Looking Beyond the Internet

The Rise of Software Defined Infrastructure

SwitchOn Workshop – Brazil / US at FIU

January 2015, Miami

Chip Elliott, BBN
celliott@bbn.com

My thesis

- SDN is just an opening act
- A major transformation of the Internet has begun
- We can now catch glimpses of what lies beyond
- We can get there by a series of step by step actions

Where I am coming from - GENI

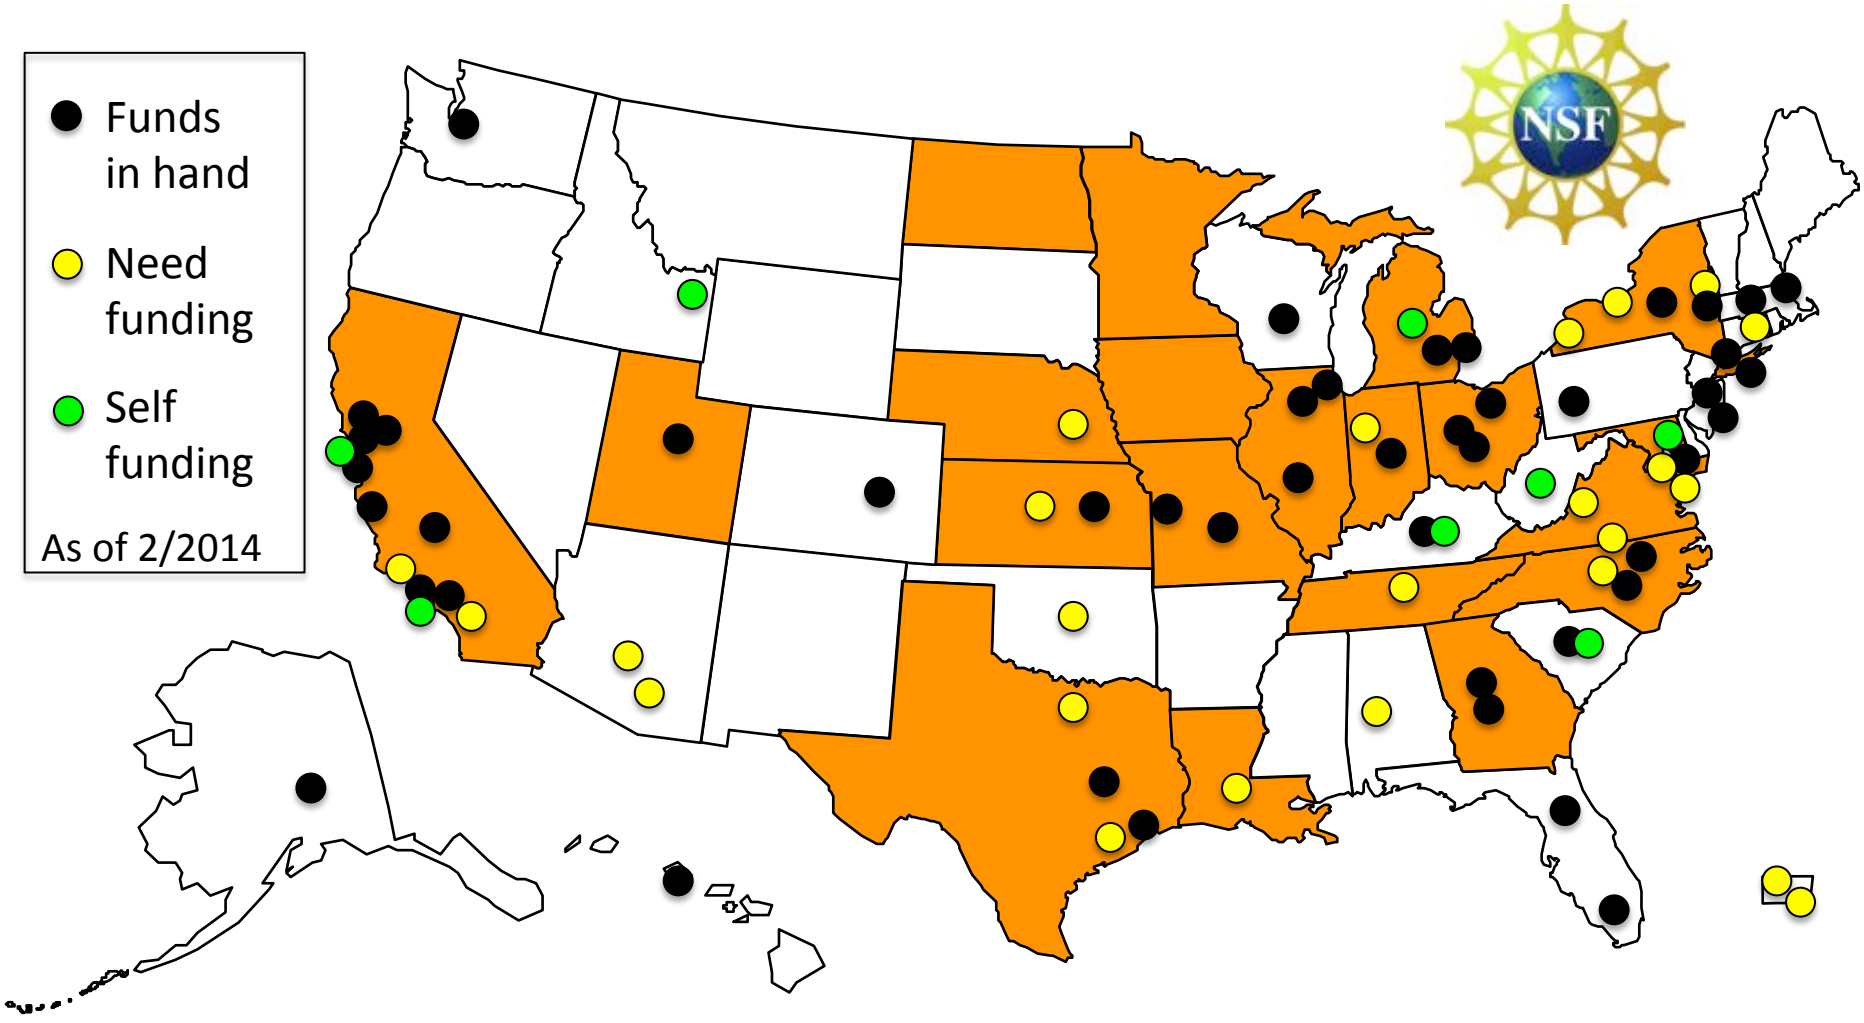

We're building out GENI through universities across the US

Slices and deep programmability

Install the software I want *throughout* my network slice
(into firewalls, routers, clouds, …)

And keep my slice isolated from your slice,
so we don't interfere with each other

