Chameleon

www. chameleoncloud.org

CHAMELEON: BUILDING A RECONFIGURABLE EXPERIMENTAL TESTBED FOR CLOUD RESEARCH

Kate Keahey

keahey@anl.gov

SWITCH ON
RESEARCH · COLLABORATION · EDUCATION

October 15th, 2015
Sao Paulo, Brazil

THE UNIVERSITY OF CHICAGO TACC NORTHWESTERN UNIVERSITY THE OHIO STATE UNIVERSITY UTSA NSF

CHAMELEON DESIGN STRATEGY

- 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: Chameleon appliances
- Complementary: "Can't do everything ourselves"
  - Complementing GENI, Grid'5000, and other experimental testbeds

CHAMELEON HARDWARE

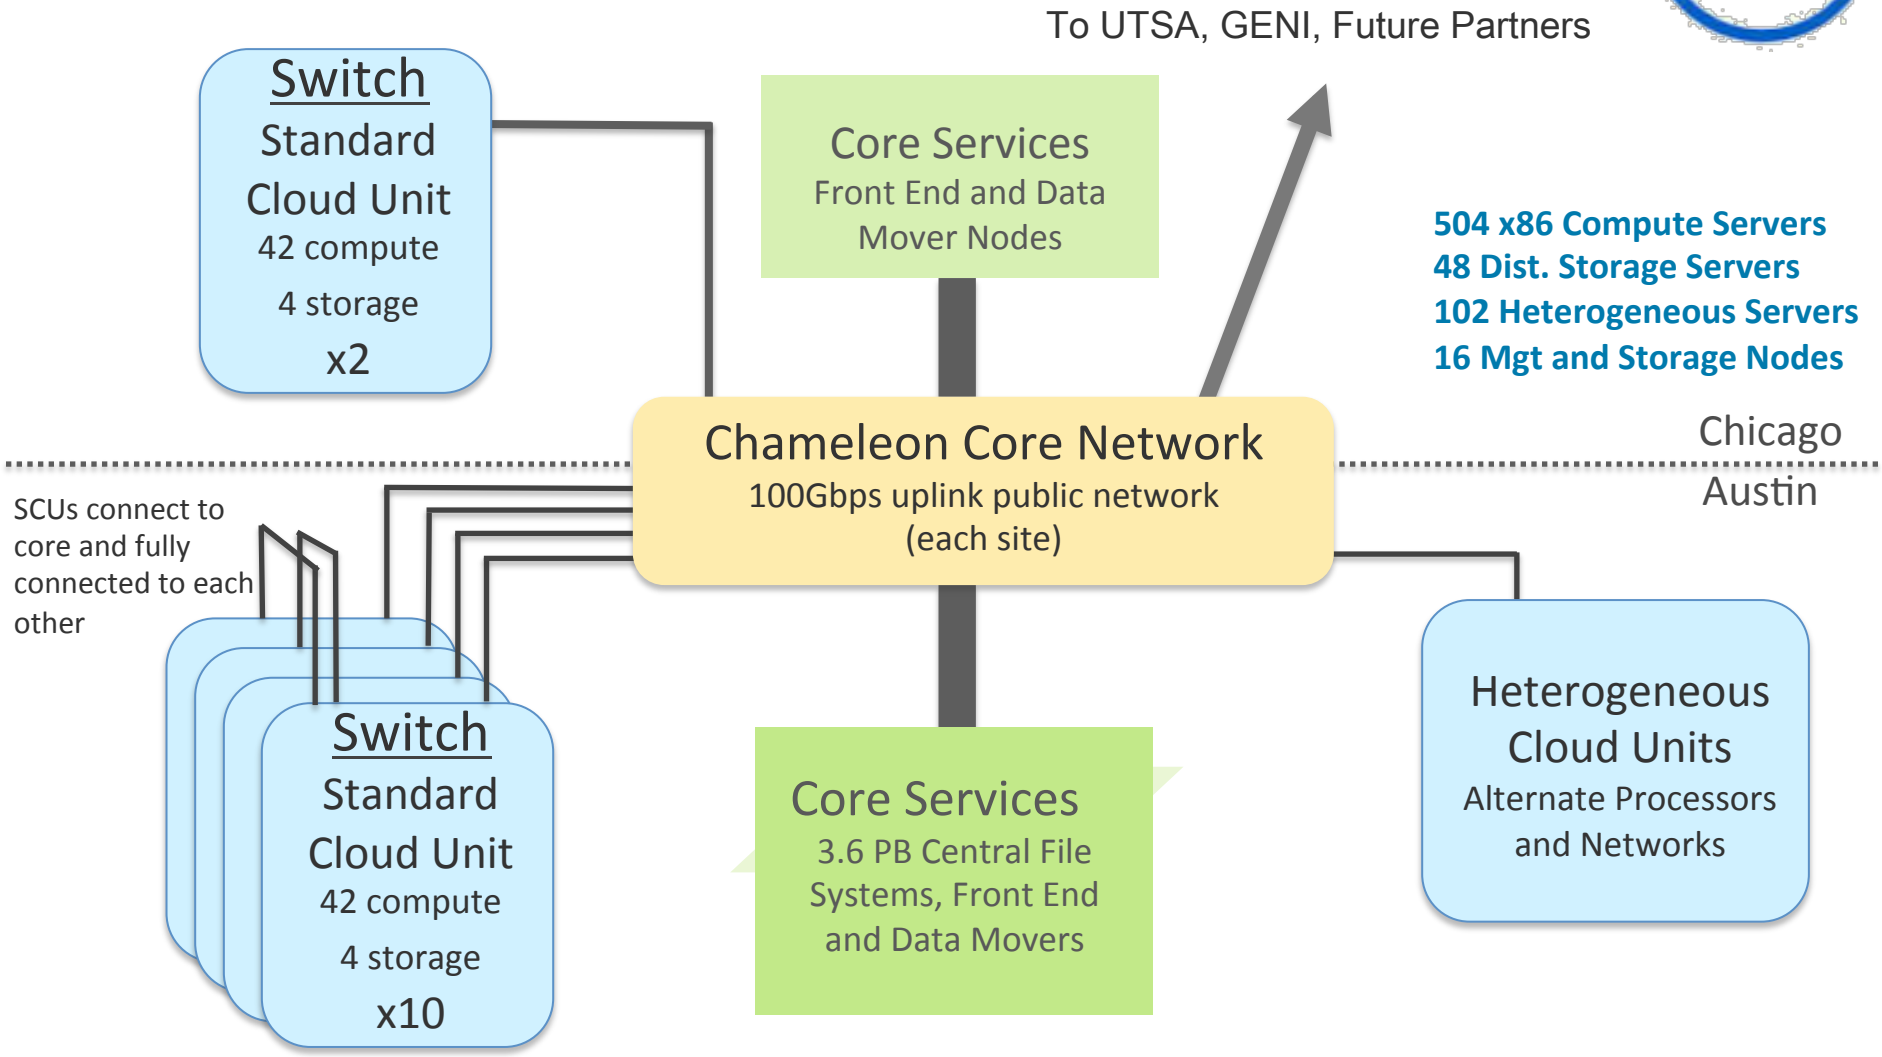

USING CHAMELEON: THE EXPERIMENTAL WORKFLOW

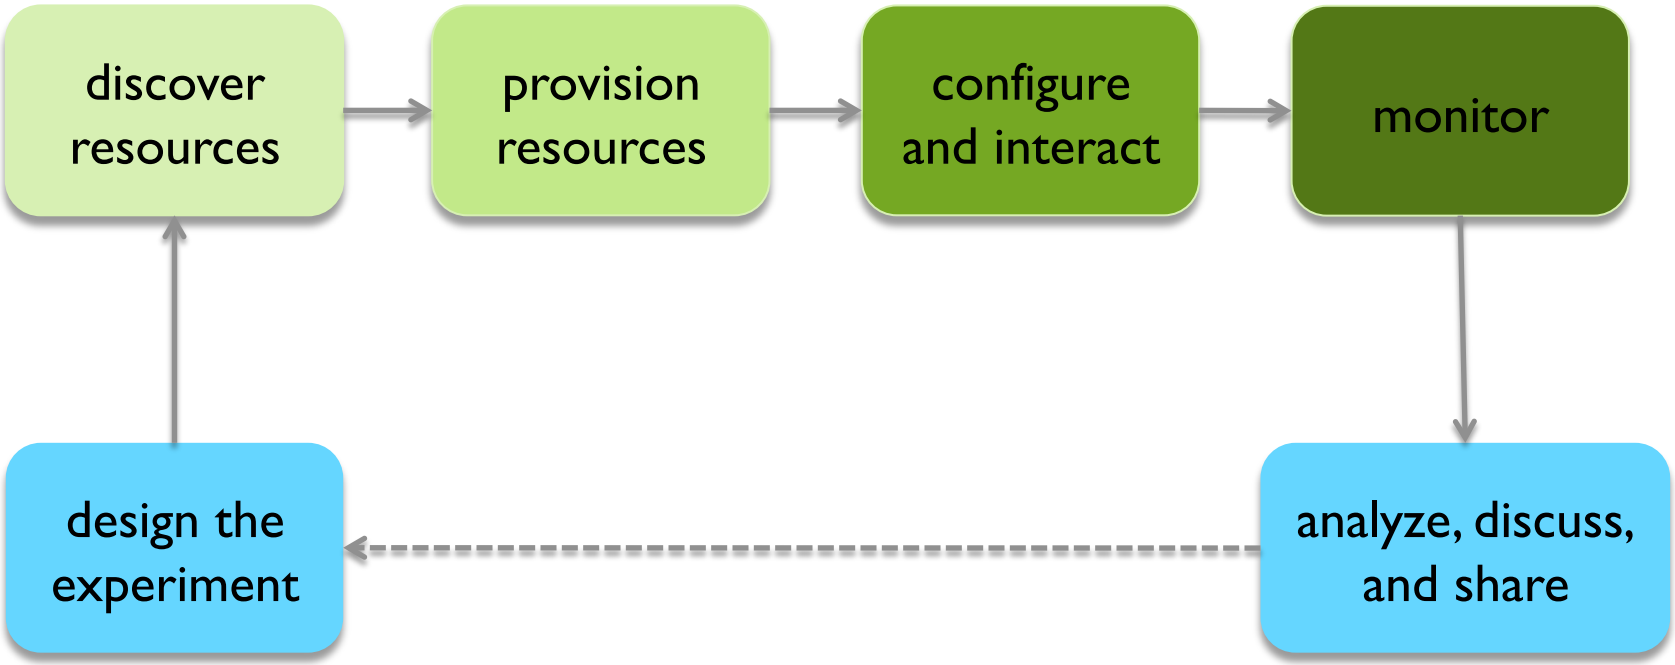

CHI: SELECTING AND VERIFYING RESOURCES

- 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 toolkit + Chameleon portal
  - Automated resource description, automated export to RM
- G5K-checks
  - Can be run after boot, acquires information and compares it with resource catalog description

CHI: PROVISIONING RESOURCES

- Resource leases
- Allocating a range of resources
  - Different node types, switches, etc.
- Multiple environments in one lease
- Advance reservations (AR)
  - Sharing resources across time
- Upcoming extensions: match making, internal management

