Chameleon

www. chameleoncloud.org

CHAMELEON: A LARGE-SCALE, RECONFIGURABLE EXPERIMENTAL ENVIRONMENT FOR CLOUD RESEARCH

Principal Investigator: Kate Keahey

Co-PIs: J. Mambretti, D.K. Panda, P. Rad, W. Smith, D. Stanzione

NB: With an Additional Brief Description of ClouldLab

14th Annual Global LambdaGrid Workshop
Queenstown, New Zealand
September 30 - October 1, 2014

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

SOLVING TODAY'S RESEARCH CHALLENGES

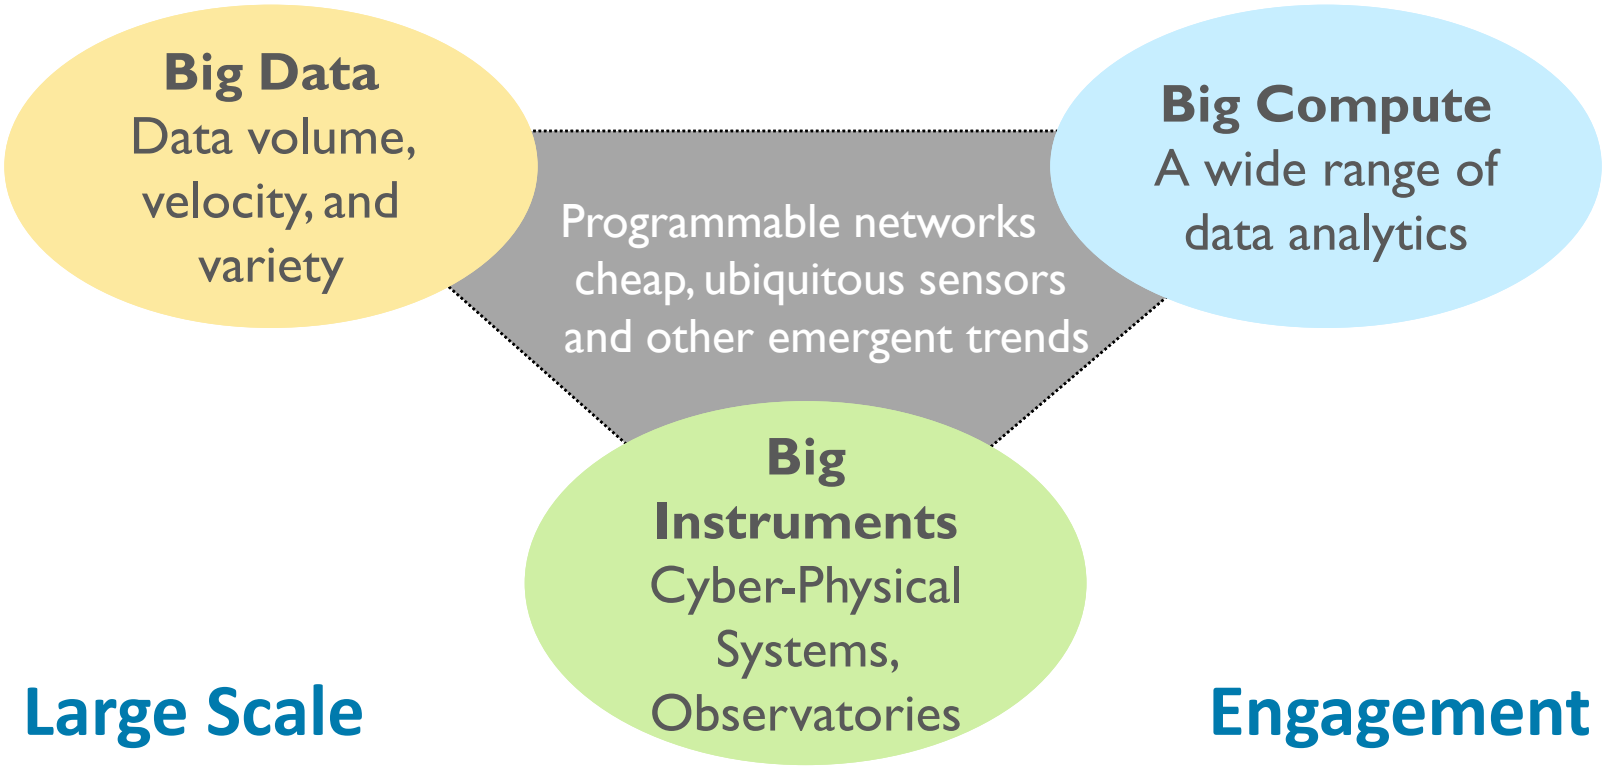

Large Scale **Engagement**

Reconfigurability **Connectedness**

NATIONAL SCIENCE FOUNDATION COMPUTER AND INFORMATION SCIENCE AND ENGINEERING (CISE) RESEARCH INFRASTRUCTURE PROGRAM

- The CISE Research Infrastructure (CRI) program drives discovery and learning in the computing disciplines by supporting the creation and enhancement of world-class computing research infrastructure.
- This infrastructure enables CISE researchers to advance the frontiers of CISE research. Through the CRI program, CISE seeks to ensure that individuals from a diverse range of academic institutions have access to such infrastructure.
- Mid-Scale Infrastructure - NSFCloud solicitation constitutes a track within the CRI program specifically supporting research infrastructure that enables the academic research community to develop and experiment with novel cloud architectures addressing emerging challenges, including real-time and high-confidence systems.
- Phase I for NSFCloud supports required infrastructure design and ramp-up activities, as well as demonstration of readiness for full-fledged execution.
- Phase II will enable the infrastructure to be fully staffed and operational, fulfilling the proposed mission of serving as a testbed that is used extensively by the research community.

CHAMELEON: A POWERFUL AND FLEXIBLE EXPERIMENTAL INSTRUMENT

- Large-scale instrument
  - Targeting Big Data, Big Compute, Big Instrument research
  - More than 650 nodes and 5 PB disk over two sites,100G network
- Reconfigurable instrument
  - Bare metal reconfiguration, operated as single instrument, graduated approach for ease-of-use
- Connected instrument
  - Workload and Trace Archive
  - Partnerships with production clouds: CERN, OSDC, Rackspace, Google, and others
  - Partnerships with users
- Complementary instrument
  - Complementing GENI, Grid'5000, and other testbeds
- Sustainable instrument
  - Strong connections to industry