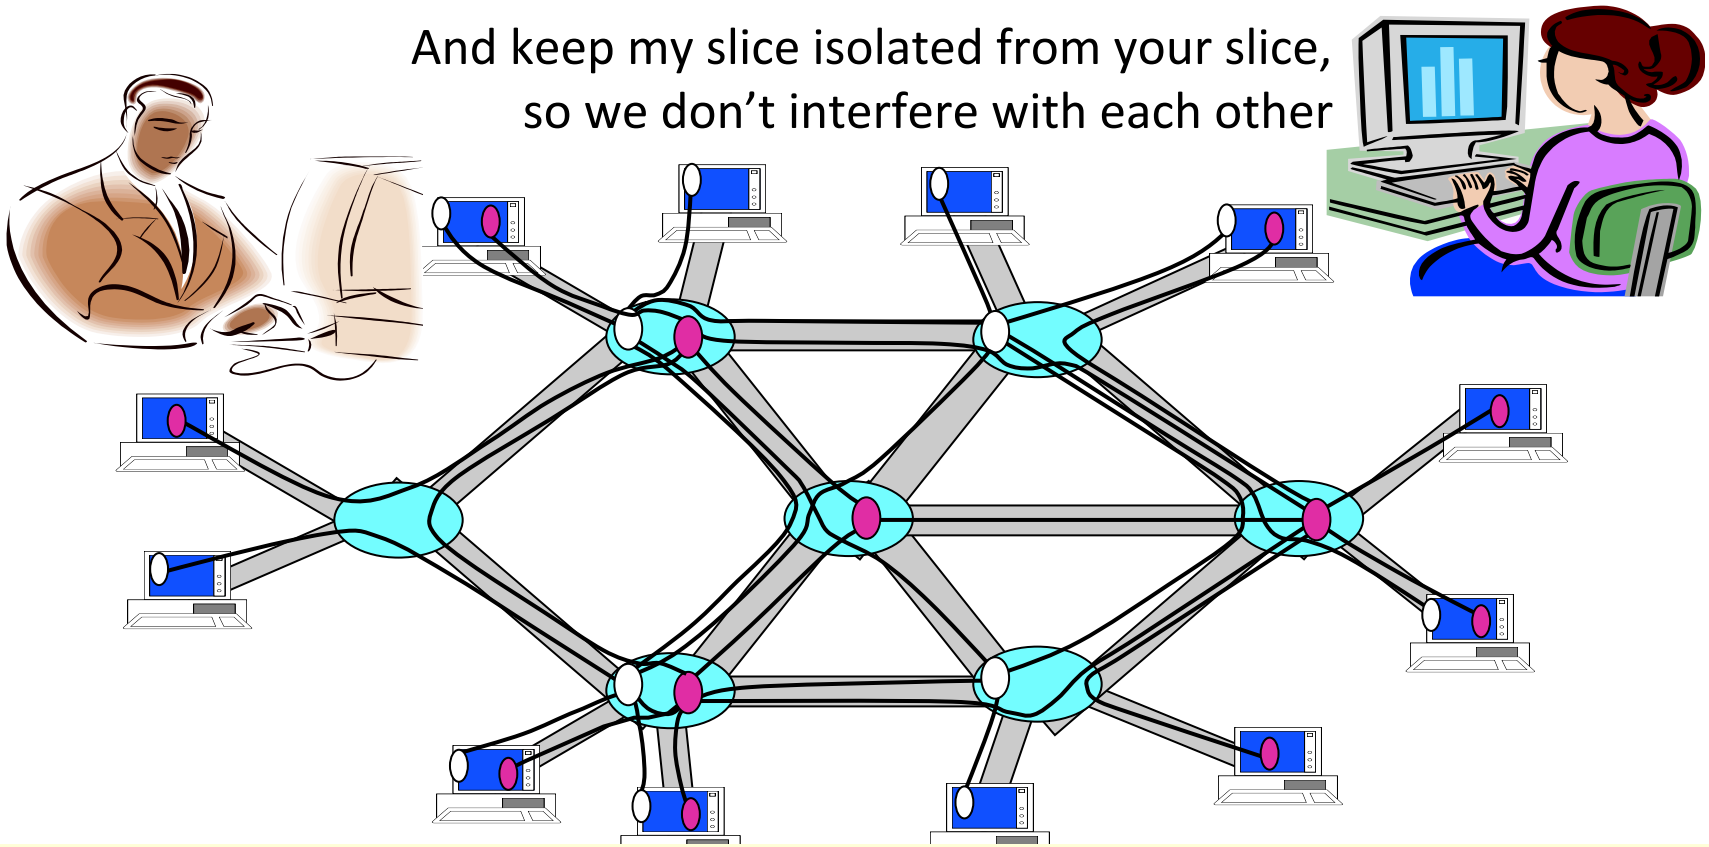

We can run many different "future internets" in parallel

Agile, deeply programmable infrastructure

- Emerging technologies that enable coherent network / processor / storage virtualization provide a great basis for agile cyber infrastructure.

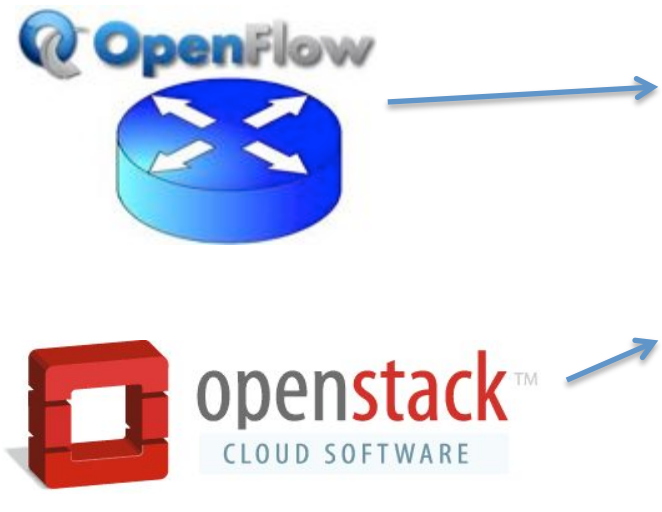

GENI Rack – OpenFlow switch
with sliced compute and storage

GENI Rack providers

Slices span many organizational boundaries

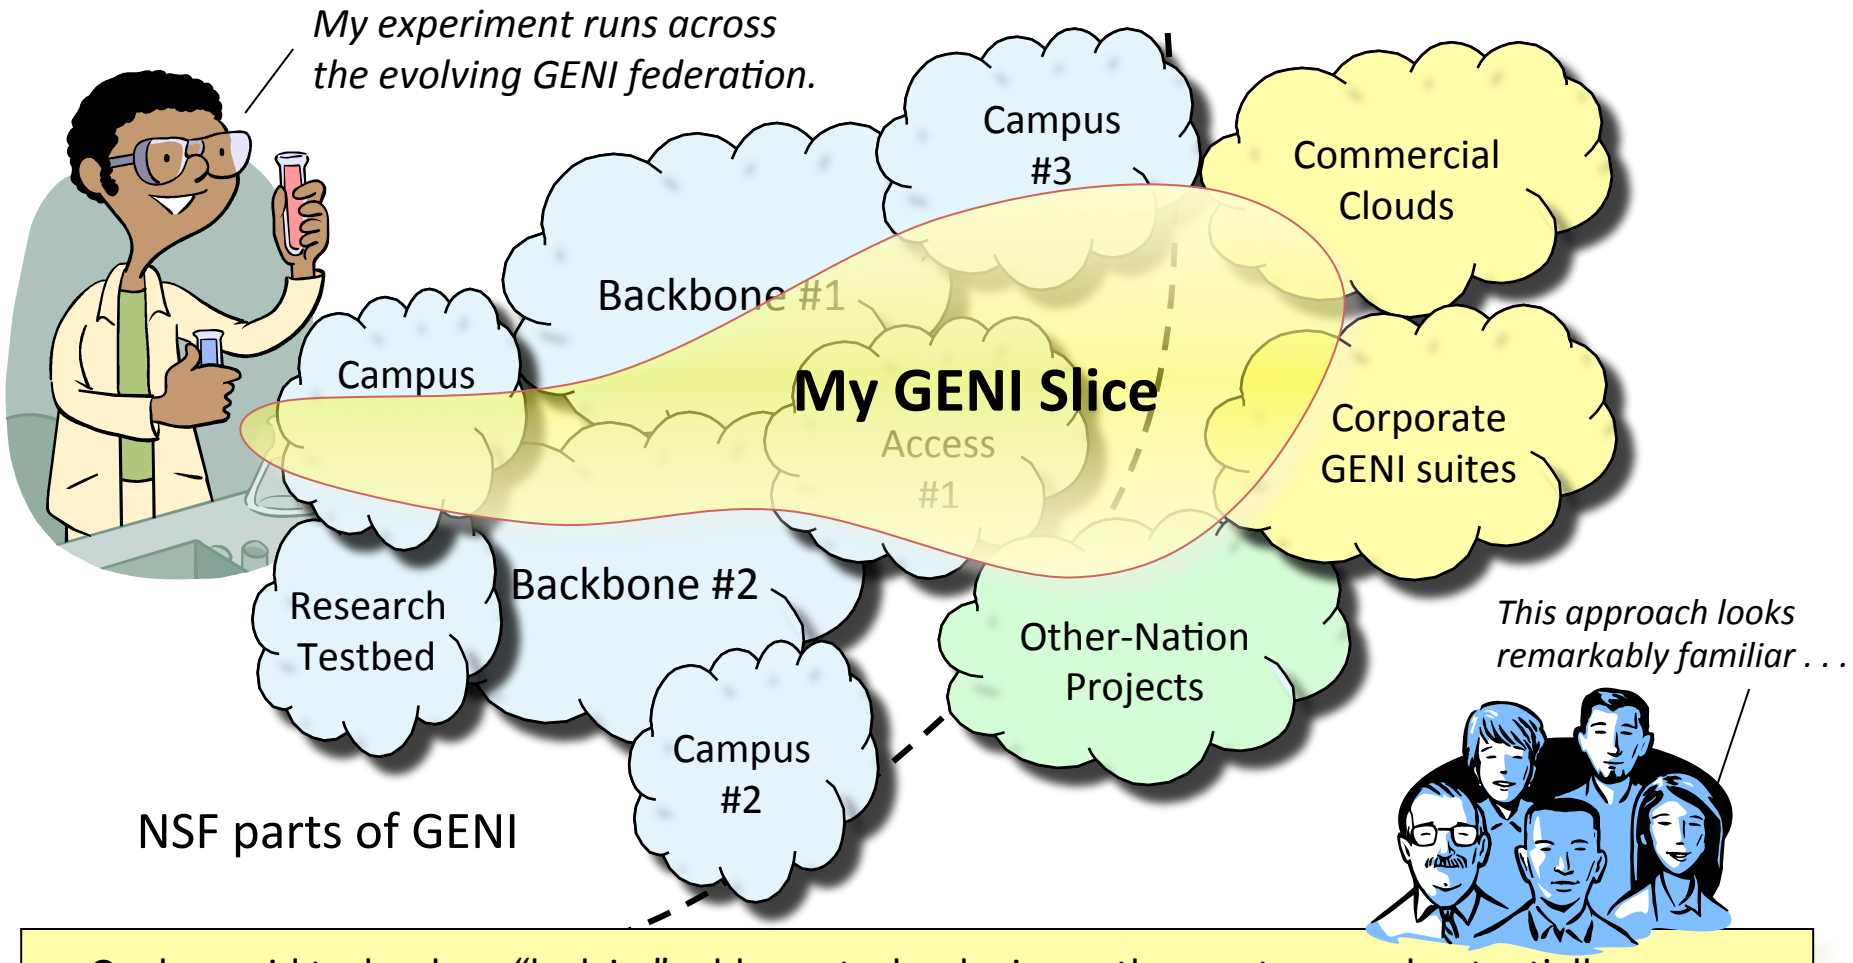

Goals: avoid technology "lock in," add new technologies as they mature, and potentially grow quickly by incorporating existing infrastructure into the overall "GENI ecosystem"