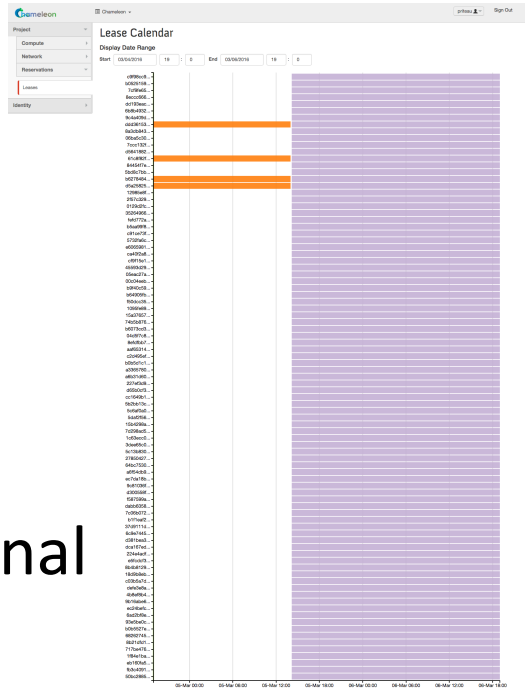

- OpenStack Nova/Blazar
- Extensions to support Gantt chart displays and other features

CHI: CONFIGURE AND INTERACT

- Map multiple appliances to a lease
- Allow deep reconfiguration (including BIOS)
- Snapshotting for image sharing
- 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

BUILDING CHI: CHAMELEON BARE METAL

- Defining requirements (proposal stage)
- Developing architecture
- Technology Evaluation and Risk Analysis
  - Rough requirements based analysis
  - Technology evaluation: Grid'5000 and OpenStack
  - Implementation proposals
- Implementing CHI
- Technology Preview deployment
- Early User and public availability

CHAMELEON STATUS AND TIMELINE

- 10/14: Project starts
- 12/14: FutureGrid@Chameleon (OpenStack KVM cloud)
- 04/15: Chameleon Technology Preview on FG hardware
- 06/15: Chameleon Early User on new homogenous hardware
- 07/15: Chameleon Public availability
- 09/15: Chameleon KVM OpenStack cloud available
- 10/15: Global storage available
- 2016: Heterogenous hardware available

