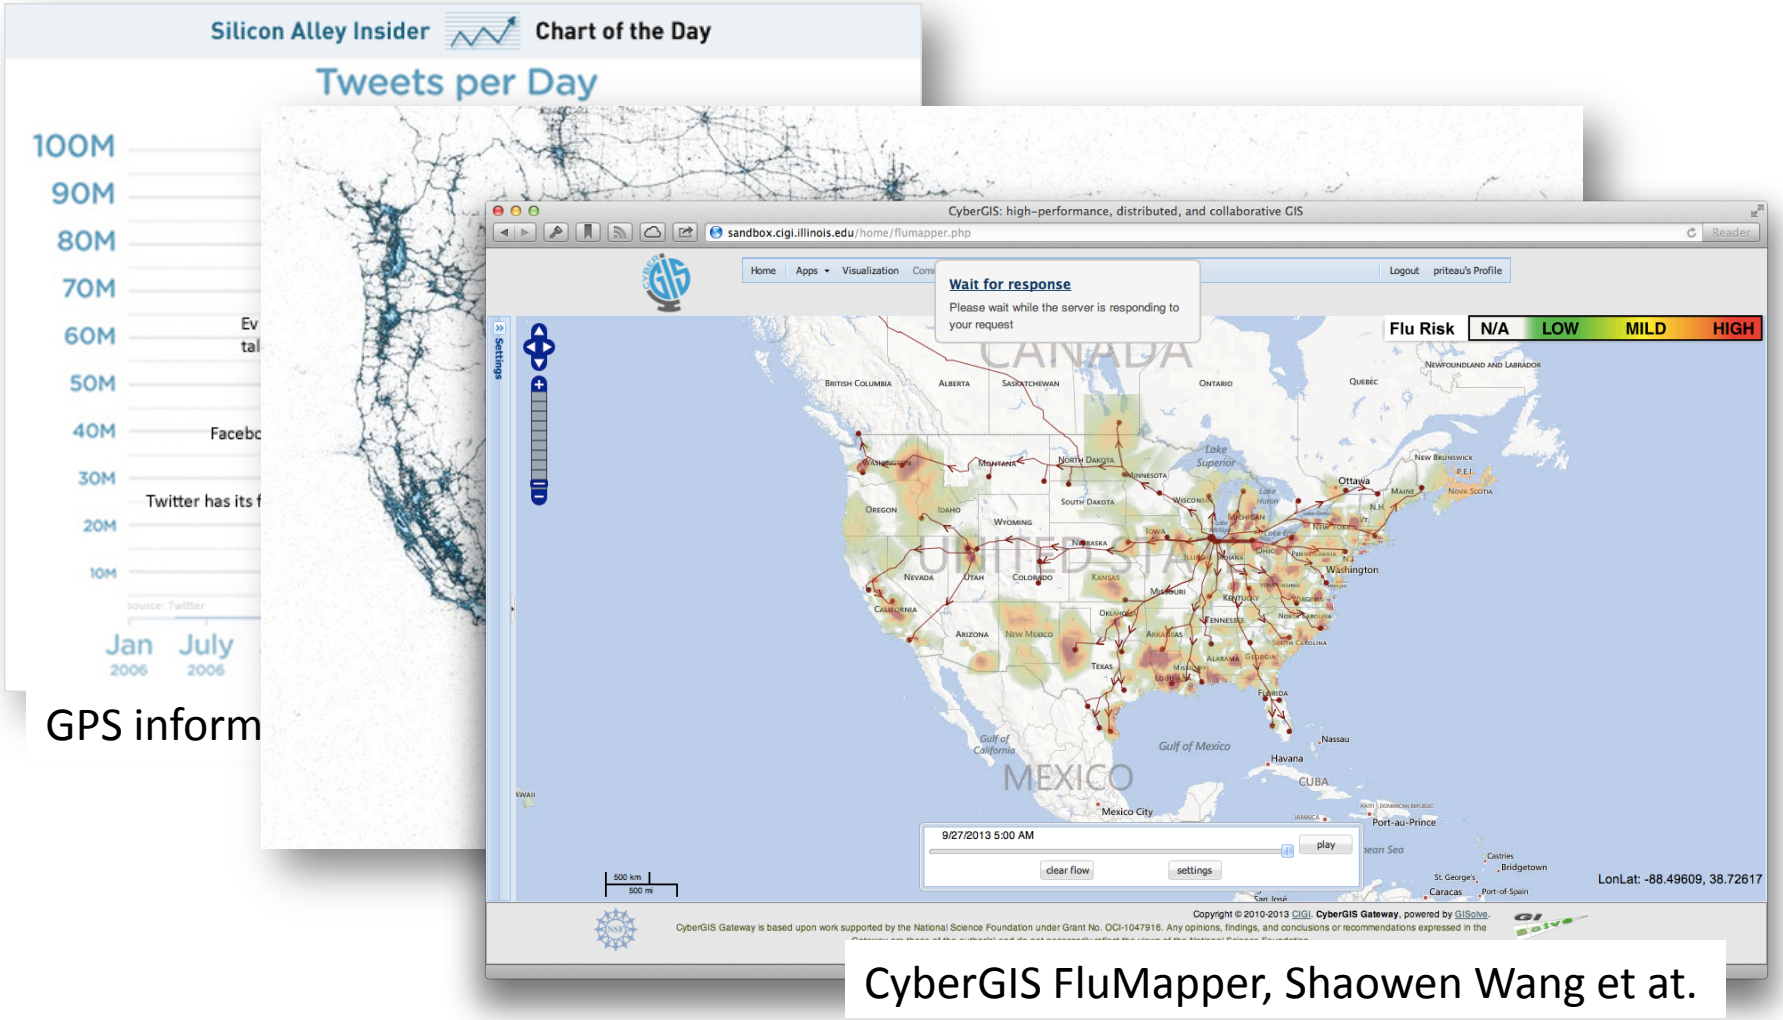

GPS inform

CyberGIS FluMapper, Shaowen Wang et at.

Virtual Observatories

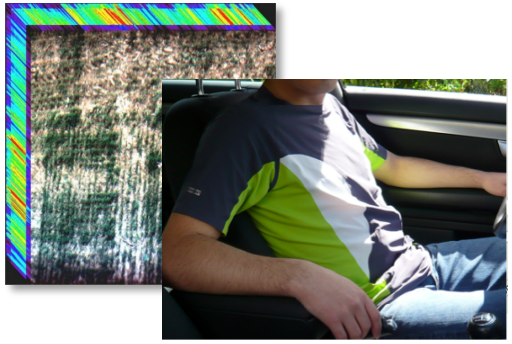

Sensor-based processing

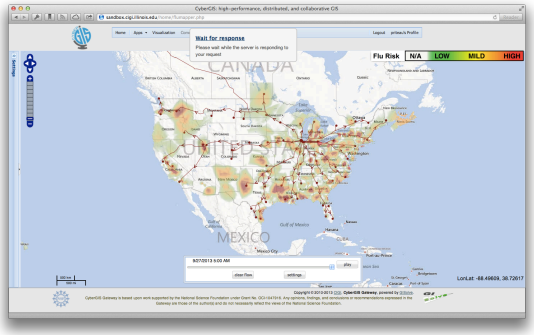

Social media

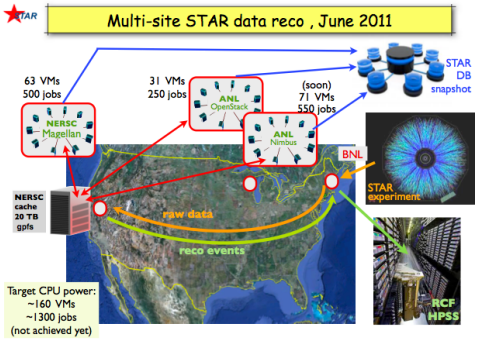

Support for experiments

Mobile devices

From data to insight:
on-demand access
to data products and processing

Citizen science

Platform Goals

- Integration platform for real-time data streams
- Publishing platform for innovative algorithms and methods
  - Well-defined interfaces for publishing operations
  - Leverage cloud technologies to package methods so that they can be run automatically
- Dynamic publish/subscribe network applying relevant methods to the data
  - Users can add or delete new processing elements as situation warrants
- Real-time response, scalable and highly available stream processing
  - Providing near real-time response time as needed via leveraging on-demand cloud resources
  - Scaling to numbers of data, simulations, and users
  - Leveraging automation for restart and HA

