Impactful Routing Research with PEERING

Combining intradomain emulation with *real* BGP connectivity

Brandon Schlinker, Ethan Katz-Bassett, Italo Cunha
University of Southern California, California, USA
Universidade Federal de Minas Gerais, Minas Gerais, Brazil

AIMS
February 2015

PEERING: The BGP Testbed

With **PEERING**, experiments can exchange

BGP routes and **traffic** at locations around the world

PEERING: The BGP Testbed

With **PEERING**, experiments can exchange

BGP routes and **traffic** at locations around the world

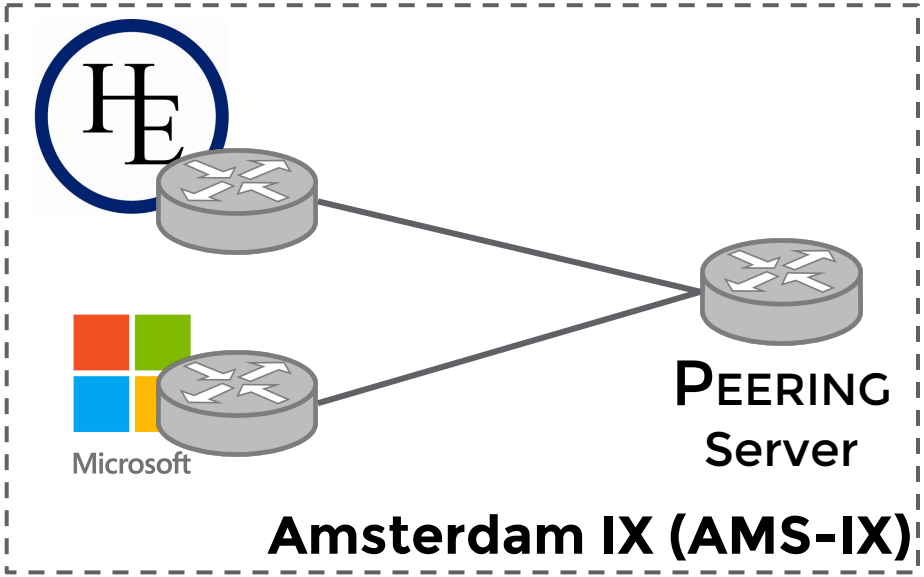