CHAMELEON HARDWARE

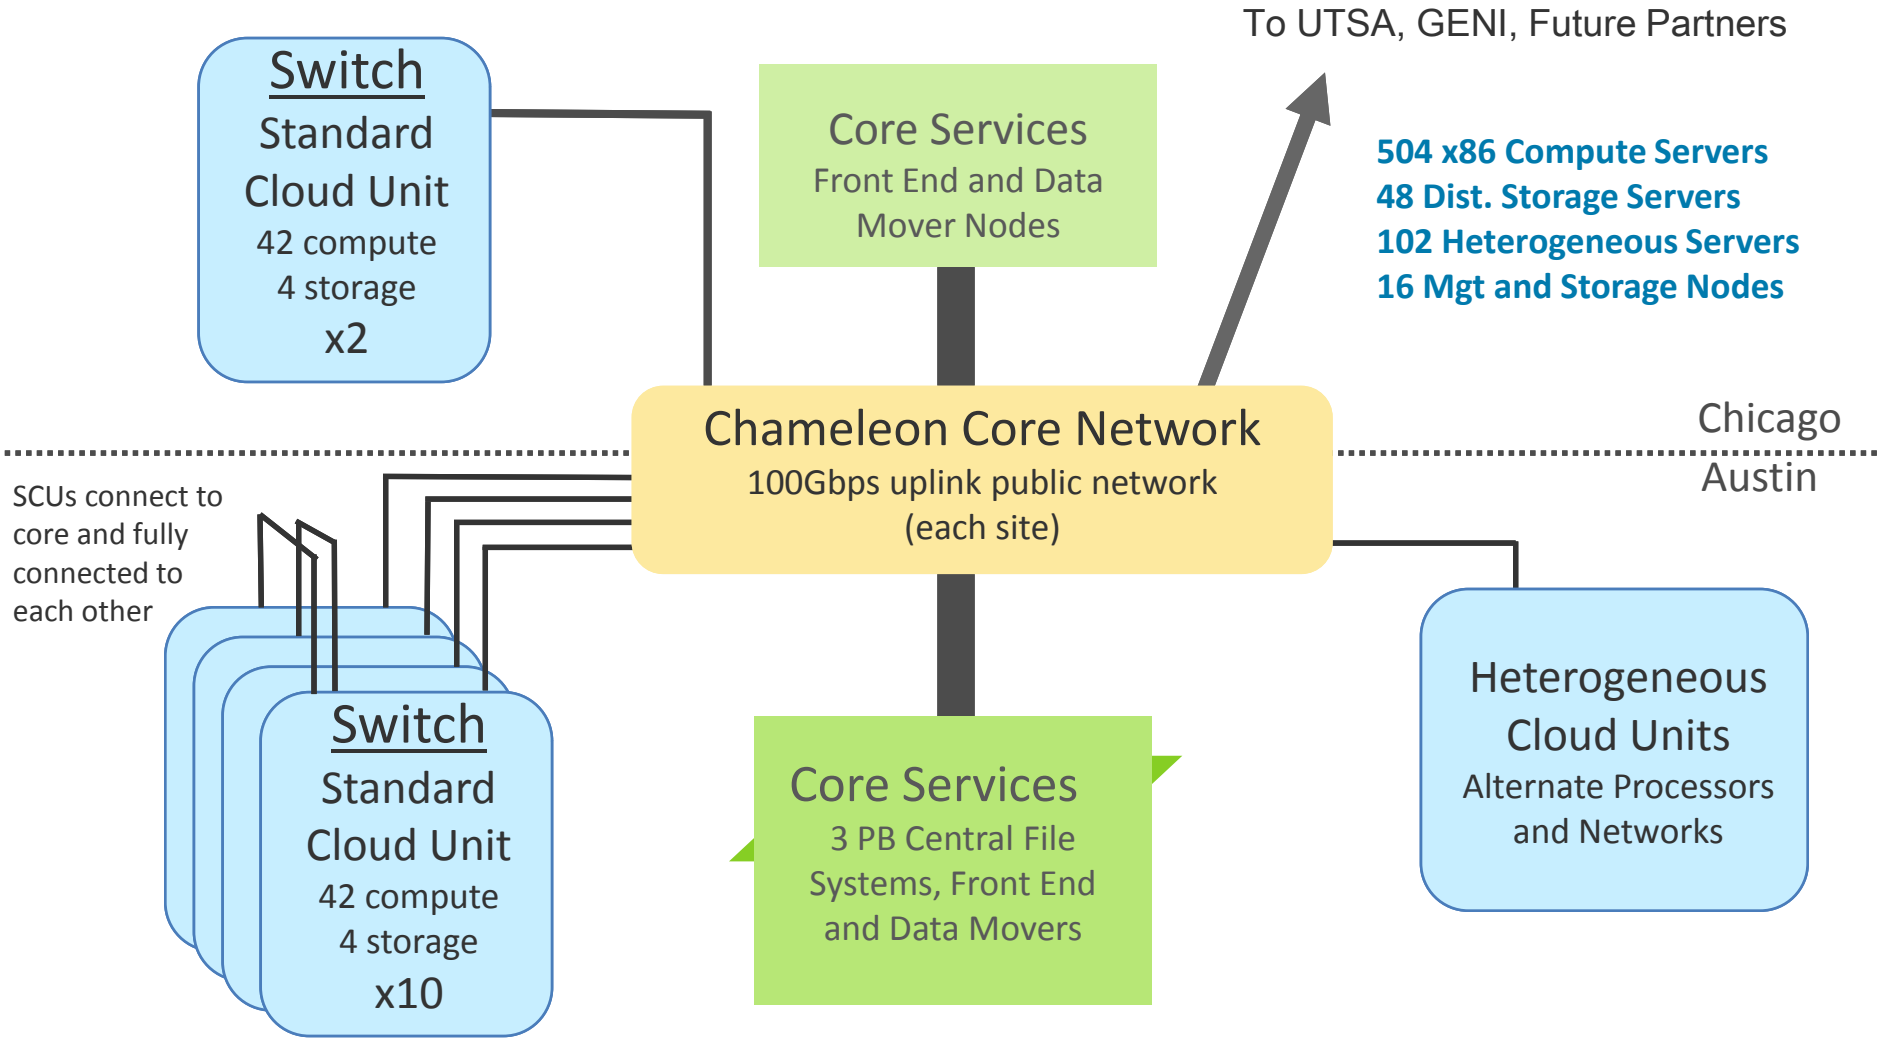

CAPABILITIES AND SUPPORTED RESEARCH

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, pre-configured images reconfiguration

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

Isolated partition, full bare metal reconfiguration

OUTREACH AND ENGAGEMENT

- Early User Program
  - Committed users, driving and testing new capabilities, enhanced level of support
- Chameleon Workshop
  - Annual workshop to inform, share experimental techniques solutions and platforms, discuss upcoming requirements, and showcase research
- Advisory Bodies
  - Research Steering Committee: advise on capabilities needed to investigate upcoming research challenges
  - Industry Advisory Board: provide synergy between industry and academia

PROJECT SCHEDULE

- Fall 2014: Hit the ground running: FutureGrid resources at UC and TACC available as OpenStack clouds
- Spring 2015: Maintain the momentum: Initial bare metal reconfiguration available on FutureGrid UC&TACC resources for Early Users
- Summer 2015: New hardware: large-scale homogenous partitions available to Early Users
- Fall 2015: Large-scale homogenous partitions generally available
- 2015/2016: Refinements to experiment management capabilities
- Fall 2016: Heterogeneous hardware available

INTERNATIONAL PARTNERSHIPS

- Multiple cloud testbeds are being implemented around the world
- This projects will engage with the communities developing those cloud testbeds
- A special focus will be interoperability among such testbeds
- These processes will build on existing international research testbeds such as the international SDN/OpenFlow testbed based on the GLIF