And it works! GENI is seeing heavy use

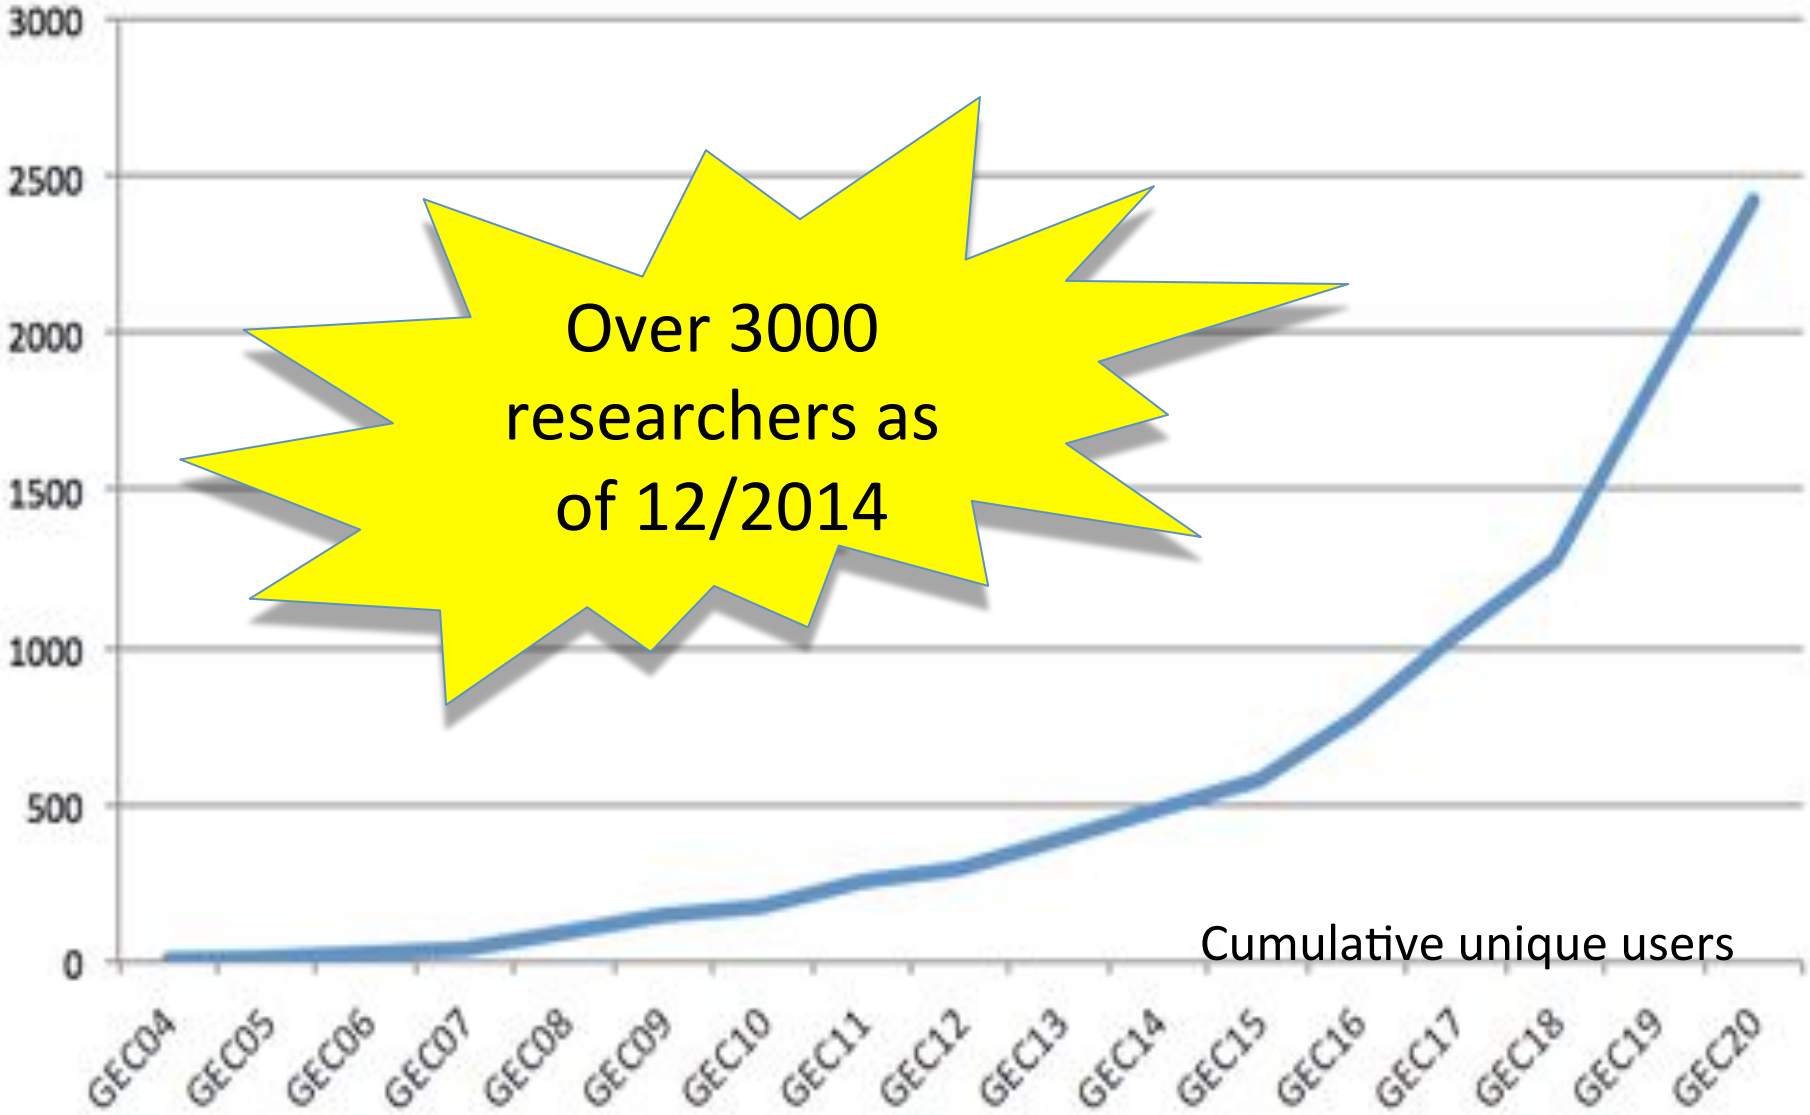

Macro-scale: the Rise of Global Interoperability

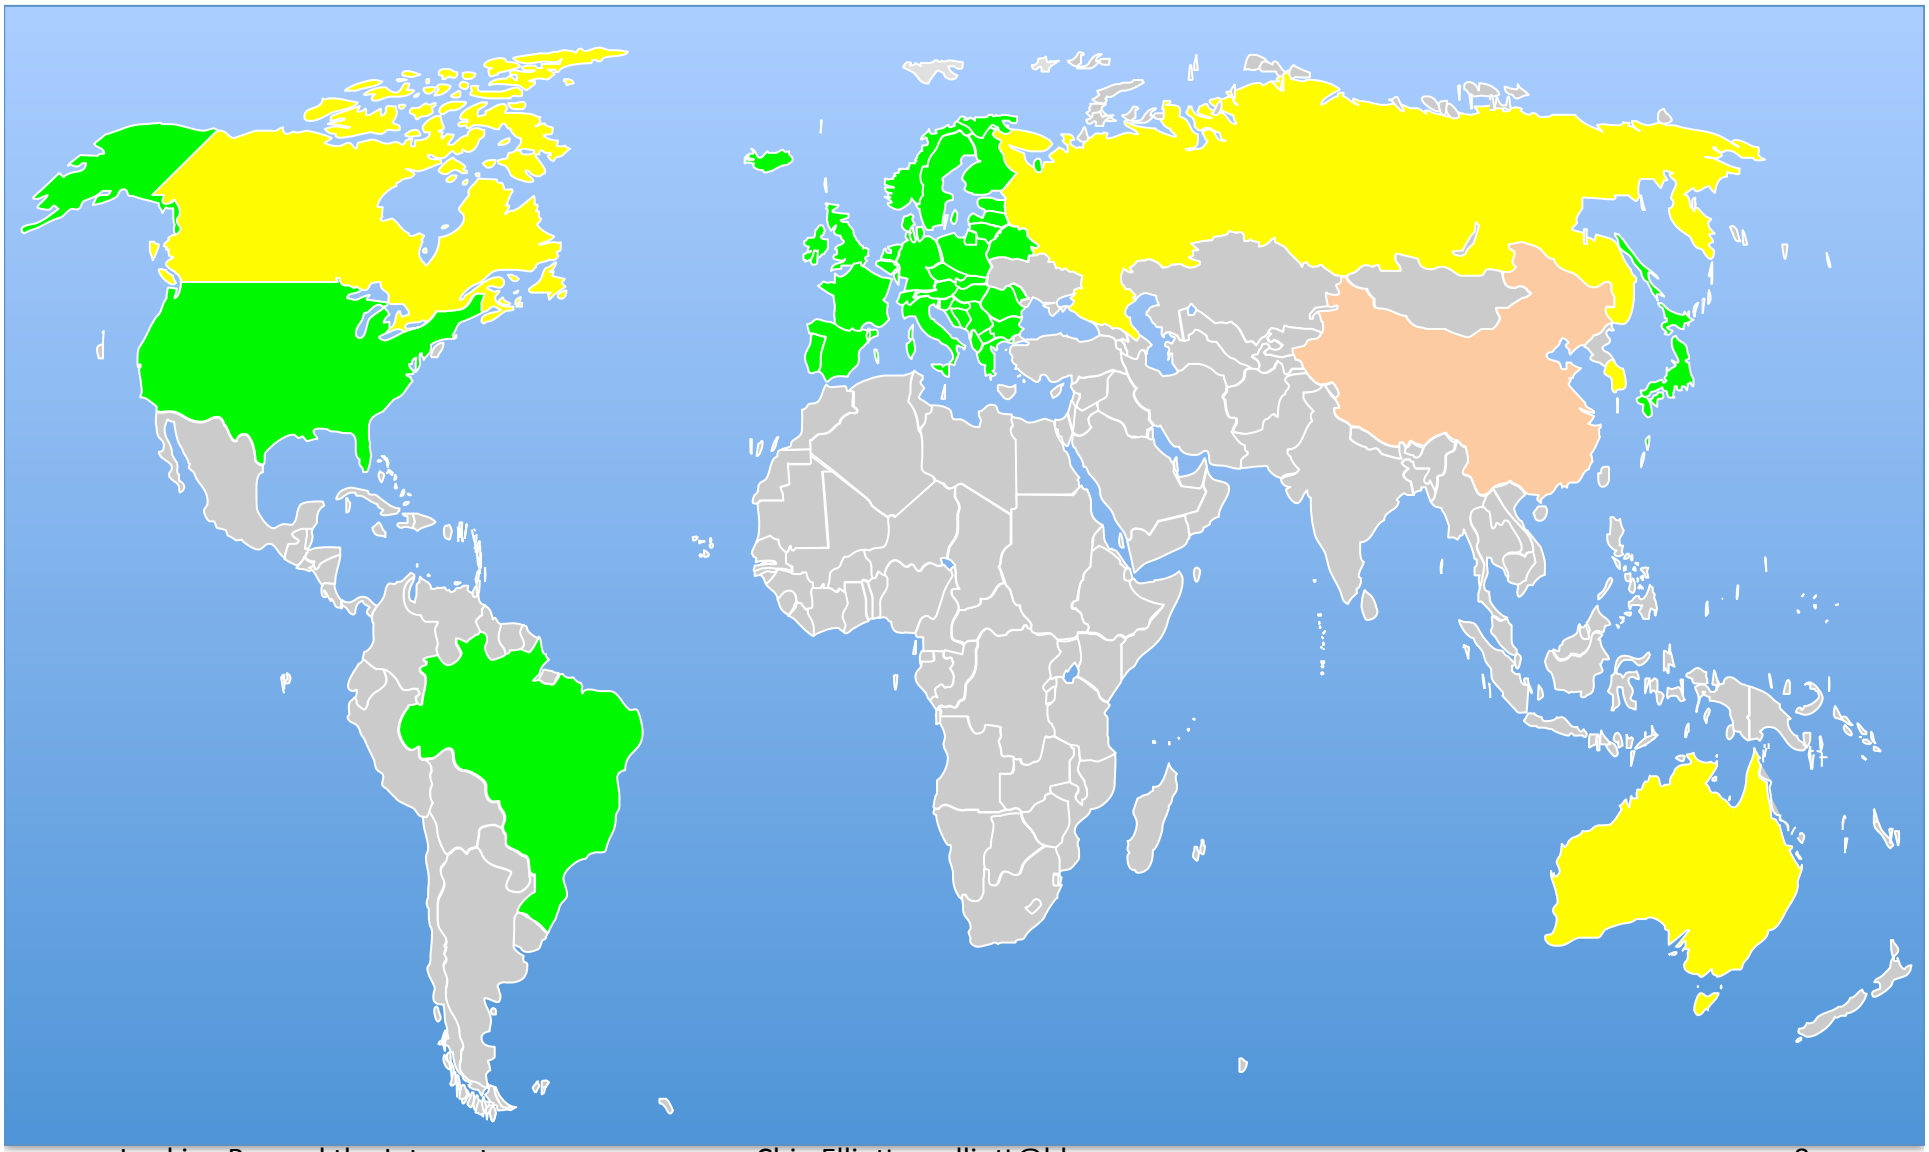

And now just starting up - **NSFCloud**

CloudLab

Chameleon

Rob Ricci

NSF

Kate Keahey

CloudLab


THE UNIVERSITY OF UTAH®
WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON
CLEMSON
UNIVERSITY
Raytheon
BBN Technologies
UMASS
AMHERST
usignite

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

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

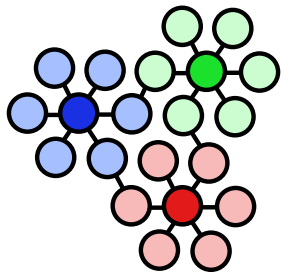

emulab

What Is CloudLab?