PEERING: The BGP Testbed

With **PEERING**, experiments can exchange

BGP routes and **traffic** at locations around the world

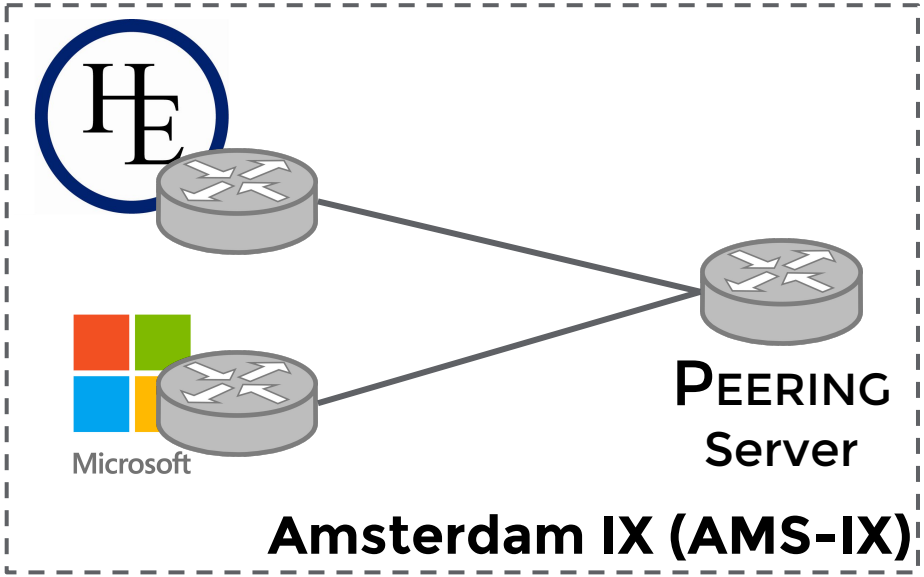

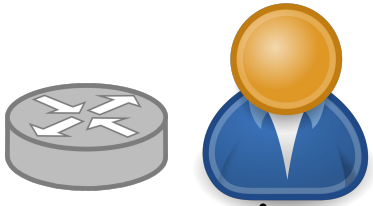

Experimentor
Allocated:
184.164.224.0/23

PEERING: The BGP Testbed

With **PEERING**, experiments can exchange

BGP routes and **traffic** at locations around the world

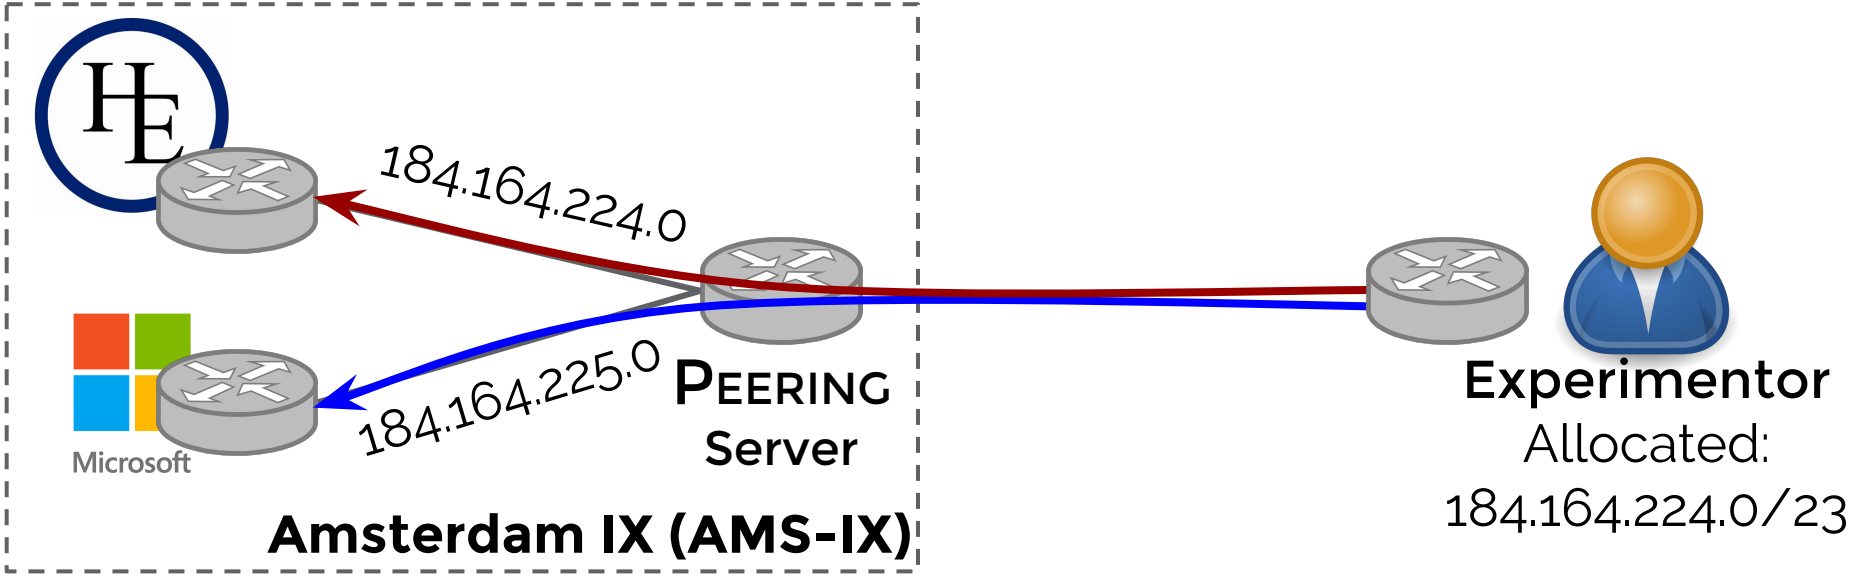

PEERING provides unprecedented control

Route monitors / traceroutes only measure *existing* routes
Simulations and emulations lack realism

