



CloudLab





Why We're Building 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



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

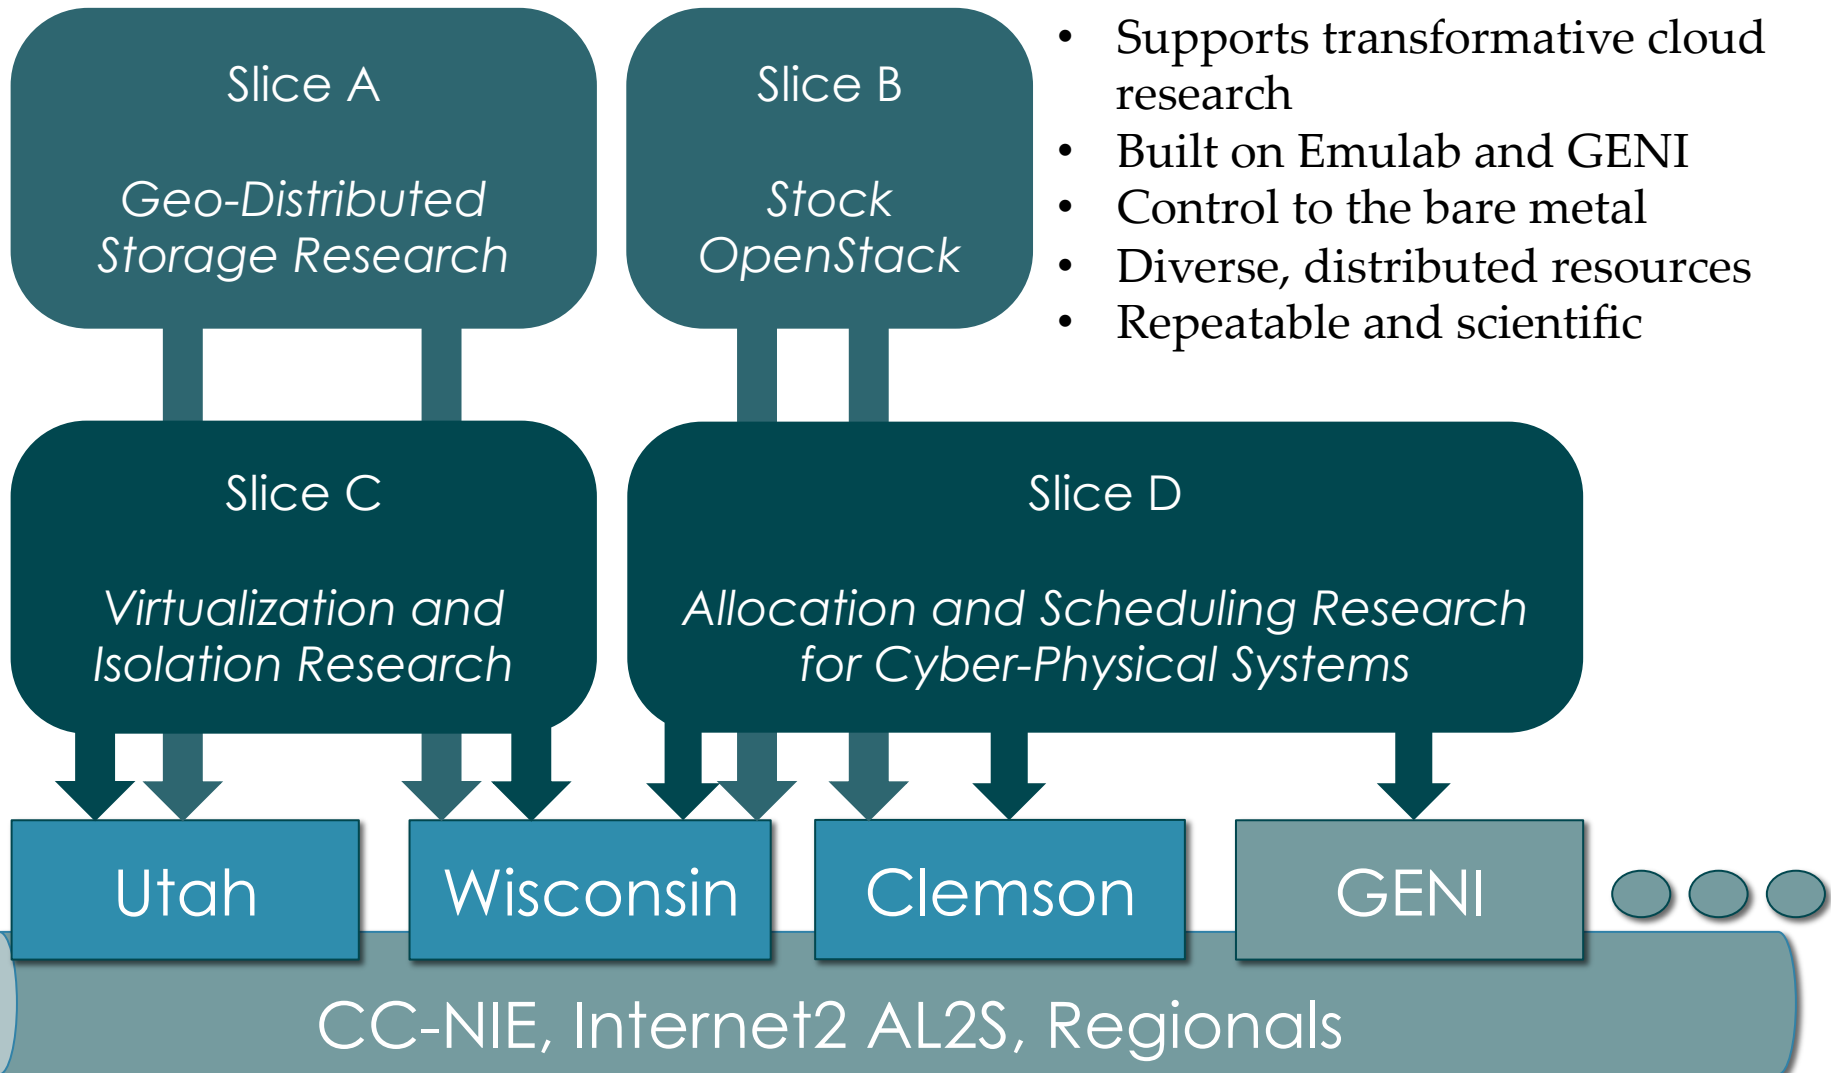


Application Research Questions

- 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




What Is CloudLab?





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)



emulab



geni
Exploring Networks
of the Future



protogeni



CloudLab's Hardware

One facility, one account, three locations

- About 5,000 cores each (15,000 total)
- 8-16 cores per node
- Baseline: 4GB 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*

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*



Hardware Schedule

Half this year

Refreshes in future years



Utah/HP: It Lives!





Utah/HP: Low-power ARM64

1.3

2 switches

315 nodes
2,520 cores
8.5 Tbps

8 cores


120 GB Flash

45 cartridges

64 GB RAM





Utah/HP Network: Core switch



320 Gb uplink

-
-
-

x7






Utah - Suitable for experiments that:

- ... explore power/performance tradeoffs
- ... want instrumentation of power and temperature
- ... want large numbers of nodes and cores
- ... want to experiment with RDMA via RoCE
- ... need bare-metal control over switches
- ... need OpenFlow 1.3
- ... want tight ARM64 platform integration
- ... want to use FPGAs (possibly)




Wisconsin/Cisco





Are we there yet?





