



Nimbus Platform: How Does Your Cluster Grow?

Kate Keahey

keahey@mcs.anl.gov

Argonne National Laboratory

Computation Institute, University of Chicago


7/1/13



www.nimbusproject.org

1

Ocean Observatory Initiative (OOI)



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



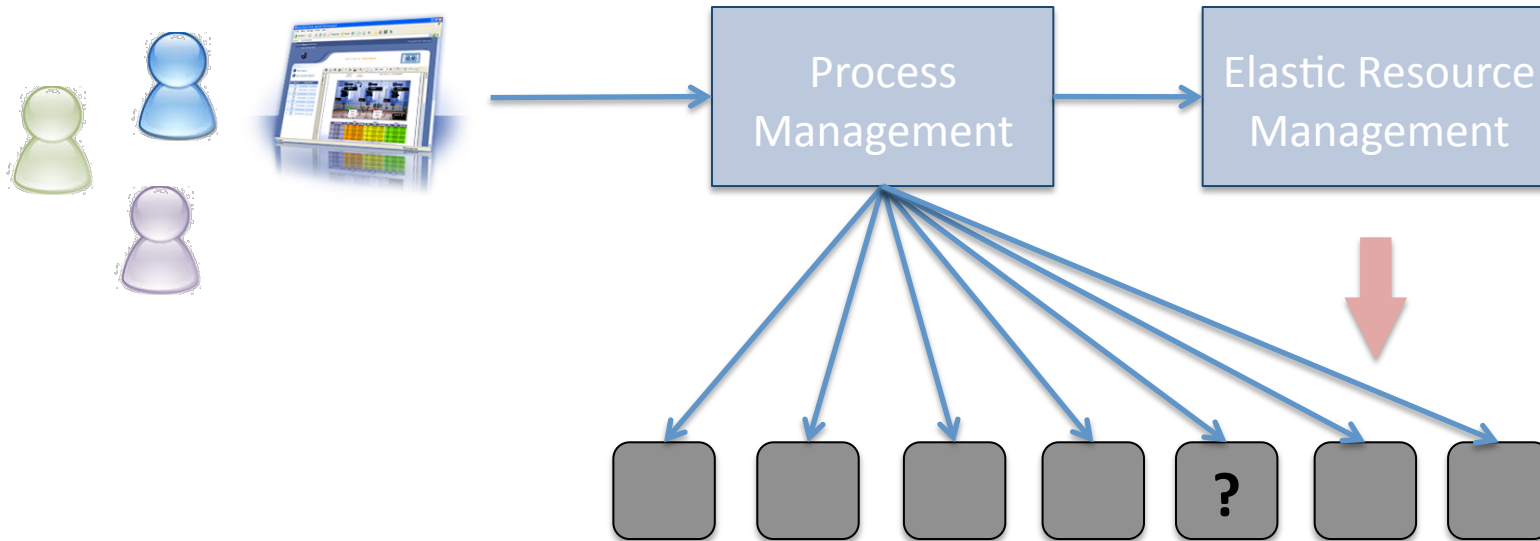
www.nimbusproject.org



What's New About Clouds

- Cloud computing benefits...
 - On-demand access (manage peaks)
 - Pay-as-you-go (manage valleys)
 - Virtualization (custom environments)
 - Outsourcing (convenience, articulation)
- ...enabling scenarios
 - Support of sensor-based processing
 - Support for “brainstorming” scenarios
 - Support for community/processing growth




If Wishes Were Fishes...