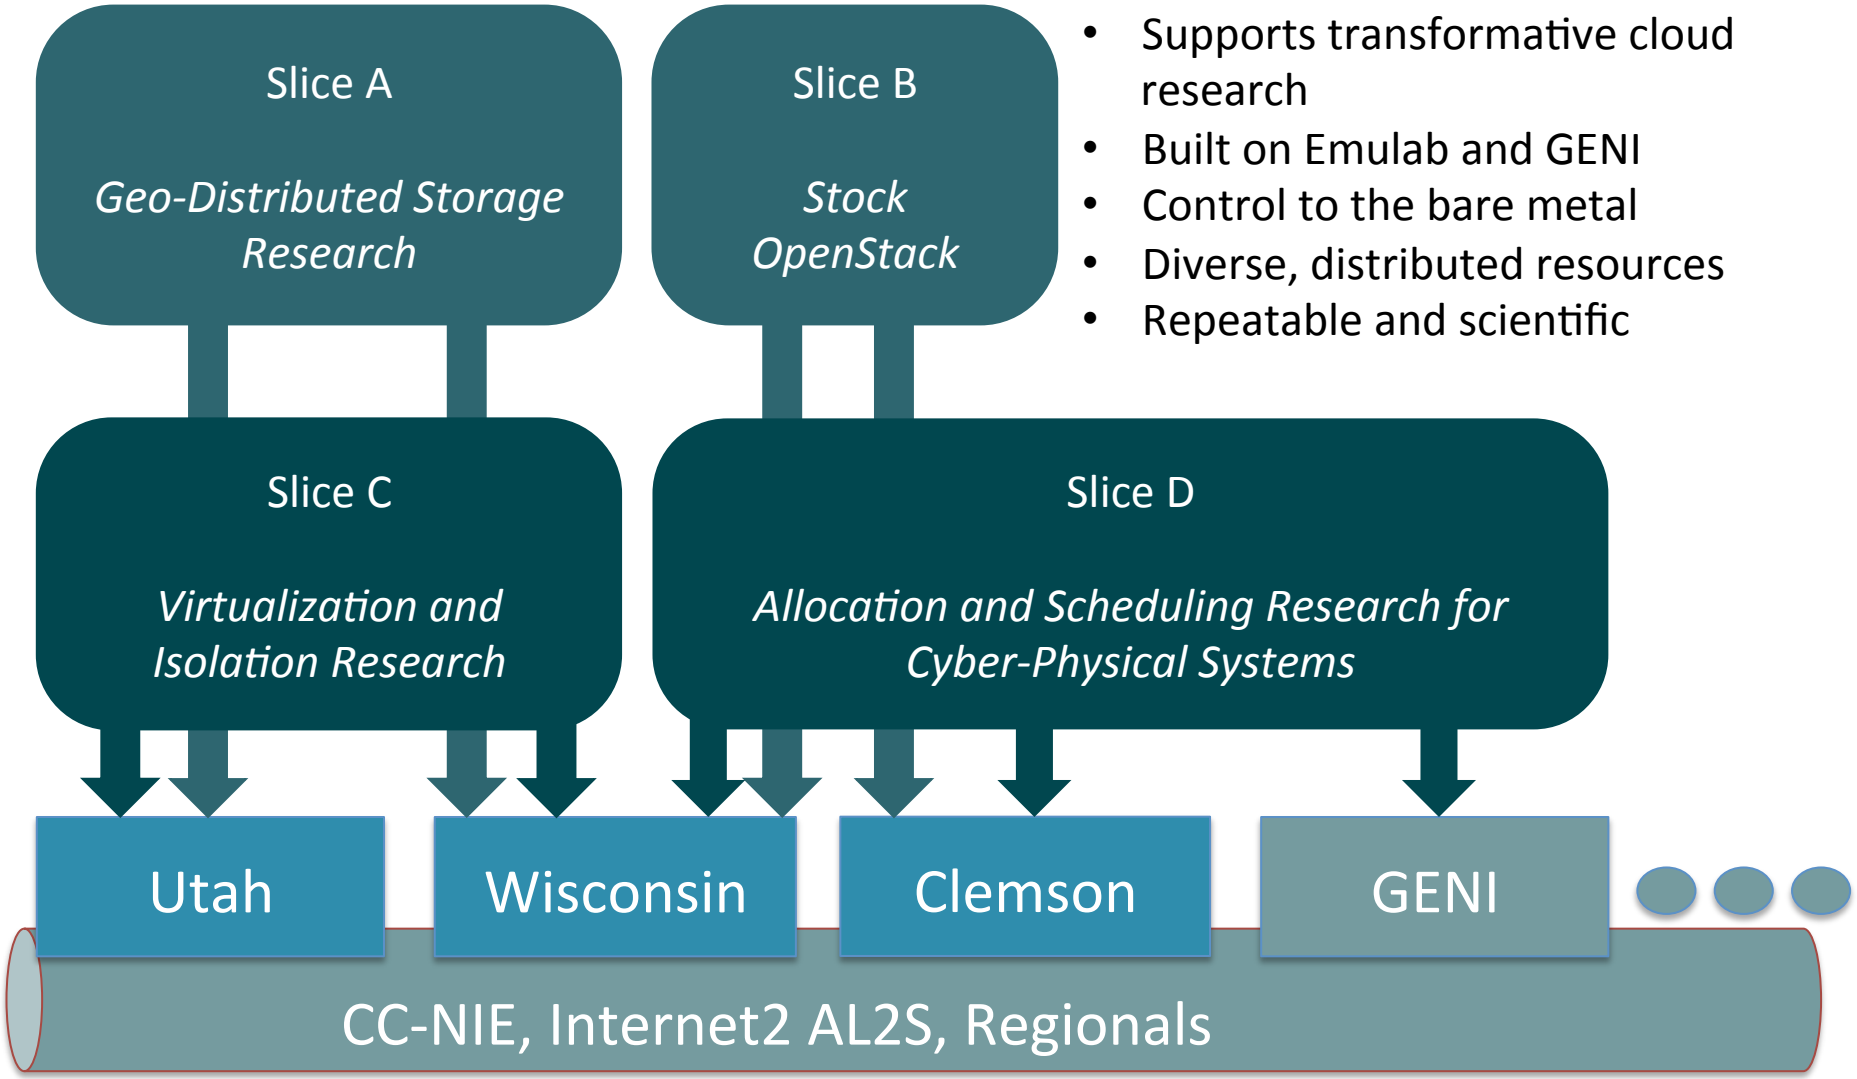

- Supports transformative cloud research
- Built on Emulab and GENI
- Control to the bare metal
- Diverse, distributed resources
- Repeatable and scientific

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 ARMs
- Disk on x86
- Very dense
- *HP*