Compute and storage

90X Cisco 220 M4



10X Cisco 240 M4



- 2X 8 cores @ 2.4GHz
 - 128GB RAM
 - 1X 480GB SSD
 - 2X 1.2 TB HDD
- 1X 1TB HDD
 - 12X 3TB HDD
(donated by Seagate)

Over the next year: ≥ 140 additional servers;

Limited number of accelerators, e.g., FPGAs, GPUs (planned)



Networking

Nexus 3132q



Nexus 3172pq



- OF 1.0 (working with Cisco on OF 1.3 support)
- Monitoring of instantaneous queue lengths
- Fine-grained tracing of control plane actions
- Support for multiple virtual router instances per router
- Support for many routing protocols



Experiments supported

Large number of nodes/cores, and bare-metal control over nodes/switches, for sophisticated network/memory/storage research

- ... Network I/O performance, intra-cloud routing (e.g., Conga) and transport (e.g., DCTCP)
- ... Network virtualization (e.g., CloudNaaS)
- ... In-memory big data frameworks (e.g., Spark/Shark)
- ... Cloud-scale resource management and scheduling (e.g., Mesos; Tetris)
- ... New models for Cloud storage (e.g., tiered; flat storage; IOFlow)
- ... New architectures (e.g., RAM Cloud for storage)