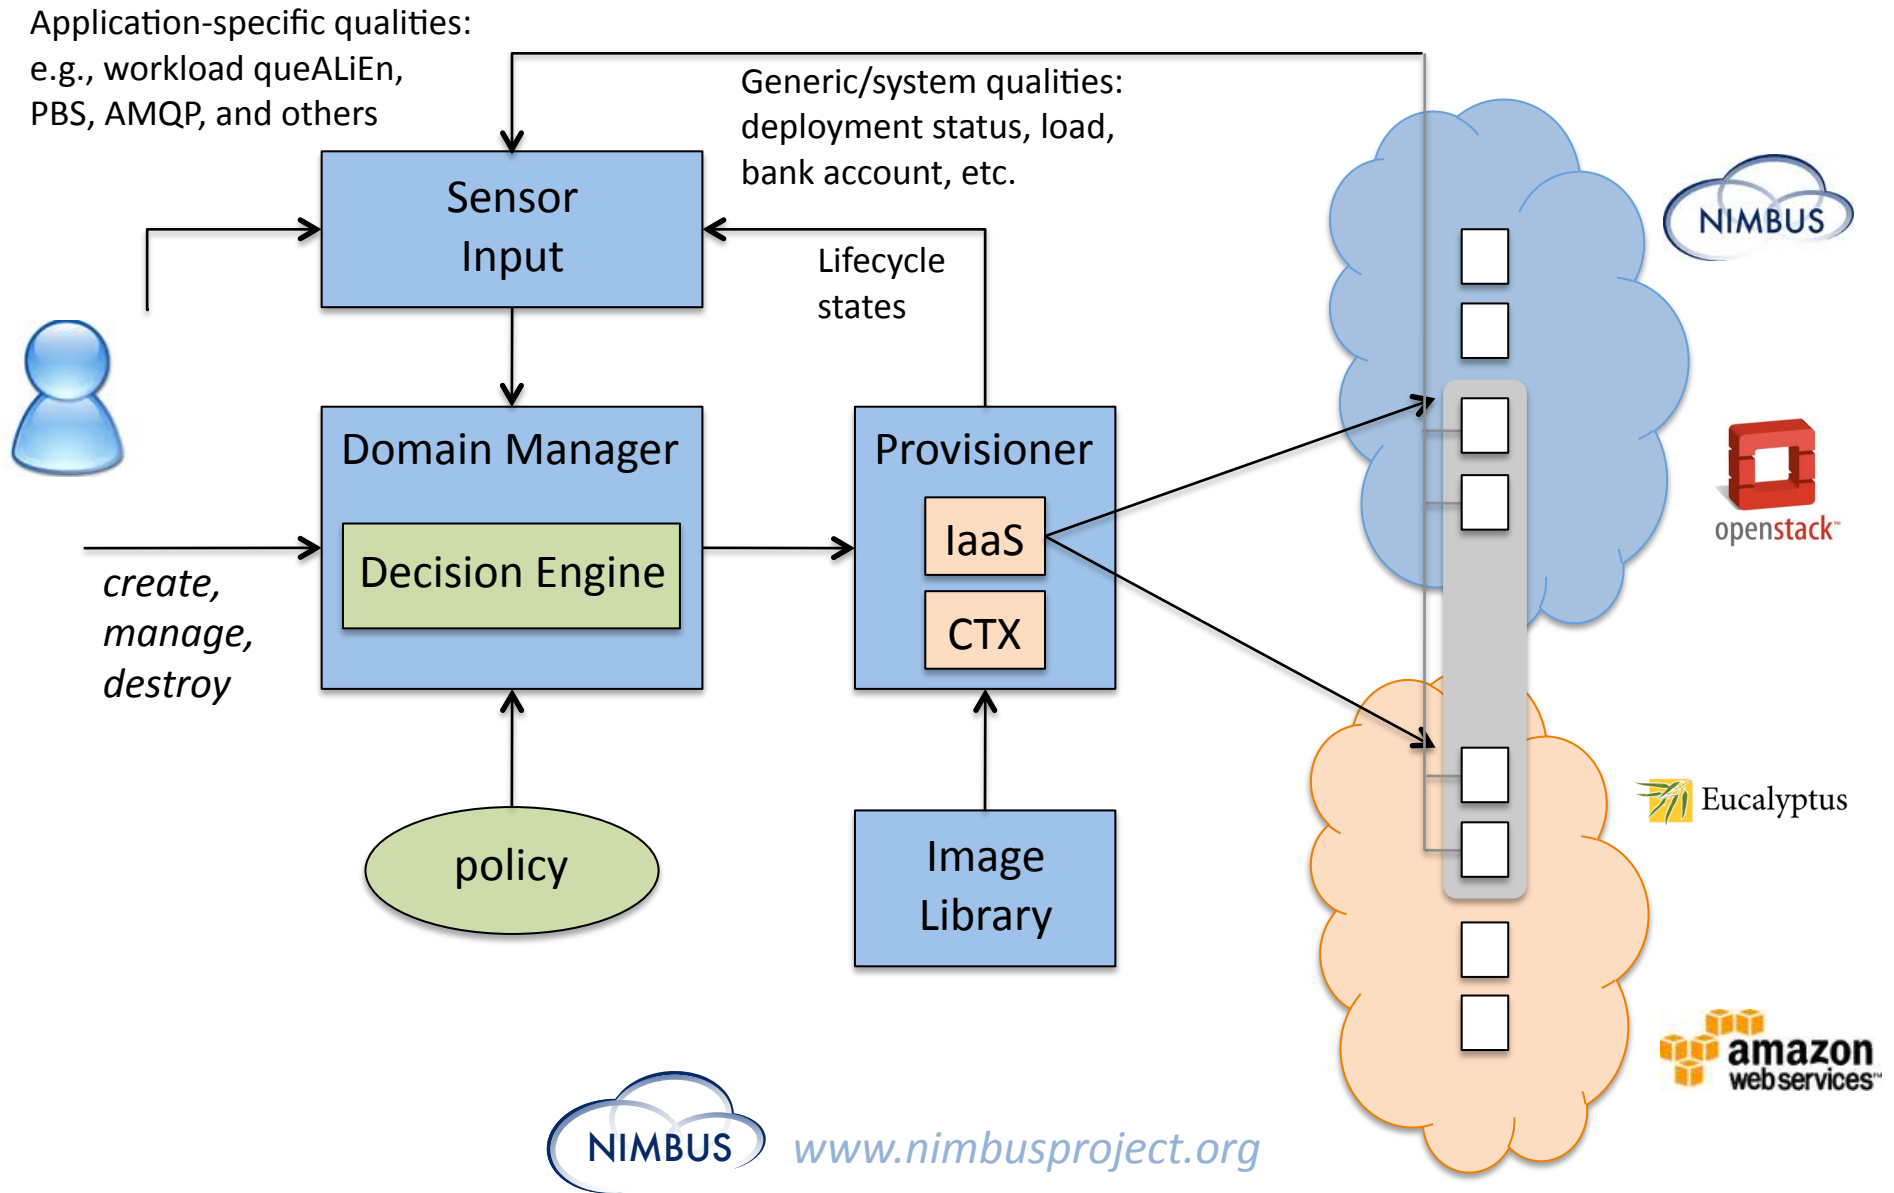
Managing Domains

Domain Management

*Monitor and regulate domain properties
based on system-specific and application-specific metrics*