PEERING provides unprecedented control

Route monitors / traceroutes only measure *existing* routes
Simulations and emulations lack realism

With **PEERING**, experiments can **make changes...**

- route poisoning to check how other networks react
- announce / withdraw routes at different PoPs / for different peers
- select their outgoing routes

Measuring ROA Filter Adoption with PEERING

Route Origin Authorizations (ROA)

- specifies which networks are allowed to announce a prefix

Existing studies have focused on the **adoption** of ROAs

- do prefixes have ROAs and do they match the observed routes?
- but **ROAs are only effective if they are *used* in routing decisions**

Measuring ROA Filter Adoption with PEERING

Route Origin Authorizations (ROA)

- specifies which networks are allowed to announce a prefix

Existing studies have focused on the **adoption** of ROAs

- do prefixes have ROAs and do they match the observed routes?
- but **ROAs are only effective if they are *used* in routing decisions**

With **PEERING**, we're measuring how / if **ROAs** are used in decisions

- coordinate BGP announcements, ROA manipulations
- observe how ASes *react* (traceroutes, BGP collectors)

Announce Anchor Prefix (184.164.224.0/24)

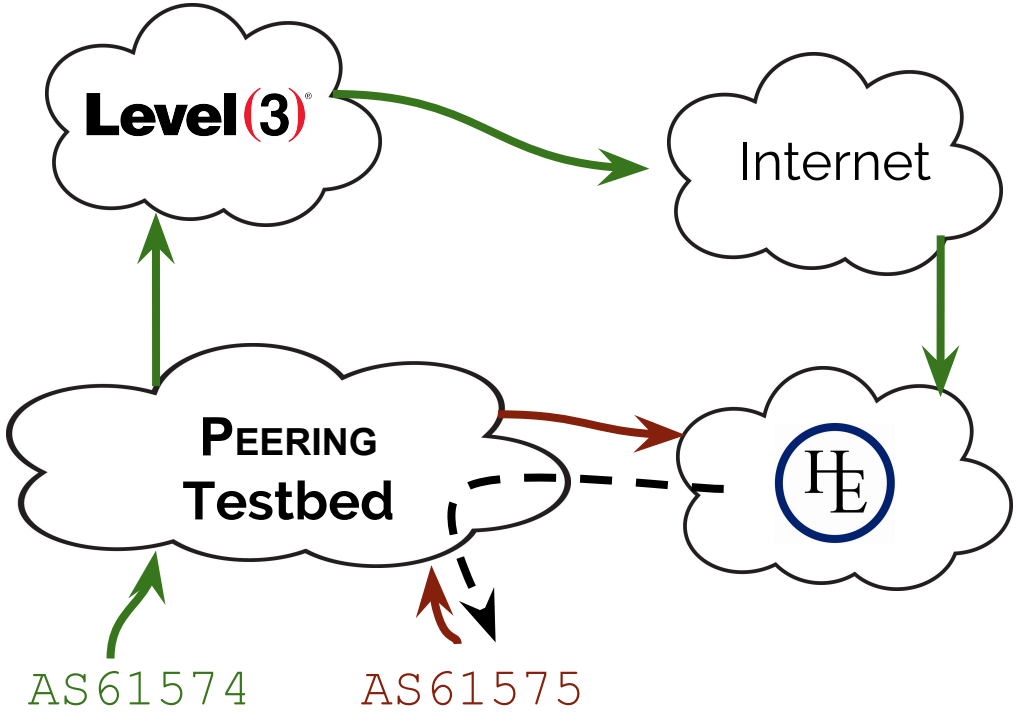

Announce *anchor* prefix
prefix originated from two ASNs

AS61574's announcement is propagated to a transit provider

AS61575's announcement is propagated directly to peer

Expected path for HE to prefix
HE -> PEERING -> AS61575
(shortest path)

Check Behavior for Test Prefix (184.164.225.0/24)