Clemson/Dell: High Memory, IB

20 cores/node

1 x 40 Gb IB/node

8 nodes/chassis

2*x 10 GbE OF/node

10 chassis/rack

2*x 1 GbE OF/node



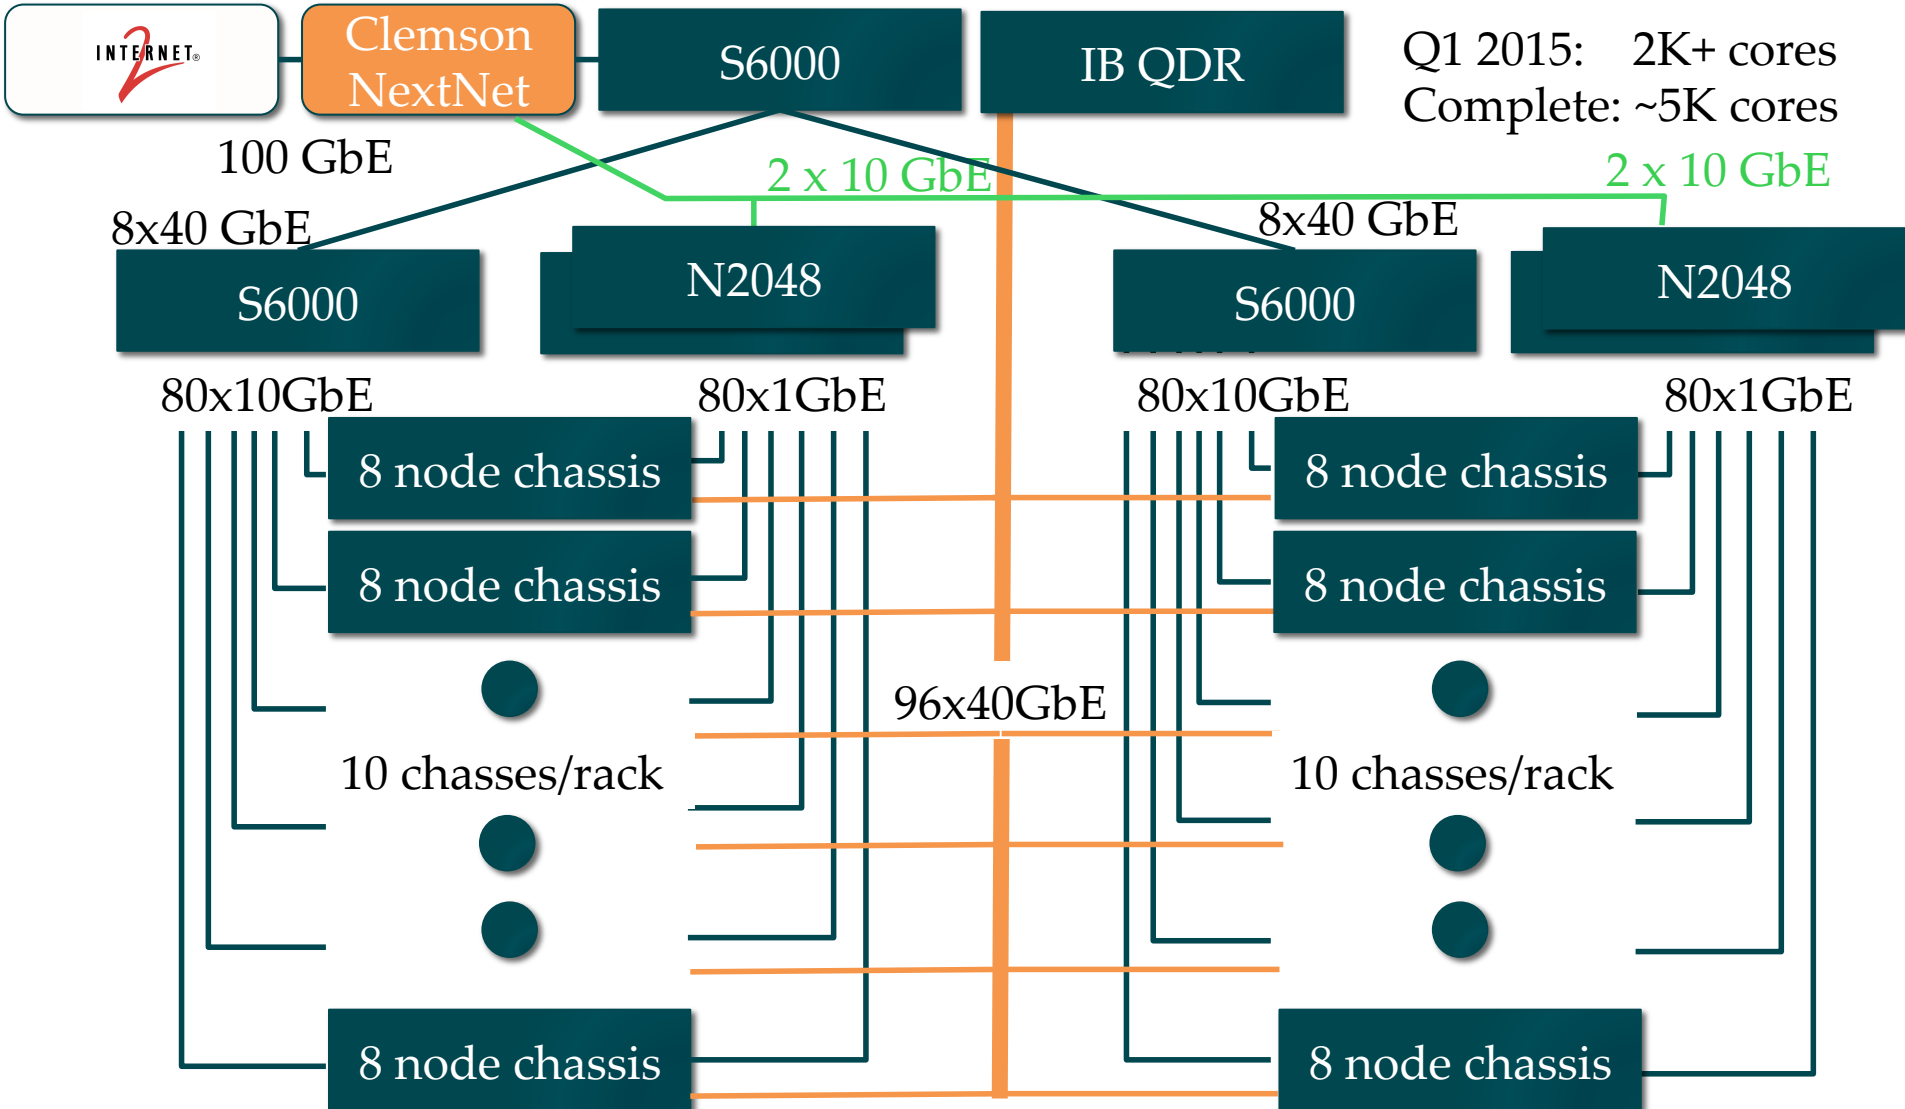
256 GB RAM/node

2 x 1 TB drive/server

* 1 NIC in 1st build



Clemson/Dell Network: IB + 10 GbE






Clemson - Suitable for experiments that:

- ... need large per-core memory
 - e.g., High-res media processing
 - e.g. Hadoop
 - e.g., Network Function Virtualization
- ... want to experiment with IB and/or GbE networks
 - e.g., hybrid HPC with MPI and TCP/IP
 - e.g., cyber physical system
- ... need bare-metal control over switches
- ... need OpenFlow 1.3



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






Early Interest in CloudLab





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






Demo

<https://www.cloudlab.us/p/emulab-ops/ARM64OpenStack>





cloudlab.us


Home Manual  Sign Up

Login

Username

Password

[Forgot Password?](#) [Geni User?](#) [Login](#)

Powered by  emulab Question or comment? Join the Help Forum Supported by NSF © 2014 The University of Utah



cloudlab.us

Select a Profile

Search

- ARM64OpenStack**
- arm64-ubuntu14-10
- arm64-ubuntu14
- OpenStack
- OnePC-Ubuntu14
- OneVM
- bstest2
- Ubuntu1404
- Docker

ARM64OpenStack

```
graph TD; compute3 --- central; compute1 --- central; compute2 --- central; central --- controller; central --- networkmanager;
```

An ARM64 OpenStack instance with a controller, network manager, and three compute nodes.

Select Profile Cancel



cloudlab.us

Home Manual Actions

ricci logged in Logout

What is CloudLab?

Run an Experiment

Selected Profile: ARM64OpenStack


An ARM64 OpenStack instance with a controller, network manager, and three compute nodes.

Change Profile

Your SSH key Update

Utah Cloudlab

Create!

Powered by  emulab Question or comment? Join the Help Forum Supported by NSF © 2014 The University of Utah



cloudlab.us

Home Manual Actions

ricci logged in Logout

Please wait while we get your experiment ready

URN: urn:publicid:IDN+emulab.net+slice+ricci-QV1031

State: created

Profile: ARM64OpenStack

Expires: Tomorrow at 11:46 AM (in 16 hours)

Extend Terminate



cloudlab.us

Home Manual Actions ricci logged in Logout

Experiment expires: Tomorrow at 11:46 AM (in 16 hours) >

Profile Instructions >

Topology View List View Manifest

Node	Shell (in-browser)	SSH command (if you provided your own key)	Console
controller		<code>ssh -p 22 ricci@ms0237.utah.cloudlab.us</code>	n/a
networkmanager		<code>ssh -p 22 ricci@ms0212.utah.cloudlab.us</code>	n/a
compute1		<code>ssh -p 22 ricci@ms0228.utah.cloudlab.us</code>	n/a
compute2		<code>ssh -p 22 ricci@ms0226.utah.cloudlab.us</code>	n/a
compute3		<code>ssh -p 22 ricci@ms0219.utah.cloudlab.us</code>	n/a

Powered by emulab Question or comment? Join the Help Forum Supported by NSF © 2014 The University of Utah



controller.ricci-qv1031.emulab-net.utah.cloudlab.us

ubuntu[®] OpenStack Dashboard

Log In

User Name

admin

Password

.....



Sign In



controller.ricci-qv1031.emulab-net.utah.cloudlab.us

Admin

System

Overview

Hypervisors

Host Aggregates

Instances

Flavors

Images

Networks

Routers

Defaults

System Information

Identity

Hypervisor Summary



VCPU Usage
Used 0 of 24



Memory Usage
Used 1.5GB of 188.9GB



Disk Usage
Used 0Bytes of 327GB

Hypervisor Compute Host

Hypervisors

Hostname	Type	VCPUs (used)	VCPUs (total)	RAM (used)	RAM (total)	Storage (used)	Storage (total)	Instances
compute2.ricci-qv1031.emulab-net.utah.cloudlab.us	QEMU	0	8	512MB	63GB	0Bytes	109GB	0
compute3.ricci-qv1031.emulab-net.utah.cloudlab.us	QEMU	0	8	512MB	63GB	0Bytes	109GB	0
compute1.ricci-qv1031.emulab-net.utah.cloudlab.us	QEMU	0	8	512MB	63GB	0Bytes	109GB	0