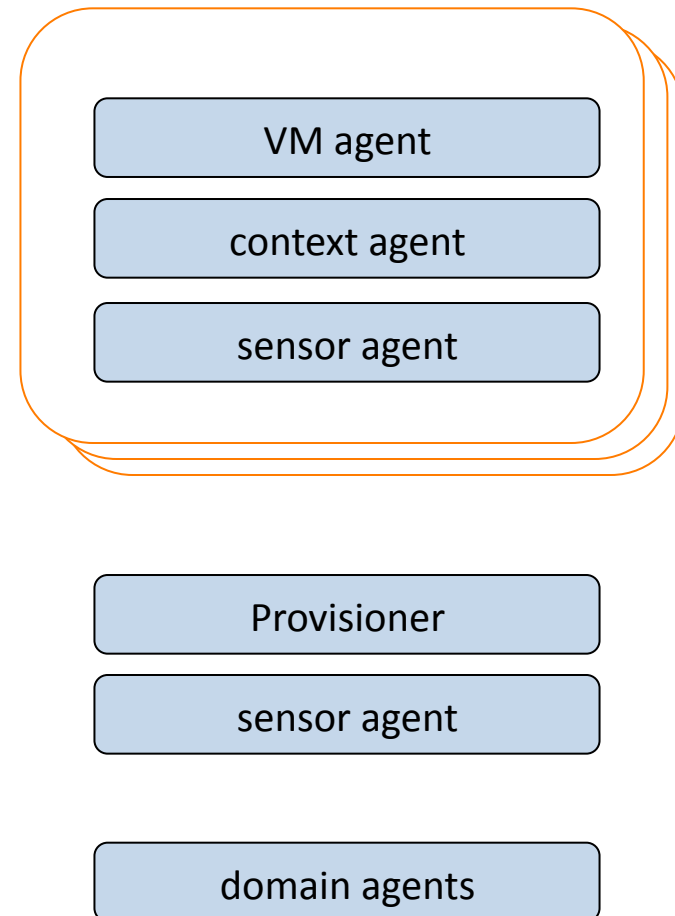


Phantom: Domain Management



Sensor Information

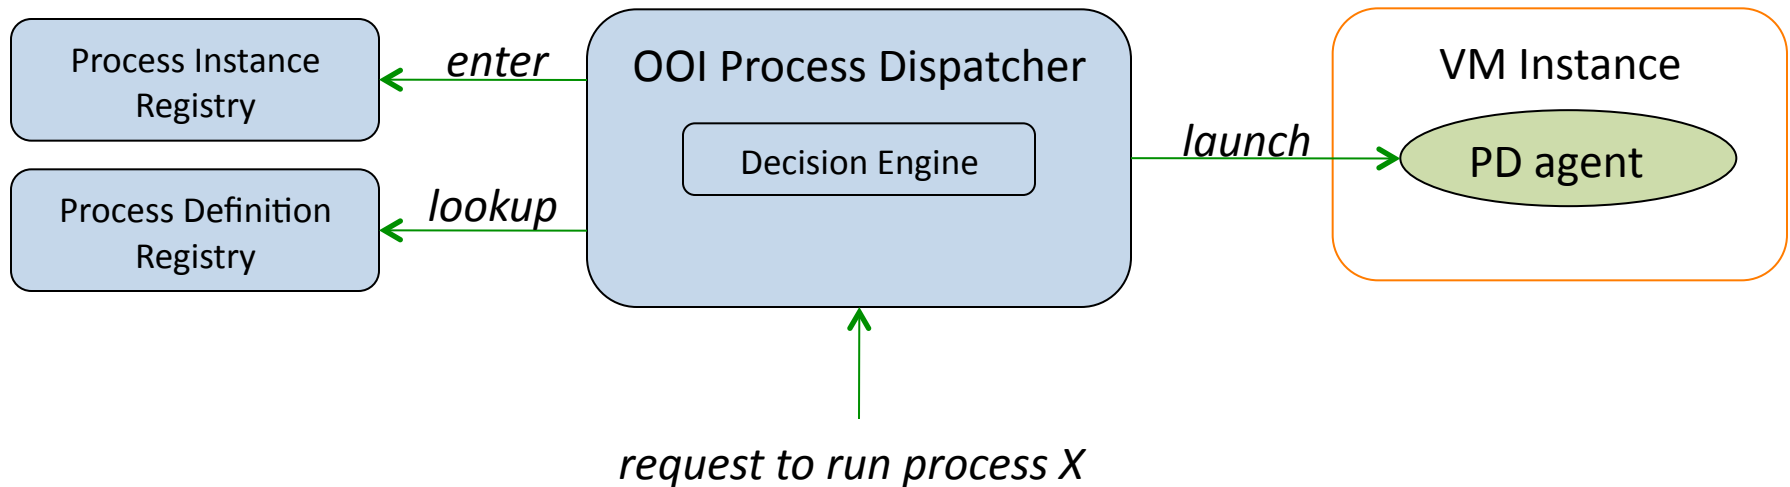
- Sensor Information:
 - Lifecycle states
 - Contextualization
 - Heartbeat
 - Special process monitoring
 - System sensors
 - TSDB
 - CloudWatch
 - Application-specific sensors
- Extensible set