Architecture

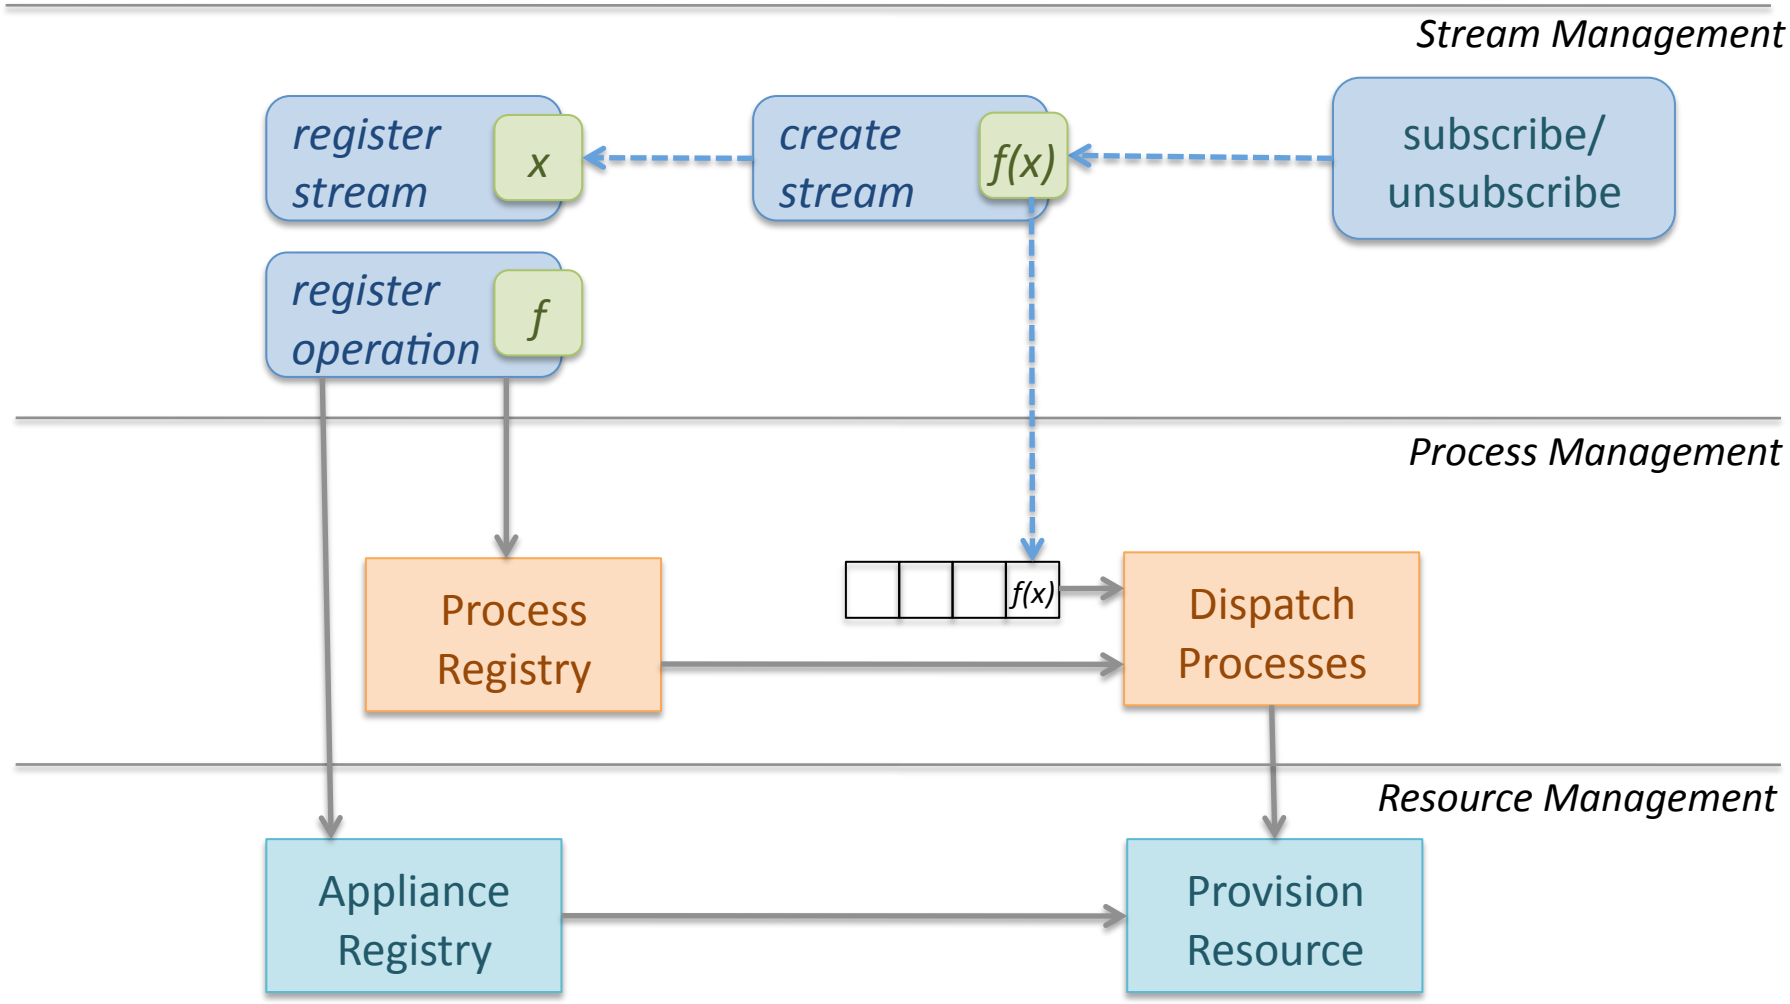

Resource Configuration and Management: the Nimbus Platform

Appliances: Configuration

- 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

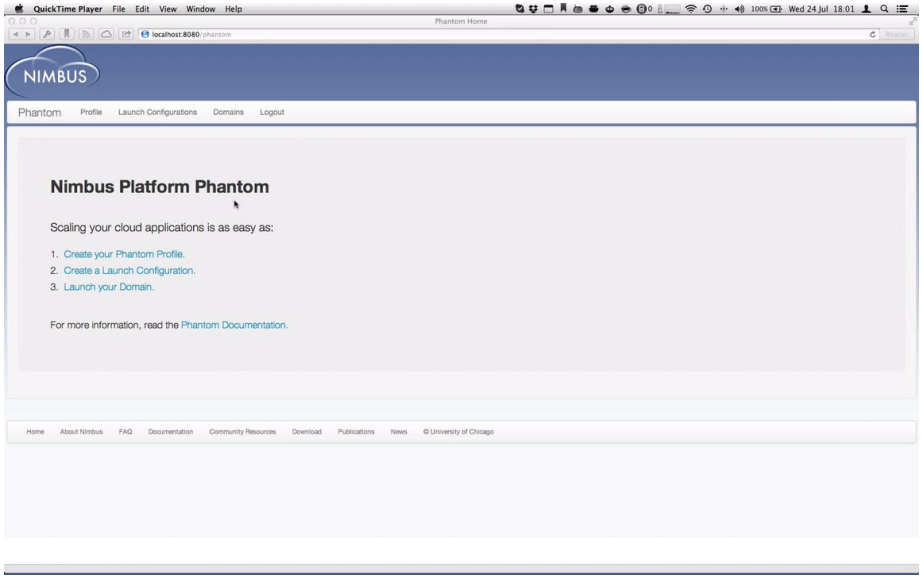

Appliances: Contextualization

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

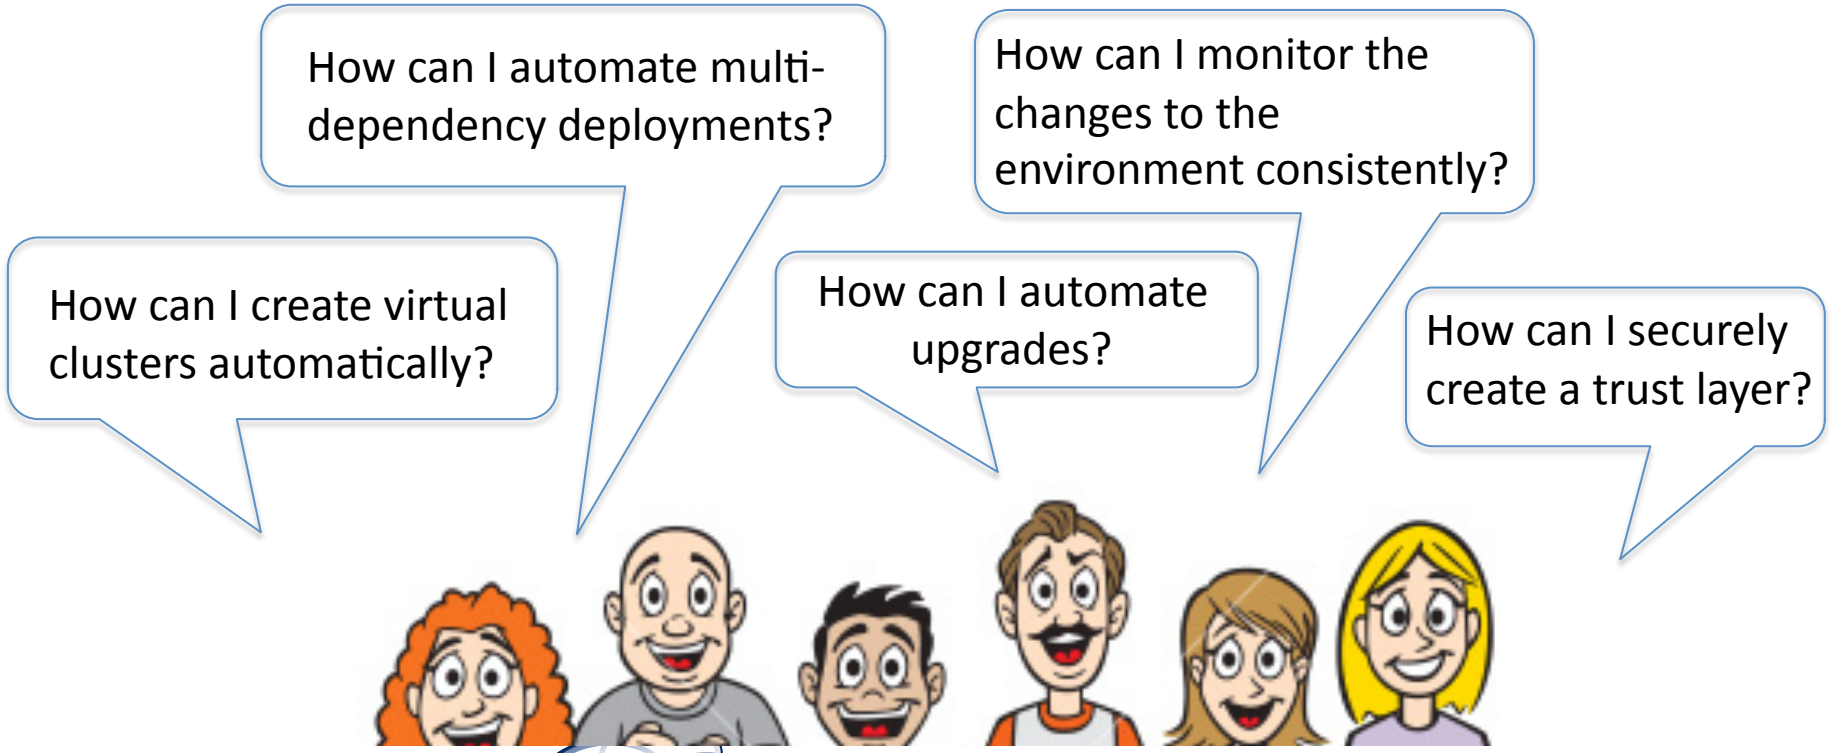

Nimbus Phantom: "The SmartPlayer"

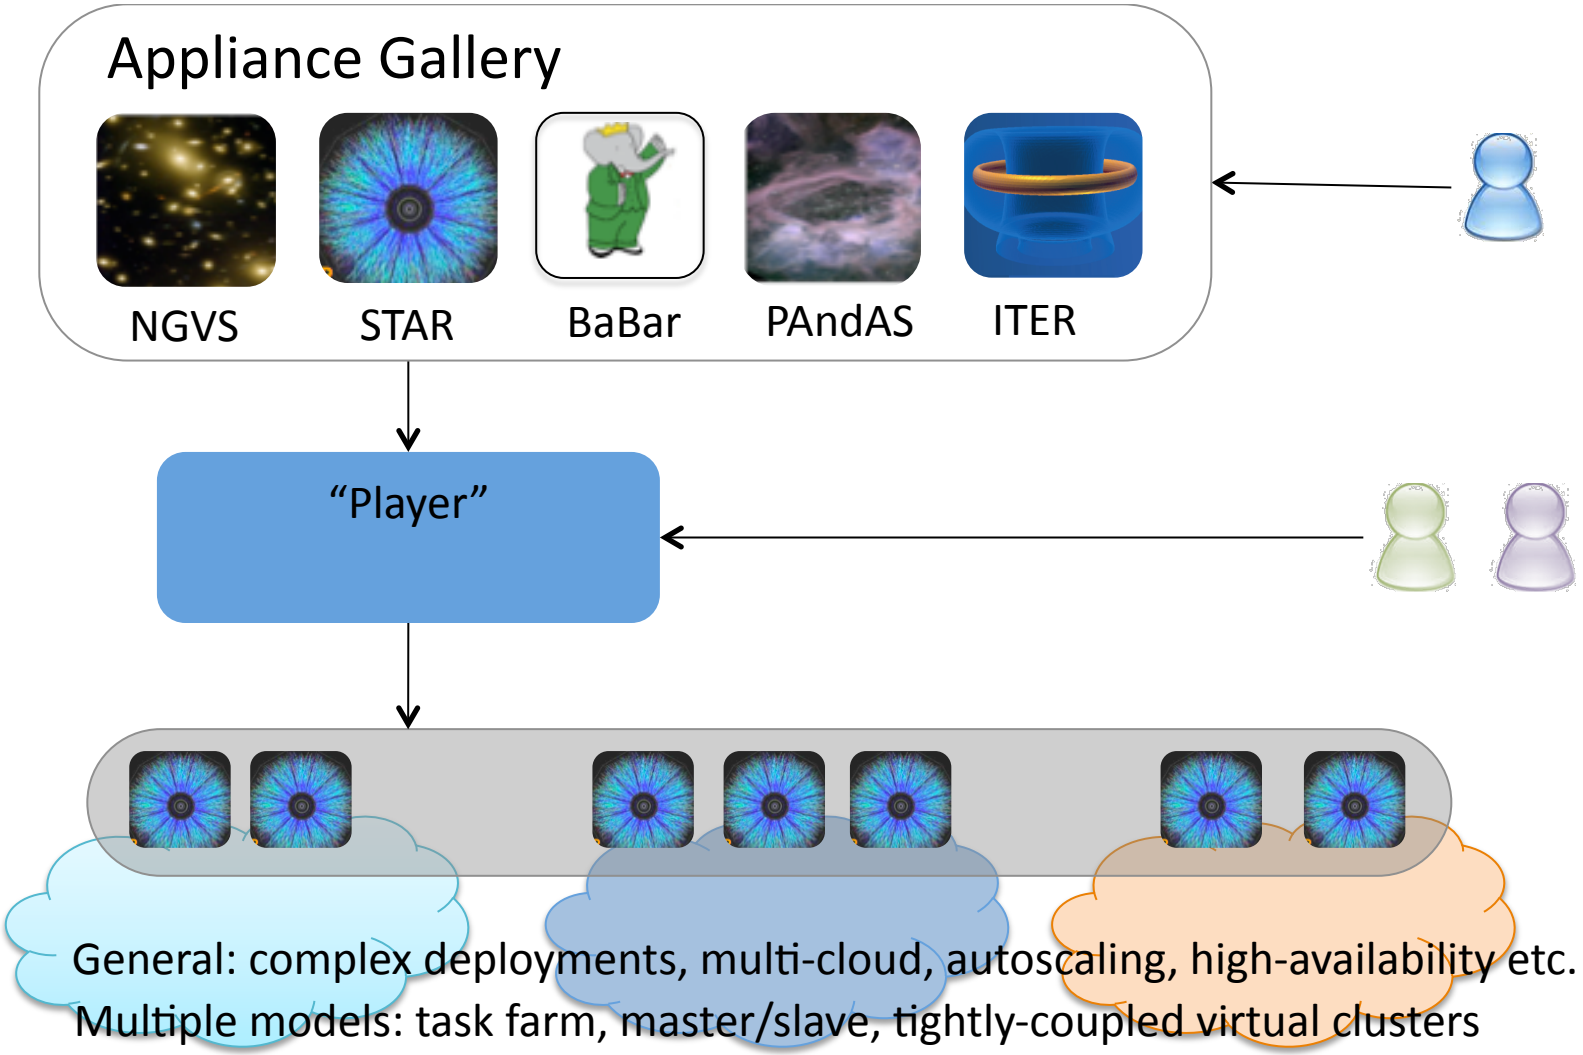