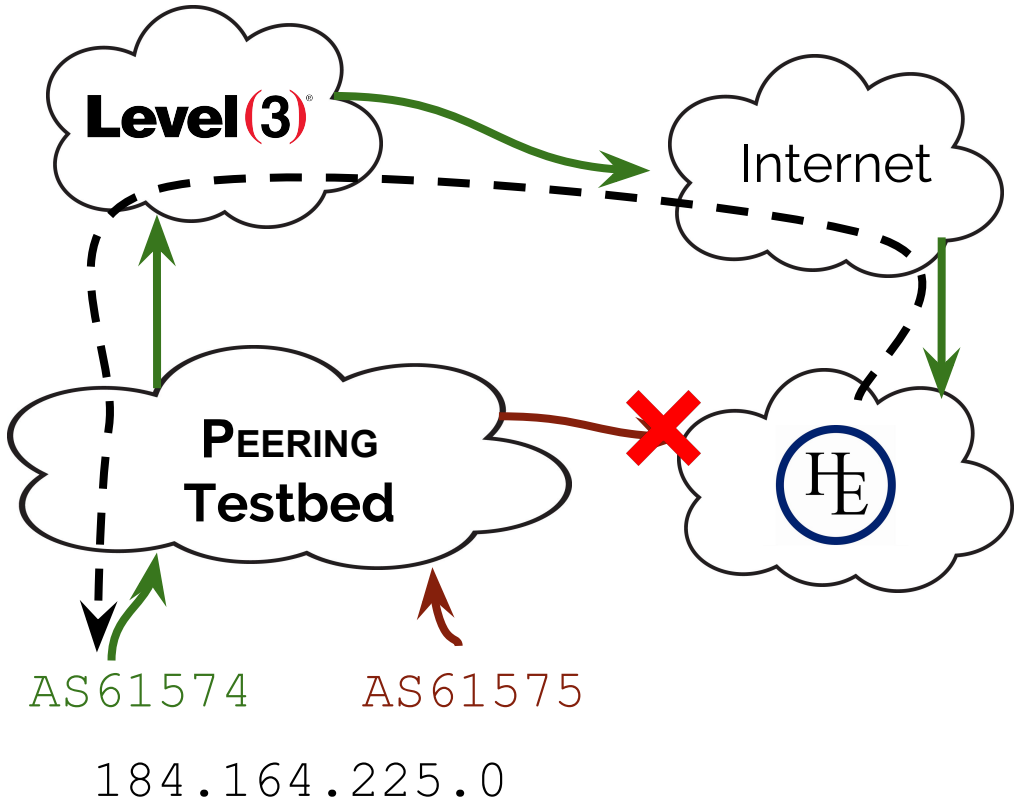

Add ROA for test prefix
AS61574 is valid origin

Announce test prefix
originated from same ASNs
propagated to same peers

If HE's filters account for ROAs
will reject route from 61575
prefer longer route from 61574

Measuring Path Performance with PEERING

Large content and cloud providers have many paths to destination

- result of lots of peering at IXPs and backbones between PoPs

What's the value of this rich interdomain connectivity?

- can it help improve end-user experience? (bypass congestion?)
- what's the relative value of different IXP connections?

Measuring Path Performance with PEERING

Large content and cloud providers have many paths to destination

- result of lots of peering at IXPs and backbones between PoPs

What's the value of this rich interdomain connectivity?

- can it help improve end-user experience? (bypass congestion?)
- what's the relative value of different IXP connections?

With **PEERING**, we can:

- direct traffic across different paths and measure performance
- build and evaluate systems that leverage this rich connectivity

Research Supported by PEERING

- LIFEGUARD: route around failures	**SIGCOMM 12**
- PECAN: joint content & network routing	**SIGMETRICS 13**
- PoiRoot: locate root cause of path changes	**SIGCOMM 13**
- ARROW: deployable fix to routing problems	SIGCOMM 14
- SDX: software-defined Internet exchange	SIGCOMM 14
- Measuring Internet routing policies	**IMC 15**
- Sprite: SDN-based inbound traffic engineering	SOSR 15
- RAPTOR: Routing attacks on TOR	USENIX Security 15

(**bold** = **PEERING** required)

