USC Viterbi
School of Engineering
Department of Computer Science

Towards Impactful Routing Research:

Running Your Own (Emulated) AS on the (Real) Internet

By Brandon Schlinker*, Kyriakos Zarifis*, Italo Cunha♮, Nick Feamster †, Ethan Katz-Bassett*, and Minlan Yu*

University of Southern California, California, USA*
Universidade Federal de Minas Gerais, Minas Gerais, Brazil♮
Georgia Institute of Technology, Georgia, USA†

CoNEXT Student Workshop
Santa Barbara, California, USA
December 2013

Recent Innovation in Networking

Focused Primarily on Intradomain Routing / SDN

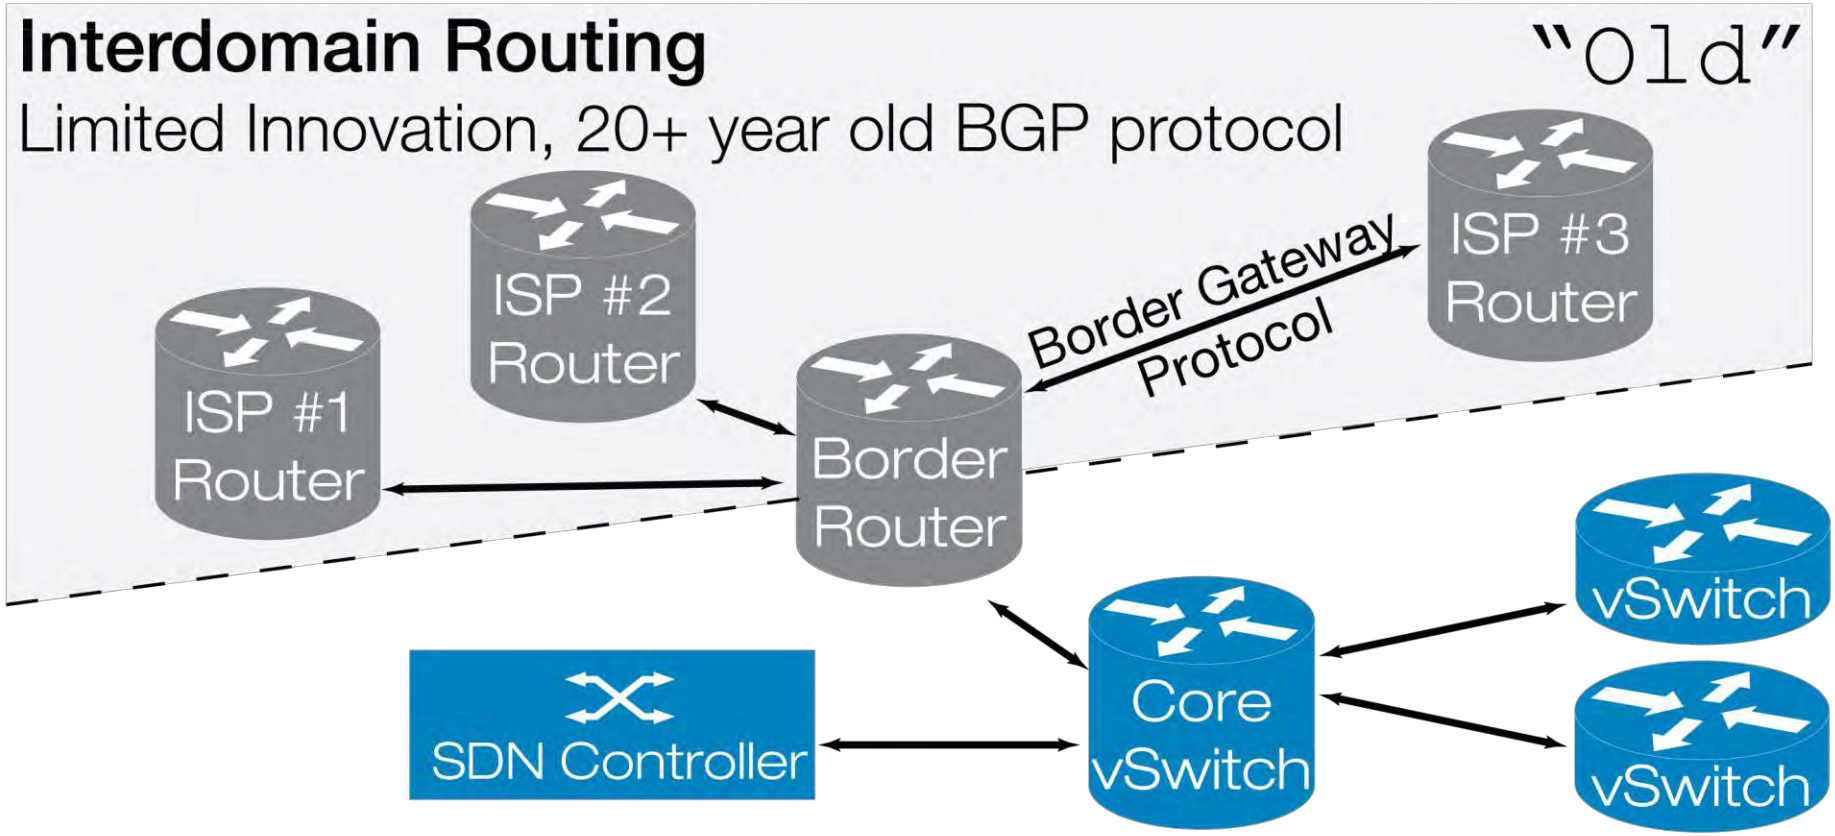

"New"

Recent Innovation in Networking

Islands of Innovation using BGP for communication

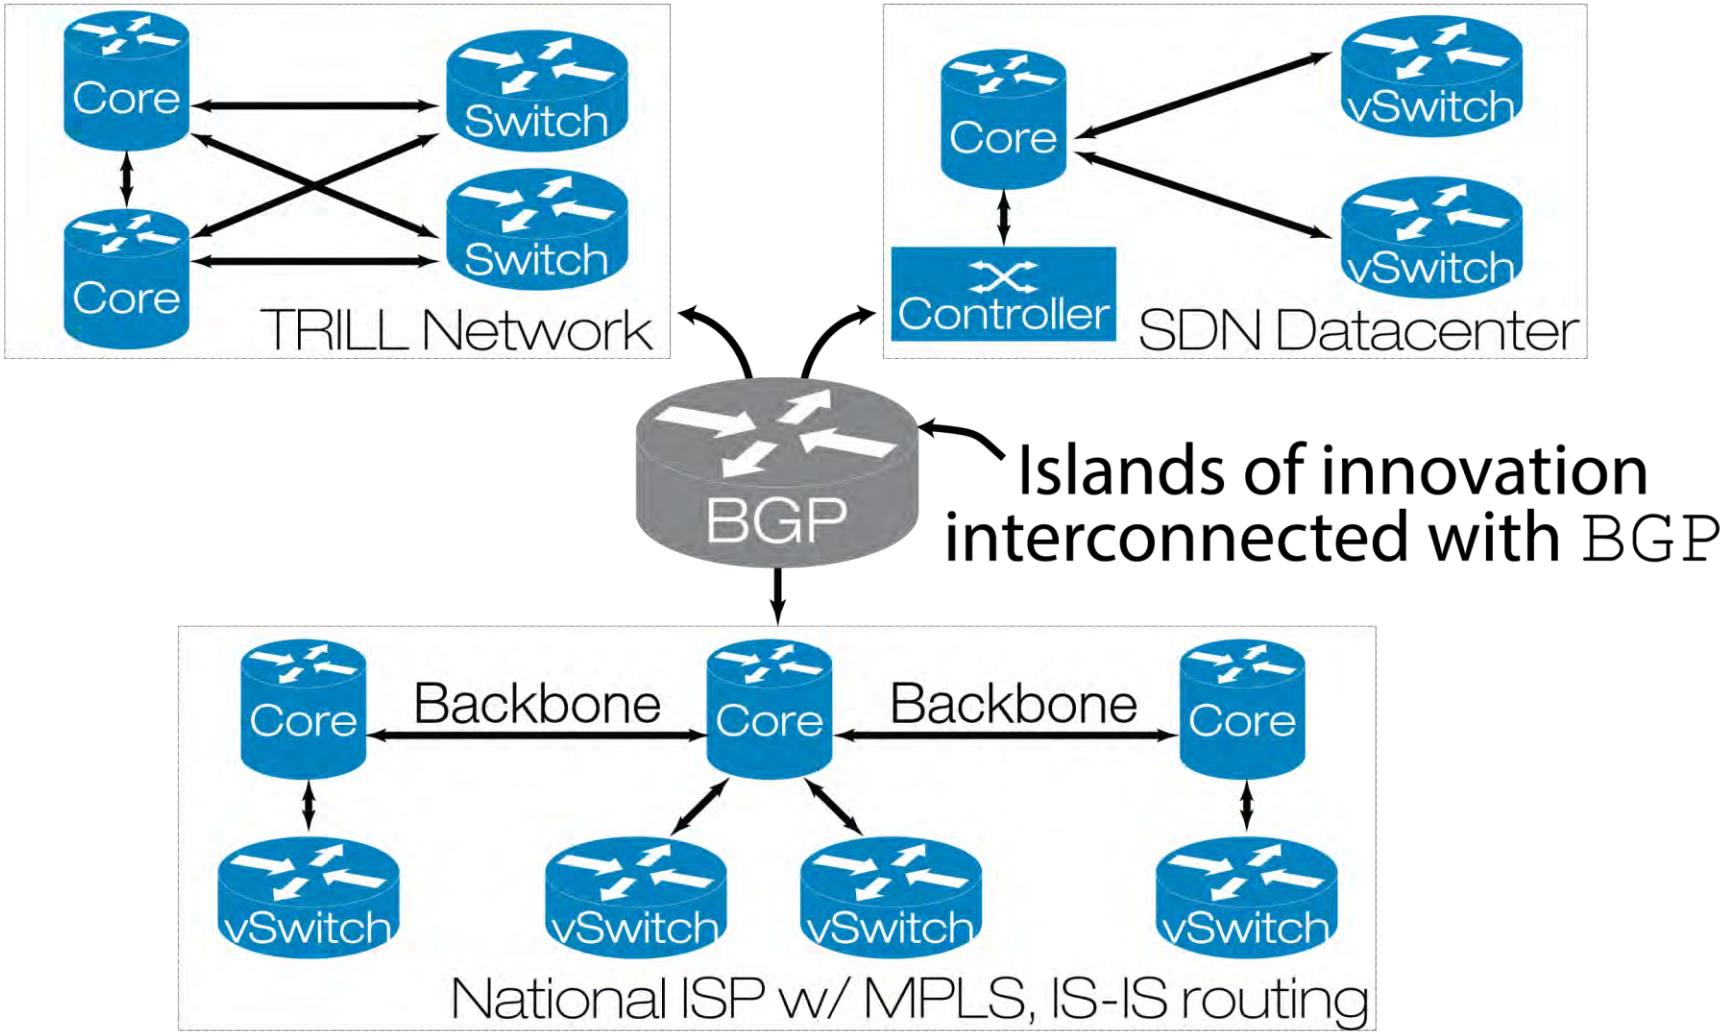

Border Gateway Protocol

Designed for a past era, BGP struggles to keep up

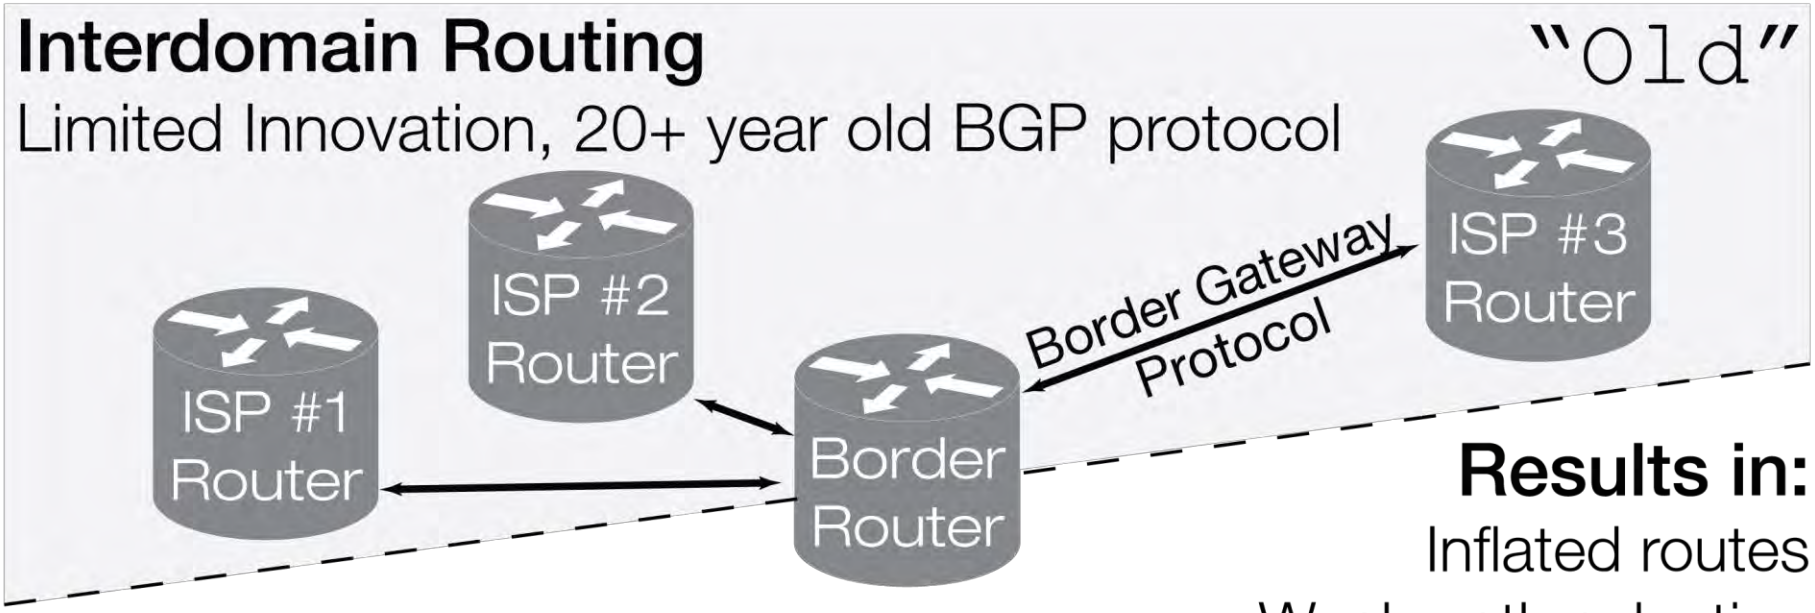

Results in:

Inflated routes

Weak path selection

Security vulnerabilities

Lengthy convergence time

Failures related to route redistribution

QoS problems due to path oscillations

Network isolation instead of collaboration

Recent Innovation in Networking

Islands of Innovation using BGP for communication

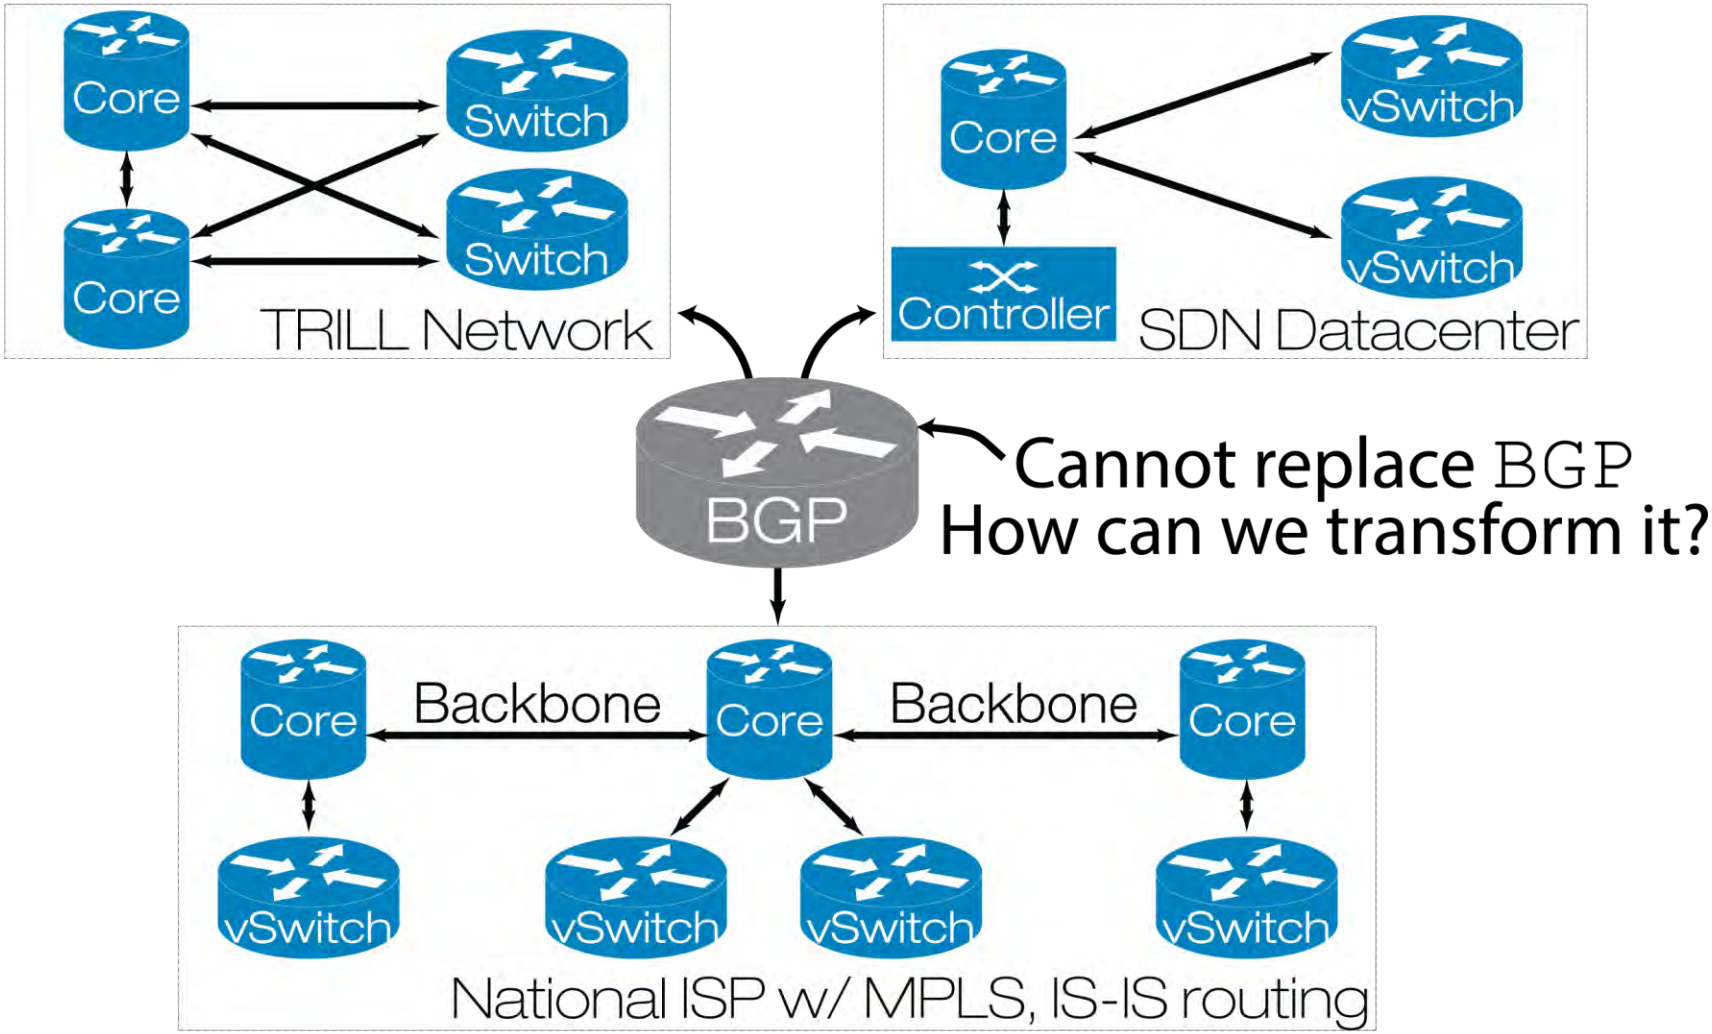

Where are the interesting problems?

At the intersection of Inter and Intradomain Routing

Interdomain Routing

Limited Innovation, 20+ year old BGP protocol

Interactions = Interesting problems

Solving these problems can improve internet routing

Intradomain Routing

SDN, Enhanced IGP & Fabric Protocols

Anycast Experiment

Involves Both Intradoman and Interdomain Technologies

Example Experiment:
Intercontinental Anycast Network

Anycast requires:

ⓐ BGP connectivity to advertise & exchange traffic
ⓑ Intradomain routing to direct internal traffic

Want to understand how traffic flows given conditions

Experiment involves *intradomain* and *interdomain* routing

Anycast Experiment Wishlist

Intradoman and Interdomain Technologies

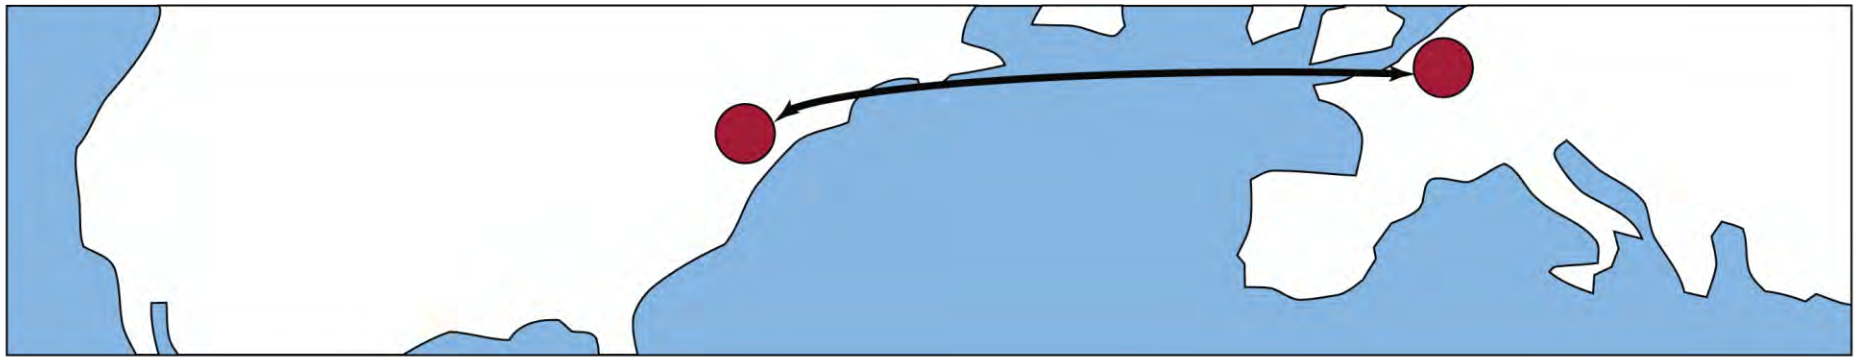

Wish List for Anycast Experiment:

① Intercontinental private wide-area network (WAN) between US and EU

Anycast Experiment Wishlist

Intradoman and Interdomain Technologies

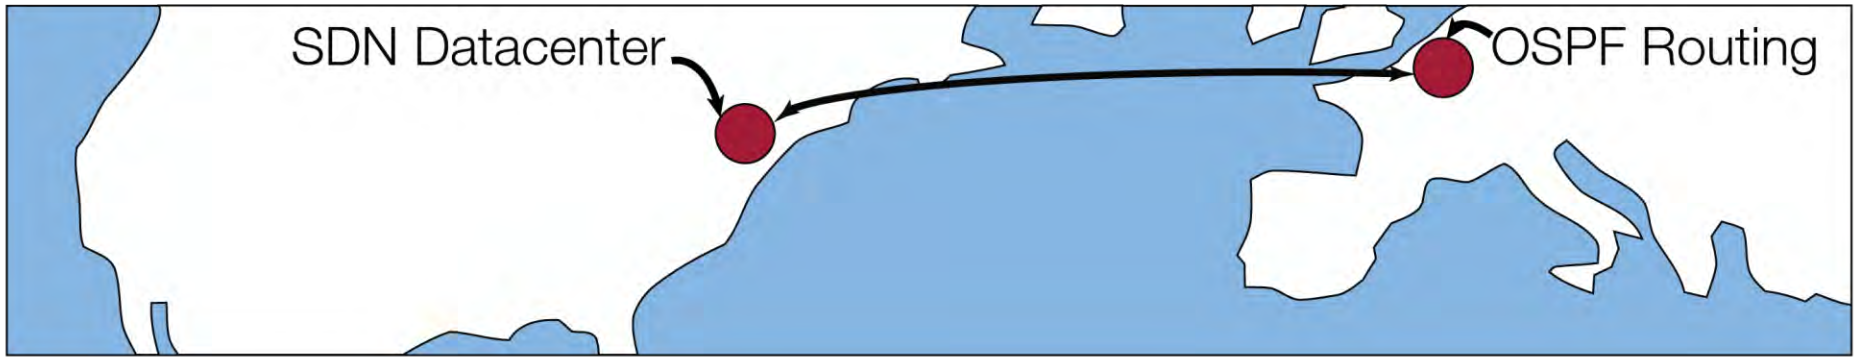

Wish List for Anycast Experiment:

① Intercontinental private wide-area network (WAN) between US and EU

② Different routing domains in each datacenter (SDN & OSPF)

Anycast Experiment Wishlist

Intradoman and Interdomain Technologies

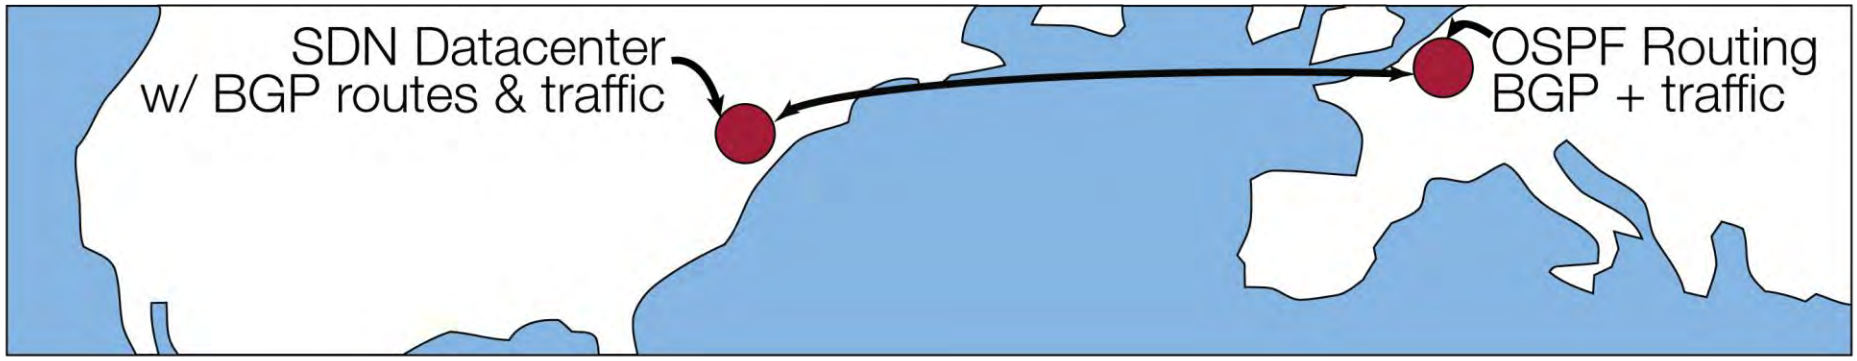

Wish List for Anycast Experiment:

1. Intercontinental private wide-area network (WAN) between US and EU
2. Different routing domains in each datacenter (SDN & OSPF)
3. Ability to exchange *routes and traffic* via BGP with other peers

Anycast Experiment Wishlist

Intradoman and Interdomain Technologies

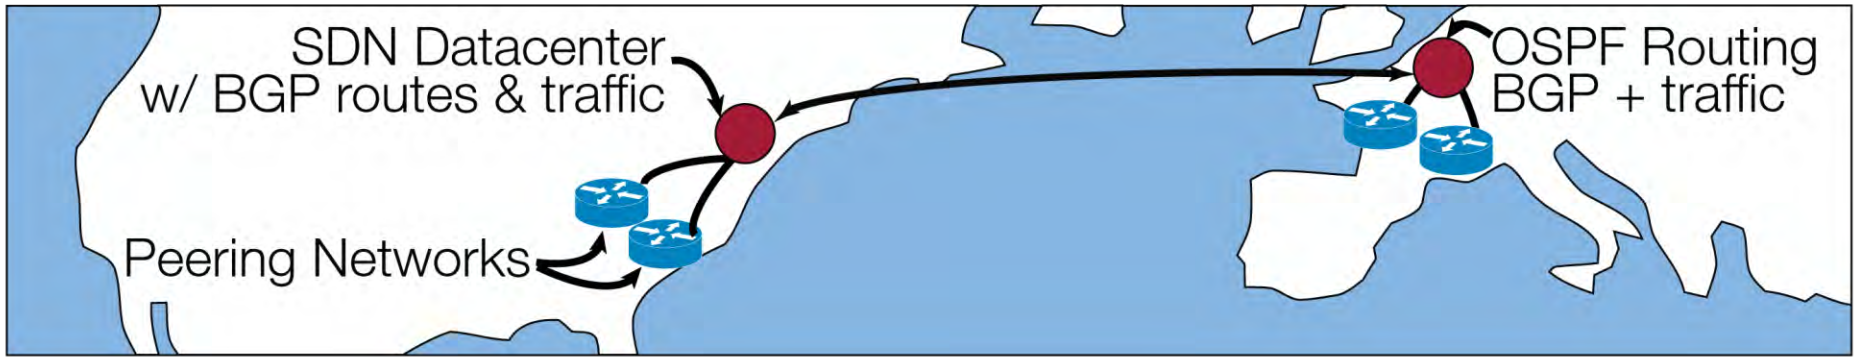

Wish List for Anycast Experiment:

1. Intercontinental private wide-area network (WAN) between US and EU
2. Different routing domains in each datacenter (SDN & OSPF)
3. Ability to exchange *routes and traffic* via BGP with other peers
4. Rich peering connectivity with *REAL* providers in US and EU

Anycast Experiment

Experiment has substantial requirements and overhead

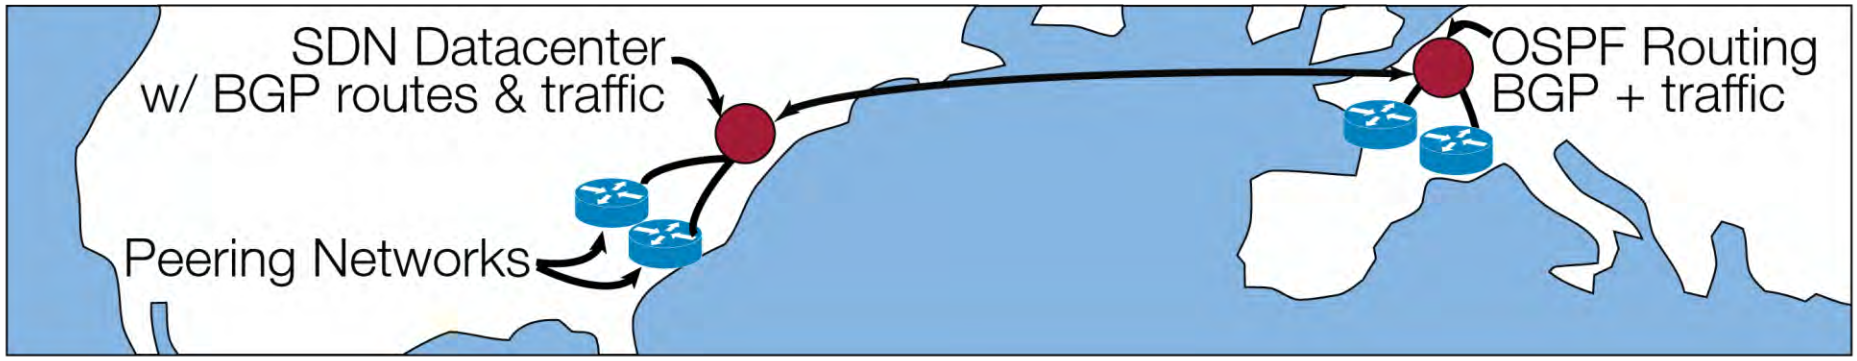

To conduct this experiment, you need to:

1. Acquire ASN, IP space, equipment, colo, diverse connectivity
2. Setup and maintain supporting infrastructure
3. Build control infrastructure to execute experiment(s)

Anycast Experiment

Experiment has substantial requirements and overhead

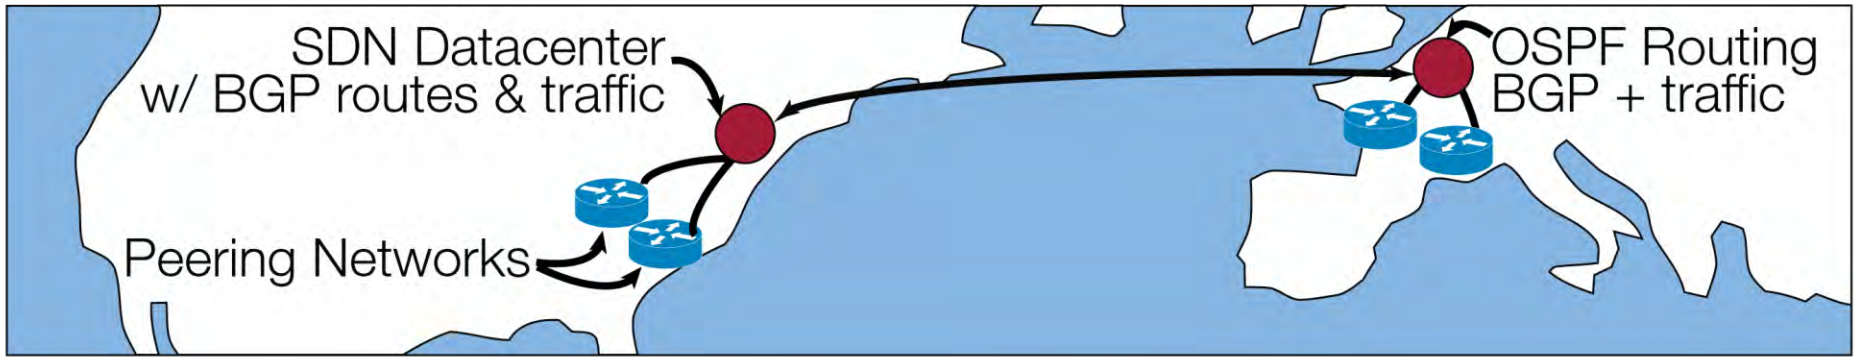

To conduct this experiment, you need to:

1. Acquire ASN, IP space, equipment, colo, diverse connectivity
2. Setup and maintain supporting infrastructure
3. Build control infrastructure to execute experiment(s)

