



CloudLab





The Need Addressed by CloudLab

- Clouds are changing the way we look at a lot of problems
 - Impacts go far beyond Computer Science
- ... but there's still a lot we don't know, from perspective of
 - Researchers (those who will transform the cloud)
 - Users (those who will use the cloud to transform their own fields)
- To investigate these questions, we need:
 - Flexible, scalable **scientific infrastructure**
 - That enables exploration of **fundamental** science in the cloud
 - Built **by** and **for** the research community



The CloudLab Vision

- A “meta-cloud” for building clouds
- Build your own cloud on our hardware resources
- Agnostic to specific cloud software
 - Run existing cloud software stacks (like OpenStack, Hadoop, etc.)
 - ... or new ones built from the ground up
- Control and visibility all the way to the bare metal
- “Sliceable” for multiple, isolated experiments at once

With CloudLab, it will be as easy to get a cloud tomorrow as it is to get a VM today



What Is CloudLab?

Slice A

*Geo-Distributed
Storage Research*

Slice B

*Stock
OpenStack*

- Supports transformative cloud research
- Built on Emulab and GENI
- Control to the bare metal
- Diverse, distributed resources
- Repeatable and scientific

Slice C

*Virtualization and
Isolation Research*

Slice D

*Allocation and Scheduling Research
for Cyber-Physical Systems*

Utah

Wisconsin

Clemson

GENI

CC-NIE, Internet2 AL2S, Regionals



CloudLab's Hardware

One facility, one account, three locations

- About 5,000 cores each (15,000 total)
- 8-16 cores per node
- Baseline: 8GB RAM / core
- Latest virtualization hardware
- TOR / Core switching design
- 10 Gb to nodes, SDN
- 100 Gb to Internet2 AL2S
- *Partnerships with multiple vendors*

Wisconsin

- **Storage and net.**
- Per node:
 - 128 GB RAM
 - 2x1TB Disk
 - 400 GB SSD
- Clos topology
- *Cisco and HP*

Clemson

- **High-memory**
- 16 GB RAM / core
- 16 cores / node
- Bulk block store
- Net. up to 40Gb
- High capacity
- *Dell*


Utah

- **Power-efficient**
- ARM64 / x86
- Power monitors
- Flash on ARM64s
- Disk on x86
- Very dense
- *HP*



Technology Foundations

- Built on Emulab and GENI (“ProtoGENI”)
- In active development at Utah since 1999
- Several thousand users (incl. GENI users)
- Provisions, then gets out of the way
 - “Run-time” services are optional
- Controllable through a web interface and GENI APIs
- *Scientific instrument for repeatable research*
 - Physical isolation for most resources
 - *Profiles* capture everything needed for experiments
 - Software, data, and hardware details
 - Can be shared and published (eg. in papers)





Who can use CloudLab?

- US academics and educators
 - Researchers in cloud architecture and novel cloud applications
 - Teaching classes, other training activities
- No charge: free for research and educational use
- International federations expected
- Apply on the website at www.cloudlab.us

CloudLab Users So Far

May 2016:
300 projects
1,250 users
21,000 experiments





Cloud Architecture Research

- Exploring **emerging and extreme** cloud architectures
- Evaluating **design choices** that exercise hardware and software capabilities
- Studying **geo-distributed** data centers for low-latency applications
- Developing different **isolation** models among tenants
- Quantifying **resilience** properties of architectures
- Developing new **diagnostic** frameworks
- Exploring cloud architectures for **cyber-physical systems**
- Enabling **realtime** and near-realtime compute services
- Enabling data-intensive computing (“**big data**”) at high performance in the cloud



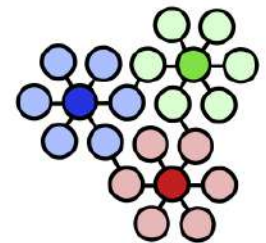
Application Research Questions

- Experiment with **resource allocation** and scheduling
- Develop enhancements to **big data frameworks**
- Intra- and inter-datacenter **traffic engineering** and routing
- New tenant-facing **abstractions**
- New **mechanisms** in support of cloud-based services
- Study adapting **next-generation stacks** to clouds
- New troubleshooting and **anomaly detection** frameworks
- Explore different degrees of **security** and isolation
- **Composing** services from heterogeneous clouds
- **Application-driven** cloud architectures