Resource Management Goals

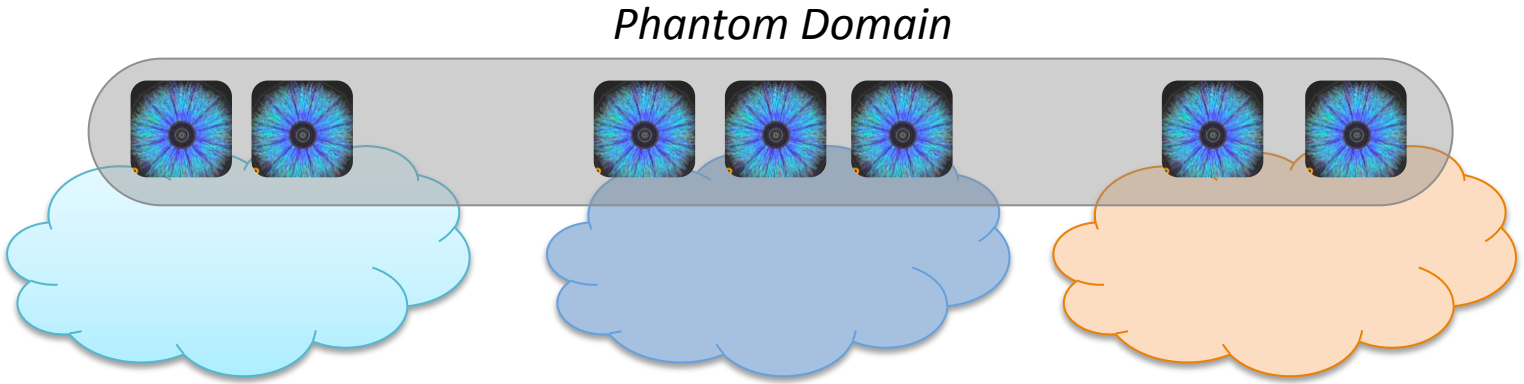

- Scalability and availability: regulates domain properties (compute, storage) using system and application metrics
- Multi-cloud: works with multiple providers
- Brokering: Finding the best resource fit
- Extensible monitoring: VM-based (OpenTSDB, traffic sentinel), provider-based (CloudWatch) and custom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