Experiments often impossible given requirements

Anycast Experiment

Using Testbeds to Lower Barriers

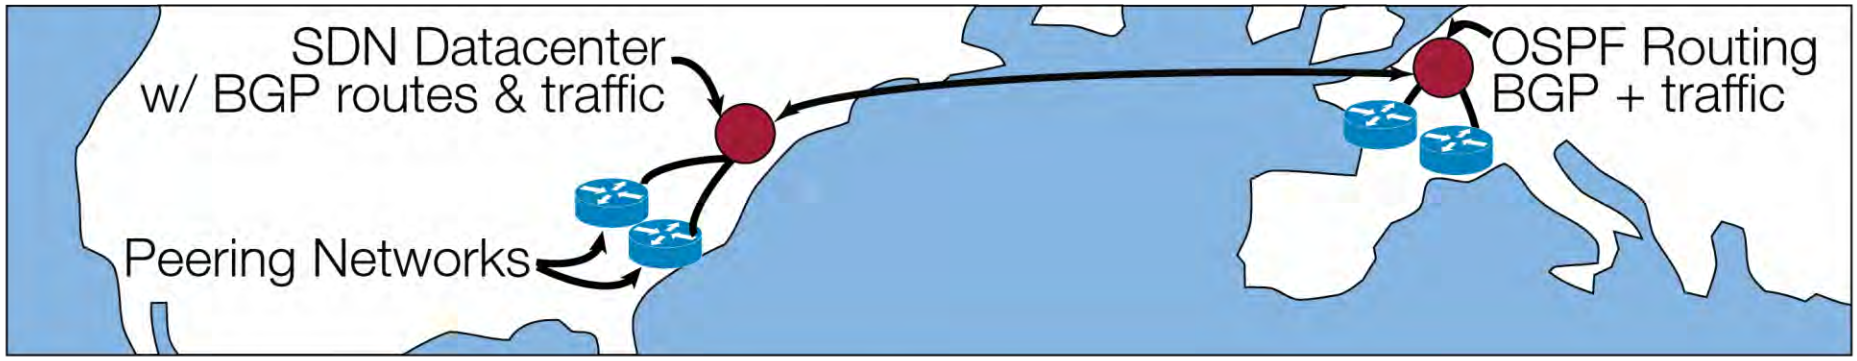

No testbeds support this experiment

Anycast Experiment

Using Testbeds to Lower Barriers

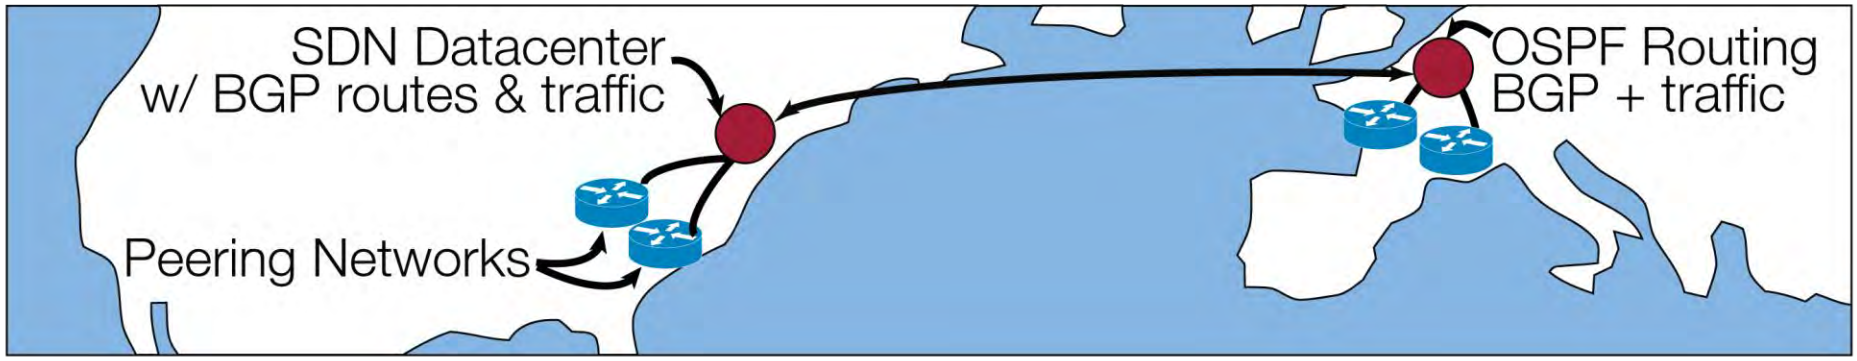

No testbeds support this experiment

However, we have pieces which bring us closer...

We'll improve these pieces and then merge them to create a testbed for this Anycast experiment

Anycast Experiment

Intradoman and Interdomain Pieces of Experiment

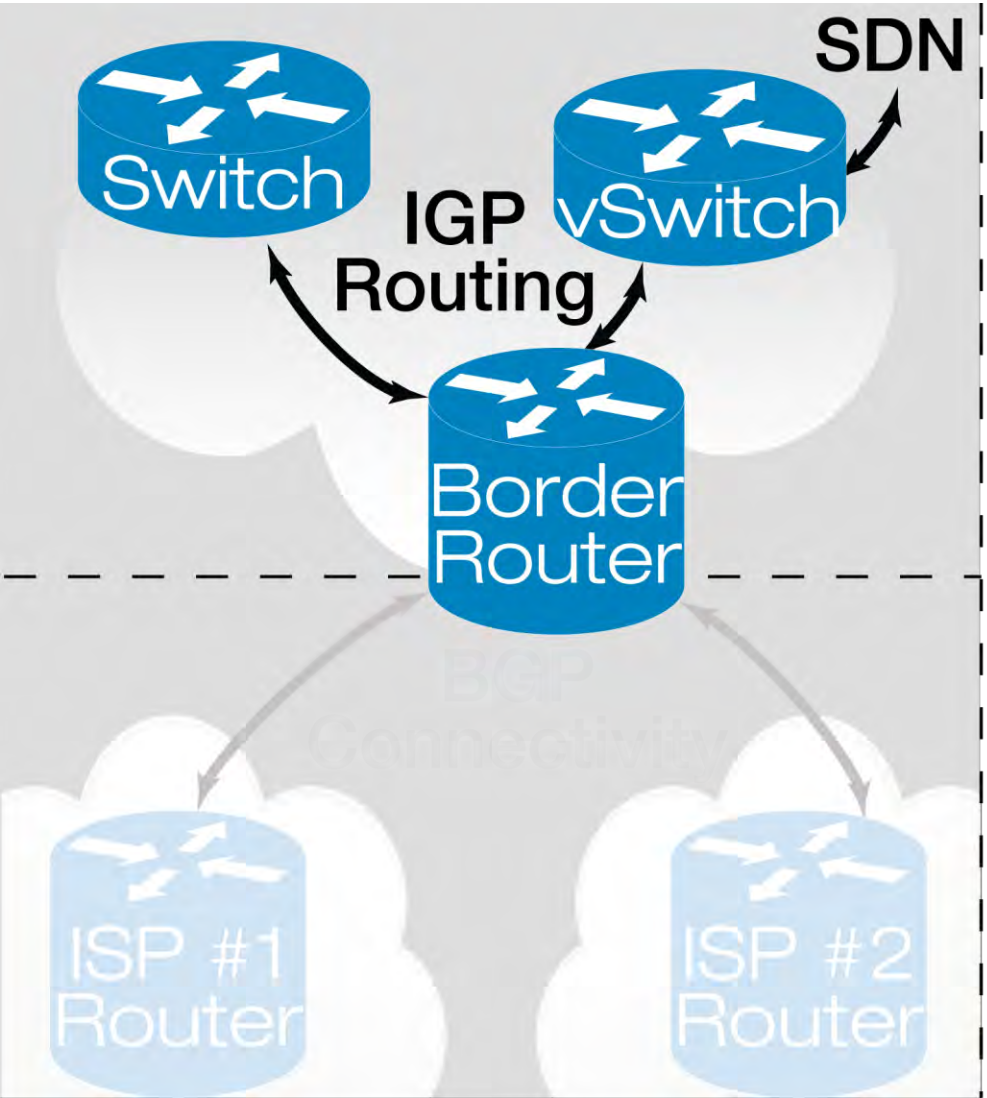

Mininet *Can* Emulate

SDN components

Network topologies

Anycast Experiment

Intradoman and Interdomain Pieces of Experiment

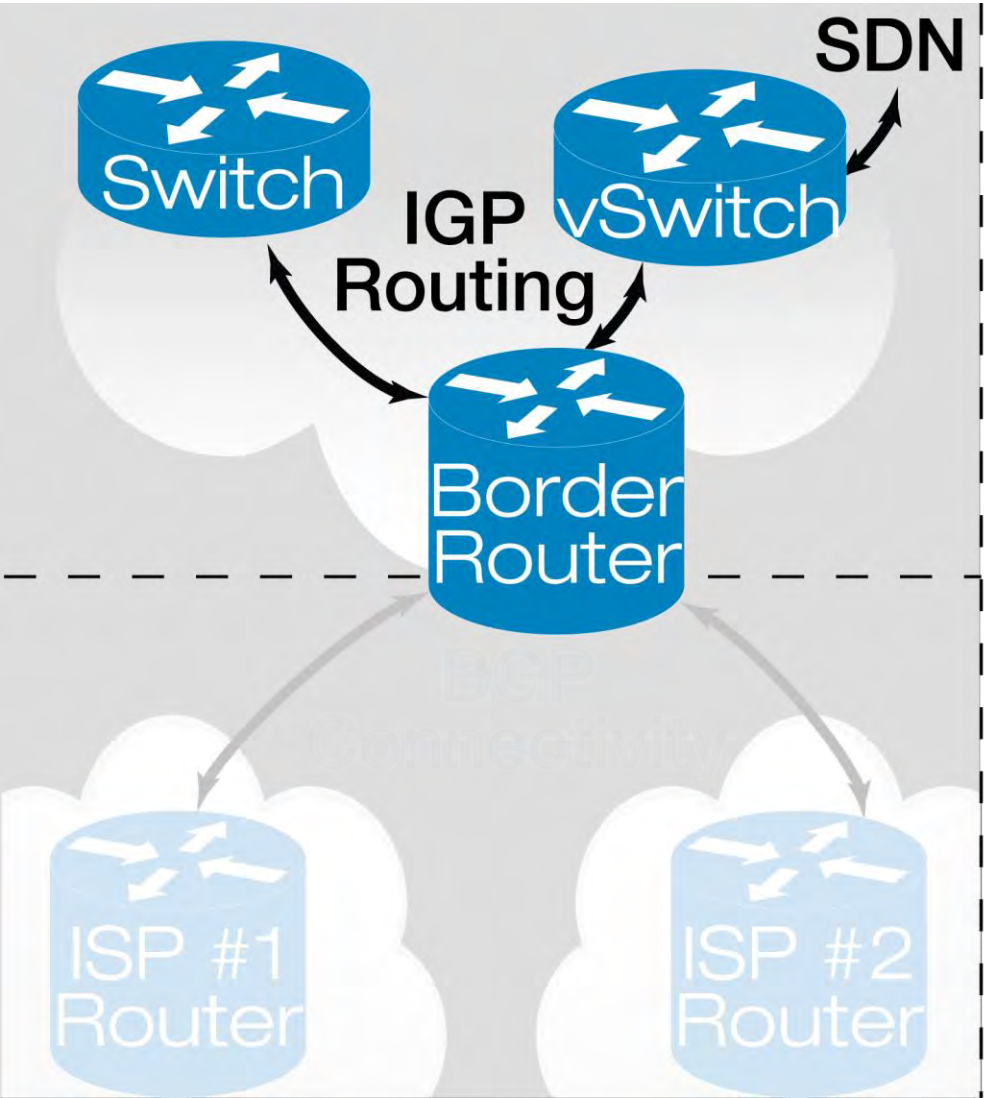

Mininet *Can* Emulate

SDN components

Network topologies

Mininet *Cannot* Emulate

Inter & intra domain routing

Routing engines (Quagga)

Anycast Experiment

Intradoman and Interdomain Pieces of Experiment

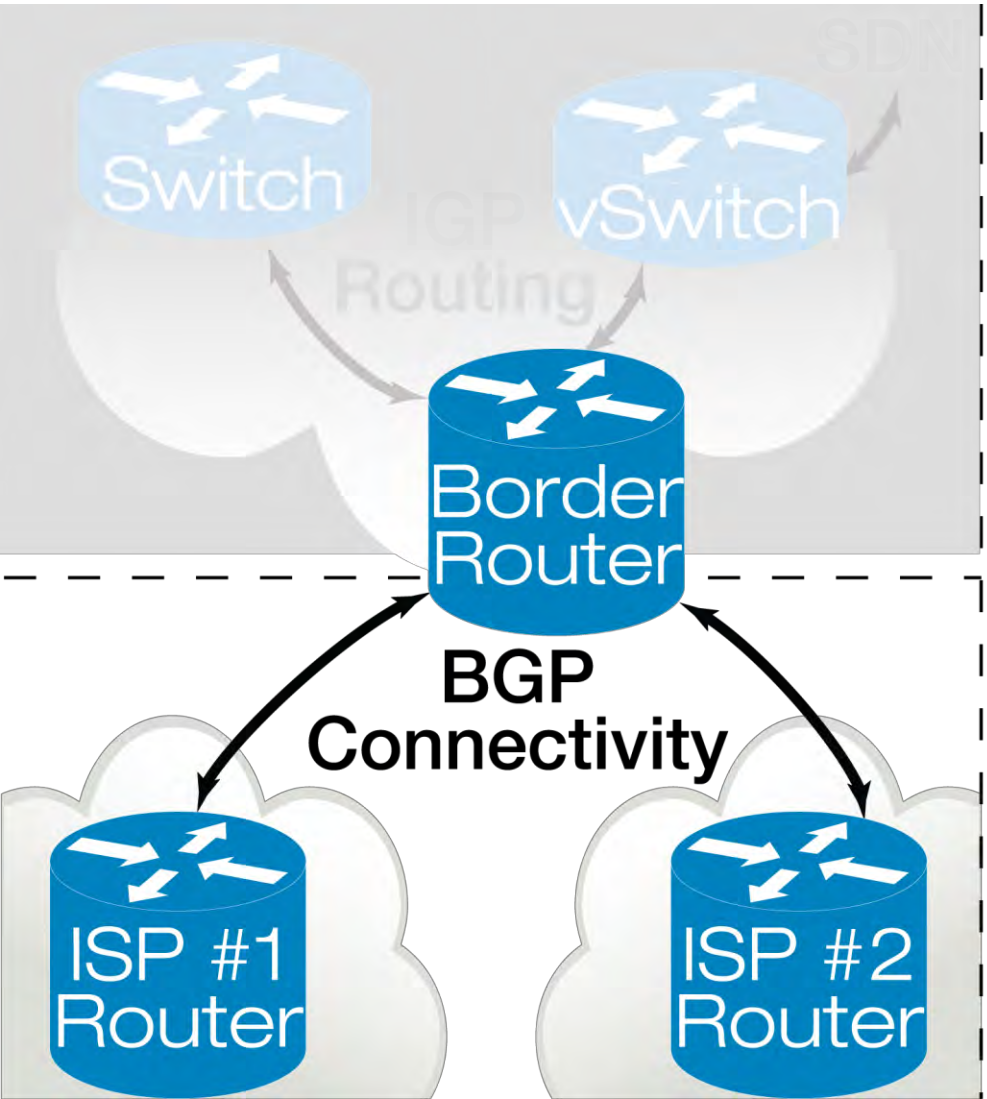

Mininet *Can* Emulate

SDN components

Network topologies

Mininet *Cannot* Emulate

Inter & intra domain routing

Routing engines (Quagga)