Displaying 3 items



Building Your Own Profiles





Copy an Existing Profile

The screenshot shows the CloudLab web interface. At the top, there is a navigation bar with "Home", "Manual", and "Actions" buttons, a central logo, and a "Logout" button. The user is logged in as "rpruser". A green notification box states "Your experiment is ready!" with a right-pointing arrow. Below this, the experiment details are listed: URN: urn:publicid:IDN+emulab.net+slice+rpruser-QV992, State: ready, Profile: arm64-ubuntu14, and Expires: 12-07T21:24Z (in 16 hours). At the bottom right of this box, there are three buttons: "Clone" (blue), "Extend" (green), and "Terminate" (red). The "Clone" button is circled in orange. Below the notification box is a blue "Profile Instructions" button with a right-pointing arrow. At the bottom of the page, there is a tabbed interface with "Topology View" selected, and other tabs for "List View", "Manifest", and "node" (with a close icon).

cloudlab.us

Home Manual Actions

rpruser logged in Logout

Your experiment is ready! >

URN: urn:publicid:IDN+emulab.net+slice+rpruser-QV992
State: ready
Profile: arm64-ubuntu14
Expires: 12-07T21:24Z (in 16 hours)

Clone Extend Terminate

Profile Instructions >

Topology View List View Manifest node ^x



Use a GUI

The screenshot shows a web browser window at `cloudlab.us` displaying the **Topology Editor** interface. The interface includes a sidebar on the left with configuration options and a main workspace on the right showing a network topology diagram.

Configuration Sidebar:

- Custom Type
- Hardware Type**
Dropdown menu: (any)
- Custom Hardware
- Disk Image**
Dropdown menu: Ubuntu 12.04 LTS 64-bit
- Custom Disk Image
- Install Scripts** Add
 - URL:
Input field: ex: `http://example.com/mystuff.tai`

Topology Diagram:

- A central square node is connected to five peripheral server nodes: `cloud-controller`, `name-node`, `worker-1`, and `worker-5`.
- The `cloud-controller` node is highlighted with a green border.
- Buttons for **Tidy View** and **Delete All** are located in the top right of the workspace.



Write Python Code

```
two-vm.py (~/Desktop) - VIM
import geni.rspec.pg as pg
rspec = pg.Request()

# Create XenVM nodes
node1 = pg.XenVM("node1")
node2 = pg.XenVM("node2")

# Create interfaces for each node.
iface1 = node1.addInterface("if1")
iface2 = node2.addInterface("if2")

rspec.addResource(node1)
rspec.addResource(node2)


# Create a link with the type of LAN.
link = pg.LAN("lan")

# Add both node interfaces to the link.
link.addInterface(iface1)
link.addInterface(iface2)
```




Build From Scratch

cloudlab.us


Home Manual Actions  rpruser logged in Logout


Create Profile

Name 

Project

Your rspec

Description 

Instructions 

List on the home page for anyone to view.

Who can instantiate your profile?

Anyone on the internet (guest users)

Only registered users of the website



Availability

Technology preview since August

Available for use now!





Sign Up At CloudLab.us

The screenshot shows a web browser window with the URL "cloudlab.us". The page has a dark teal header with navigation links for "Home" and "Manual" on the left, the CloudLab logo in the center, and "Sign Up" and "Login" buttons on the right. The main content area is titled "Start Project" and is divided into two columns: "Personal Information" and "Project Information".

Start Project

Personal Information

Username

Full Name

Email

Institutional Affiliation

Please Select Country

Please Select State

City

Project Information

Join Existing Project Start New Project

Project Name

Project Title (short sentence)

Project Page URL

Project Description (details)



The CloudLab Team



Robert Ricci (PI)
Eric Eide
Steve Corbató
Kobus Van der Merwe



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



KC Wang (co-PI)
Jim Bottum
Jim Pepin
Amy Apon



Chip Elliott (co-PI)
Larry Landweber



Mike Zink (co-PI)
David Irwin



Glenn Ricart (co-PI)





...or use your GENI or Emulab account

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.