Phantom: Domain Management

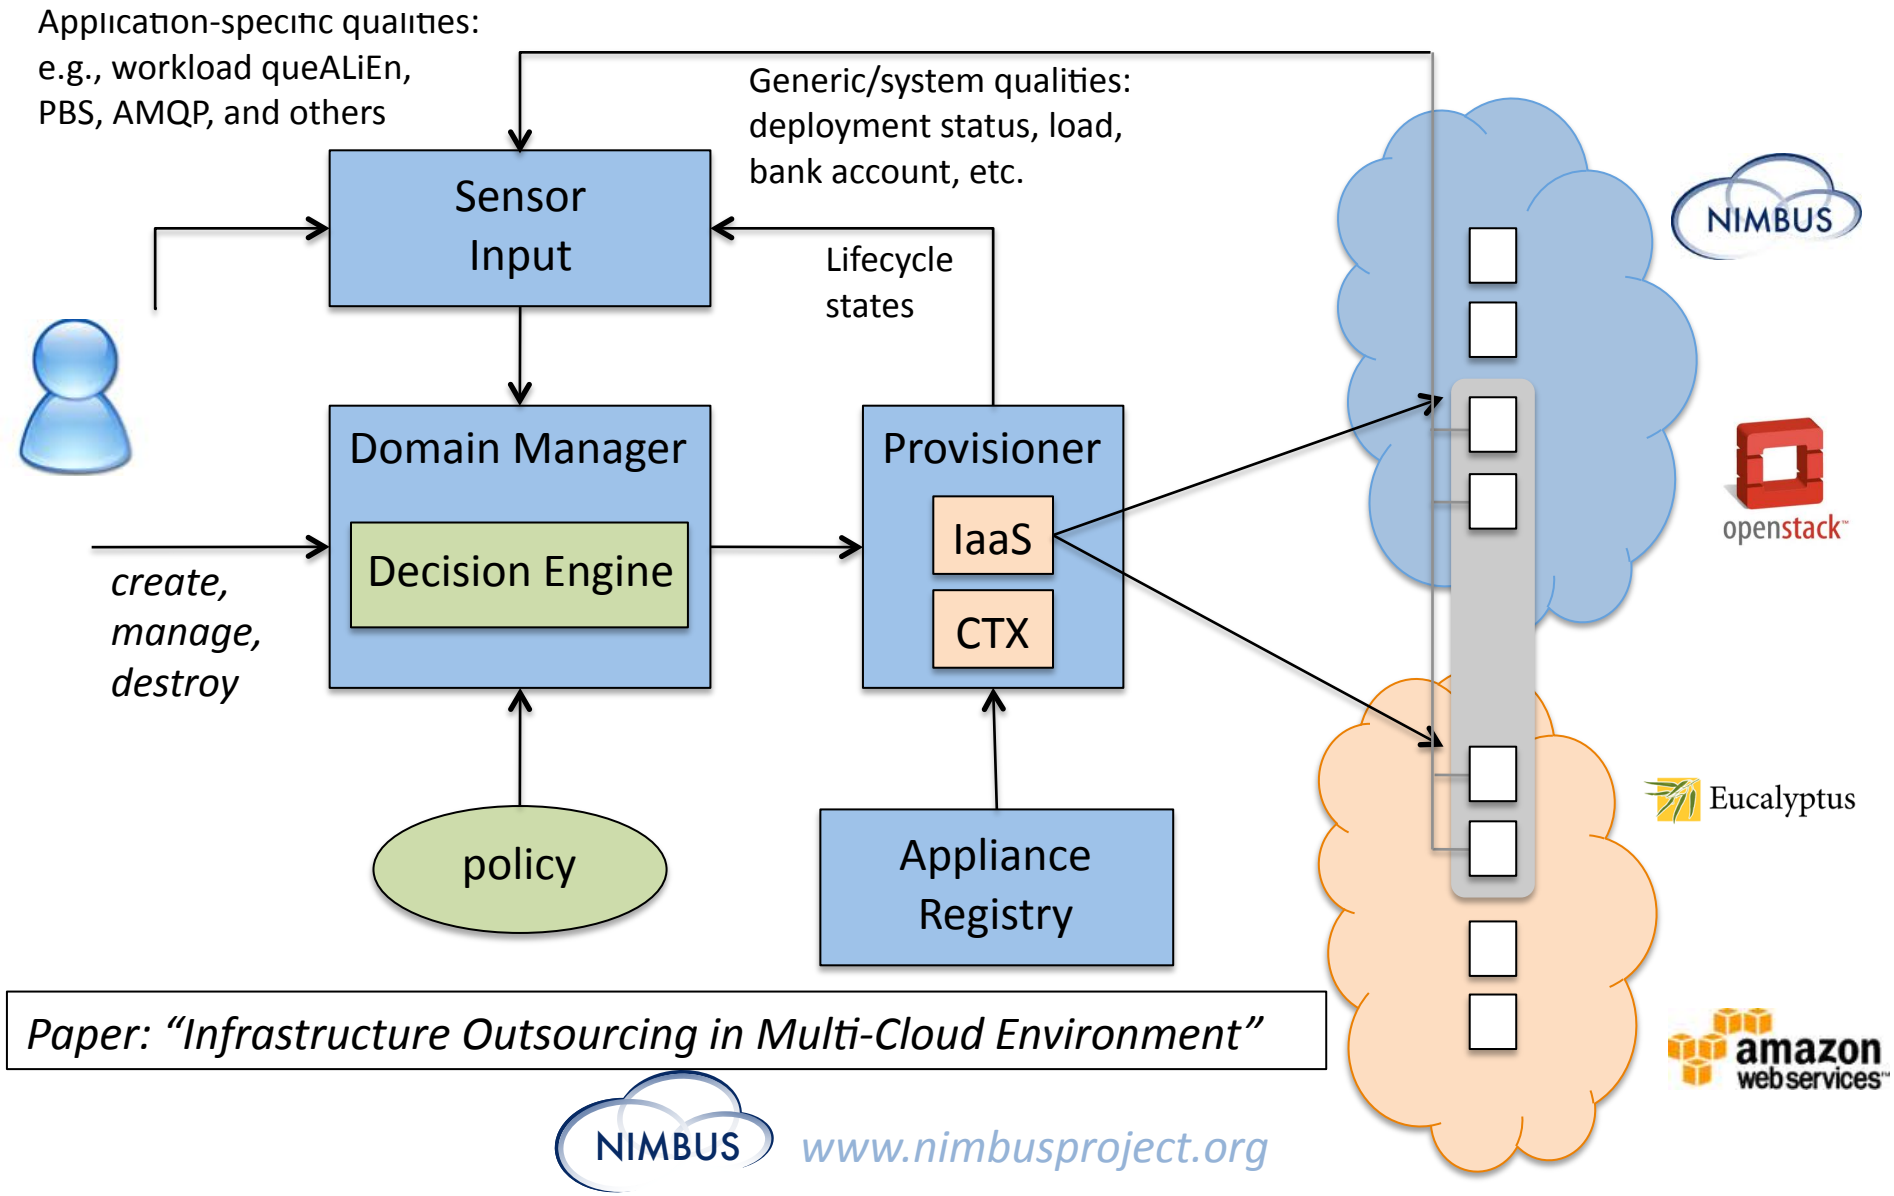

Paper: "Infrastructure Outsourcing in Multi-Cloud Environment"

NIMBUS *www.nimbusproject.org*

Extensible Monitoring

- Sensor Information:
  - **External:** lifecycle, contextualization, system sensors (CloudWatch)
  - **Internal:** Contextualization, heartbeat, system sensors (OpenTSDB, traffic sentinel)
  - **Custom:** job queue length, special process monitoring, GPS, and other application-specific sensors
- Extensible set

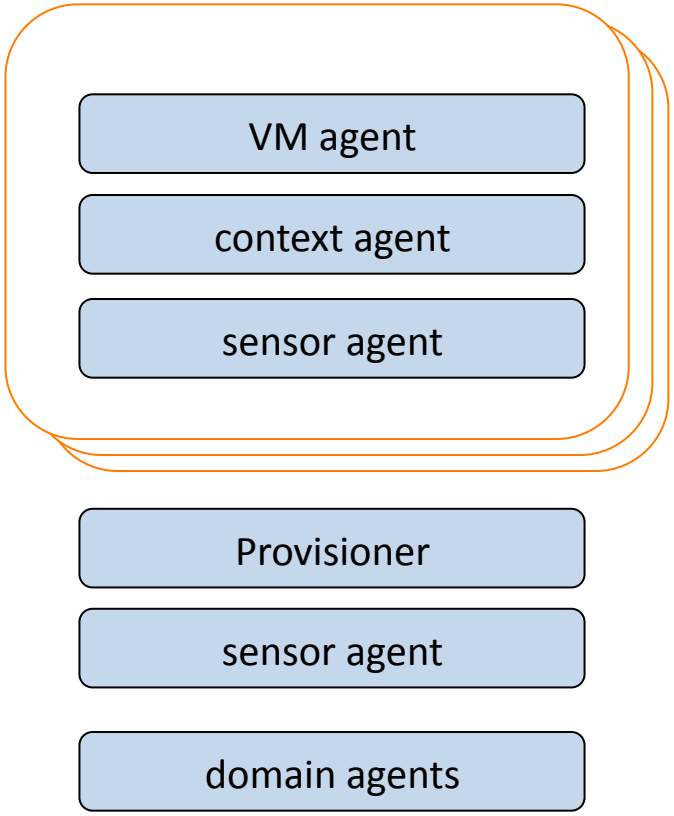

Phantom Interfaces and Clients

- Web application
  - Easy to use – but pre-defined policies
- Scripting
  - REST Phantom API
  - AWS autoscale API (boto-compatible)
- AMQP interface