KUAS
NetFPGA
NOX
OpenFlow
NCKU
NetFPGA
NOX
OpenFlow
ETRI
NOX
OpenFlow
NetFPGA
KiSTi
University of Victoria
TOMATO
OpenFlow
InstaGENI
CRC
NOX
OpenFlow
NetFPGA
cenic
InstaGENI
OpenFlow
JGN-X
NSI
東京大学
THE UNIVERSITY OF TOKYO
VNode
PACIFIC GIGAPOP
aarnet
NICTA
ExoGENI
OpenFlow
NCHC
NetFPGA
NOX
OpenFlow
ofNSI
KRLight
NSI
canarie
APPLIED RESEARCH CENTER FOR COMPUTER NETWORKS
OpenFlow
GLORIAD
NSI
SingLight
SingAREN
OpenFlow
NetFPGA
Northern Wave
SoX
Southern Crossroads
GENI AM
OpenFlow
InstaGENI
renci
ExoGENI
OpenFlow
MREN
STARLIGHT
The Optical STAR TAP
GENI AM
ExoGENI
OpenFlow
NSI
ofNSI
INTERNET2
GENI AM
OpenFlow
BEN
Breakable Experimental Network
ExoGENI
Duke
UNIVERSITY
ExoGENI
OpenFlow
NORTHWESTERN UNIVERSITY
NetFPGA
NOX
OpenFlow
TOMATO
InstaGENI
AMPATH
NSI
OpenFlow
MAGPI
uen
UTAH EDUCATION NETWORK
GENI AM
OpenFlow
THE UNIVERSITY OF UTAH
VNode
OpenFlow
InstaGENI
RNP
NSI
NLight
NSI
NOX
OpenFlow
NetFPGA
UNIVERSITY OF AMSTERDAM
NSI
ExoGENI
Princeton University
OpenFlow
InstaGENI
USP
InstaGENI
TOMATO
TECHNISCHE UNIVERSITÄT KAISERSLAUTERN
DFN
Deutsches Forschungsnetz
GEANT2
NSI
ja.net
University of BRISTOL
OpenFlow
University of Essex
OpenFlow
grnet
NSI
Belnet
dedicated connectivity
iMinds
CONNECT.INNOVATE.CREATE
InstaGENI

TEAM

Kate Keahey
Chameleon PI
Science Director

Paul Rad
Industry Liason

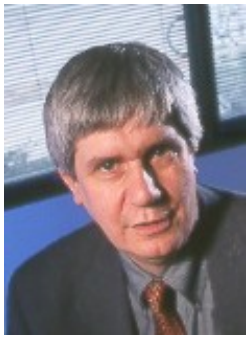

Joe Mambretti
Programmable networks

Warren Smith
Director of Operations

DK Panda
High-performance networks

Dan Stanzione
Facilities Director

CLOUDLAB: A SECOND NSFCLOUD PROJECT

- CloudLab is a large-scale distributed infrastructure based at the University of Utah, Clemson University and the University of Wisconsin, on top of which researchers will be able to construct many different types of clouds.
- Each site will have unique hardware, architecture and storage features, and will connect to the others via 100 gigabit-per-second connections on I2's advanced platform, supporting OpenFlow (an open standard that enables researchers to run experimental protocols in campus networks) and other software-defined networking technologies.
- CloudLab will be a facility where researchers can build their own clouds and experiment with new ideas with complete control, visibility and scientific fidelity.
- CloudLab will help researchers develop clouds that enable new applications with direct benefit to the public in areas of national priority
- In total, CloudLab will provide approximately 15,000 processing cores and in excess of 1 petabyte of storage at its three data centers. Each center will comprise different hardware, facilitating additional experimentation. In that capacity, the team is partnering with three vendors: HP, Cisco and Dell to provide diverse, cutting-edge platforms for research. Like Chameleon, CloudLab will feature bare-metal access. Over its lifetime, CloudLab is expected to run dozens of virtual experiments simultaneously and to support thousands of researchers.
- Other partners on CloudLab include Raytheon BBN Technologies, the University of Massachusetts Amherst and US Ignite, Inc.

PARTING THOUGHTS

- Chameleon is a large-scale, responsive experimental testbed
  - Targeting critical research problems at scale
  - Evolve with the community input
- Reconfigurable environment
  - Support use cases from bare metal to production clouds
  - Support for repeatable and reproducible experiments
- One-stop shopping for experimental needs
  - Trace and Workload Archive, user contributions, requirement discussions
- Engage the community
  - Network of partnerships and connections with scientific production testbeds and industry
  - Partnerships with existing testbeds
  - *Outreach activities, including to the international community*
- Come visit us at www.chameleoncloud.org!