Utah/HP: Low-power ARM64

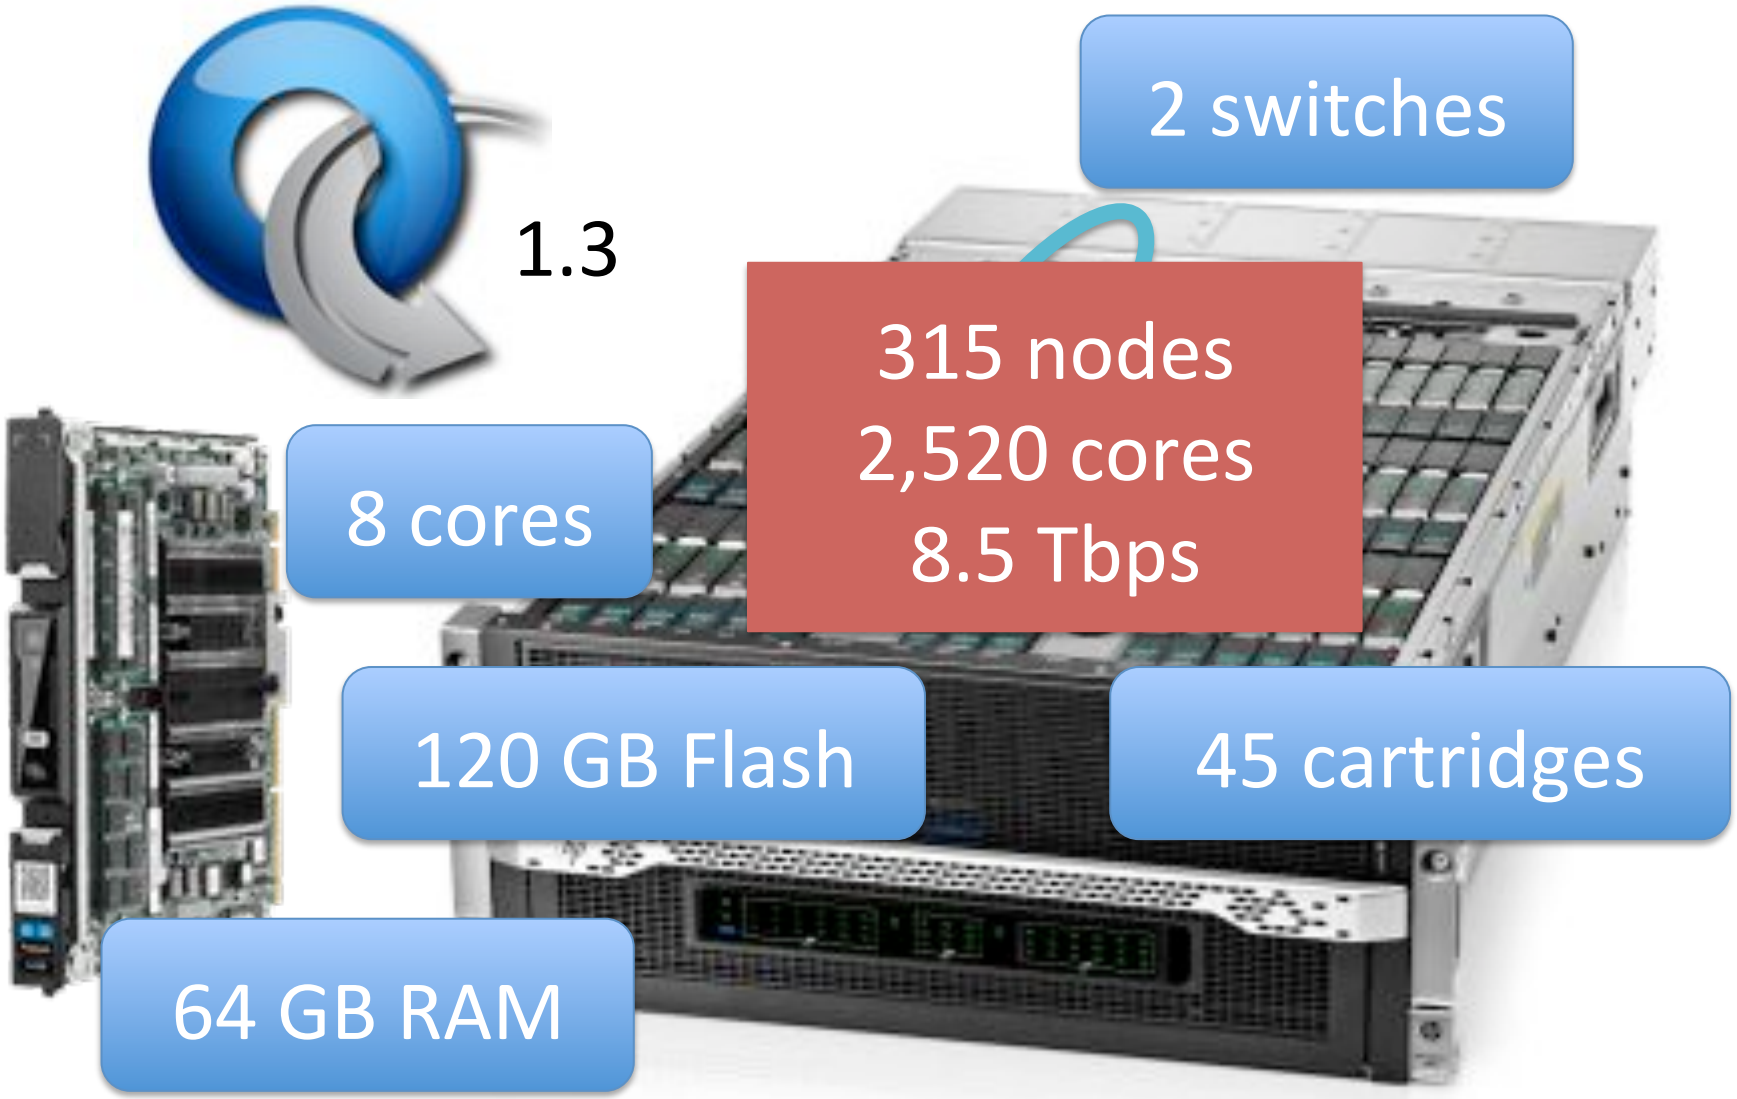

Wisconsin/Cisco

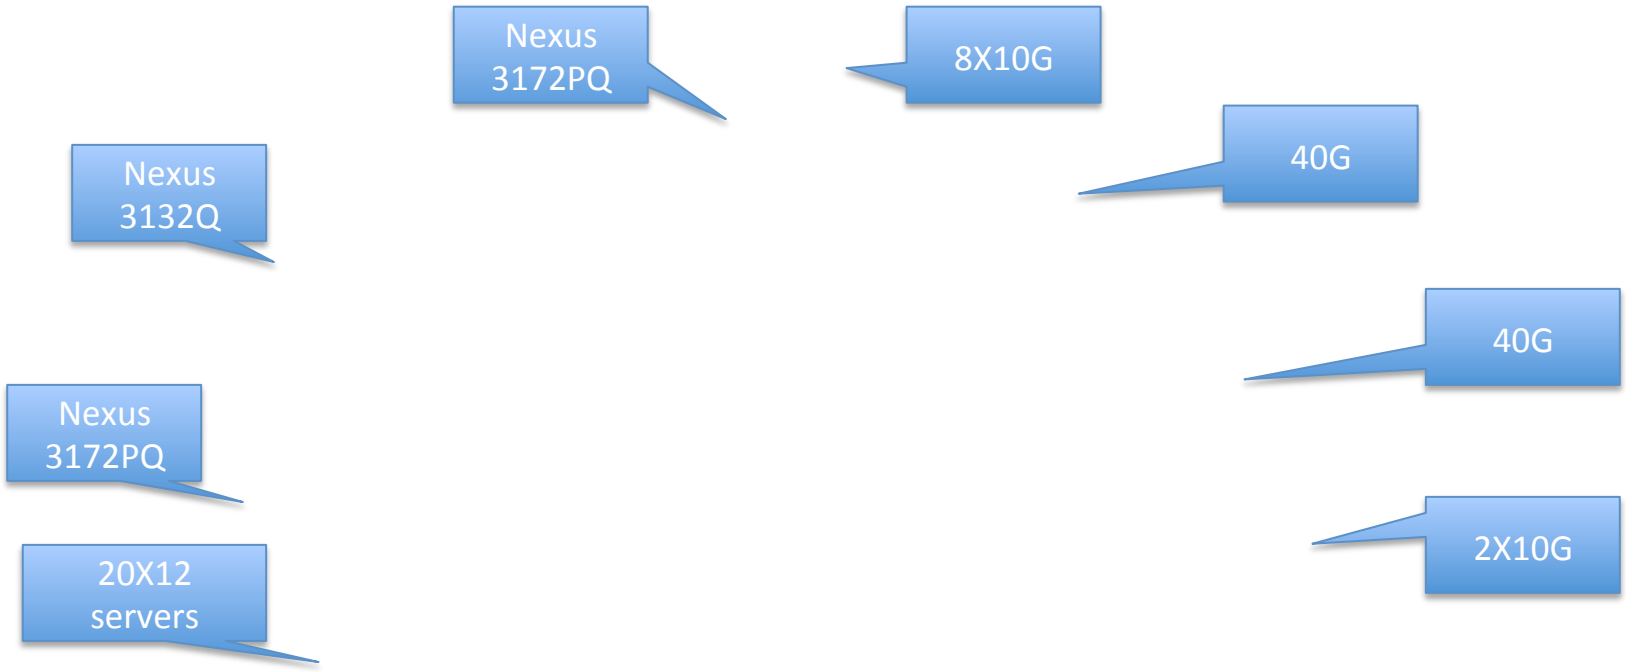

Clemson/Dell: High Memory, IB

20 cores/node

1 x 40 Gb IB/node

8 nodes/chassis

2*x 10 GbE OF/node

10 chasses/rack

2*x 1 GbE OF/node

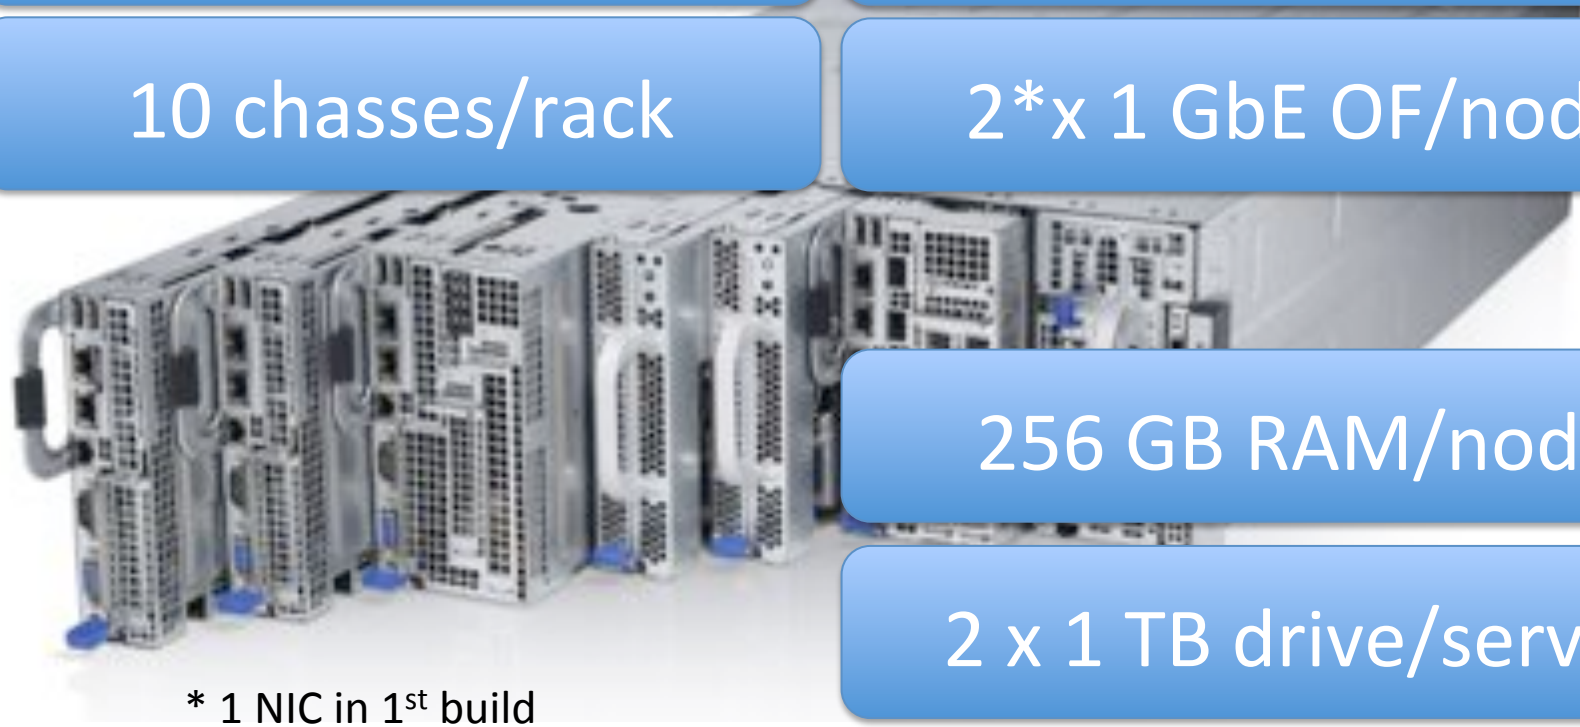

256 GB RAM/node

2 x 1 TB drive/server

* 1 NIC in 1st build

Use your GENI account
or sign up now at CloudLab.us

My thesis

- SDN is just an opening act
- A major transformation of the Internet has begun
- We can now catch glimpses of what lies beyond
- We can get there by a series of step by step actions

Major trends are converging

Multi-tenant Datacenters

Software Defined Infrastructure

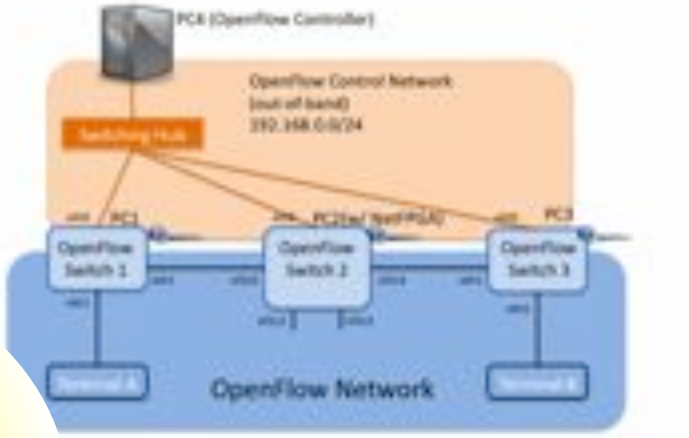

Software Defined Networks