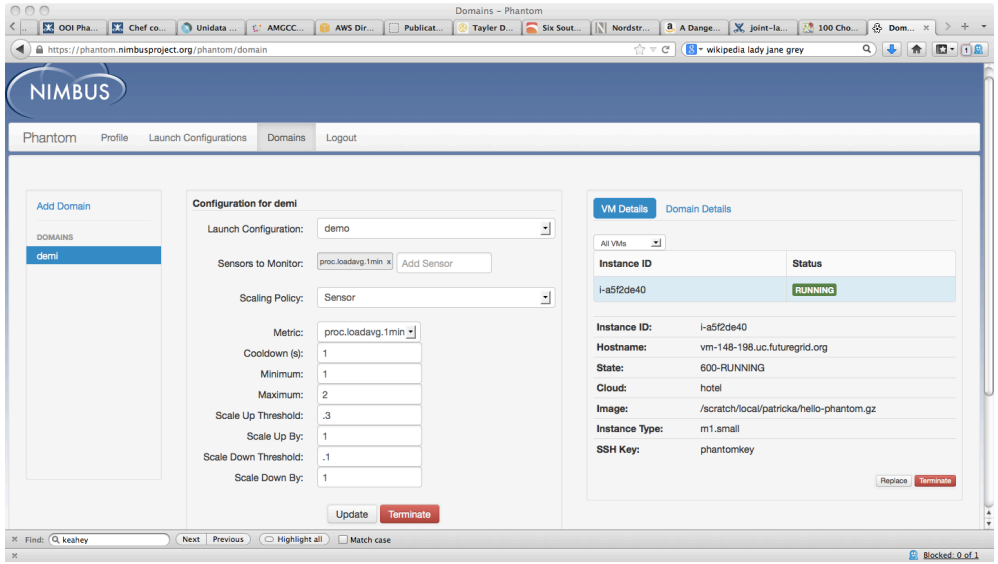

Adventures in Availability

Mean time between failures

$$A = \frac{MTBF}{MTBF + MTTR}$$

Mean time to repair

- Time to repair (TTR)
  - Diagnosis
  - Time to scale (TTS)
    - PENDING (request)
    - STARTED (deployment)
    - RUNNING (contextualization)

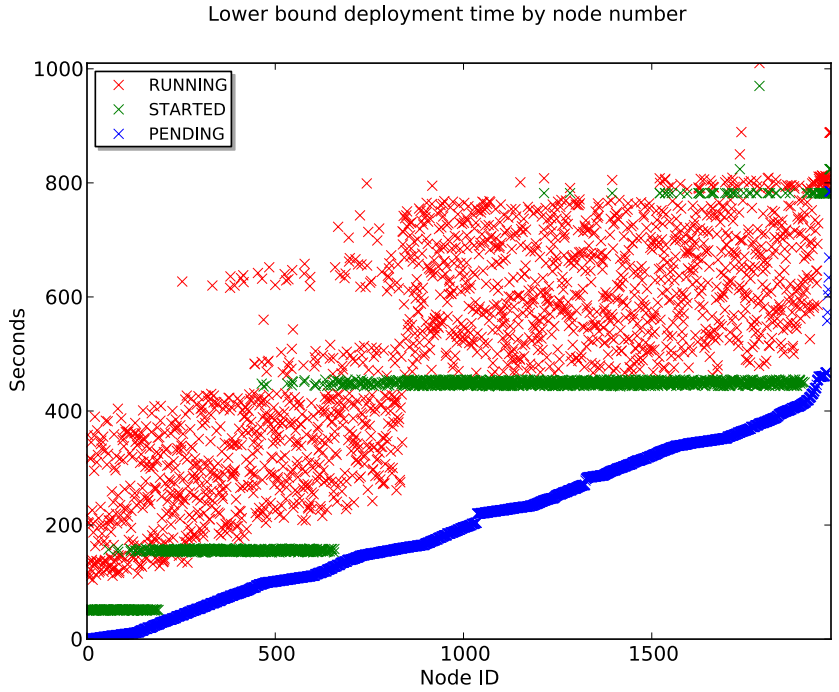

TTS: preliminary results for 2,000 VMs provisioned on AWS EC2

See also scalability articles at: www.nimbusproject.org

Managing Processes

Process Management Goals

- Extend process scheduling to integrate an independent resource manager
- Support a variety of existing process schedulers
- Work across cloud and traditional resources

Example: OOI Process Management

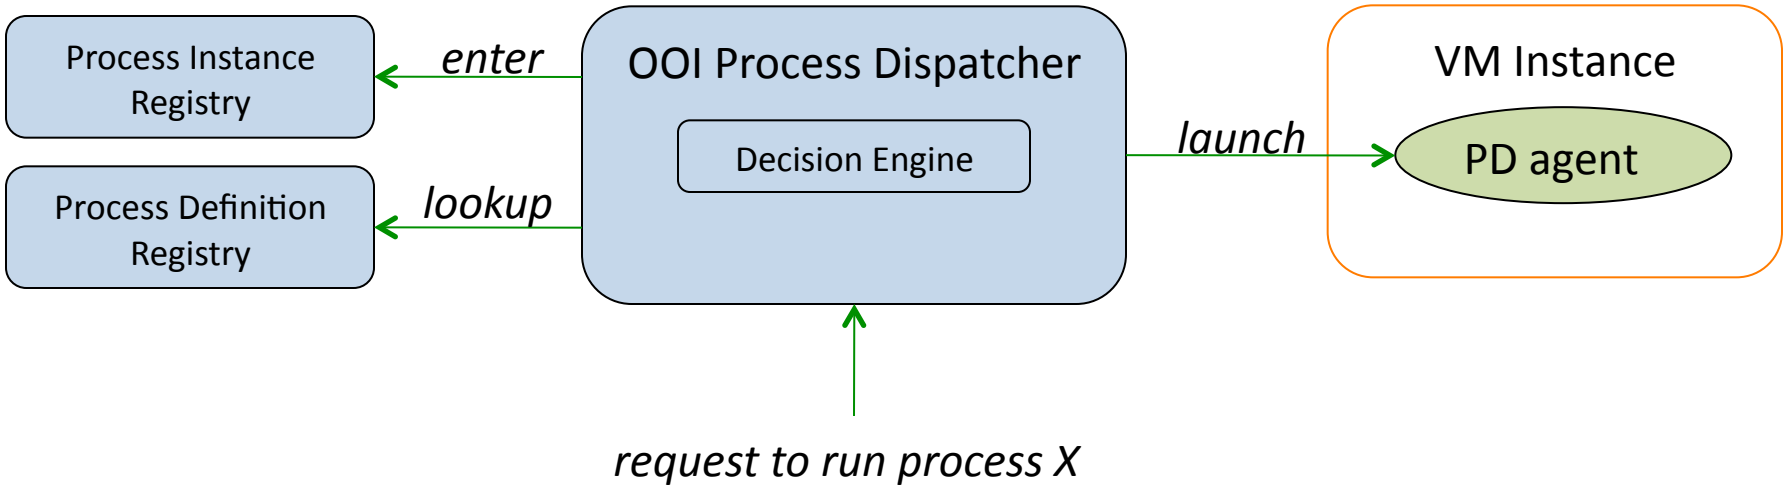

OOI Process Management (cntd)