CHAMELEON PROJECTS

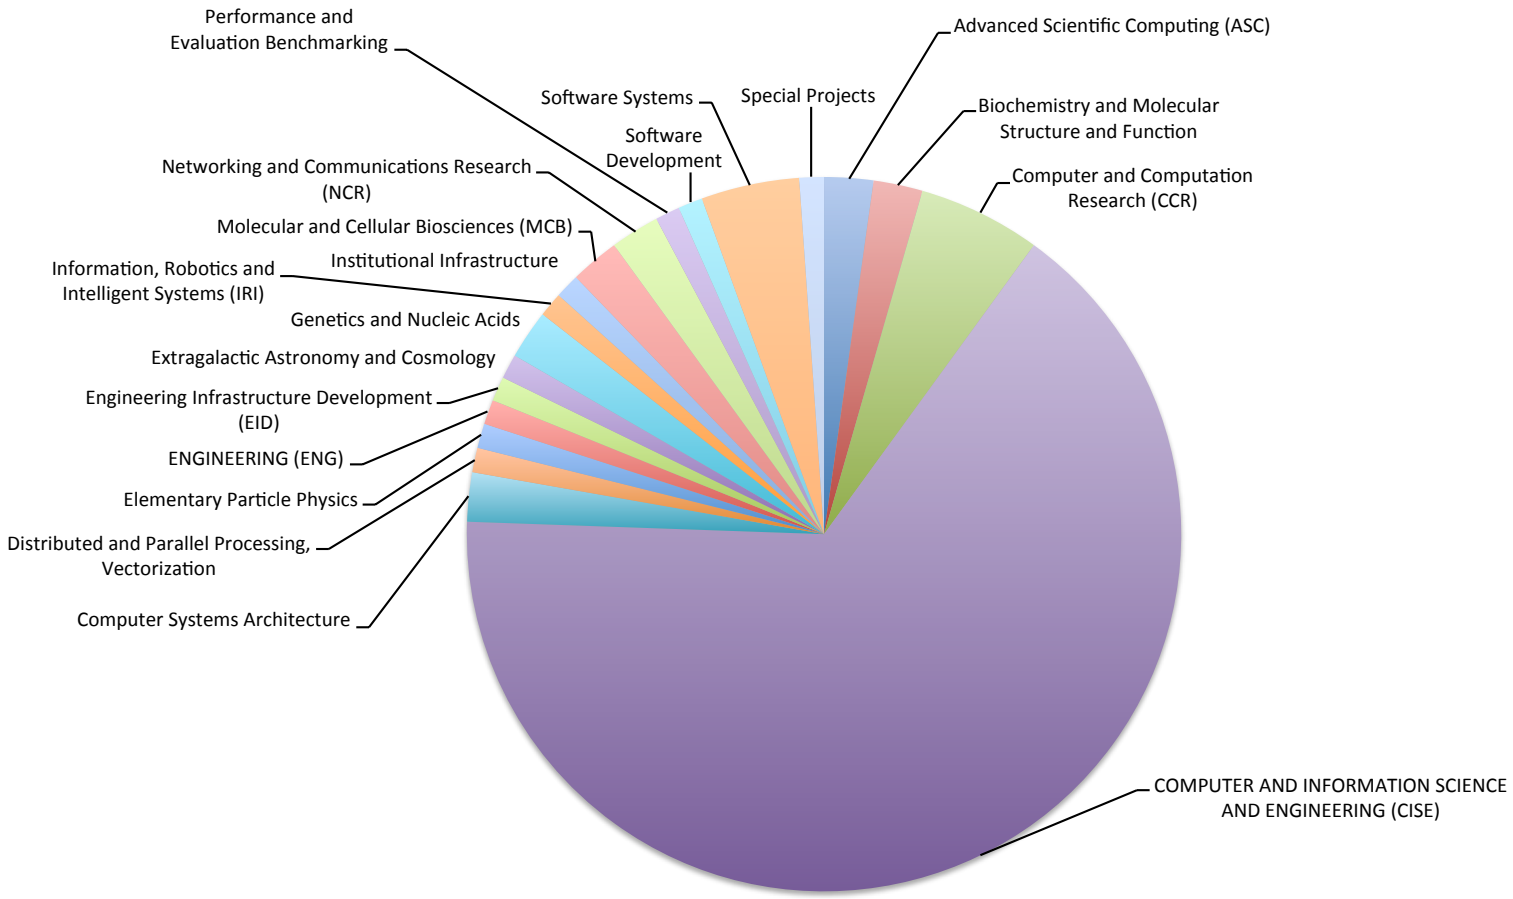

Overall: 101 projects, 187 users, 66 institutions

CHAMELEON TEAM

Kate Keahey
Chameleon PI
Science Director
Architect
University of Chicago

Paul Rad
Industry Liason
Education and training
UTSA

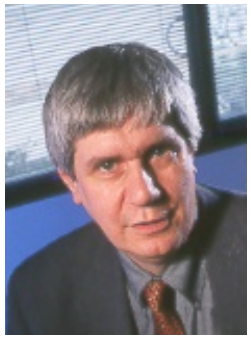

Joe Mambretti
Programmable networks
Federation activities
Northwestern University

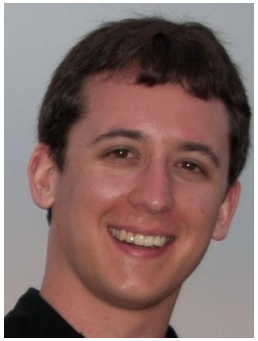

Pierre Riteau
Devops Lead
University of Chicago

DK Panda
High-perf networking
Ohio State University

Dan Stanzione
Facilities Director
TACC

PARTING THOUGHTS

- Work on your next research project @ www.chameleoncloud.org!

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

- Platform for collaboration
  - With Chameleon team: from innovative ways of extending the testbed to infrastructure research
  - With other users: sharing Chameleon appliances
  - With broader community: sharing traces, insights on CS experimentation, reproducibility, methodology