Transit Portal Provides

Network transit

BGP multiplexing

Connectivity with *real* ISPs

Mininet

Designed for Emulation of SDN Networks

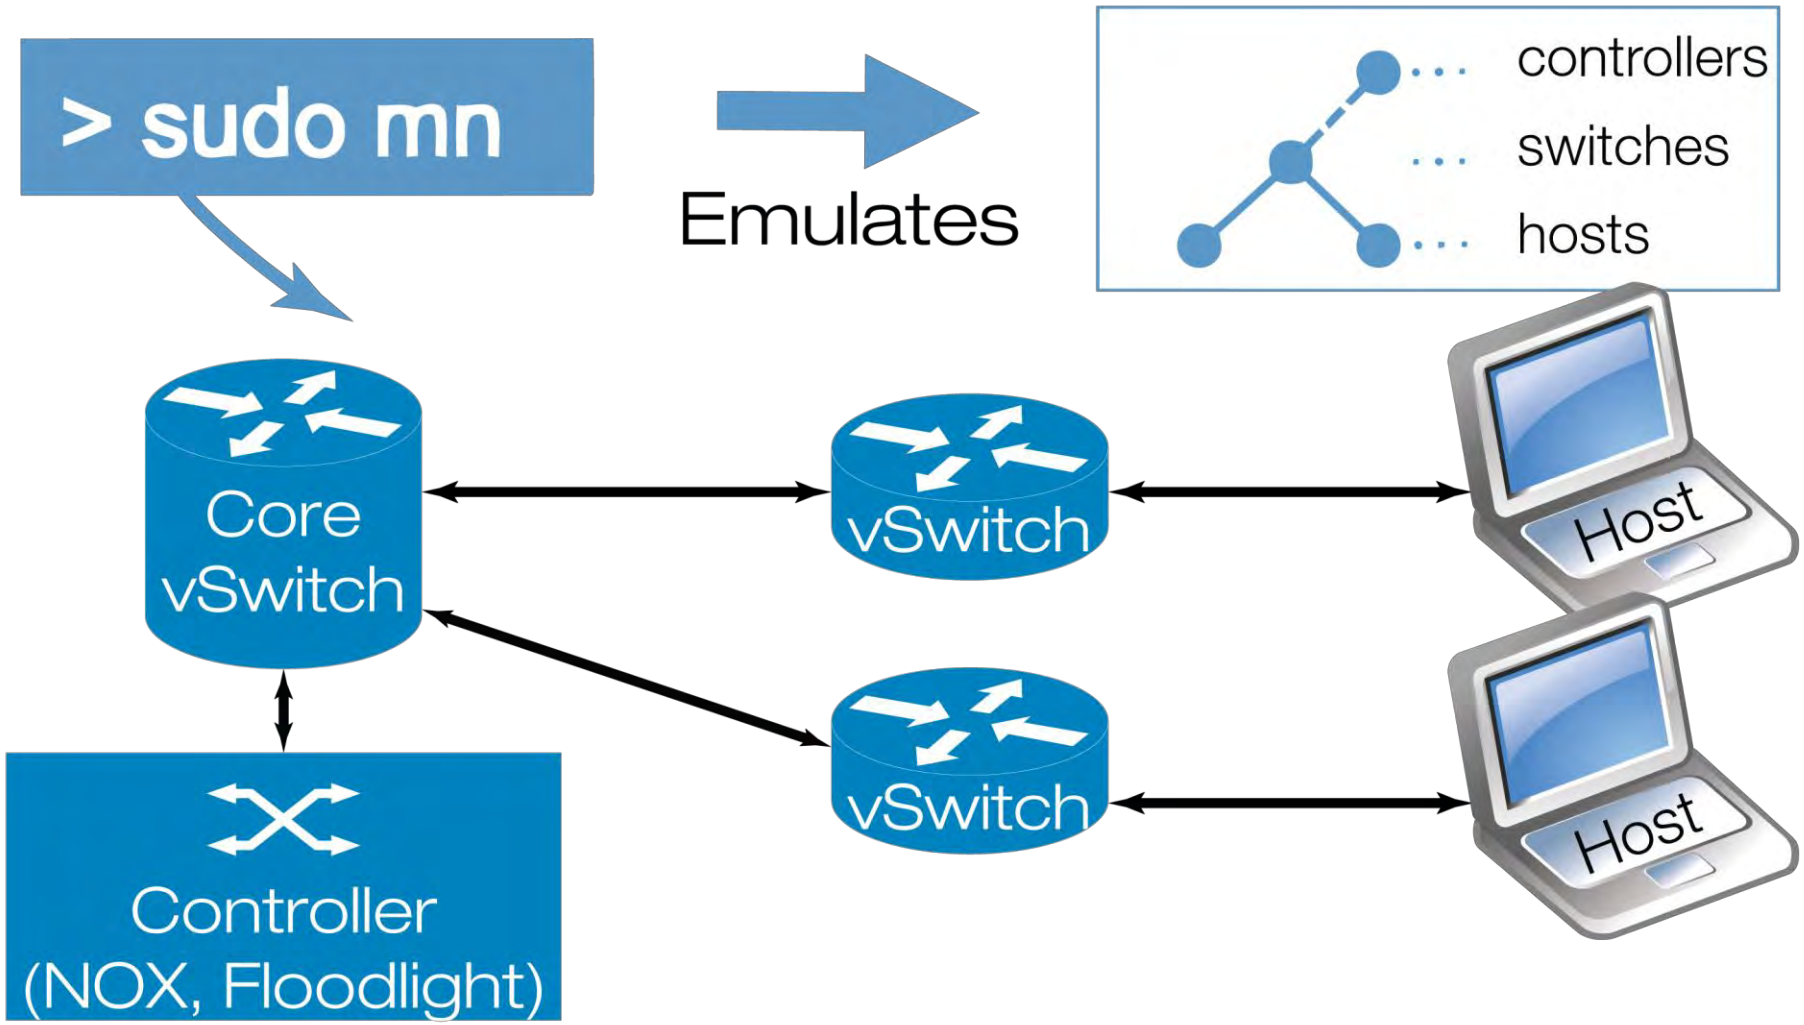

Mininet

Doesn't Support Intra and Interdomain Routing


```
> sudo mn
```

By default, cannot emulate:

interdomain routing (BGP)
intradomain routing (OSPF, RIP)
routing engines (Quagga, BIRD)

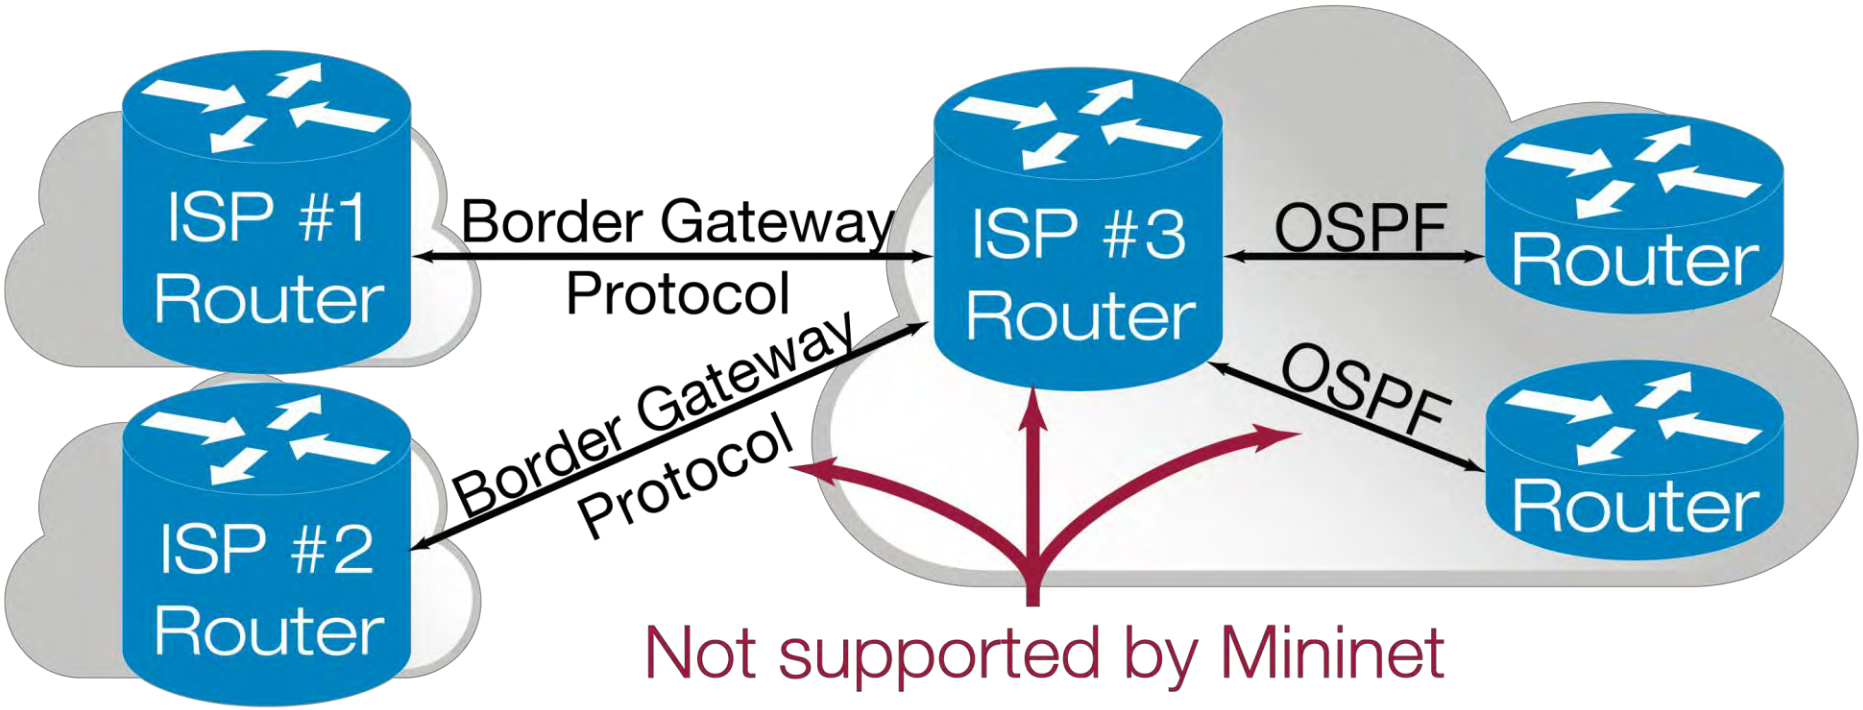

Enhanced Mininet

Supports Intra and Interdomain Routing


```
> sudo mn
```

Enhanced Mininet Adds Support For:

- topologies with **IGP + BGP** support
- **isolated routing engine instances** (Quagga, BIRD) in Mininet
- requires **significantly less resources** than virtual machines

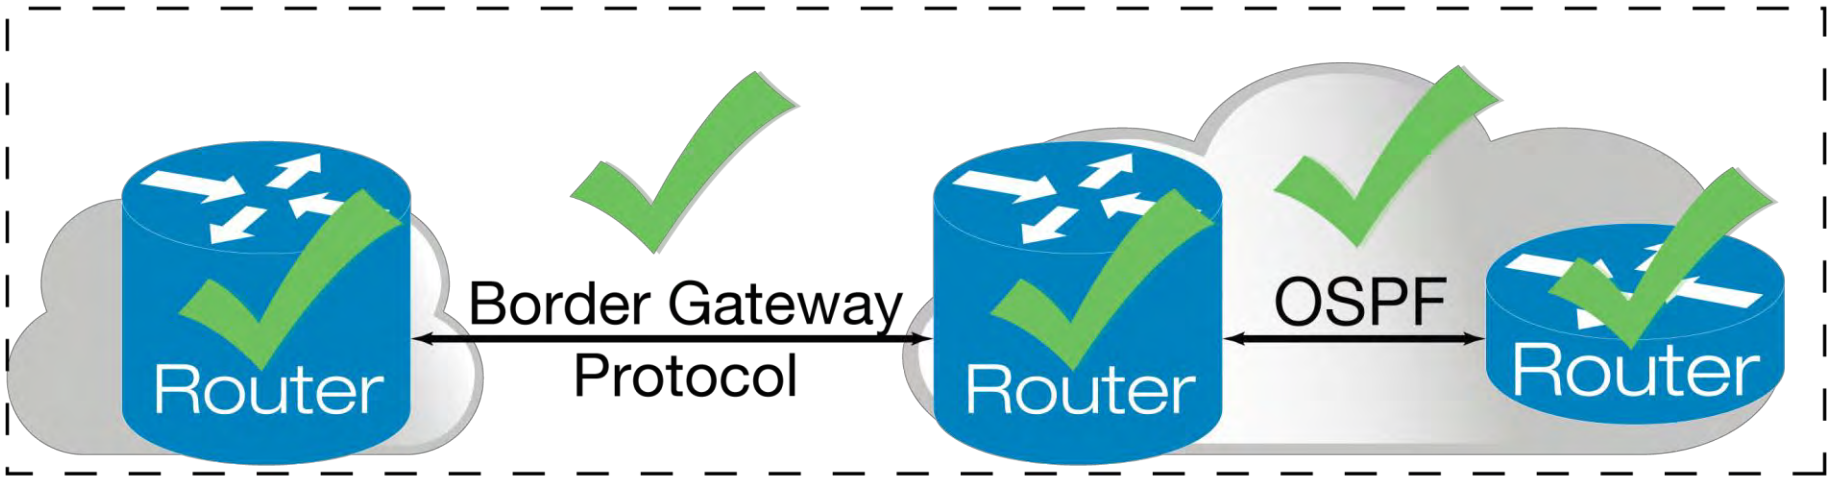

Anycast Experiment

Emulating Intradoman and Interdomain Pieces

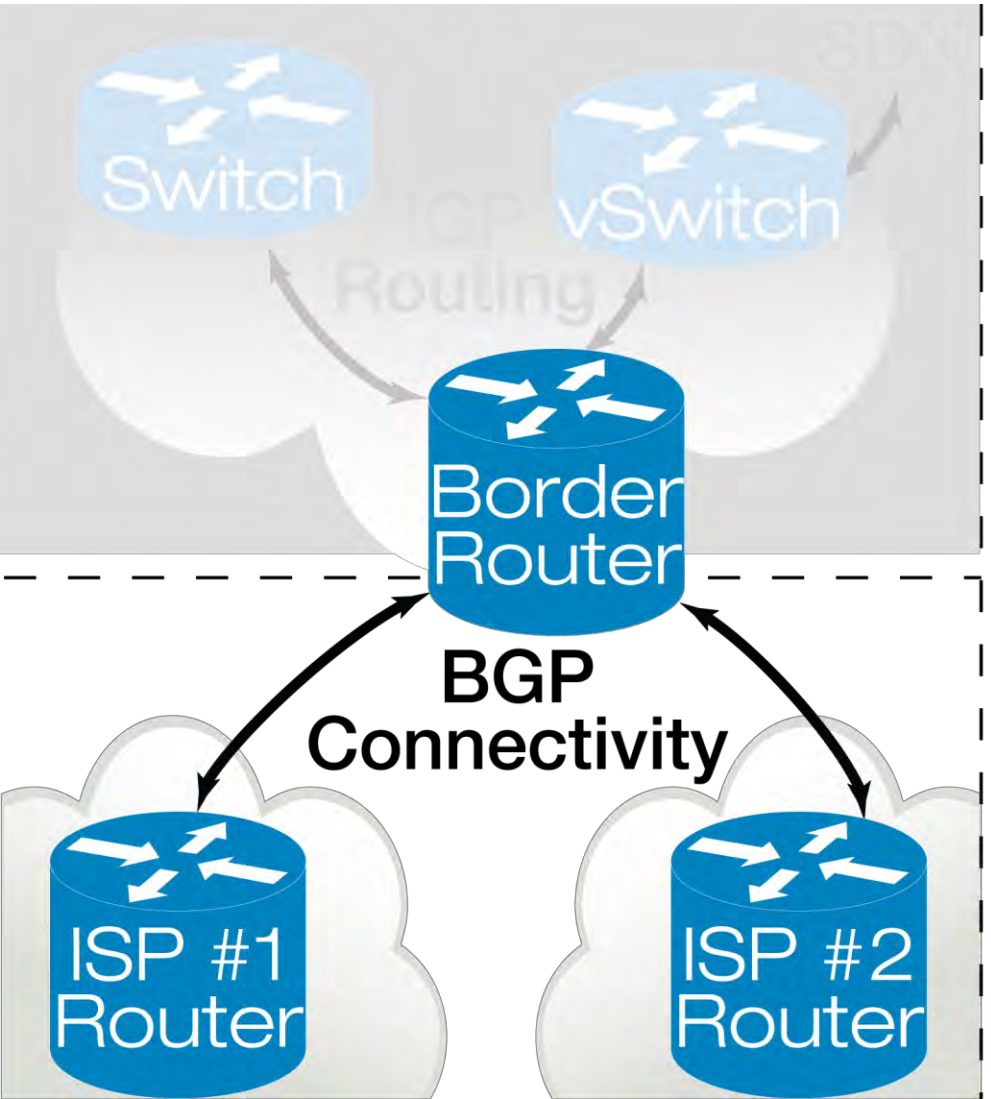

Our Enhanced Mininet Can Emulate

SDN components

Intradomain routing

Network topologies

Routing engines (Quagga)

But We Cannot Emulate Interdomain Routing

Anycast Experiment

Why don't we Emulate Peer ISPs / The Internet?

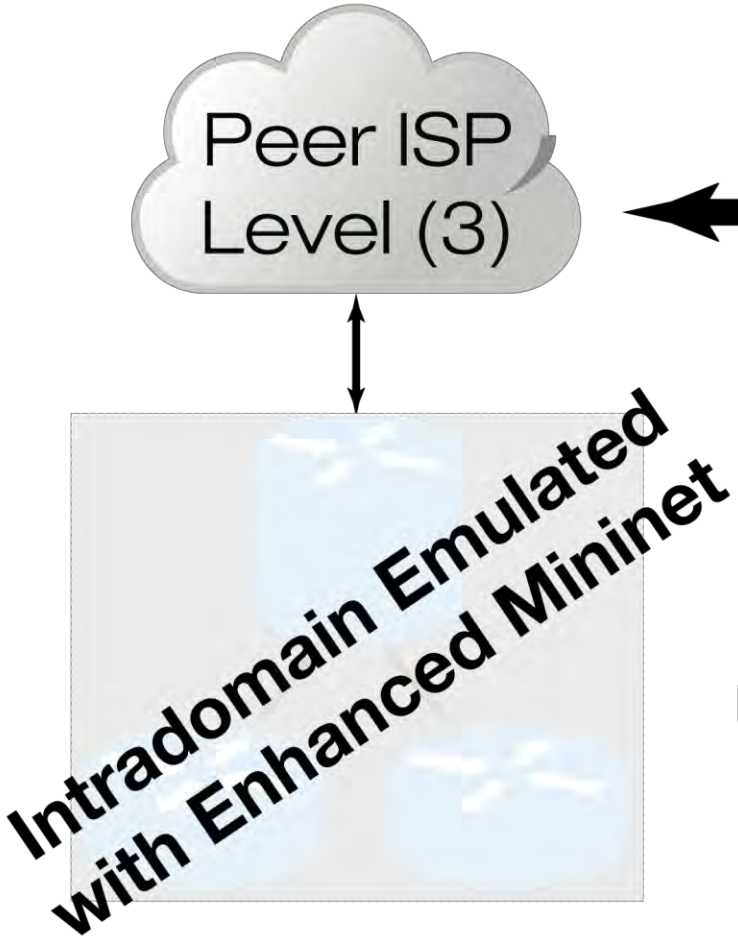

Cannot model / emulate peer ISP due to:

Unpredictable routes

Transient & persistent failures

Impact of peering relationships

Impact of neighboring networks

Unknown policies, topologies, goals

Anycast Experiment

Using Transit Portal for Interdomain and Emulating Intradoman

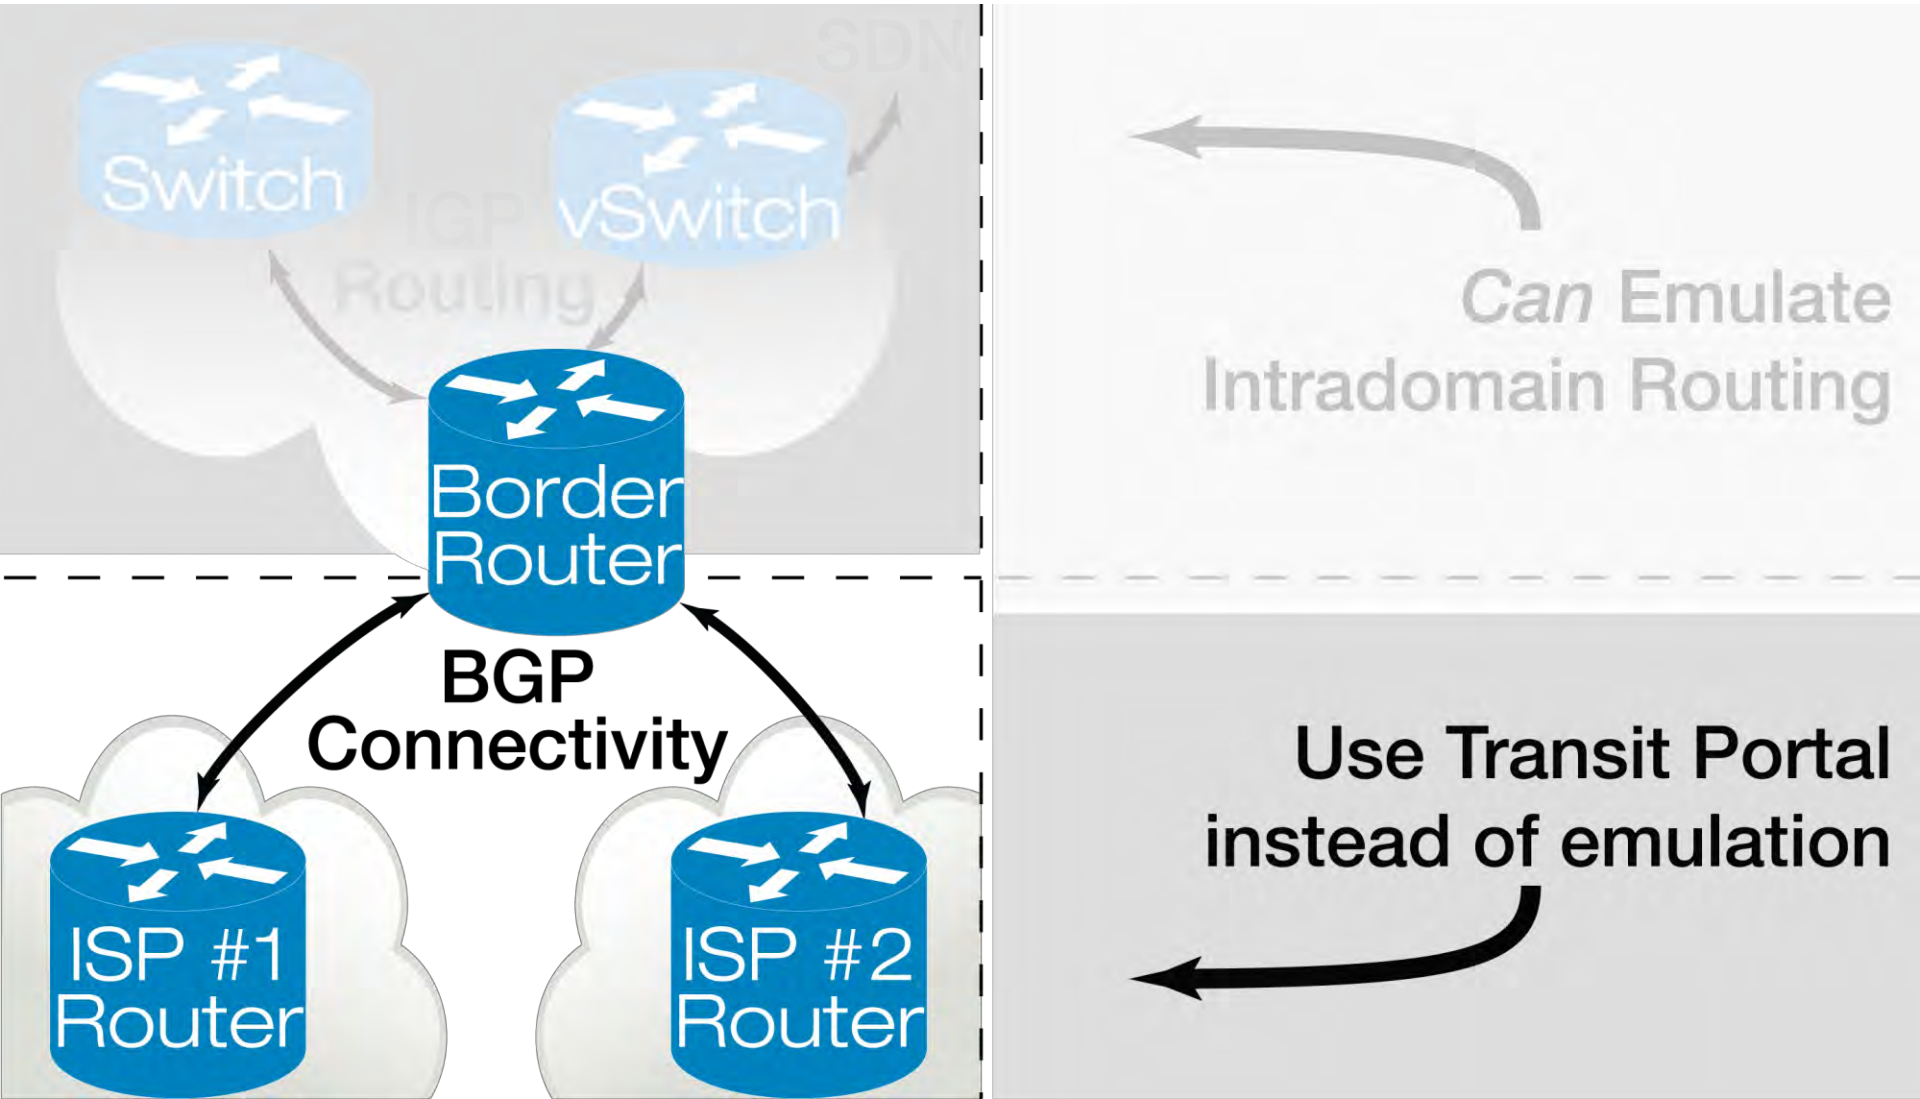

Transit Portal

BGP Multiplexing and Traffic Exchange Service

ASN: *47065* | **IP allocation:** *184.164.224.0/19*

BGP multiplexing provides private, isolated BGP sessions

6 points of presence across two continents
Connectivity to over 600 *REAL* ISPs via BGP

Infrastructure is constantly growing

Transit Portal

BGP Multiplexing and Traffic Exchange Service

Strengths of Transit Portal

Can send / receive *routes* and *traffic* via BGP

Connections to diverse set of **REAL** networks and ISPs

Exposure to unpredictable routes, policies and failures

Previous Use in Research [SIGCOMM 12']