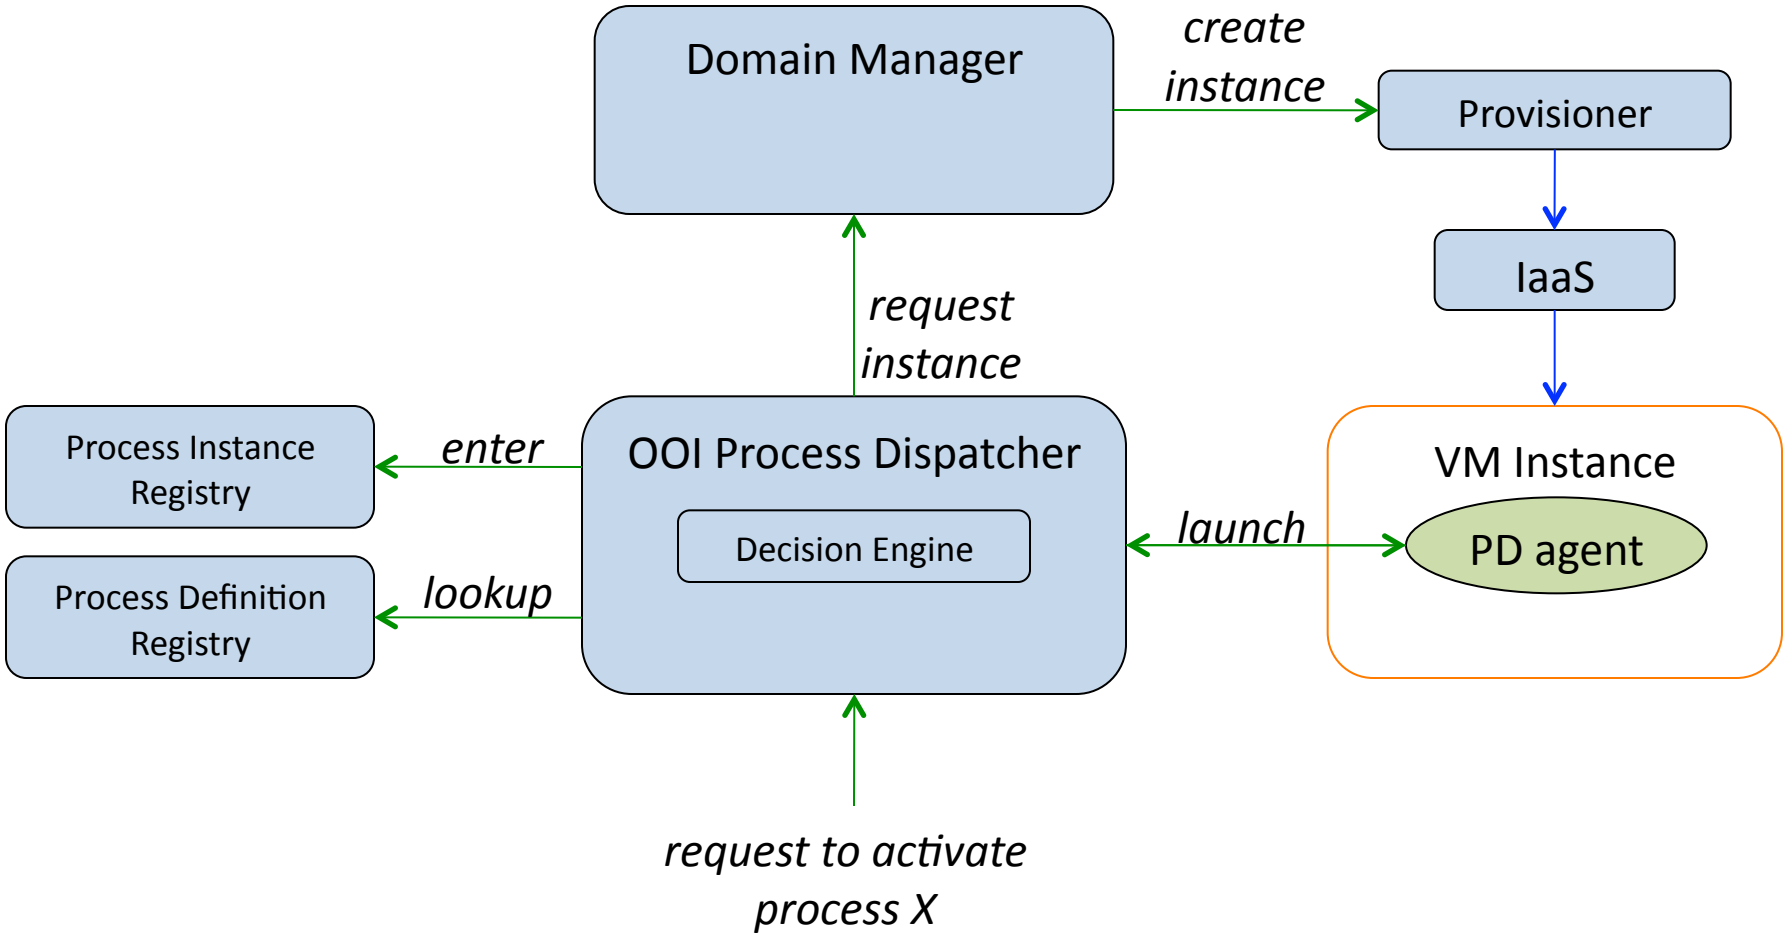

Process Management with Swift

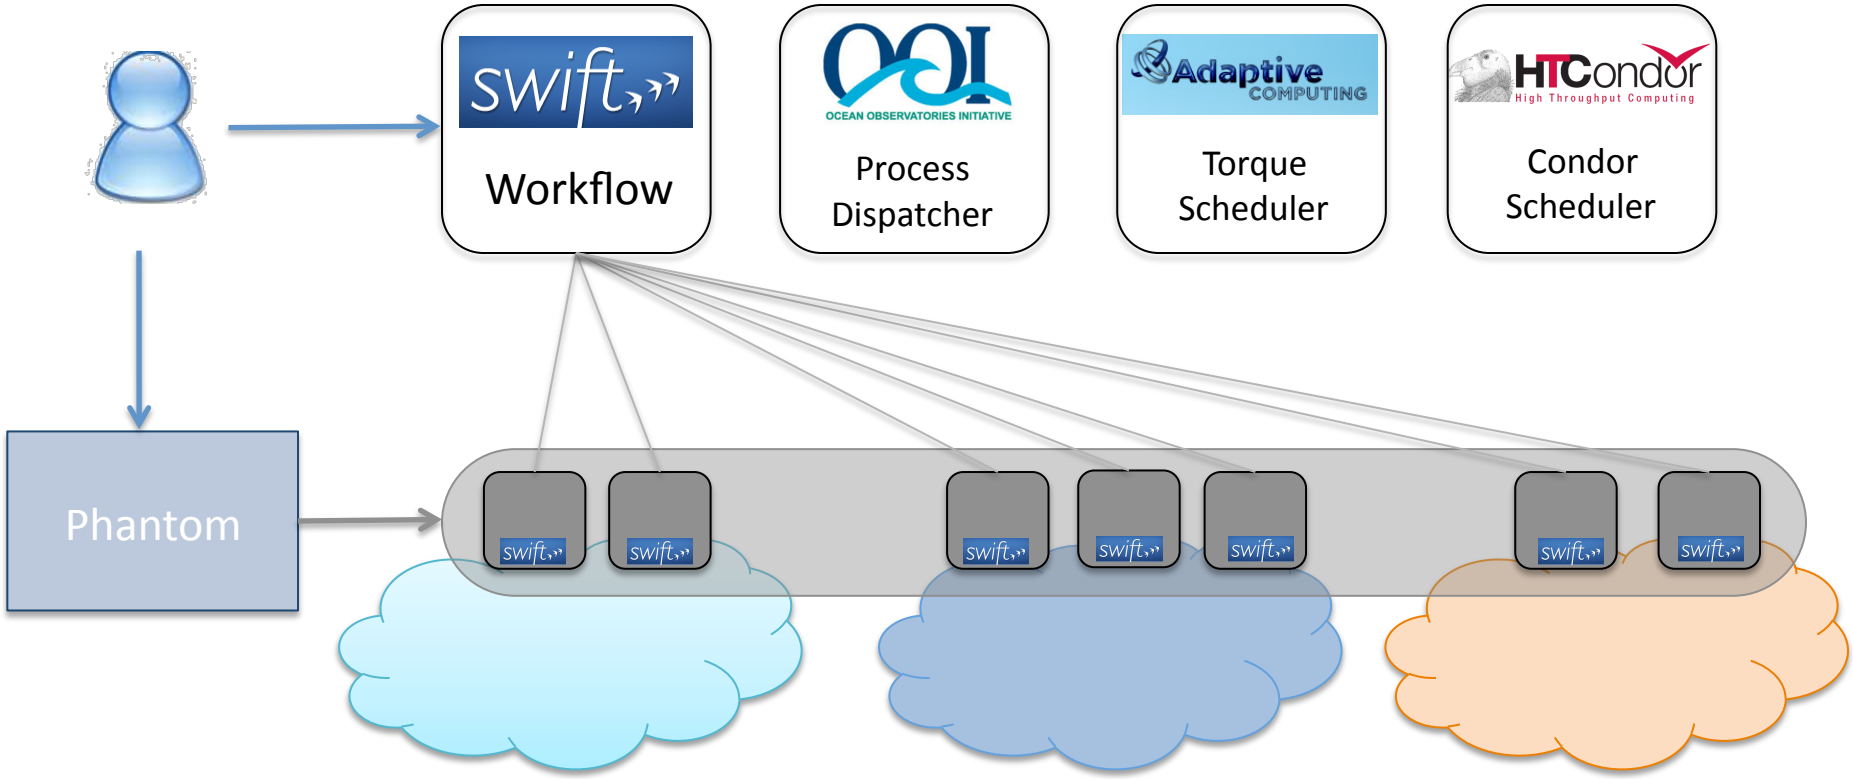

Check out the Swift and Torque appliances at www.sciencecloud.org

Using Nimbus Domains

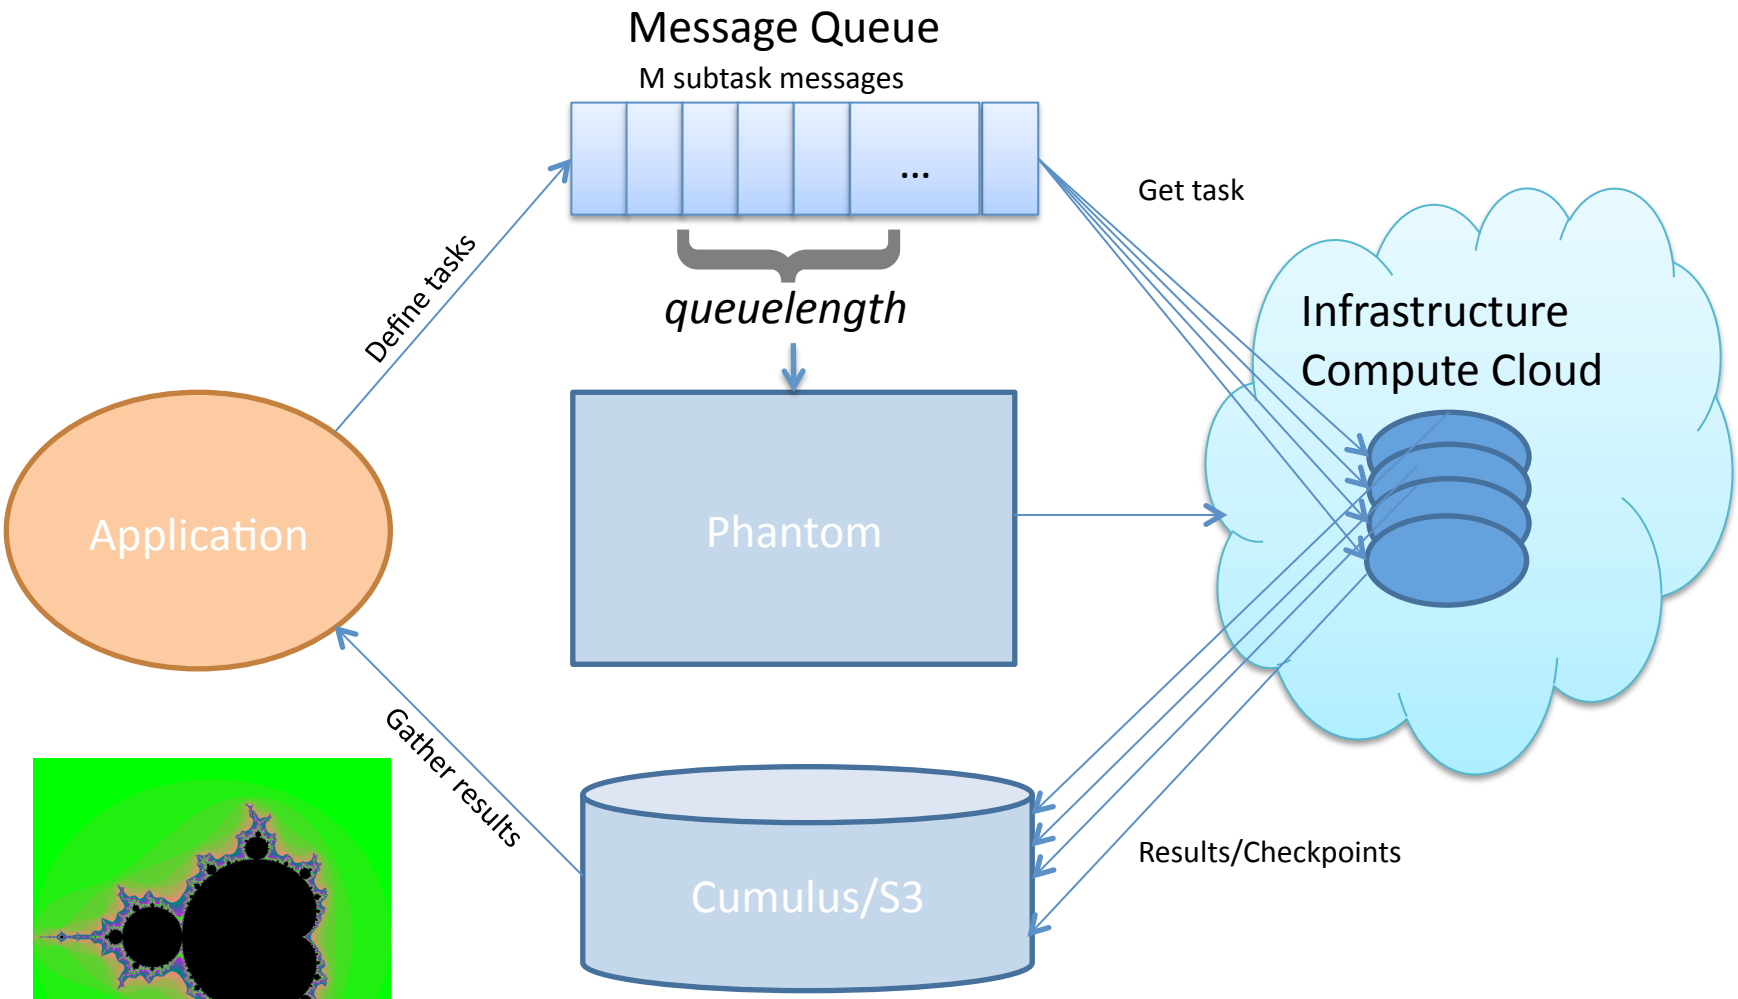

NIMBUS www.nimbusproject.org