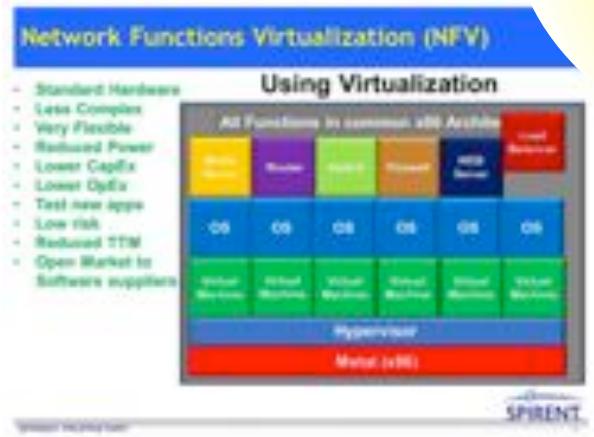

Network Functions Virtualization (NFV)

Distributed Datacenters

Driving the transformation - A radical change in "router" economics

Economics now favor pervasive computation and storage

ARPANET Imp (1969)

1 core, clock ~ 1.1 MHz
64 Kbytes RAM
No disk

Today's cost: ~ $650,000

Commodity GENI rack

Each 1U=
32 cores, 2.1 GHz
16 Gbyte, 4 Tbyte

Today's cost: $200,000
for full rack (50 x 1U)

Disk + controller (IBM 1302)
Today's cost: ~ $2,545,000

Disks were too expensive in 1969

1/3 the IMP's price, but with 1500 cores and 200 Tbytes of local storage

Instantiating services into slices

- Soon each switching point will be able to sustain 10,000 – 100,000 slices
- Can run arbitrary software in each slice
- Decoupling of "service" from infrastructure