PoiRoot: Investigating the Root Cause of Interdomain Path Changes

LIFEGUARD: Practical Repair of Persistent Route Failures

Transit Portal

A Closer Look at the AMS-IX Multiplexer

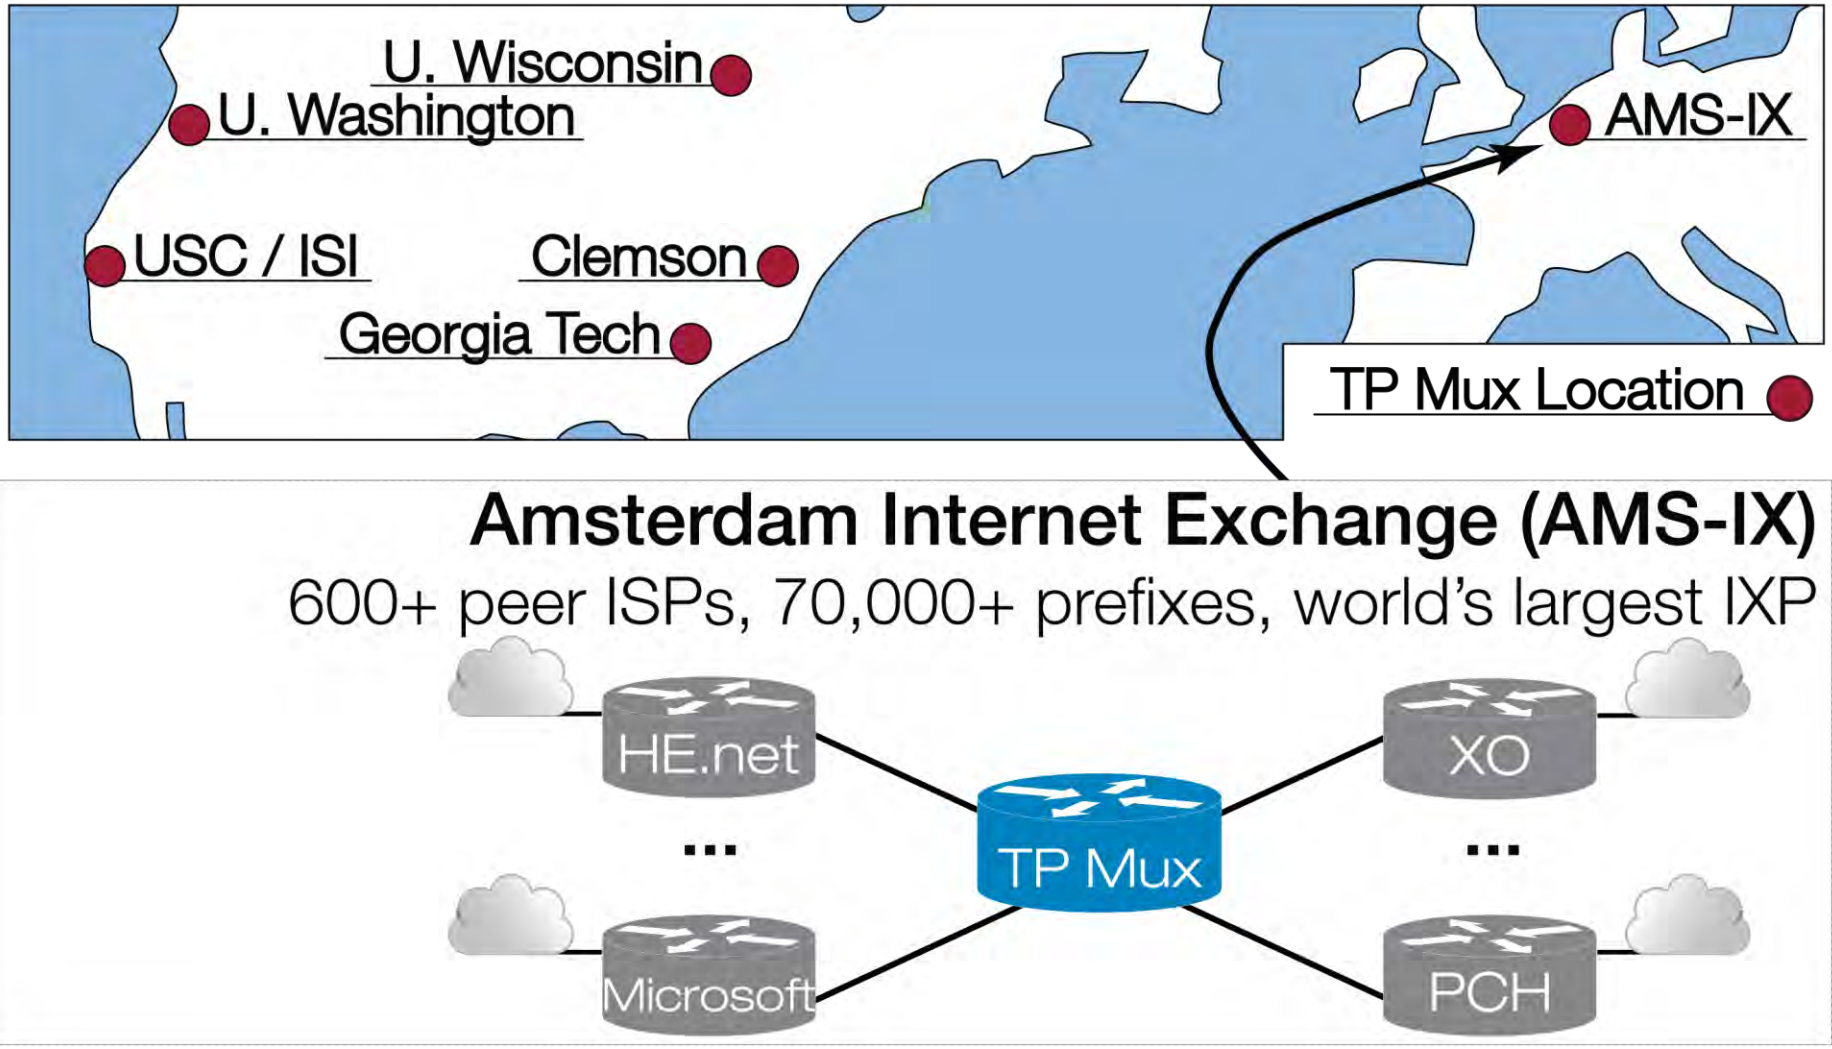

Transit Portal

A Closer Look at the AMS-IX Multiplexer

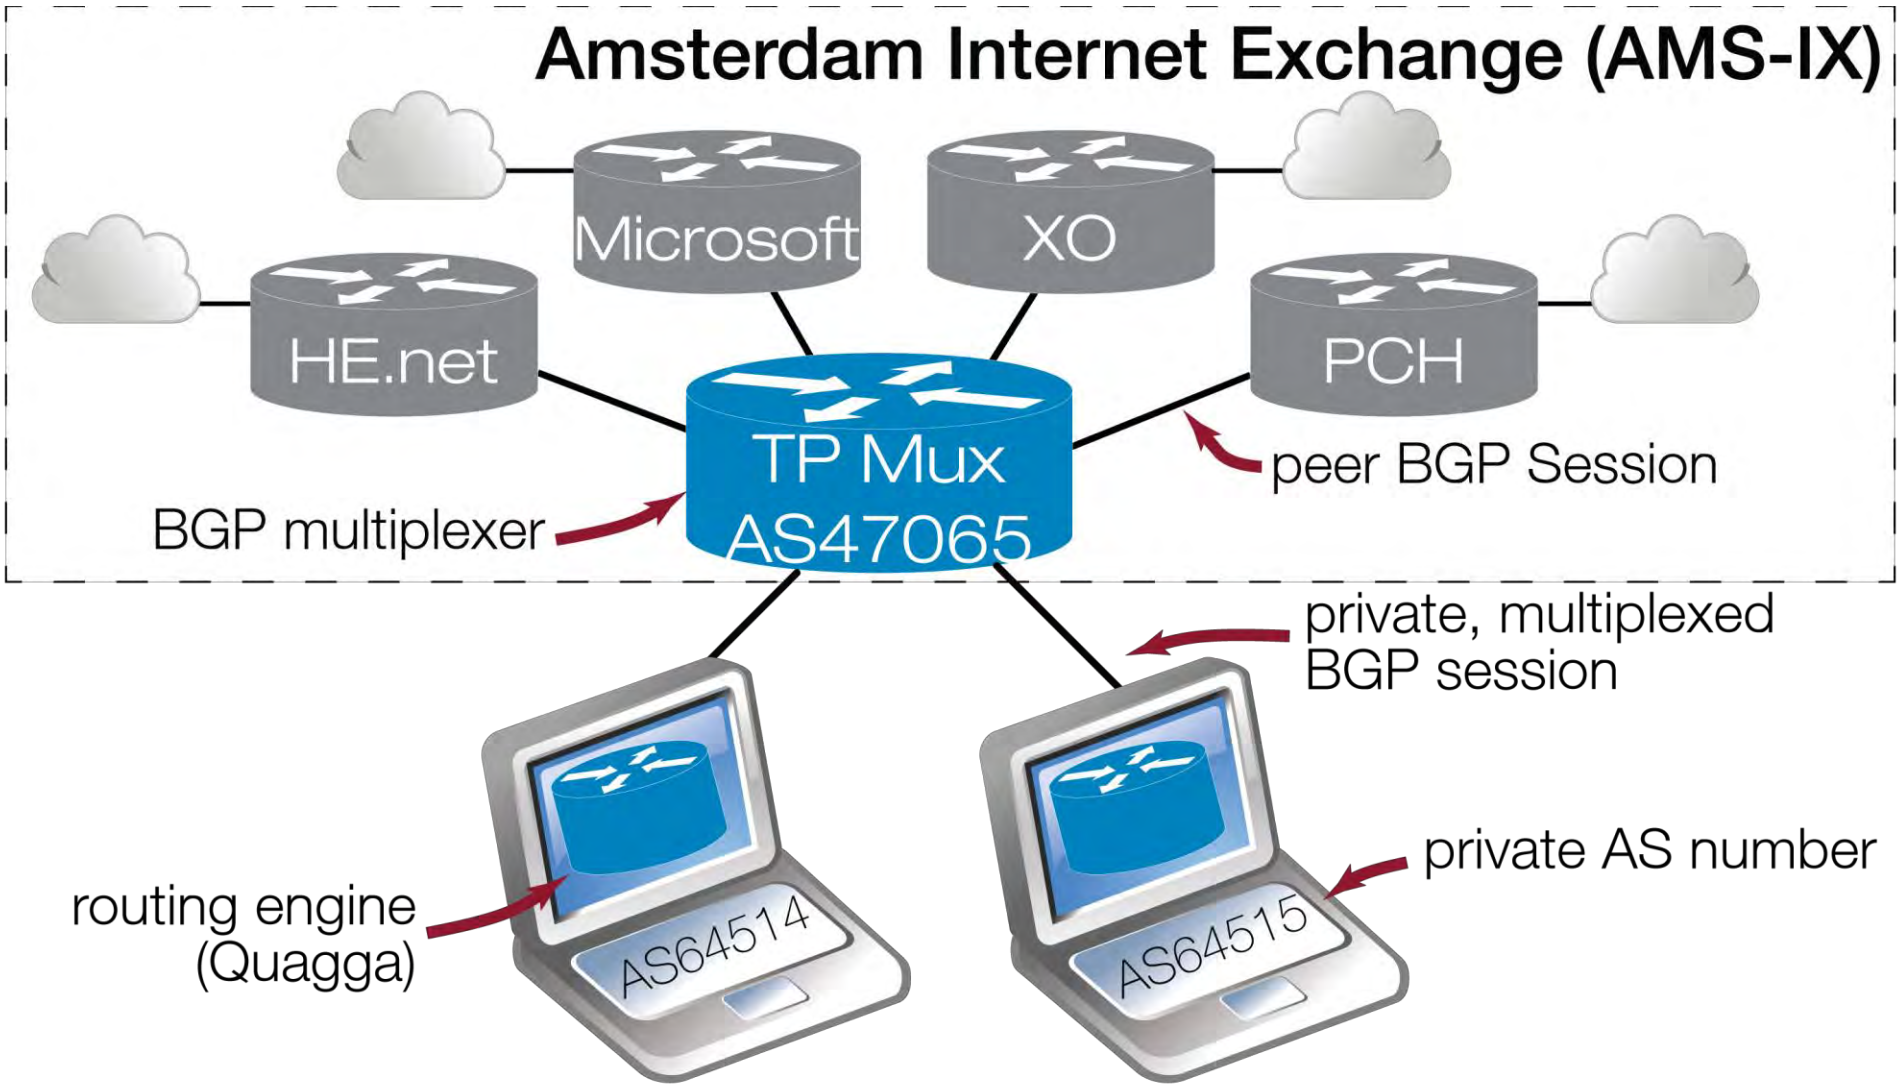

Transit Portal

A Closer Look at the AMS-IX Multiplexer