CloudBurst!

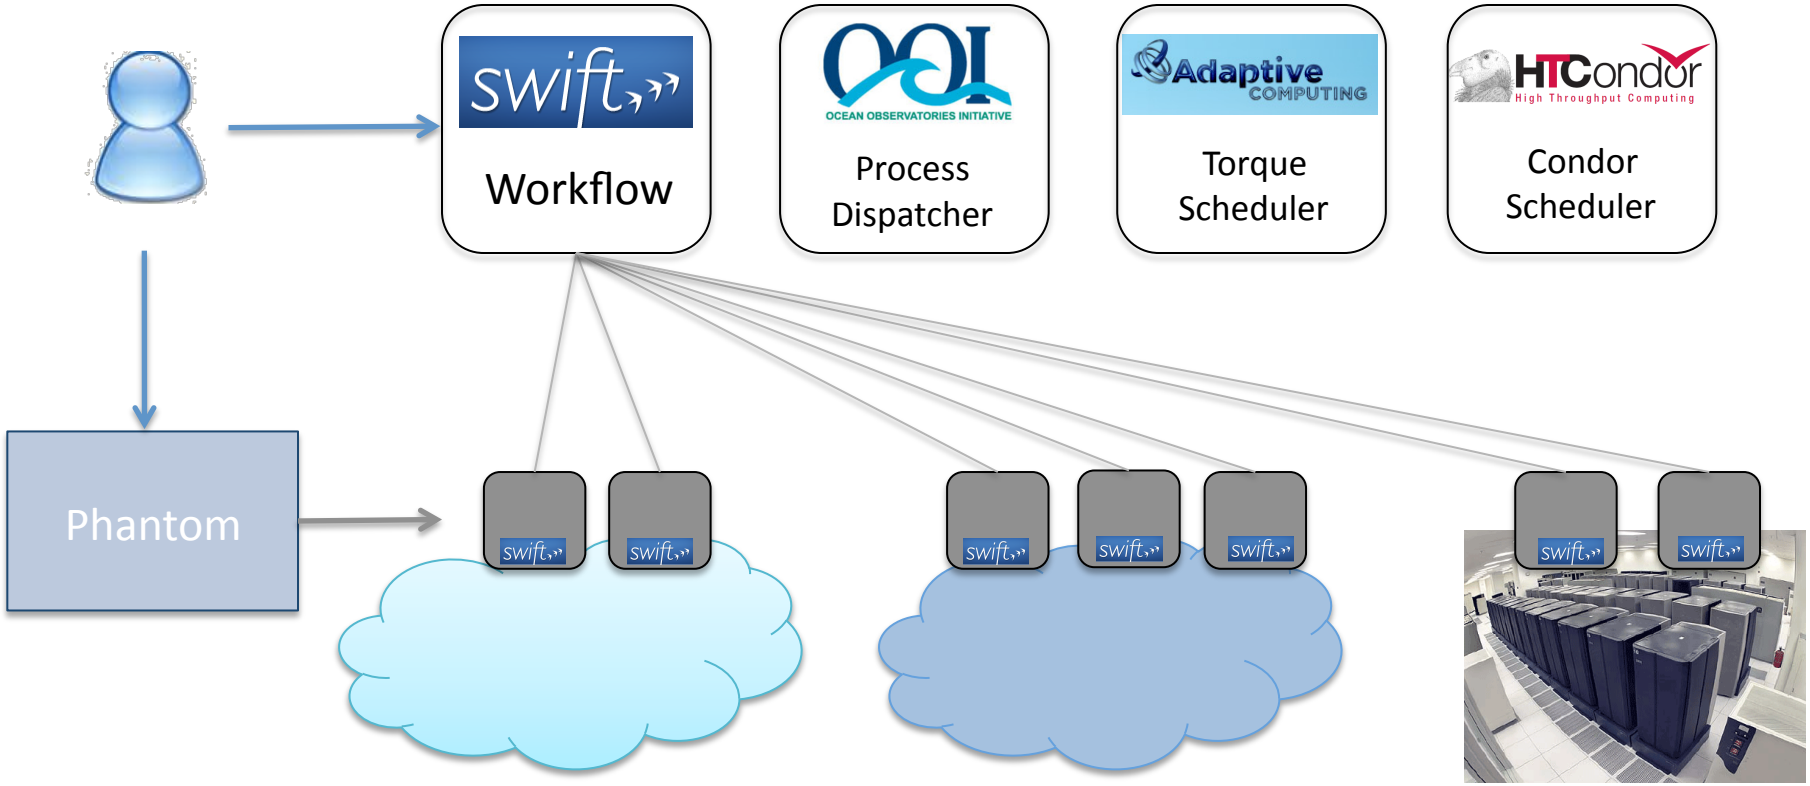
swift
Workflow
OCEAN OBSERVATORIES INITIATIVE
Process Dispatcher
Adaptive COMPUTING
Torque Scheduler
HTCondor
High Throughput Computing
Condor Scheduler
Phantom
swift
swift
swift
swift
swift
swift
swift
swift

