



[www.chameleoncloud.org](http://www.chameleoncloud.org)

## CHAMELEON: BUILDING A LARGE SCALE RECONFIGURABLE EXPERIMENTAL TESTBED FOR CLOUD RESEARCH

Kate Keahey

[keahey@anl.gov](mailto:keahey@anl.gov)

*3<sup>rd</sup> Joint Laboratory for Extreme-Scale Computing Workshop  
June 29- July 1  
Barcelona, Spain*

SEPTEMBER 15, 2015


1




# CHAMELEON: A FLEXIBLE AND POWERFUL EXPERIMENTAL INSTRUMENT

- ▶ **Large-scale:** “Big Data, Big Compute, Big Instrument research”
  - ▶ ~650 nodes (~14,500 cores), 5 PB disk over two sites, 2 sites connected with 100G network
- ▶ **Reconfigurable:** “As close as possible to having it in your lab”
  - ▶ From bare metal reconfiguration to clouds
  - ▶ Support for repeatable and reproducible experiments
- ▶ **Connected:** “One stop shopping for experimental needs”
  - ▶ Workload and Trace Archive
  - ▶ Partnerships with production clouds: CERN, OSDC, Rackspace, Google, and others
  - ▶ Partnerships with users
- ▶ **Complementary:** “Can’t do everything ourselves”
  - ▶ Complementing GENI, Grid’5000, and other experimental testbeds

# CHAMELEON HARDWARE



To UTSA, GENI, Future Partners



# CAPABILITIES AND SUPPORTED RESEARCH

Development of new models, algorithms, platforms, auto-scaling HA, etc., innovative application and educational uses

*Persistent, reliable, shared clouds*

Repeatable experiments in new models, algorithms, platforms, auto-scaling, high-availability, cloud federation, etc.

*Isolated partition, Chameleon Appliances*


Virtualization technology (e.g., SR-IOV, accelerators), systems, networking, infrastructure-level resource management, etc.

*Isolated partition, full bare metal reconfiguration*

# BUILDING CHAMELEON INFRASTRUCTURE (CHI)

- ▶ Defining requirements
- ▶ Developing architecture
- ▶ Technology Evaluation and Risk Analysis
  - ▶ Technology evaluation: Grid'5000 and OpenStack
  - ▶ Implementation proposals
- ▶ Implementing CHI
- ▶ Technology Preview (since 04/15)
- ▶ Early User (since 2 weeks ago)

# EXPERIMENTAL WORKFLOW

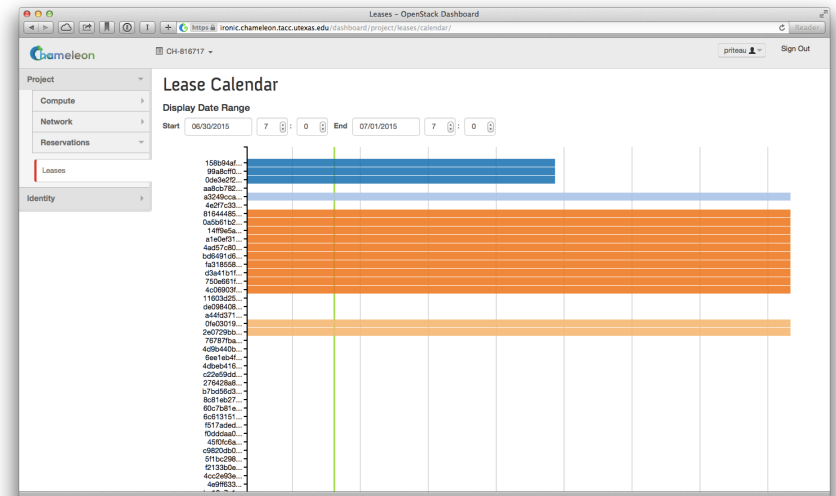


# CHI: SELECTING AND VERIFYING RESOURCES

- ▶ Complete, fine-grained and up-to-date representation
  - ▶ Machine parsable, enables match making
  - ▶ Versioned
    - ▶ “What was the drive on the nodes I used 6 months ago?”
  - ▶ Dynamically Verifiable
    - ▶ Does reality correspond to description? (e.g., failures)
- 
- ▶ Grid’5000 Registry
    - ▶ Automated resource description, automated export to RM
  - ▶ G5K-checks
    - ▶ Run at boot, acquire information, compare with resource catalog description

# CHI: PROVISIONING RESOURCES

- ▶ Resource leases
- ▶ Allocating a range of resources
  - ▶ Different node types, etc.
- ▶ Multiple environments
- ▶ Advance reservations (AR)
  - ▶ Sharing resources across time
- ▶ Extensions: match making, Gantt chart displays



- ▶ OpenStack Nova/Blazar
- ▶ Extensions to support working with more resources, match making, and displays



# CHI: CONFIGURE AND INTERACT

- ▶ Map multiple appliances to a lease
- ▶ Allow deep reconfiguration (incl. BIOS)
- ▶ Snapshotting
- ▶ Efficient appliance deployment
- ▶ Handle complex appliances
  - ▶ Virtual clusters, cloud installations, etc.
- ▶ Interact: reboot, power on/off, access to console
- ▶ Shape experimental conditions

- 
- ▶ OpenStack Ironic, Glance, and meta-data servers

# CHI: MONITORING

- ▶ Enables users to understand what happens during the experiment
- ▶ Types of monitoring
  - ▶ User resource monitoring
  - ▶ Infrastructure monitoring (e.g., PDUs)
  - ▶ Custom user metrics
- ▶ High-resolution metrics
- ▶ Easily export data for specific experiments

- 
- ▶ OpenStack Ceilometer

# PROJECT TIMELINE

- ▶ Started 09/2014
- ▶ Currently:
  - ▶ Since 12/14: FutureGrid@Chameleon (OpenStack KVM cloud)
  - ▶ Since 04/15: Chameleon Technology Preview (CHI: bare metal)
  - ▶ Since 06/17/15: Early Users on homogenous hardware
  - ▶ Overall: 57 projects, 102 users, 40 institutions
- ▶ Fall 2015: Large-scale homogenous partitions and bare metal reconfiguration generally available
- ▶ 2015/2016: Refinements to experiment management capabilities, higher level capabilities
- ▶ Fall 2016: Heterogeneous hardware available

# WORKING TOGETHER!

- ▶ Work on your next research project @ [www.chameleoncloud.org](http://www.chameleoncloud.org)!

*The most important element of any experimental testbed is users and the research they work on*

- ▶ Sign up for Early User at [chameleoncloud.org](http://chameleoncloud.org)
- ▶ Get engaged in the discussion:
  - ▶ Grid'5000 API, Blazar, representing appliances
  - ▶ Contribute appliances, research, ideas
  - ▶ Federation efforts: GENI & Grid'5000
- ▶ Come and work with us!