Updates on PEERING Testbed

More Points of Presence, More IXP Connectivity

Now control 8 ASNs, multiple IPv4 and v6 prefixes

- Officially transferred our primary ASN (47065) from GENI

PEERING projected to have 17 points of presence by mid-year

- adding 10 or 40G connectivity at CloudLab sites (3 sites)

Highest priority = Internet Exchange Points (7 sites)

- Seattle Internet Exchange (connected)
- Amsterdam Internet Exchange (connected)
- Phoenix Internet Exchange (connected)
- Equinix facilities in Dallas and Asburn (equipment at facilities)
- Brazil Internet Exchange in São Paulo (shipping soon)
- One Wilshire in Los Angeles (planning)

Measurements as a Service

Beacon service continuously issues announcements
- cycles any unallocated prefixes through announcement loop

Experiments can control beacon service
- request announcement through web interface UI
- lower overhead than setting up infrastructure locally (VPNs, BIRD)
- investigating programmatic interface (hackathon feedback)

Regular measurements and announcement
- regular traceroutes from RIPE Atlas towards all /24s (every 20 min)
- announcements archived in BGPMon

Colocating Experiments at PoPs

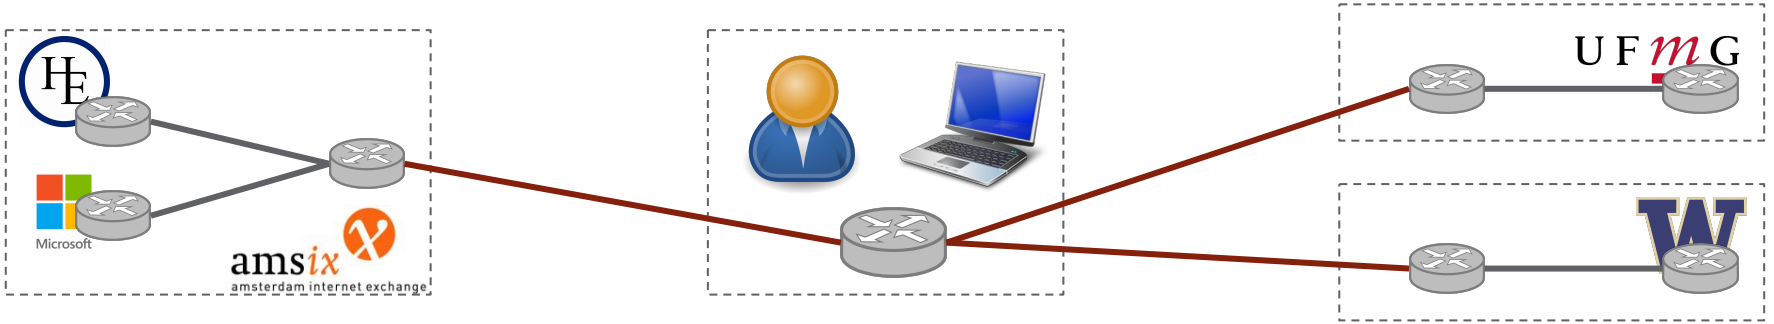

Today, clients establish VPN connections to **PEERING** PoPs

- control and data-plane traffic is relayed to their system
- sufficient for almost any control-plane experiment
- difficult to run services, conduct performance measurements