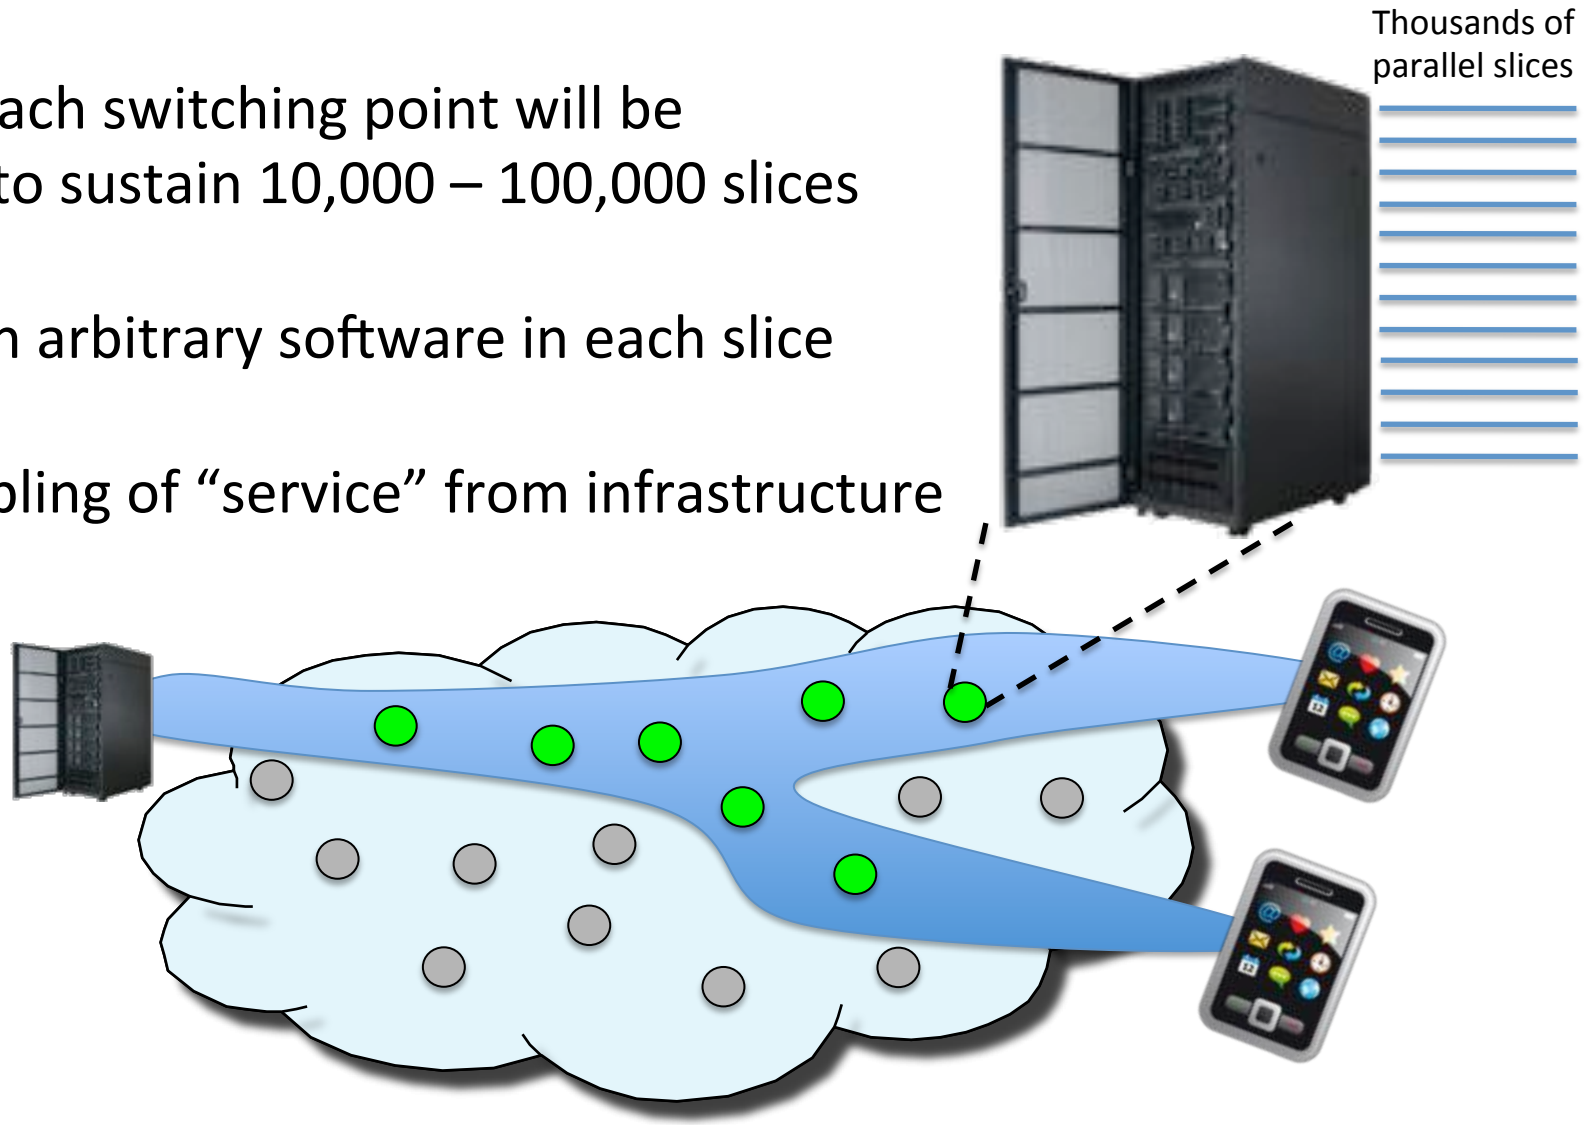

Software Defined Infrastructure

We're all heading to the same place

Rapidly create entire "sliced" cyberinfrastructure / networks on demand

Fast spin new protocols, switching strategies, virtual machines

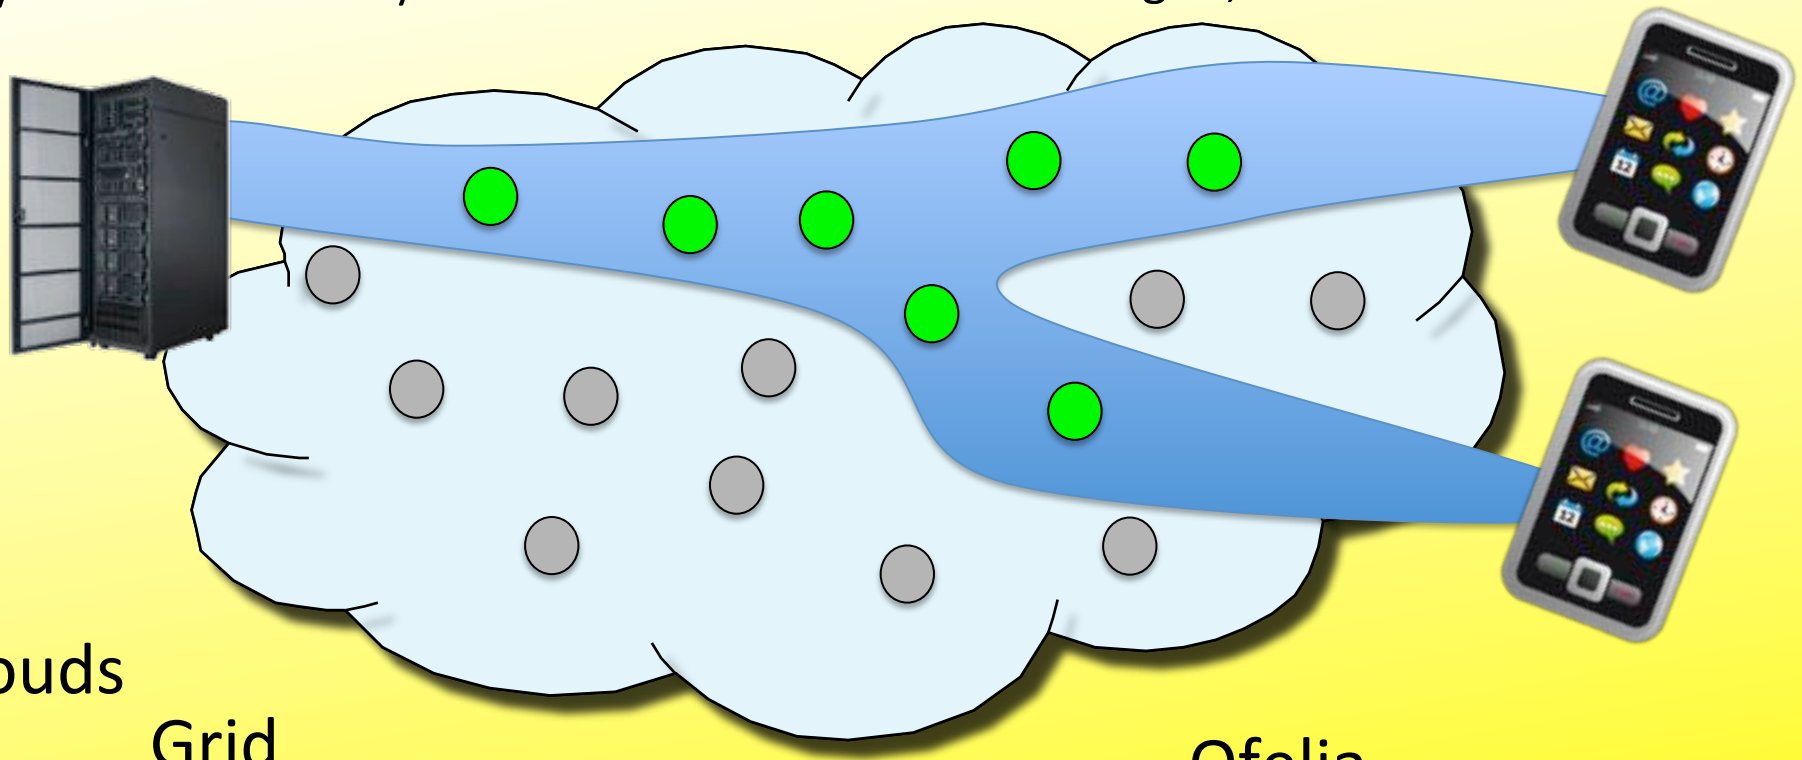

Clouds

Grid

Network functions virtualization

Software defined networks

Distributed Clouds

Vnode

Wivi

Ofelia

FLARE

US Ignite

GENI

The Rise of the "Service Store"

