# Introducing CloudLab

## Scientific Infrastructure for Advancing Cloud Architectures and Applications

ROBERT RICCI, ERIC EIDE, AND THE CLOUDLAB TEAM

Robert Ricci is a research assistant professor in the School of Computing at the University of Utah. He is the principal investigator of the CloudLab project and a co-director of the Flux Research Group. He is one of the primary architects and implementers of the Flux Group's many testbeds, including Emulab, ProtoGENI, Apt, and, now, CloudLab.
ricci@cs.utah.edu

Eric Eide is a research assistant professor in the School of Computing at the University of Utah and a co-director of the Flux Research Group. His research focuses on the design and implementation of trustworthy systems software, including the use of testbeds for repeatable, experimental computer science.
eeide@cs.utah.edu

Additional members of the CloudLab team include Steve Corbató and Jacobus Van der Merwe, University of Utah; Aditya Akella, Remzi H. Arpaci-Dusseau, and Miron Livny, University of Wisconsin-Madison; K.C. Wang, Jim Bottum, James Pepin, and Amy Apon, Clemson University; Chip Elliott and Lawrence H. Landweber, Raytheon BBN Technologies; Michael Zink and David Irwin, University of Massachusetts Amherst; and Glenn Ricart, US Ignite.

Do you have an idea for improving cloud computing? Do you need to instantiate a complete cloud stack so that you can improve part of it? Replace part of it? Tune it to better support a particular scientific workload? This article introduces CloudLab (www.cloudlab.us), a new multi-site facility that we are building to support cloud research.

Researchers and practitioners are flush with ideas for tomorrow's cloud architectures. Their proposals range from small extensions of today's popular cloud-software stacks to all-new architectures that address mobility, energy efficiency, security and privacy, specific workloads, the Internet of Things, and on and on. Many of the ideas that drive modern clouds, such as virtualization, network slicing, and robust distributed storage arose from the research community. However, today's clouds have become unsuitable for moving this research agenda forward: they have specific, unmalleable implementations of the core technologies "baked in."

To support next-generation cloud research, the community needs infrastructure that is built to support research into a wide variety of cloud architectures. CloudLab is a new, large-scale, diverse, and distributed infrastructure designed to address this need. CloudLab is not itself a cloud. Rather, it is a substrate on which researchers can build their own clouds and experiment with them in an environment that provides a high degree of realism.

Like a commercial multi-tenant cloud, CloudLab will be used by many independent experimenters at any given time. In contrast to a commercial cloud, however, CloudLab is being built as a *scientific instrument*. It will give full visibility into every aspect of the facility, and it is being designed to minimize the impact that simultaneous experiments have on each other. This means that researchers using CloudLab will be able to fully understand why their systems behave the way they do, and they can have confidence that the results that they gather are not artifacts of competition for shared hardware resources.

CloudLab is currently under construction by a team located across the University of Utah, Clemson University, the University of Wisconsin-Madison, the University of Massachusetts Amherst, Raytheon BBN Technologies, and US Ignite. Like the team, the CloudLab facility will be geographically distributed, with large clusters at three sites. Each will be a variation on a "reference architecture" that comprises approximately 5,000 cores and 300–500 TB of storage in the latest virtualization-capable hardware. The diversity between sites will help CloudLab to support many areas of research and, at the same time, help researchers test the generality of their findings.

- The University of Utah site, partnered with HP, will be a cluster with both traditional x86-64 servers and a set of low-power ARM-based servers, enabling researchers to explore power/performance tradeoffs. The cluster will be connected by a large core switch, and it will offer experimenters direct access to switch hardware.
- The University of Wisconsin-Madison site, partnered with Cisco Systems, will closely reflect the technology and architecture used in modern commercial datacenters. Its 240 servers will have a total of about 4,000 cores and SSDs, and some nodes will have large numbers of disk spindles. They will be connected with a Clos network topology.
- The Clemson site, developed in cooperation with Dell, will have three components: bulk block-storage nodes, low-density storage nodes for MapReduce/Hadoop-style computing, and generic-VM nodes for provisioning virtual machines. This cluster will focus on provisioning significant experimental environments that can be linked to other national and international resources.

Within each site, CloudLab will provide two 10 Gbps network interfaces to every node. A high-bandwidth switching infrastructure supporting software-defined networking (SDN) will let researchers instantiate a wide range of in-cluster experimental topologies. CloudLab sites will connect with each other via IP and Layer 2 links to regional and national research networks, including AL2S, the SDN-based 100 Gbps network that is part of Internet2's Innovation Platform [7]. This will enable high-speed, end-to-end SDN between all CloudLab sites.