Federated with GENI

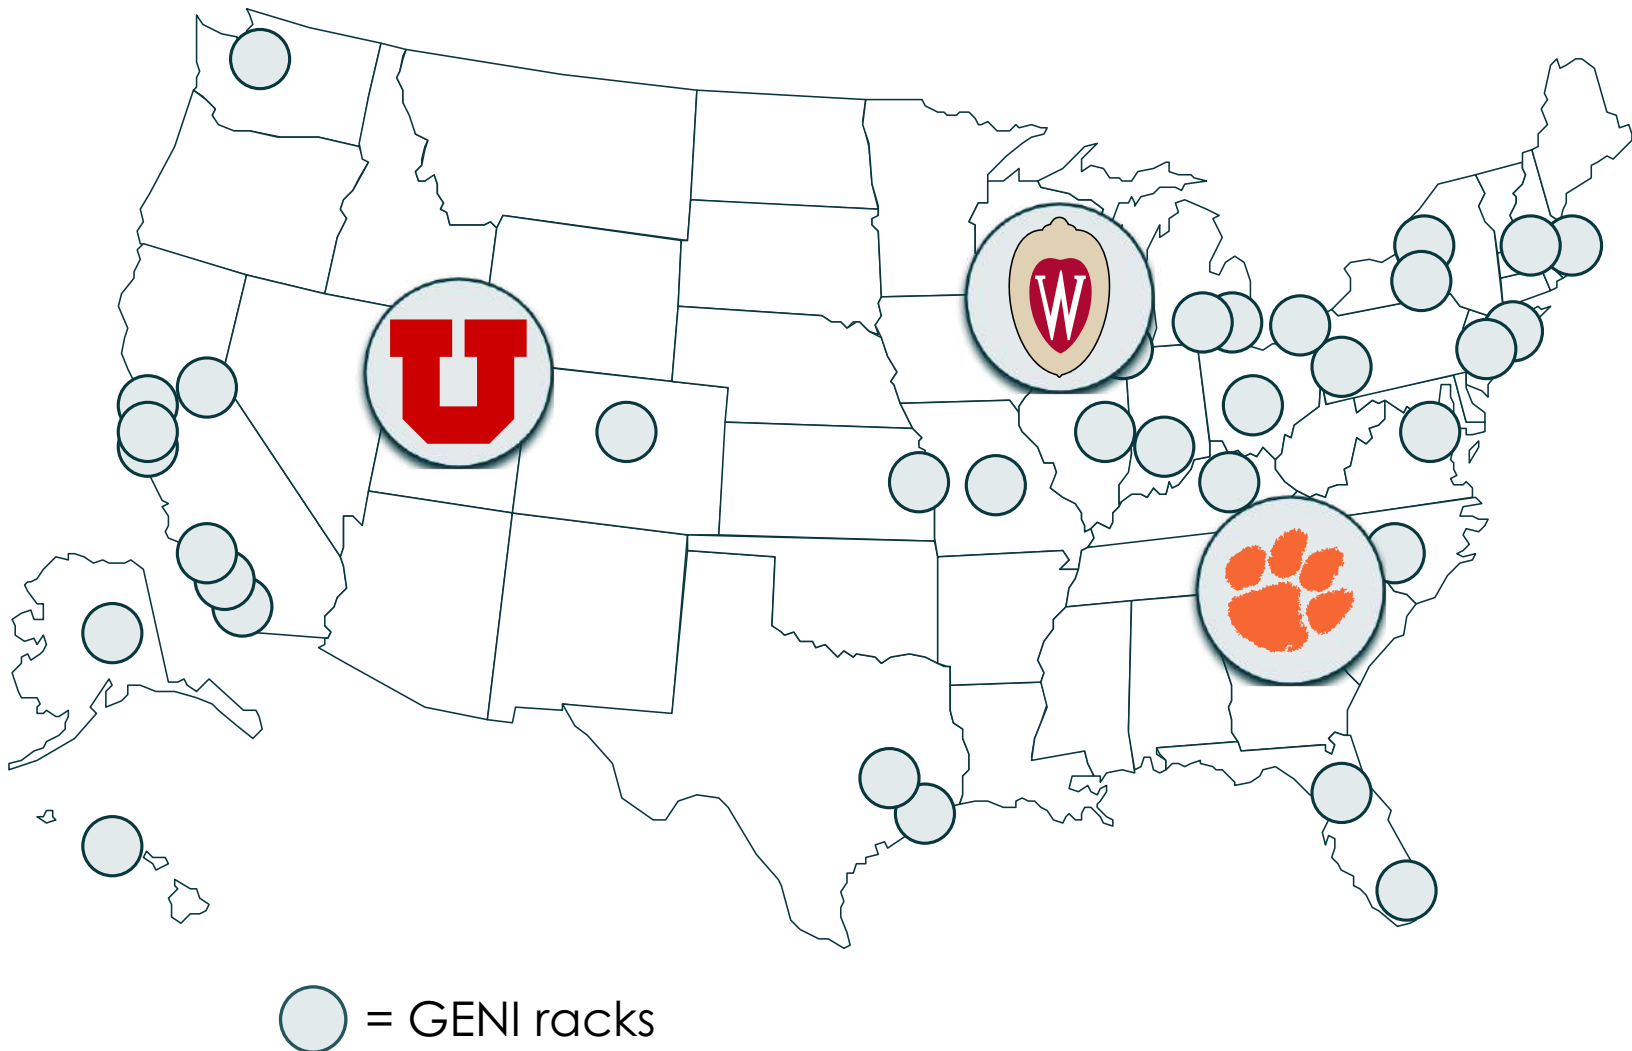
- *CloudLab can be used with a GENI account, and vice-versa*
- GENI Racks: ~ 50 small clusters around the country
- Programmable wide-area network
 - Openflow at dozens of sites
 - Connected in one layer 2 domain
- Large clusters (100s of nodes) at several sites
- Wireless and mobile
 - WiMax at 8 institutions
 - LTE / EPC testbed (“PhantomNet”) at Utah
- International partners
 - Europe (FIRE), Brazil, Japan