Colocating Experiments at PoPs

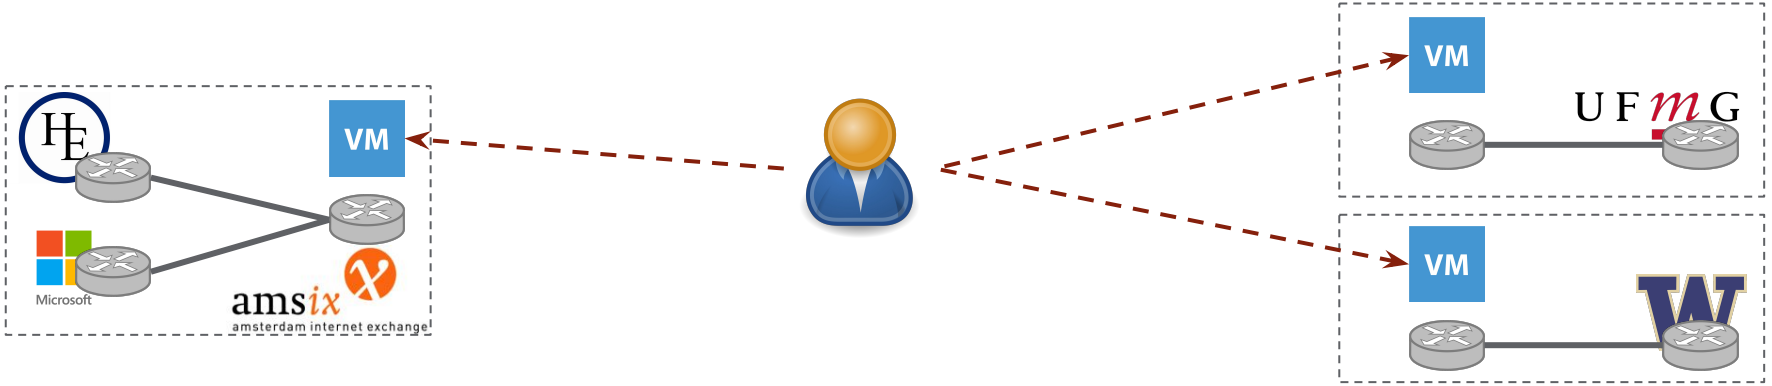

Remove backhauling by **installing user VMs at PoPs**

- supports non-resource intensive experiments
- enables hosting of Anycast content / services
- enables performance / routing experiments