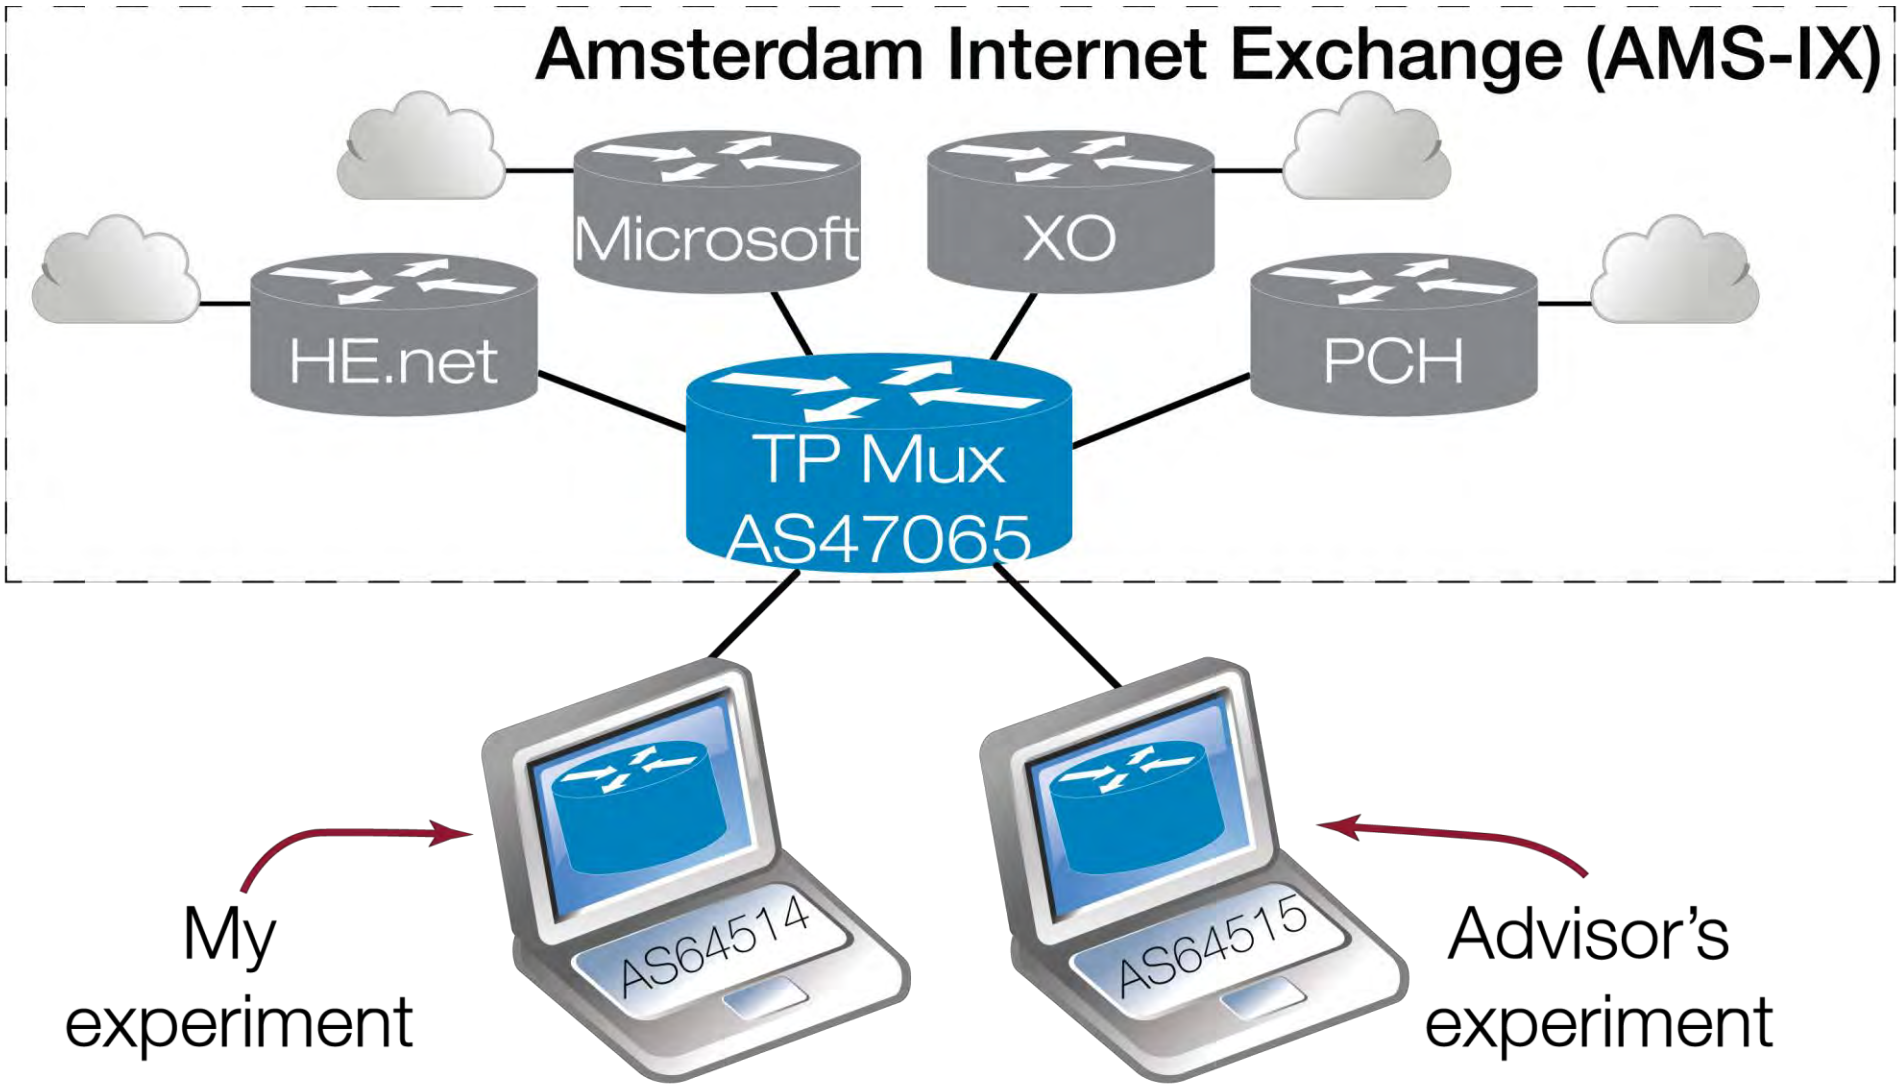

Transit Portal

A Closer Look at the AMS-IX Multiplexer

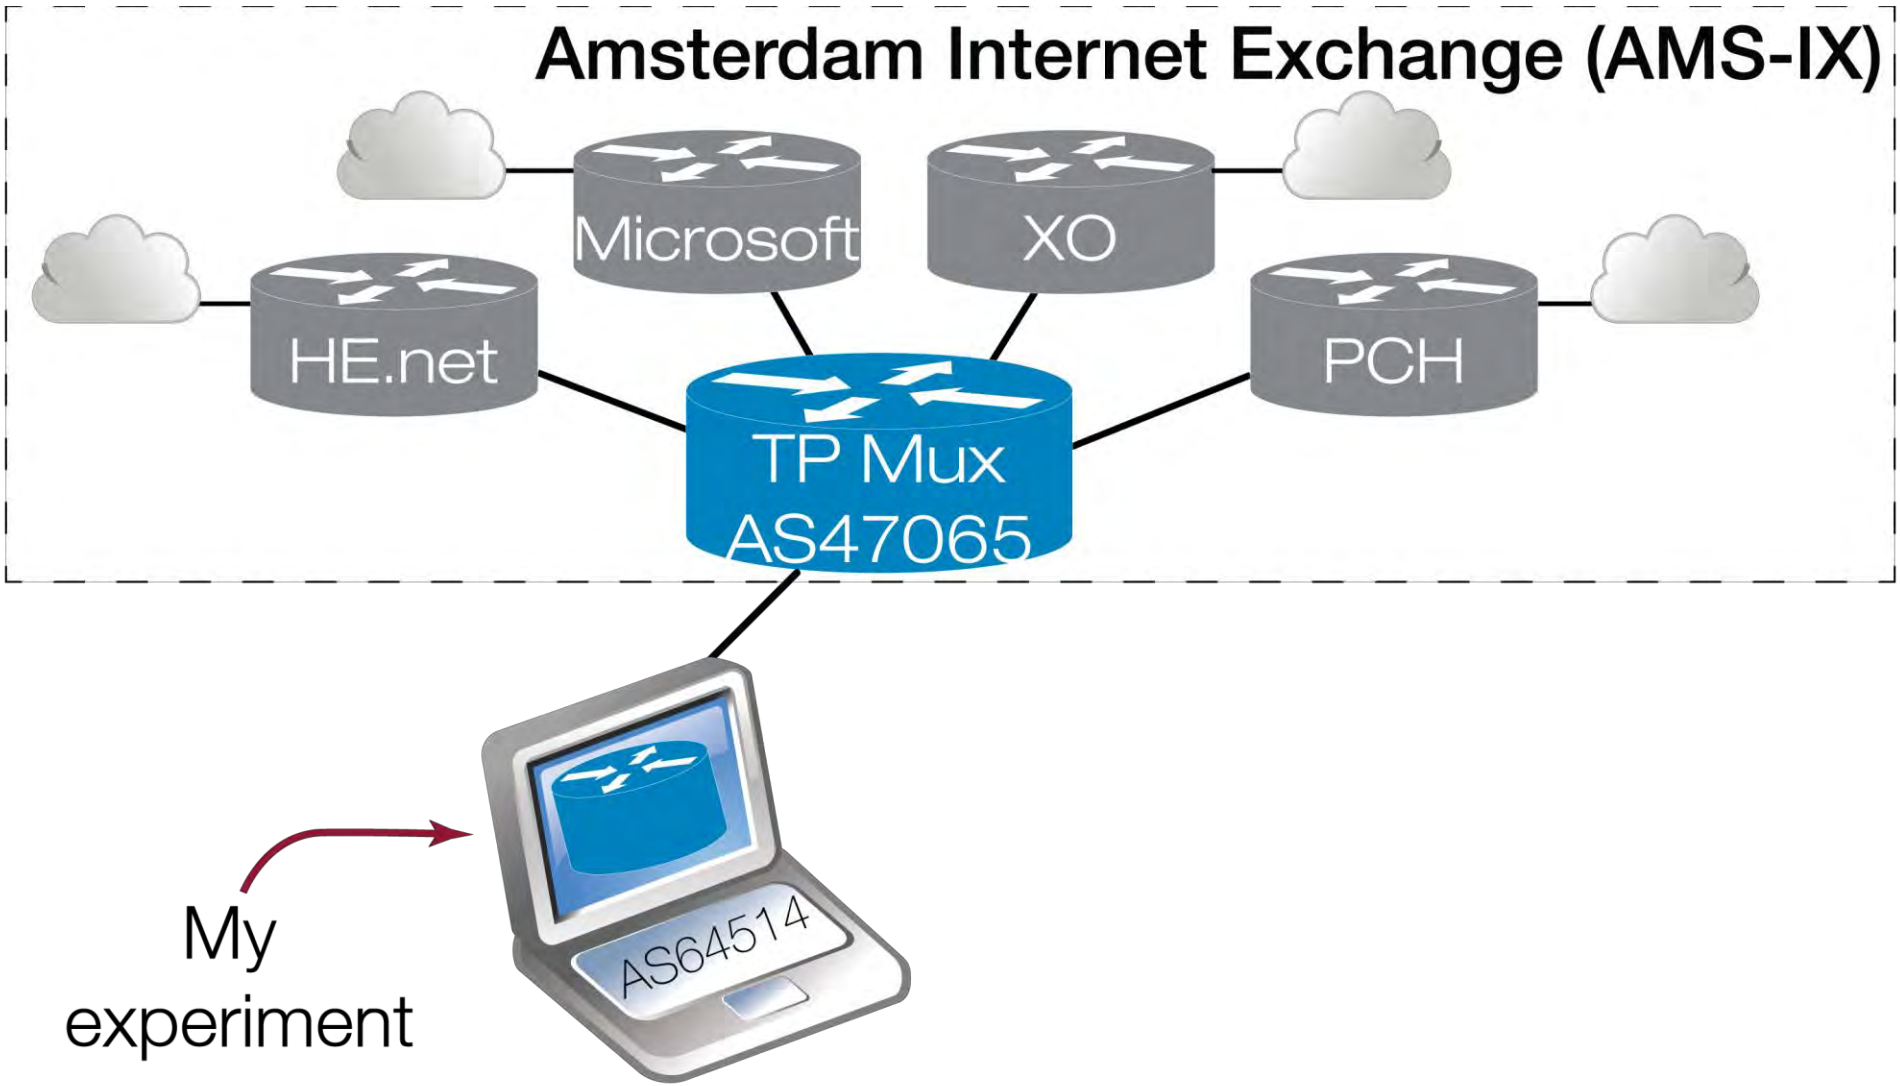

Transit Portal

A Closer Look at the AMS-IX Multiplexer

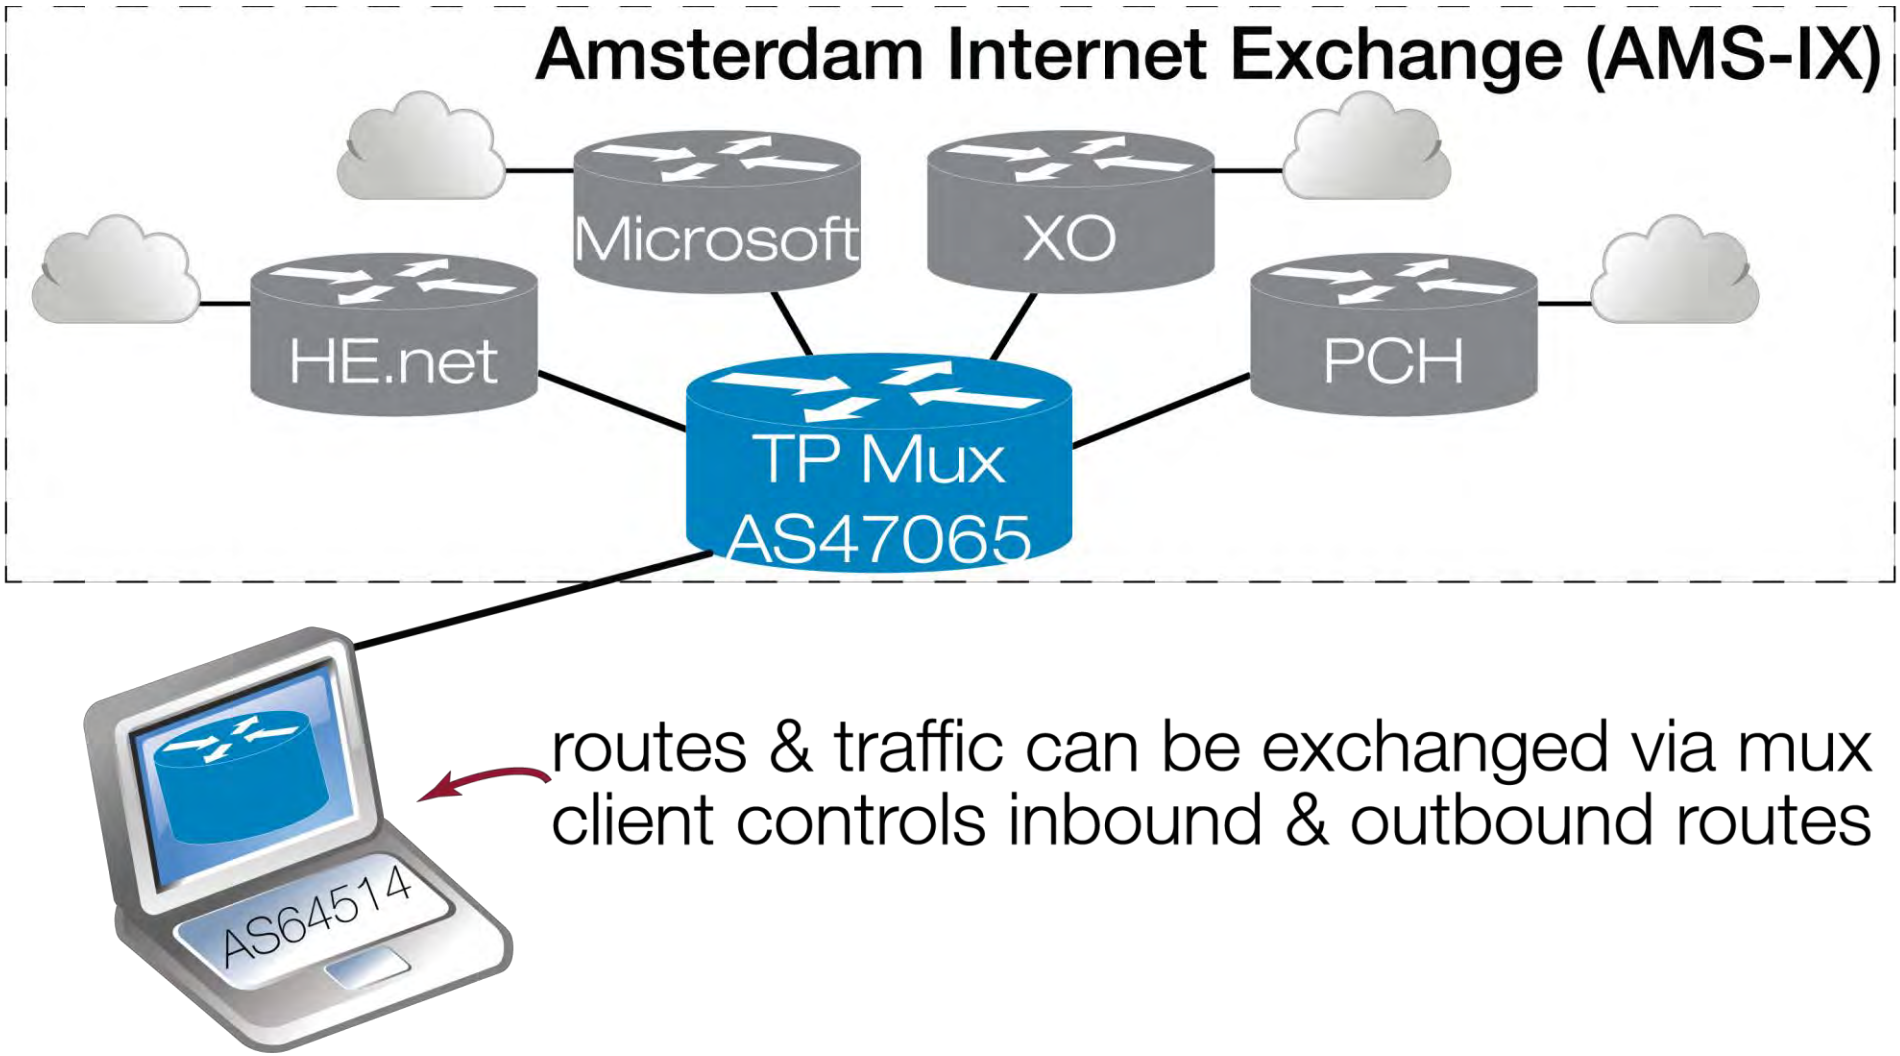

Transit Portal

A Closer Look at the AMS-IX Multiplexer

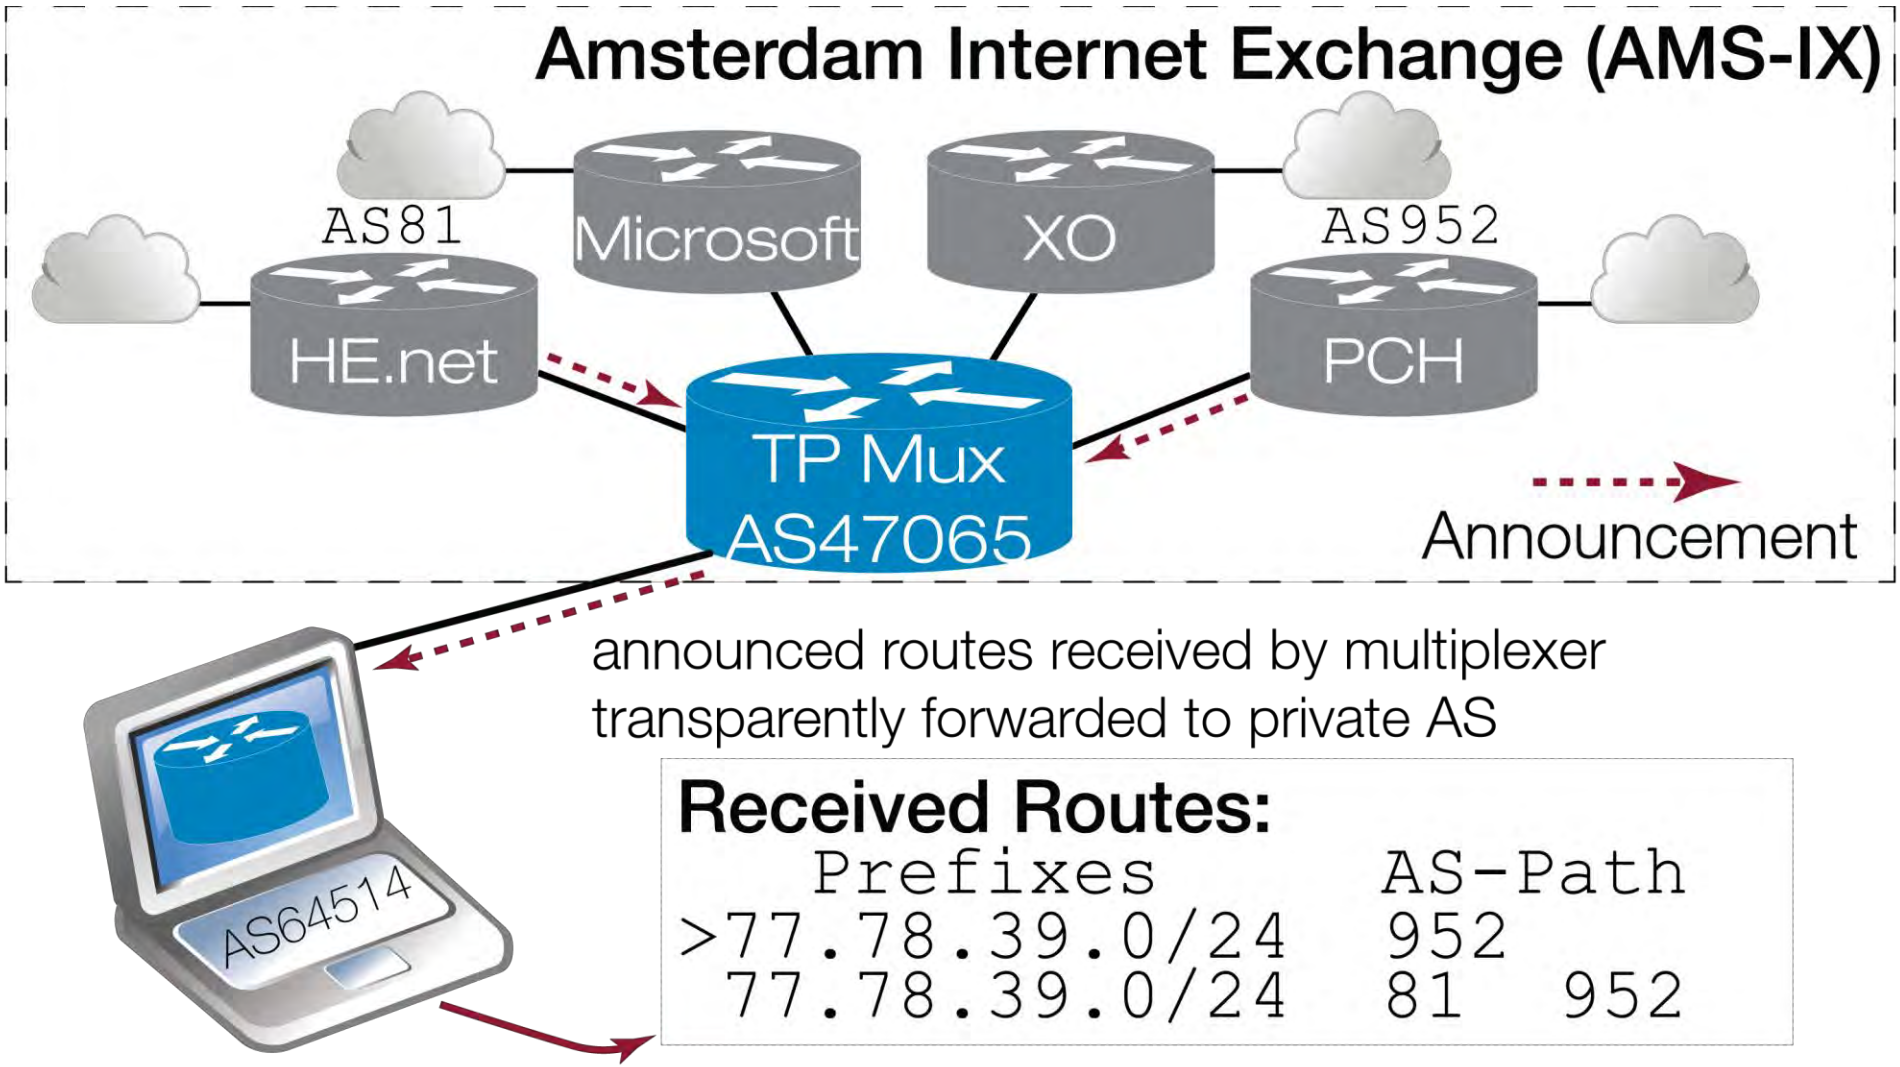

Transit Portal

A Closer Look at the AMS-IX Multiplexer