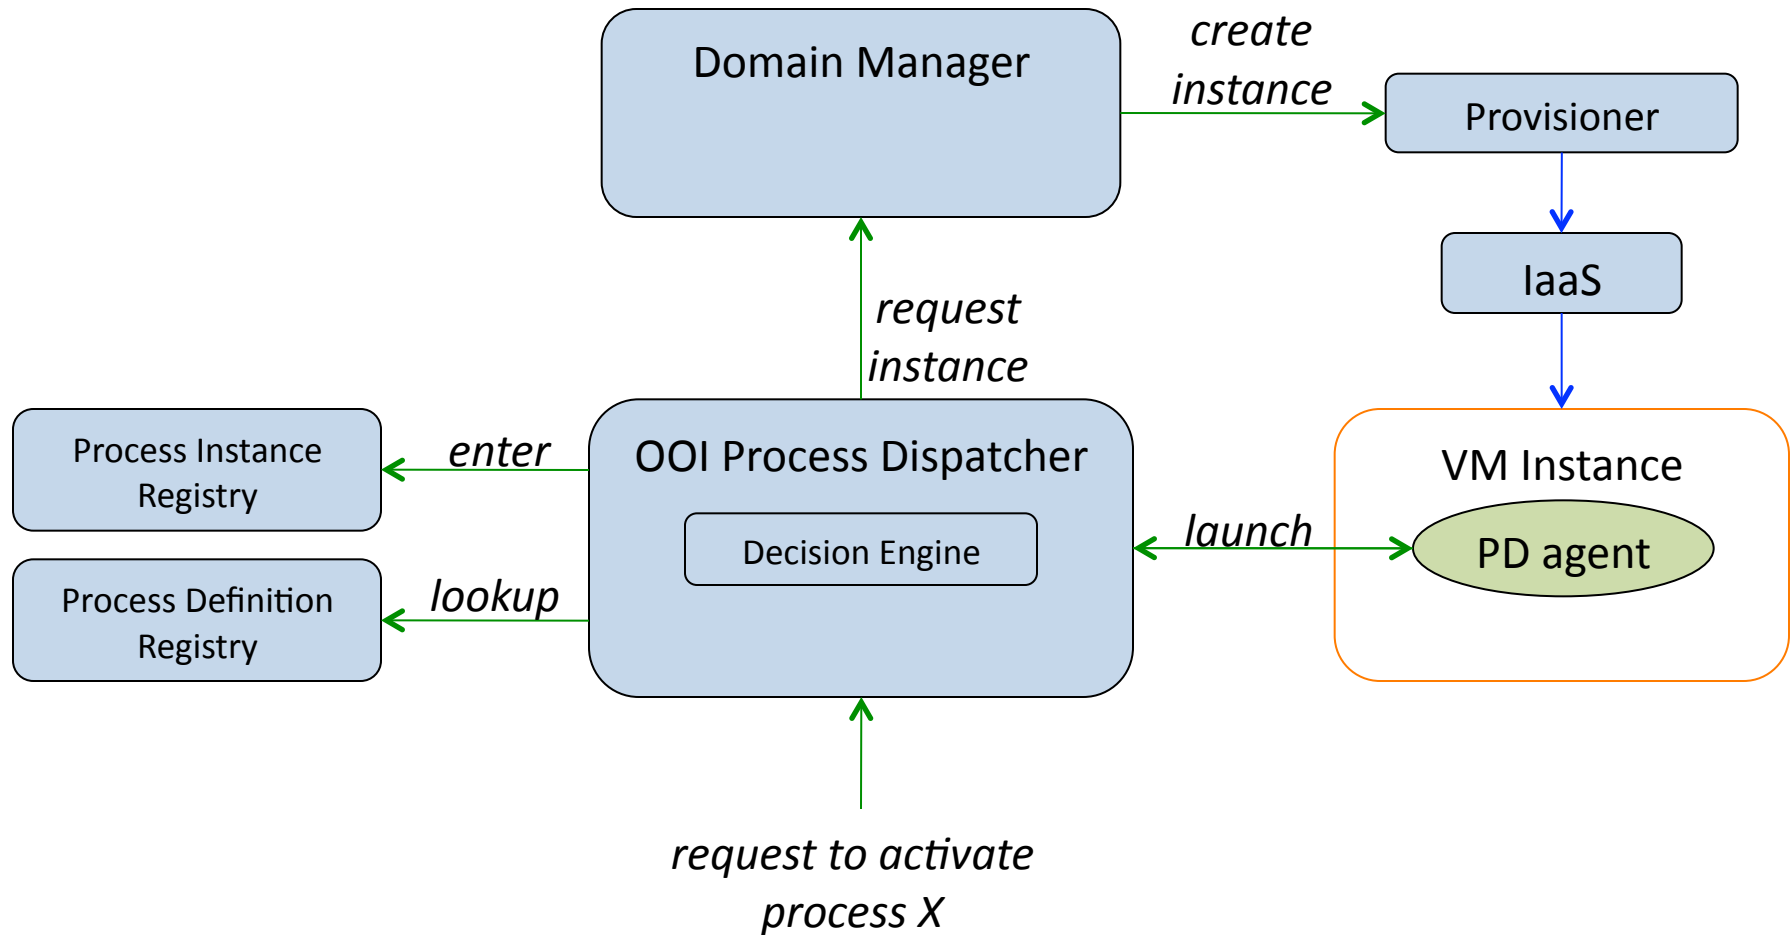
Phantom Interfaces and Clients

- Web application
- Scripting
 - Amazon autoscale API, boto-compatible
- AMQP interface