Virtualized Layer-2 Backbone Connecting Sites

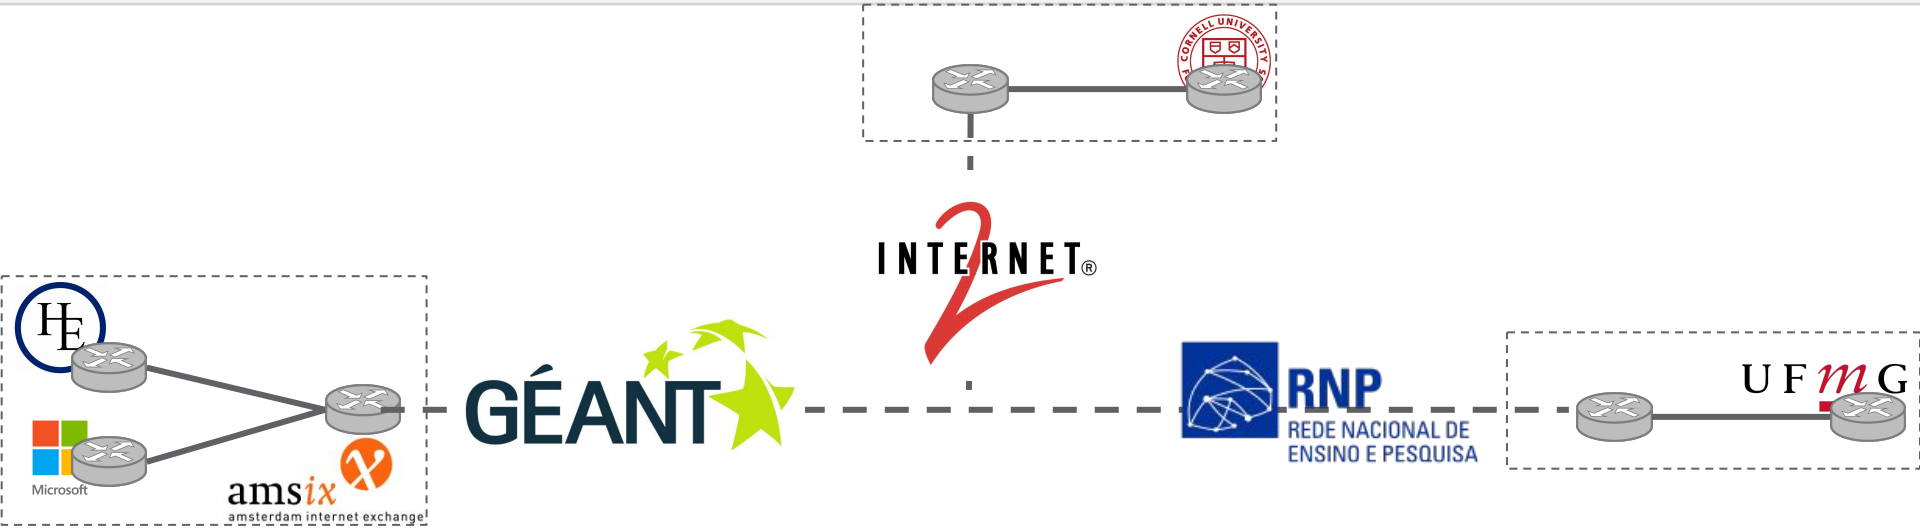

Connecting all PoPs via virtualized layer-2 interconnection

- use R&E network infrastructure
- provides performance guarantees, control over routing