- "Drag and drop" Services
- Like an App Store . . .
- . . . that instantiates
  end-to-end Services

Decoupling Service from Provider

Service Store

Tailored Service

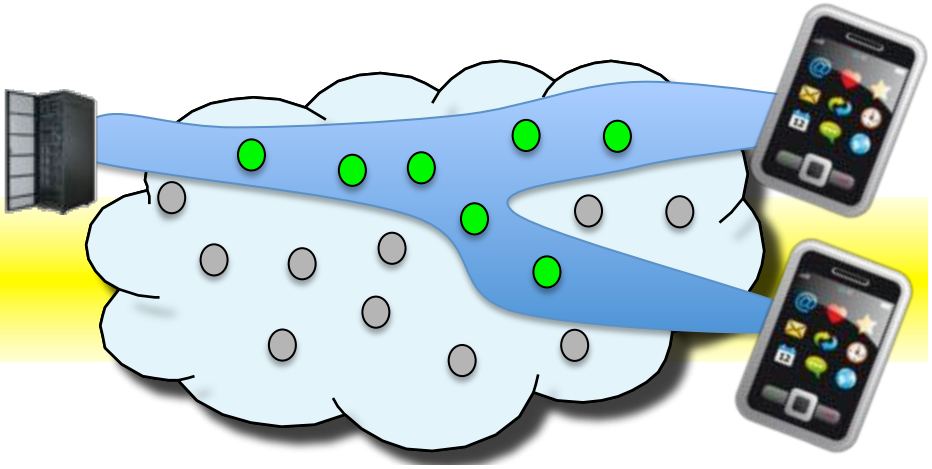

Service instantiated in a slice

Software Defined Infrastructure

Looking beyond the Internet

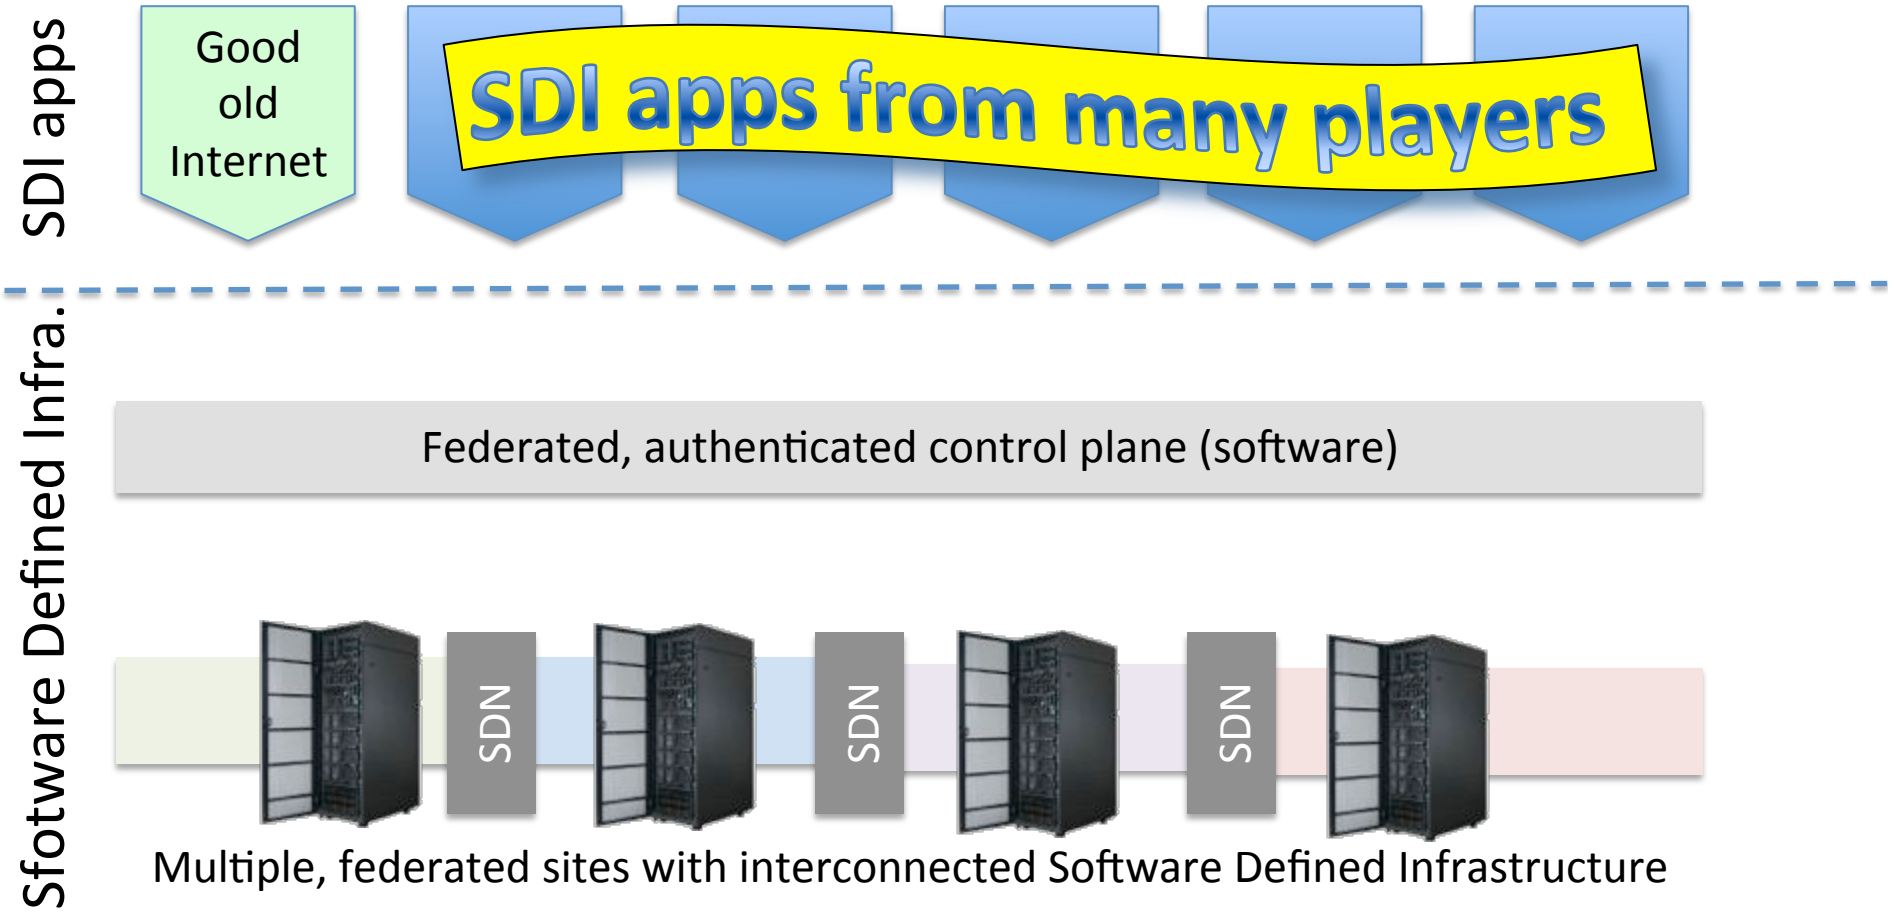

My thesis

- SDN is just an opening act
- A major transformation of the Internet has begun
- We can now catch glimpses of what lies beyond
- We can get there by a series of step by step actions

Where we are today – single-owner "islands"

Primary benefits right now: reduce cost, improve manageability

- Datacenters / clouds
- Experiments with enterprise / campus nets
- Wireless (cellular etc)

Source: NTT Communications

Early multi-domain SDN systems include Google and NTT – each datacenter is an SDN island, with a separate SDN WAN interconnecting them

(all same owner, though; no peering)

Necessary steps from today to SDI

- Step 1. Interconnect SDN islands
- Step 2. Make the network fully sliceable
- Step 3. Build out pervasive compute & storage

Step 1. Interconnect islands

SDN Workshop – Dec. 2013, Washington, DC

Inder Monga
ESnet Chief Technologist

About 100 networking experts from academia, industry, national labs and federal agencies

SDN Workshop – Key themes (1)

- Software-Defined Networking (SDN) is understood as the entire distributed infrastructure needed for next-generation commercial and/or scientific applications – i.e., **closely integrated compute, storage, and networks**.

- SDN technology has the potential and momentum to provide **game changing innovation to the entire Internet eco-system.**

- Using SDN technology, we can now envision (and in practical terms, create) **scientific "instruments on demand" or app-specific "infrastructure on demand" across multiple networks (multi-domain),** on a worldwide scale.

SDN Workshop – Key themes (2)

- The time is right for deploying prototype operational, multi-domain SDNs.
- The focus of initial SDN deployments should include **Software-Defined Exchanges (SDXs)** to enable interoperability, co-designed in close collaboration with US industry.
- These efforts should actively **engage key scientific instruments and next-generation applications as design and prototyping partners.**
- A vigorous and sustained research program should investigate the **security implications** of multi-domain/multi-layer SDNs.

Conclusions

- SDN is just an opening act
- A major transformation of the Internet has begun
- We can now catch glimpses of what lies beyond
- We can get there by a series of step by step actions