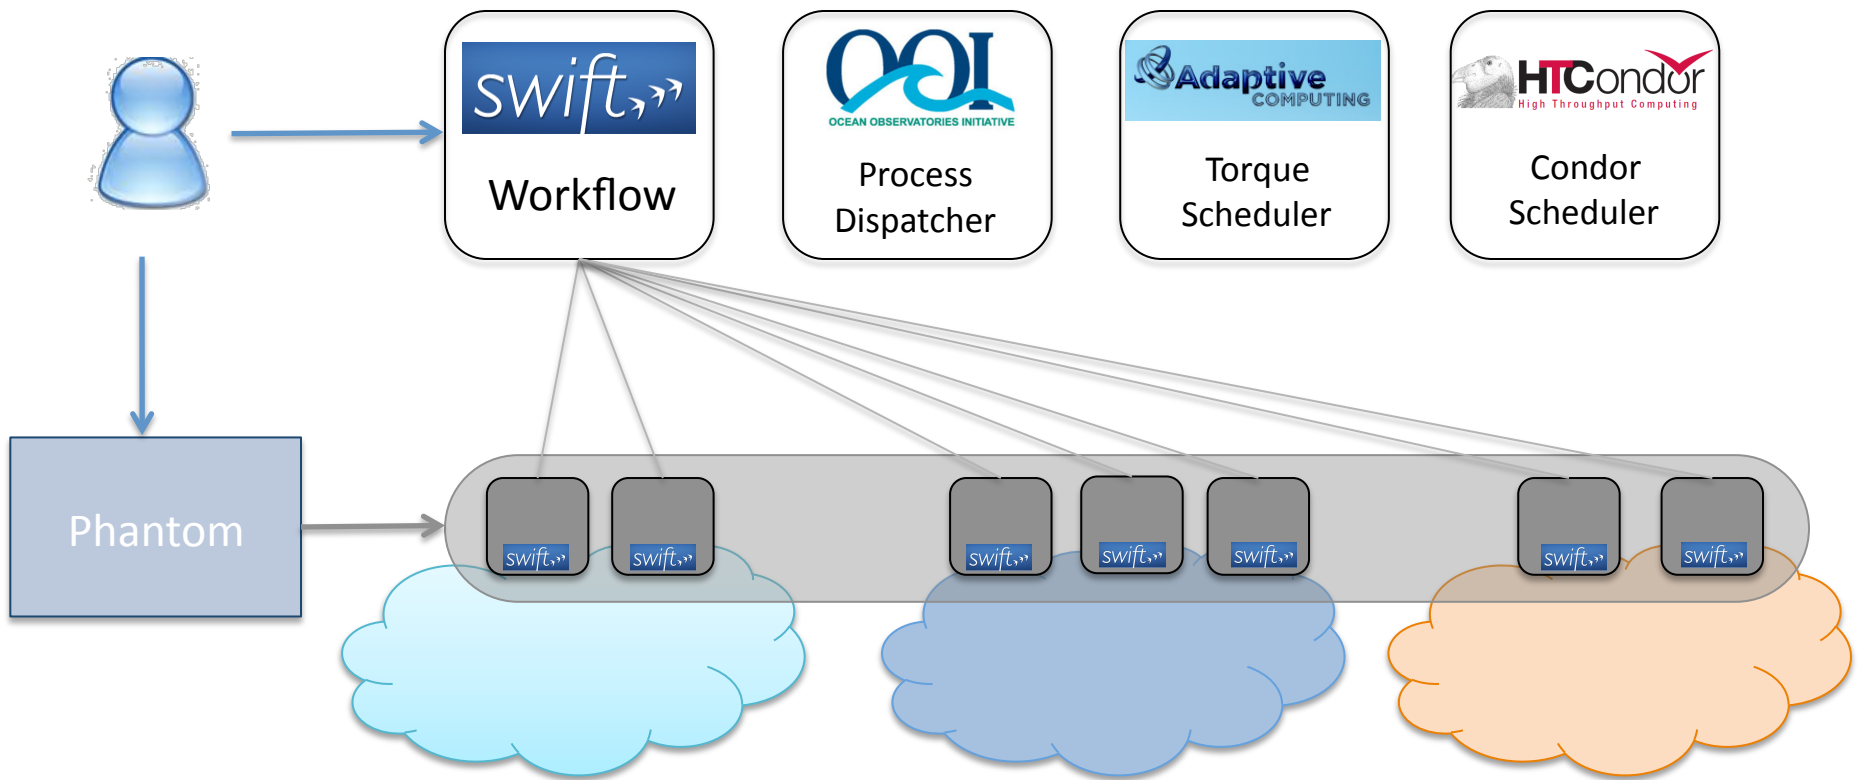
Working with Process Management I




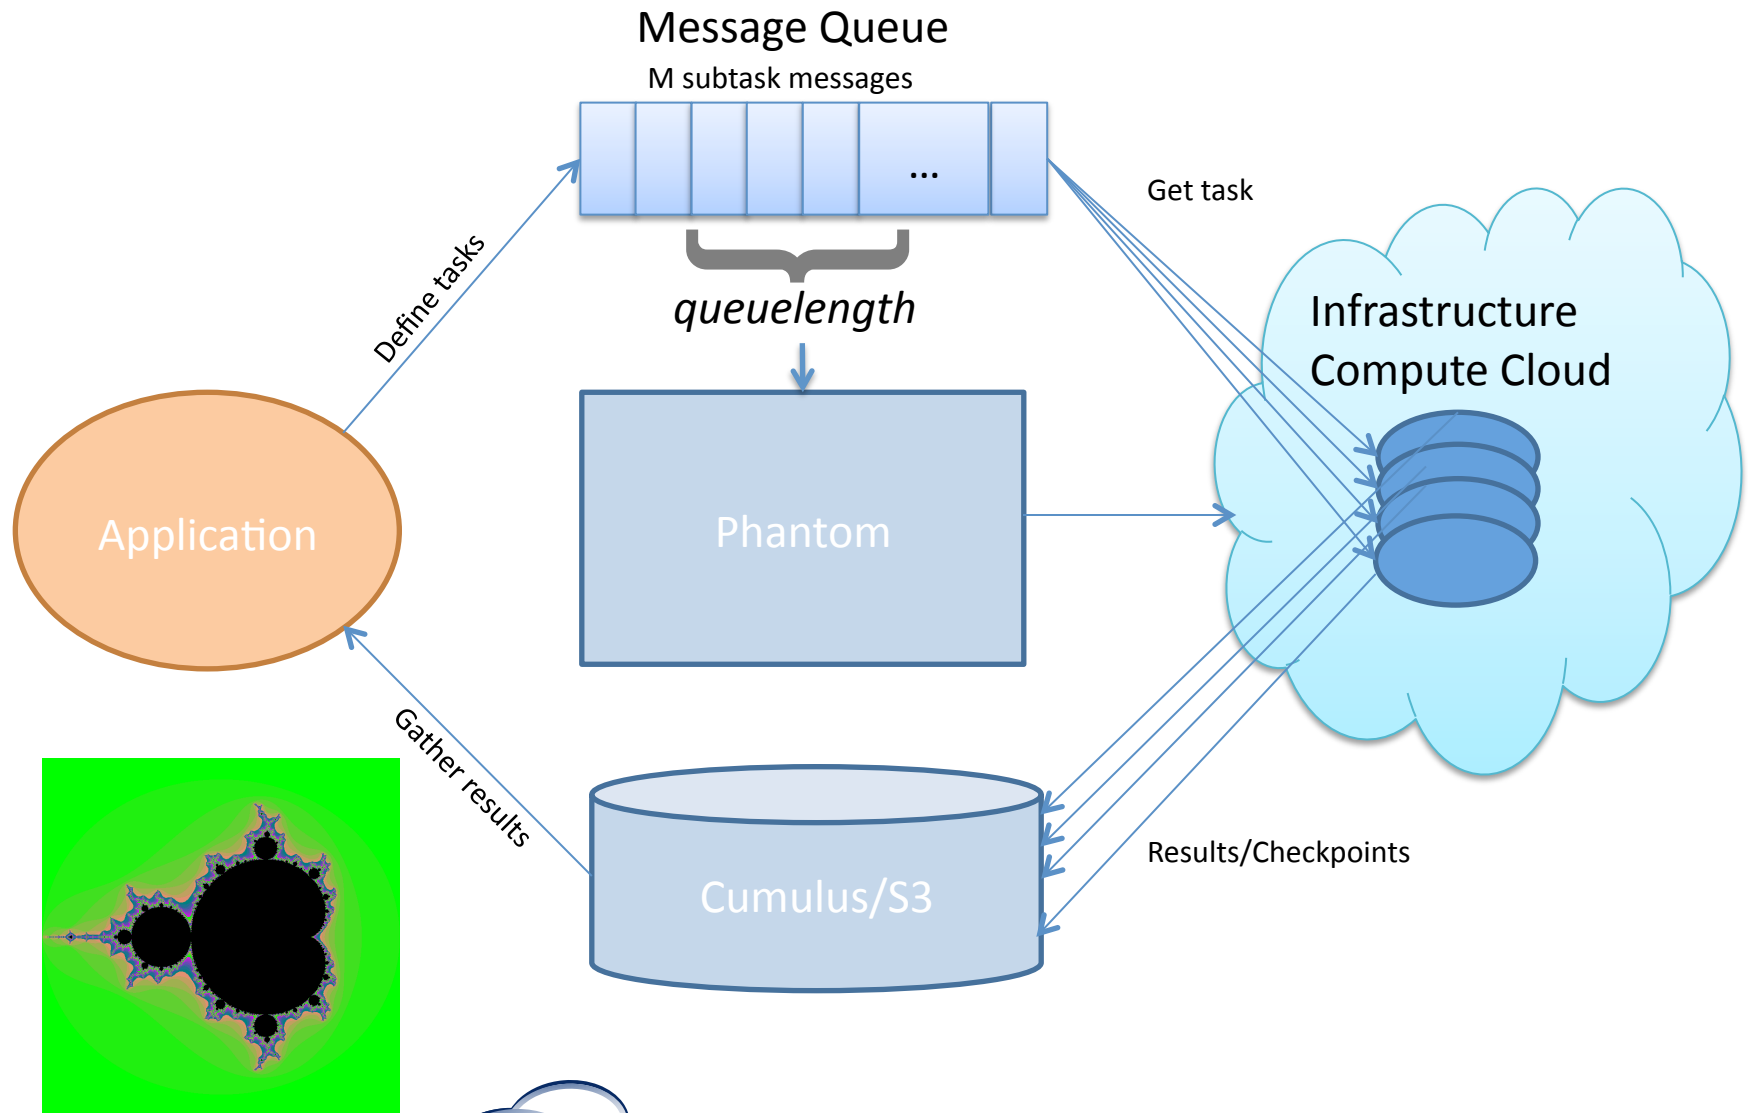
Working with Process Management II



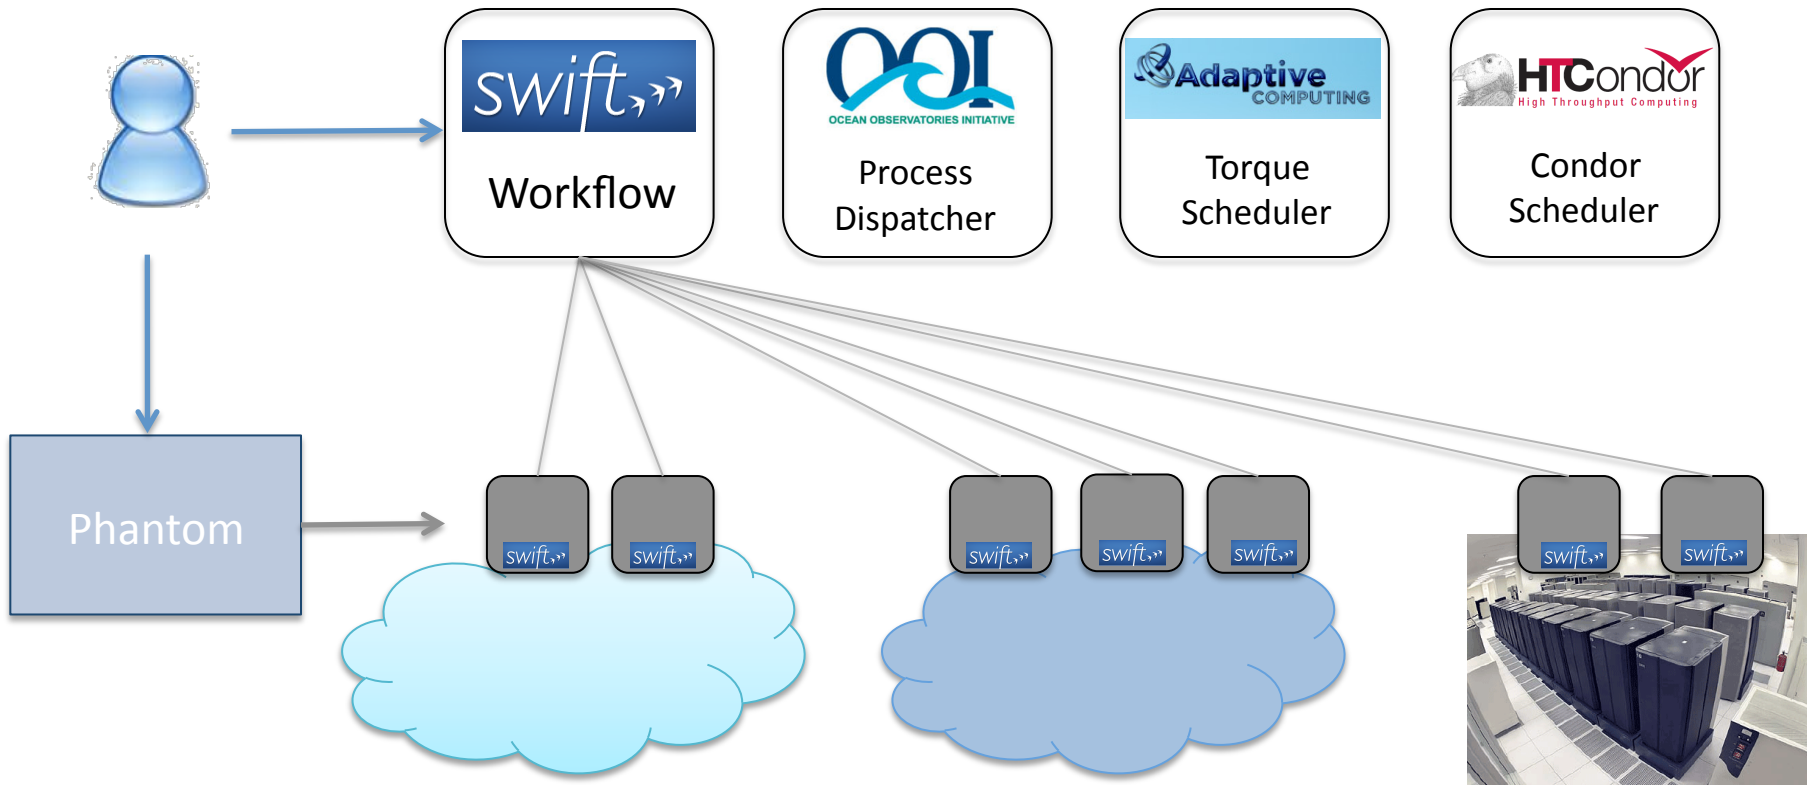
Process Management with Swift



Using Nimbus Domains



CloudBurst!




Adventures in Availability

Mean time between failures

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

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

Parting Thoughts

- Our world used to be bounded by the size of a cluster -- this is no longer the case
- Phantom is bridging the old and the new
- Tell us what you could do with an elastic cluster!



Get on FutureGrid and give it a whirl!
www.nimbusproject.org/phantom

Nimbus Team




Kate Keahey
Argonne National Lab
University of Chicago



Patrick Armstrong
University of Chicago




David LaBissoniere
University of Chicago




Pierre Riteau
University of Chicago

*Mary, Mary, quite contrary
How does your cluster grow?*



*With power bills and server kills
And racks all in a row*



*As things come up, it is a snap
It goes where I want it to go*