Experimenting with Large Cloud Networks

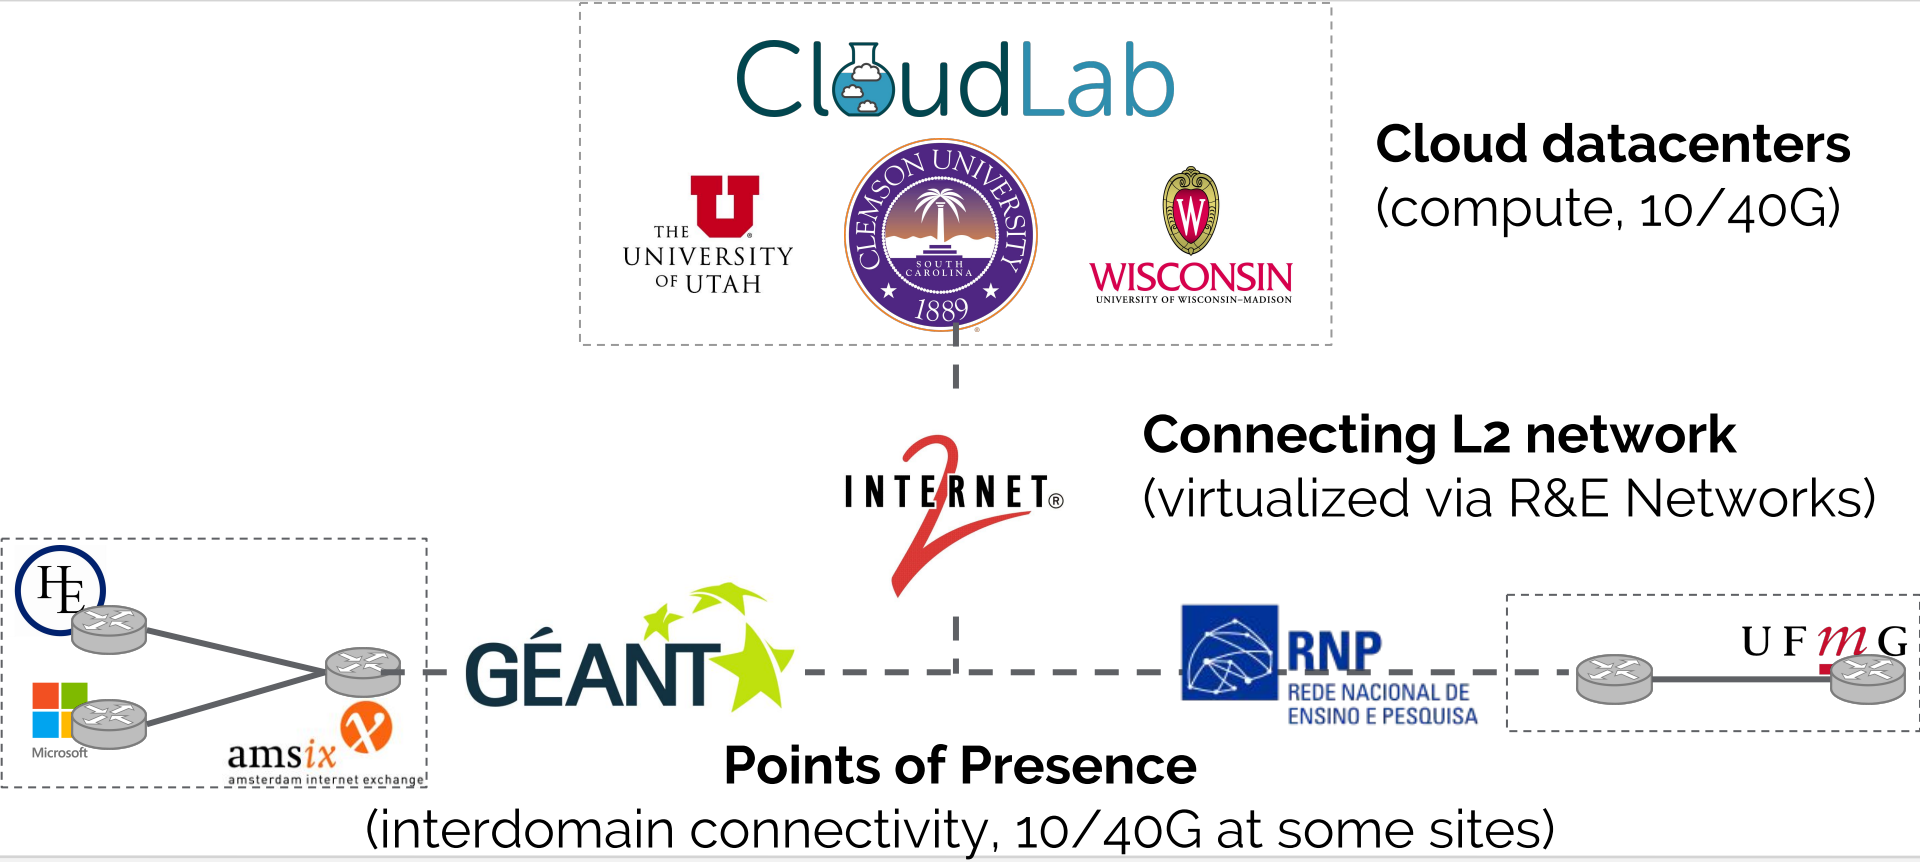

Experimenting with Large Cloud Networks

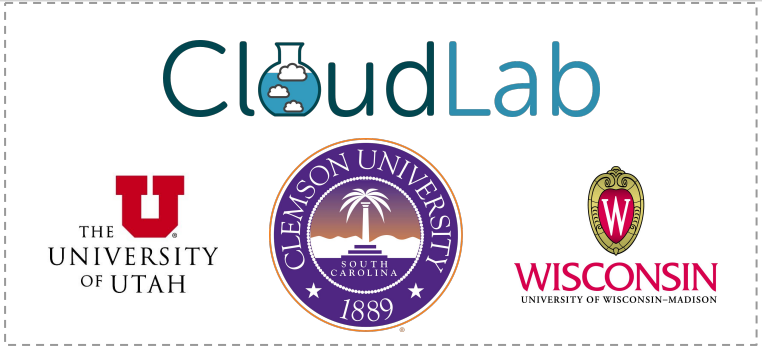

Datacenters

Backhaul

Interdomain PoPs

Qualitative representation of cloud / content provider's network

- control of intra and interdomain routing, including R&E routes
- services can be hosted for performance / route measurements
- new routing schemes can be evaluated

Improving Experiment Setup Process

New website interface for users

- experiment setup was previously manual, error-prone process
- automated majority of the steps, including allocations
- adding more visibility to website, including looking glass

Rewrote setup scripts to make them easier to use

- decide which peers at an IX receive an announcement
- decide how egress traffic is routed among available paths
- changed from Quagga to BIRD to support added functionality
- successfully supported multiple clients during hackathon

Summary

PEERING is built for the community's research:

- we've tackled the challenge of setting up this infrastructure
- deployed routers and established peerings around the world
- manage filters, traffic restrictions, peering sessions, servers

Working to expand PEERING to meet the community's needs:

- colocated experiments, backbone connectivity, CloudLab
- supporting a number of new security experiments

Contact us:

- team@peering.usc.edu