A CloudLab user will be able to provision resources from all of the CloudLab sites at once and combine them into a single experimentation environment. CloudLab's environments will also be able to connect at Layer 2 to the core GENI Network, US Ignite cities [11], and advanced HPC clusters across the United States. For example, in addition to resources from CloudLab's own clusters, a user's environment might include resources from GENI Racks [3], local fiber in a US Ignite city, or cyber-physical systems such as the U. Mass. CASA distributed weather radar system [10].

Like CloudLab's hardware, CloudLab's software is designed for diversity and flexibility in the cloud software stacks that researchers can deploy. CloudLab will be operated by a control framework that runs at a lower layer than cloud software stacks: It will directly provision and control "raw" hardware. A user will request a portion of the raw resources within CloudLab, thereby allocating a *slice* of the CloudLab facility for his or her exclusive use. The ability to allocate resources in this way is familiar to researchers who have used network testbeds such as Emulab [2, 12], GENI [4], DETER [6], and PRObE [8]. In fact, CloudLab's control software will be based on the proven software that today runs Emulab, several dozen Emulab-based sites, and also parts of GENI.

To allocate a slice of CloudLab, a user writes a *profile*, which is a description of everything needed to build a cloud: both the physical hardware (servers, disks, switches) and the software needed to transform it into a particular type of cloud. (A profile is therefore similar to the definition of an Emulab "experiment" [12] or a GENI "RSpec" [5].) Once the slice is allocated, its owner has full control over its resources. For example, the cloud stack running within the slice can create and manage virtual machines atop the physical machines that are part of the slice. CloudLab will provide canned configurations of popular cloud stacks (e.g., OpenStack [9]), storage systems (e.g., HDFS [1]), and computational frameworks (e.g., Hadoop [1]) so that experimenters can get something running quickly. Researchers will not be bound to these, however. They will be free to deploy whatever they wish on top of the resources provided by CloudLab.

Some researchers will want to create private clouds (e.g., for software development and controlled experimentation), while others will want to open their clouds to other users (e.g., to collect and evaluate real workloads). CloudLab will support both models of experimentation. In addition, researchers will be able to publish their CloudLab profiles, making it straightforward for others to reconstruct the hardware and software environments used in their studies. This will be a mechanism for repeating experiments and comparing results.

We are currently building CloudLab, but we expect that by the time you read this, one or more of the three CloudLab clusters will be up and available to early-access users who can help us to "shake out" the new hardware and software infrastructure. We have an aggressive timetable for building CloudLab, and our plan is that all of CloudLab will be open for regular use in spring 2015. The lessons learned from early adopters will drive the evolution and expansion of CloudLab going forward.

CloudLab will be available without charge to all US academic researchers and educators. In fact, if you have a GENI account or a Utah Emulab account, you can use CloudLab with your existing credentials! We encourage you to try CloudLab if you need modern infrastructure to help you invent the future of cloud computing. Visit the CloudLab Web site (www.cloudlab.us), sign up for news, and email us at contact@cloudlab.us.

This material is based upon work supported by the National Science Foundation under Grant Number 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.

***References***

[1] Apache Software Foundation: http://hadoop.apache.org/.

[2] Flux Research Group, University of Utah, Emulab—Network Emulation Testbed Home: http://www.emulab.net/.

[3] GENI Project Office, Raytheon BBN Technologies, Current GENI Rack Projects: http://groups.geni.net/geni/wiki/GENIRacksHome.

[4] GENI Project Office, Raytheon BBN Technologies: http://www.geni.net/.

[5] GENI Project Office, Raytheon BBN Technologies, Resource Specification (RSpec) Documents in GENI: http://groups.geni.net/geni/wiki/GENIExperimenter/RSpecs.

[6] Information Sciences Institute, University of Southern California, DeterLab: Cyber-Security Experimentation and Testing Facility: http://www.deterlab.net/.

[7] Internet2, Innovation Platform: http://www.internet2.edu/vision-initiatives/initiatives/innovation-platform/.

[8] New Mexico Consortium, NMC PRObE: http://www.nmc-probe.org/.

[9] OpenStack Foundation, Open Source Software for Building Private and Public Clouds: http://www.openstack.org/.

[10] University of Massachusetts Amherst, Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere: http://www.casa.umass.edu/.

[11] US Ignite, Ignite Communities: https://us-ignite.org/hub/.

[12] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar, "An Integrated Experimental Environment for Distributed Systems and Networks," in *Proceedings of the 5th Symposium on Operating Systems Design and Implementation (OSDI '02)*, December 2002, pp. 255–270.