emulab



Many Sites, One Facility





Community Outreach

- Applications in areas of national priority
 - Medicine, emergency response, smart grids, etc.
 - Through 
- “Opt in” to compute jobs from domain scientists



- Summer camps
 - Through Clemson data-intensive computing program
- Under-represented groups





Availability and Schedule

- Availability:
 - ✓ **Open and in use!**
- Hardware being deployed in stages:
 - ✓ Fall 2014: Utah / HP cluster
 - ✓ Winter 2015: Wisconsin / Cisco cluster
 - ✓ Spring 2015: Dell / Clemson cluster
- Hardware refresh in early 2016
 - ✓ Spring 2016: Clemson and Wisconsin clusters
 - Summer 2016: Utah cluster planned




Your Own Cloud in One Click

CloudLab - Instantiate a Profile

Home

Manual

Actions ▾



ricci1 logged in

Logout

1. Select a Profile

2. Parameterize

3. Finalize

Selected Profile: OpenStack

A highly-configurable OpenStack instance with a controller, network manager, and one or more compute nodes (potentially at multiple sites). This profile runs x86 or ARM64 nodes. It sets up OpenStack Kilo or Juno on Ubuntu 15.04 or 14.10, and configures all OpenStack services, pulls in some VM disk images, and creates basic networks accessible via floating IPs. You'll be able to create instances and access them over the Internet in just a few minutes. When you click the Instantiate button, you'll be presented with a list of parameters that you can change to control what your OpenStack instance will look like; **carefully** read the parameter documentation to understand the various features available to you!


Copy Profile

Show Profile

Change Profile

Previous

Next

Powered by  emulab

Question or comment? Join the Help Forum

Supported by NSF © 2015 The University of Utah



The CloudLab Team



Robert Ricci (PI)
Eric Eide
Kobus Van der Merwe



Aditya Akella (co-PI)
Remzi Arpaci-Dusseau
Miron Livny



KC Wang (co-PI)
Jim Bottum
Jim Pepin



Chip Elliott (co-PI)
Larry Landweber



Mike Zink (co-PI)
David Irwin



Glenn Ricart (co-PI)





Learn more, sign up:

www.CloudLab.us



This material is based upon work supported by the National Science Foundation under Grant No. 1419199. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.