Managing Streams

Stream Processing

- Publish/subscribe framework
- AMQP-based stream management framework for
- Data processing
  - Control channel vs data channel
- Stream registration
  - Sensors, USB, existing stream aggregation push and pull, archived streams, substreams, etc.
- Standard functions, such as archiving and displays become “sink” operations that can be added at any time

Interfaces and Clients

Register Stream | Register Operation | Create Stream | Subscriptions

Data Streams

- **Hyperspectral Camera Images**
  Spectra imagery of 128 narrow bands in the 374 to 1038 nm spectral range.
- **Weather Underground Weather Data**
  Weather Underground Rapid Fire XML data
- **ANL Weather Data**
  Weather from the ANL flux tower
- **NOAA Weather Data**
  NOAA MADIS Surface Observations
- **Amateur Radio Packet APRS Socket**

Operations

- **Hyperspectral Calibration**
  Process converting raw data values of the hyperspectral imagery to reflectance values
- **Weather Quality Control**
  Weather time consistency and plausible value check
- **Weather Generator**
  Short-term forecasts of relevant weather variables through machine learning algorithms
- **NDVI**
  Normalized Difference Vegetation Index. Spectral index corresponding to abundance of photosynthetic ally

Live Stream Output

New Stream

Stream Information | Create Stream | Create Stream and Subscribe

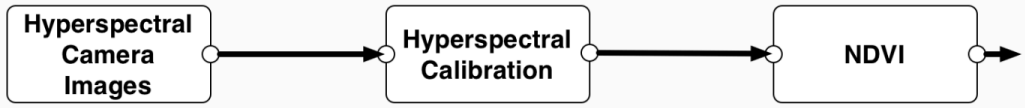

Status

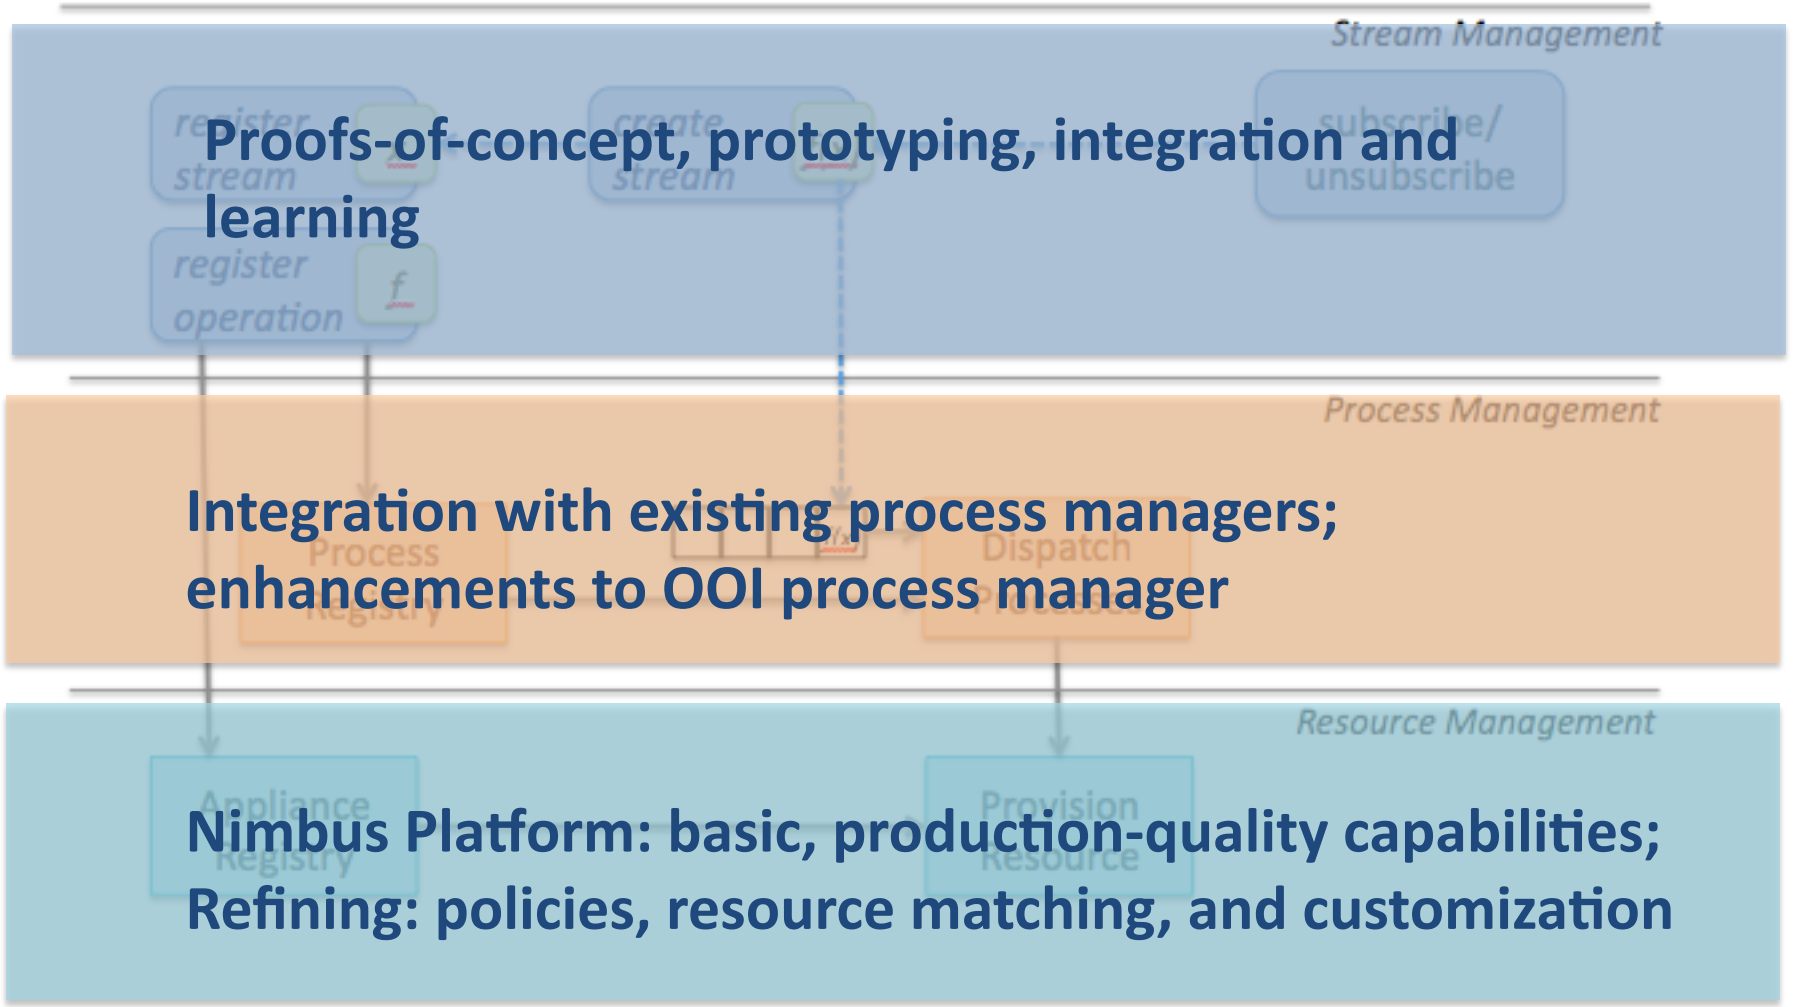
Stream Management
Proofs-of-concept, prototyping, integration and learning
Process Management
Integration with existing process managers; enhancements to OOI process manager
Resource Management
Nimbus Platform: basic, production-quality capabilities; Refining: policies, resource matching, and customization

Open Challenges

- Investigate and understand the interaction pattern
- Data management
  - Data types, transfer/placement management, etc.
- Integration and evolution of process management mechanisms
- Resource management
  - Cloud Resource models
  - SLAs and negotiation
  - Understanding and balancing response time, cost, availability, etc.
  - New repair/recovery strategies
  - Emphasis on networking and storage
- Adaptation: more efficient and reliable adaptation strategies and

Parting Thoughts

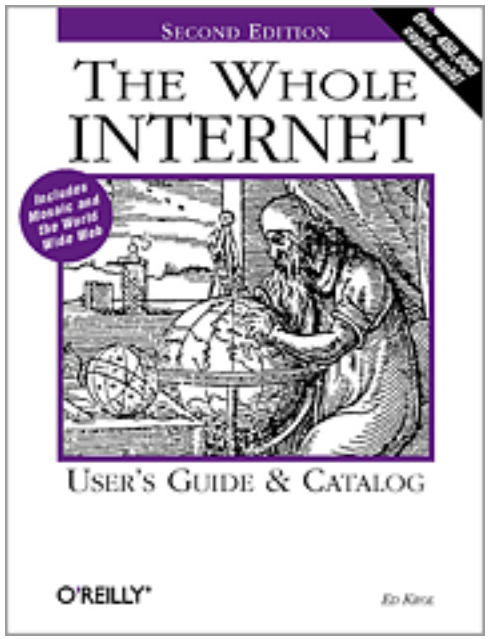

Mature Innovation: the Internet

Maturing Innovation: the cloud