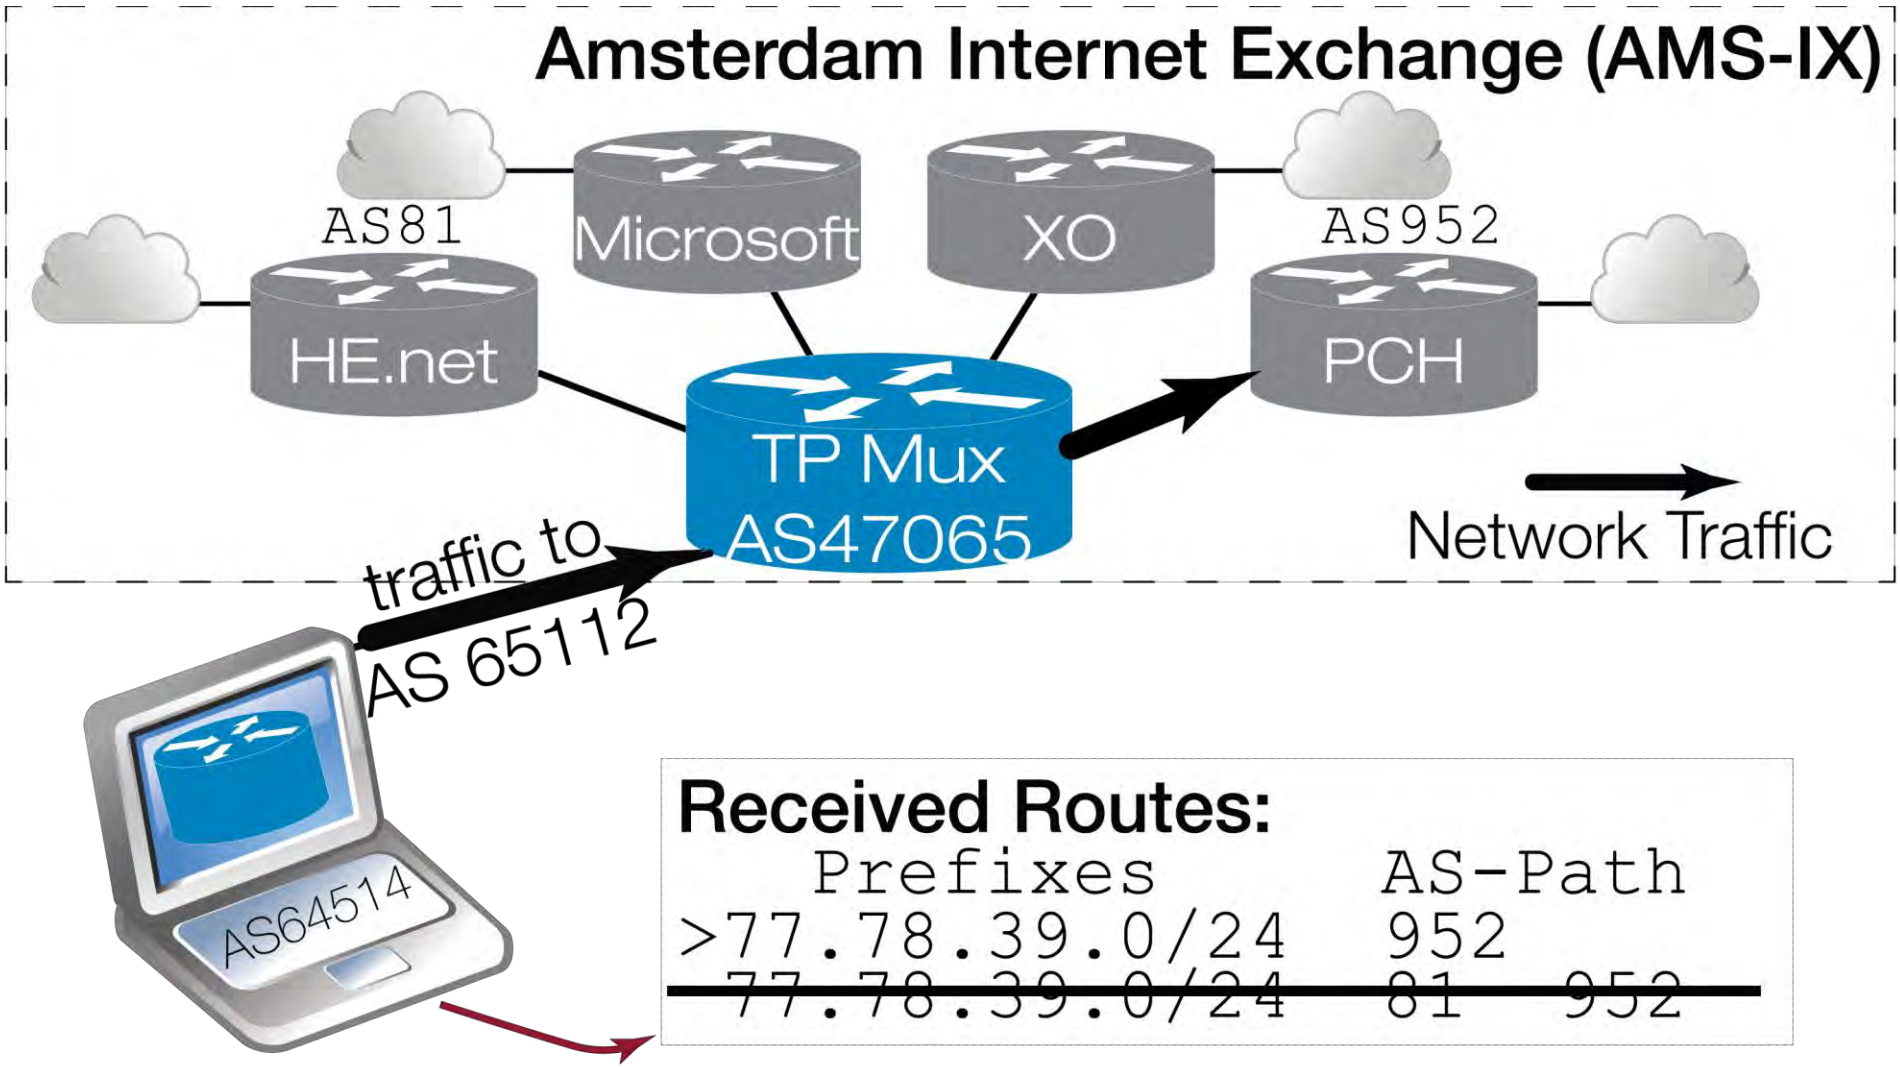

Transit Portal

A Closer Look at the AMS-IX Multiplexer

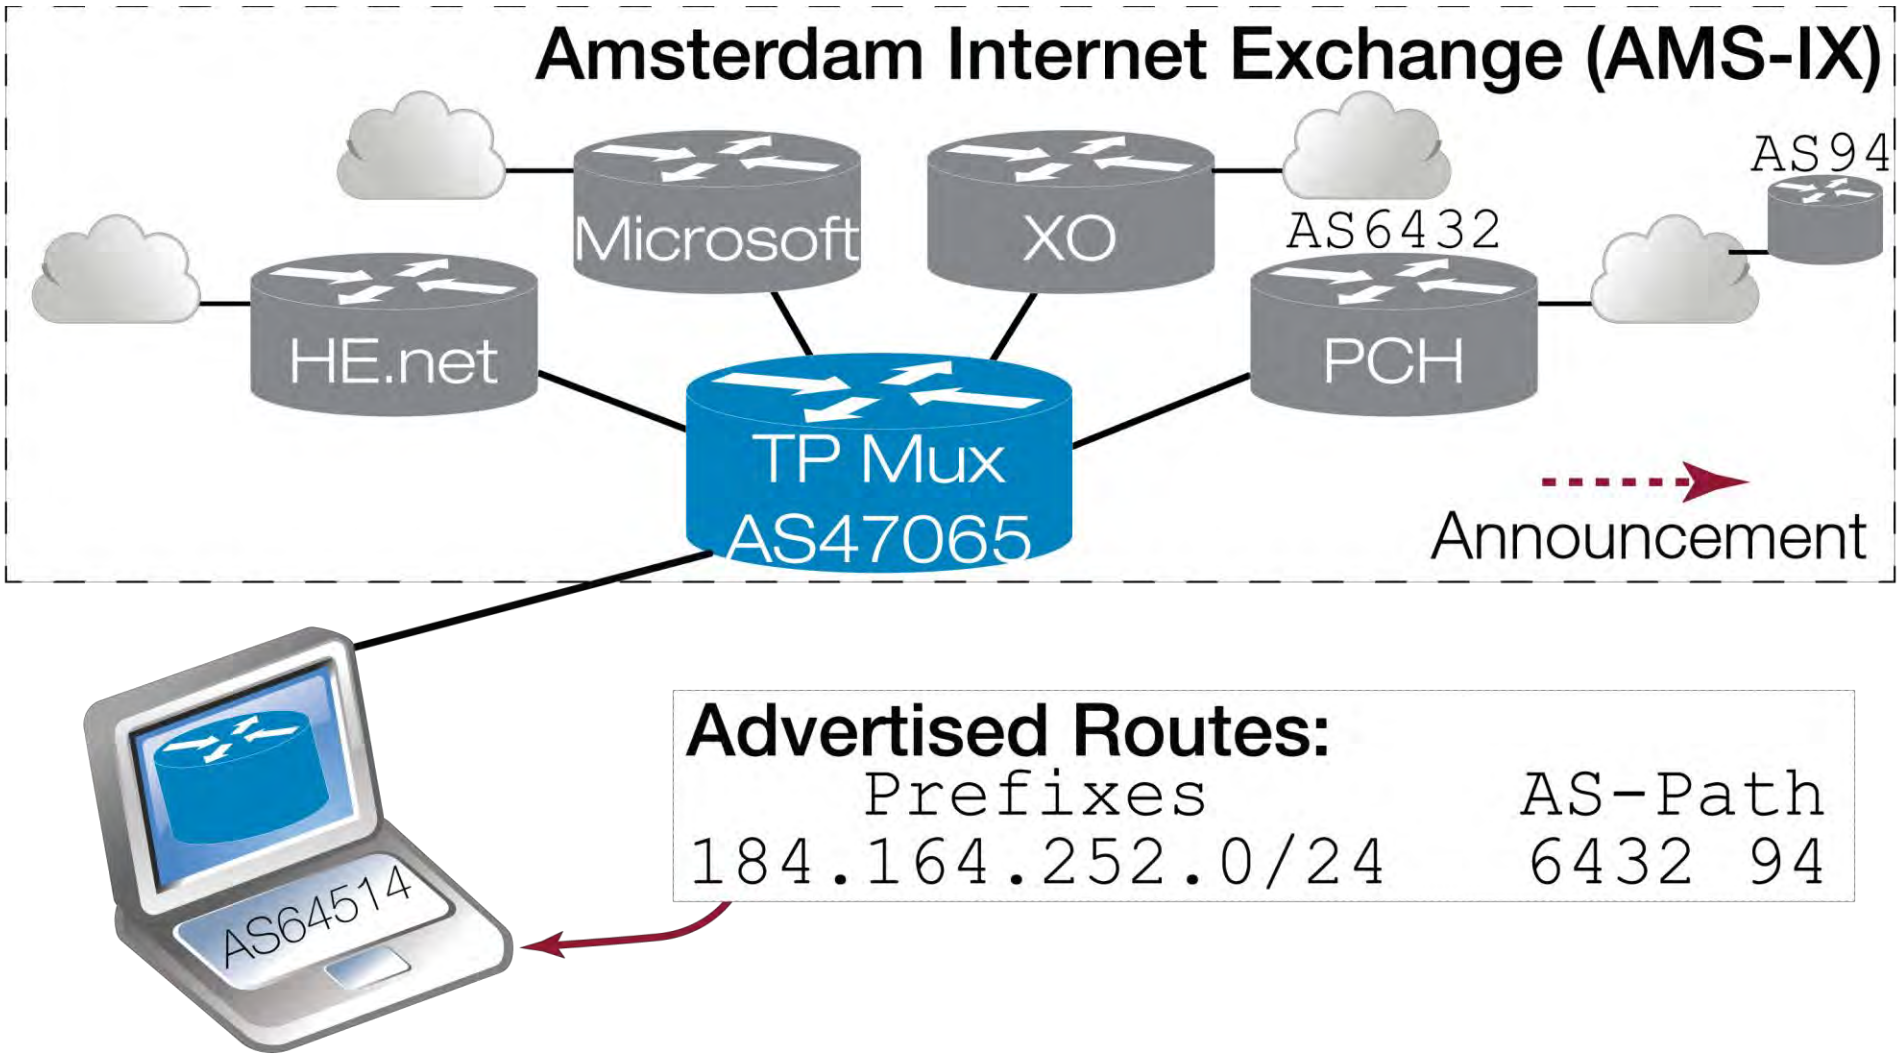

Transit Portal

A Closer Look at the AMS-IX Multiplexer

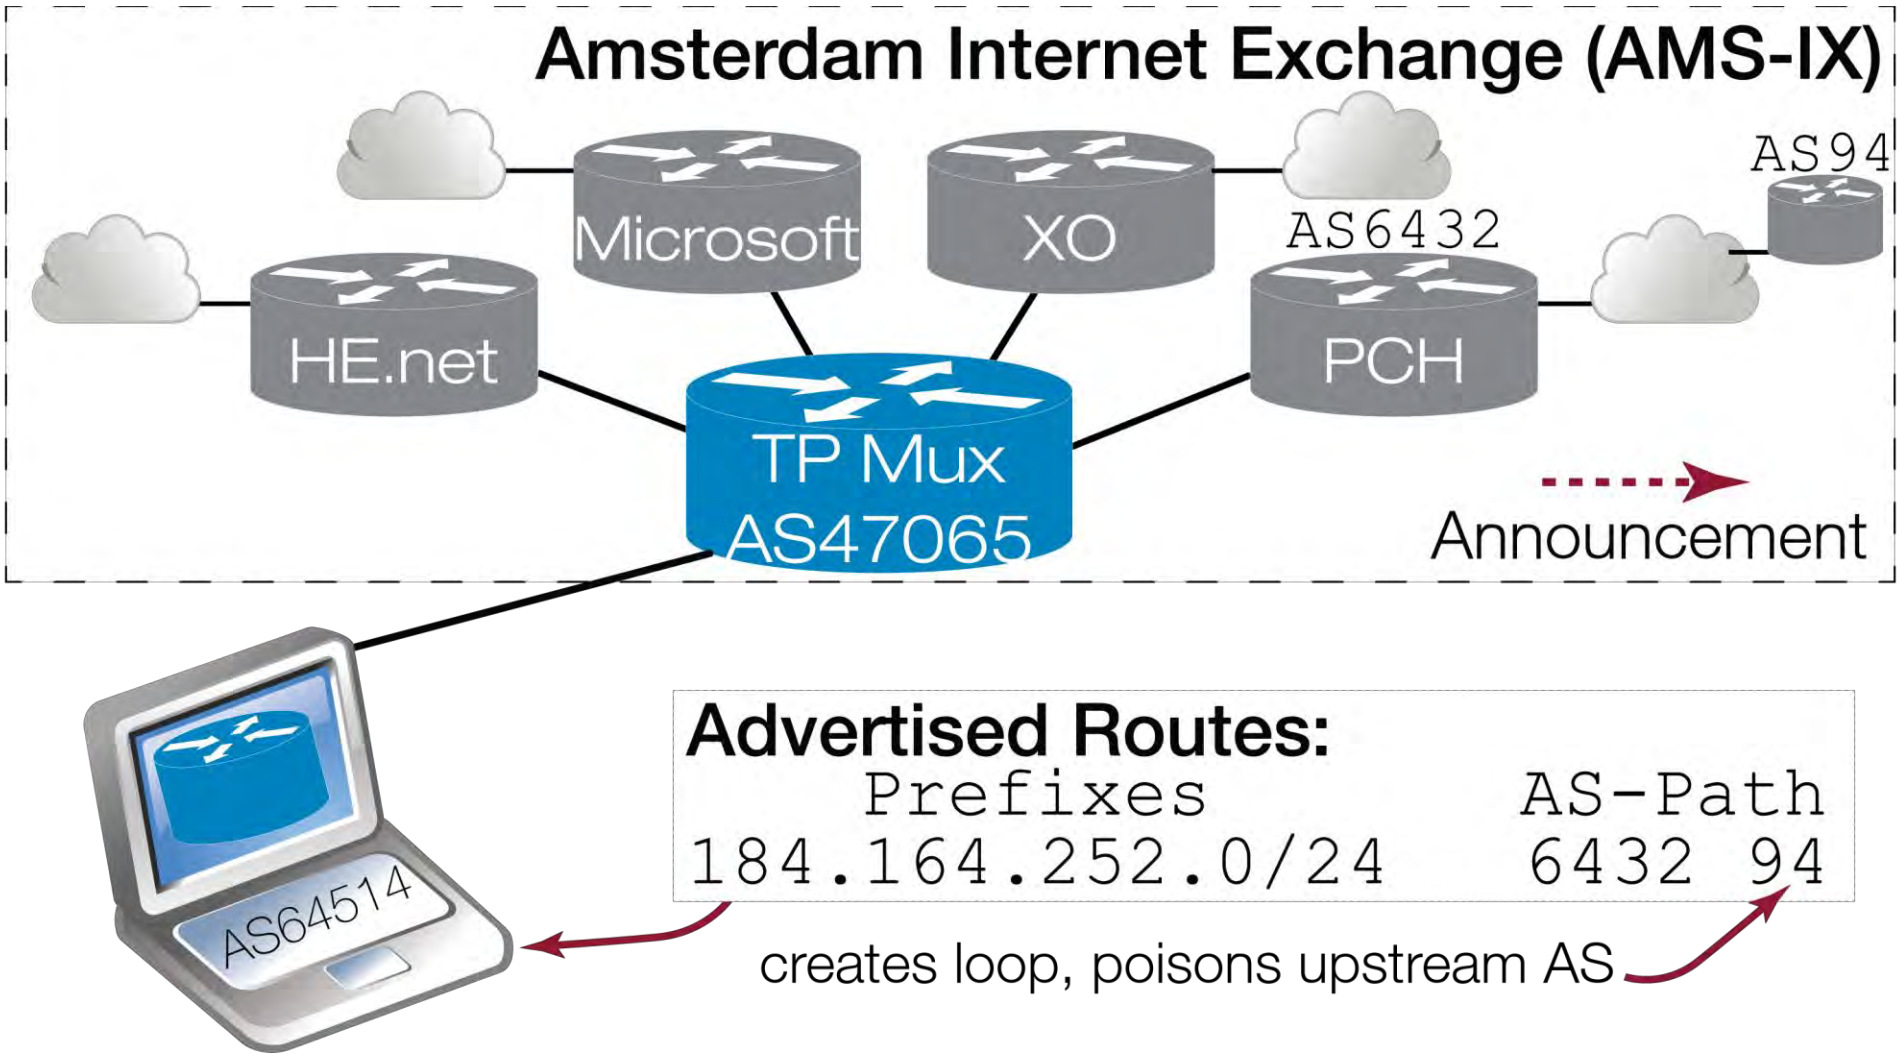

Transit Portal

A Closer Look at the AMS-IX Multiplexer

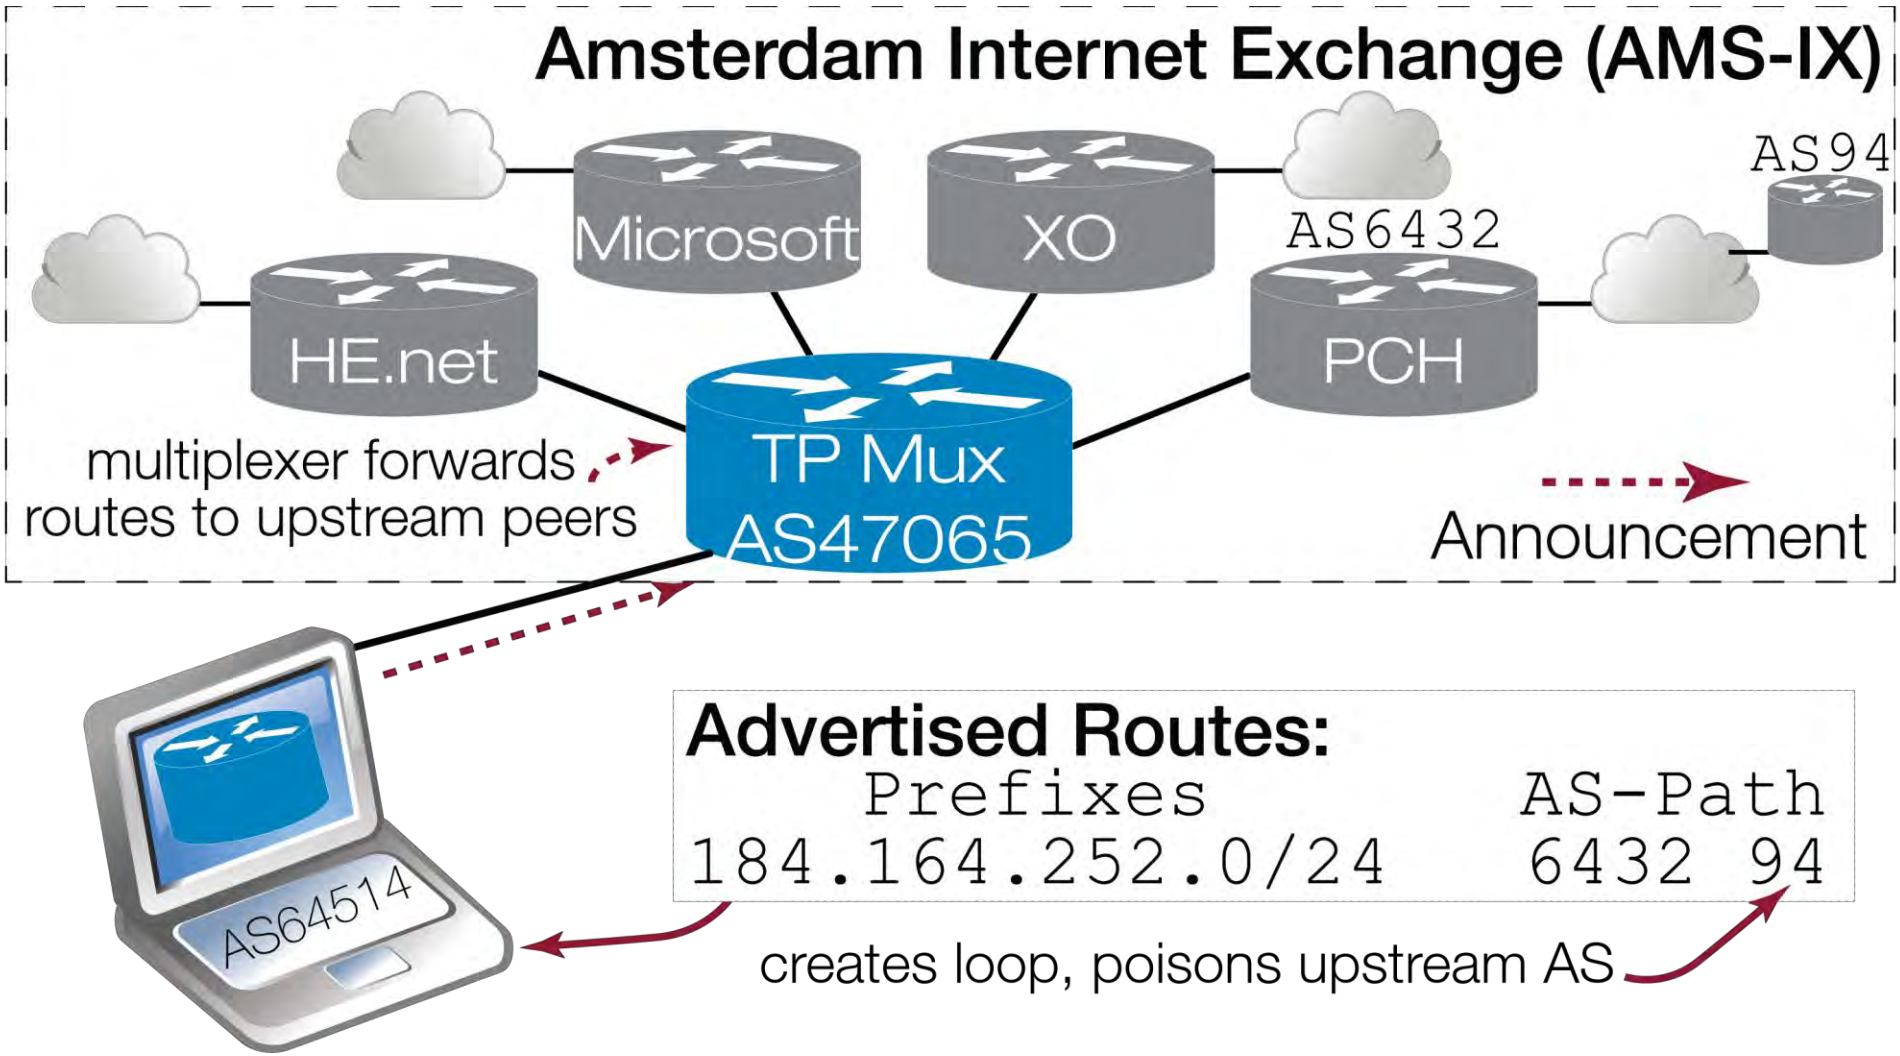

Transit Portal

A Closer Look at the AMS-IX Multiplexer

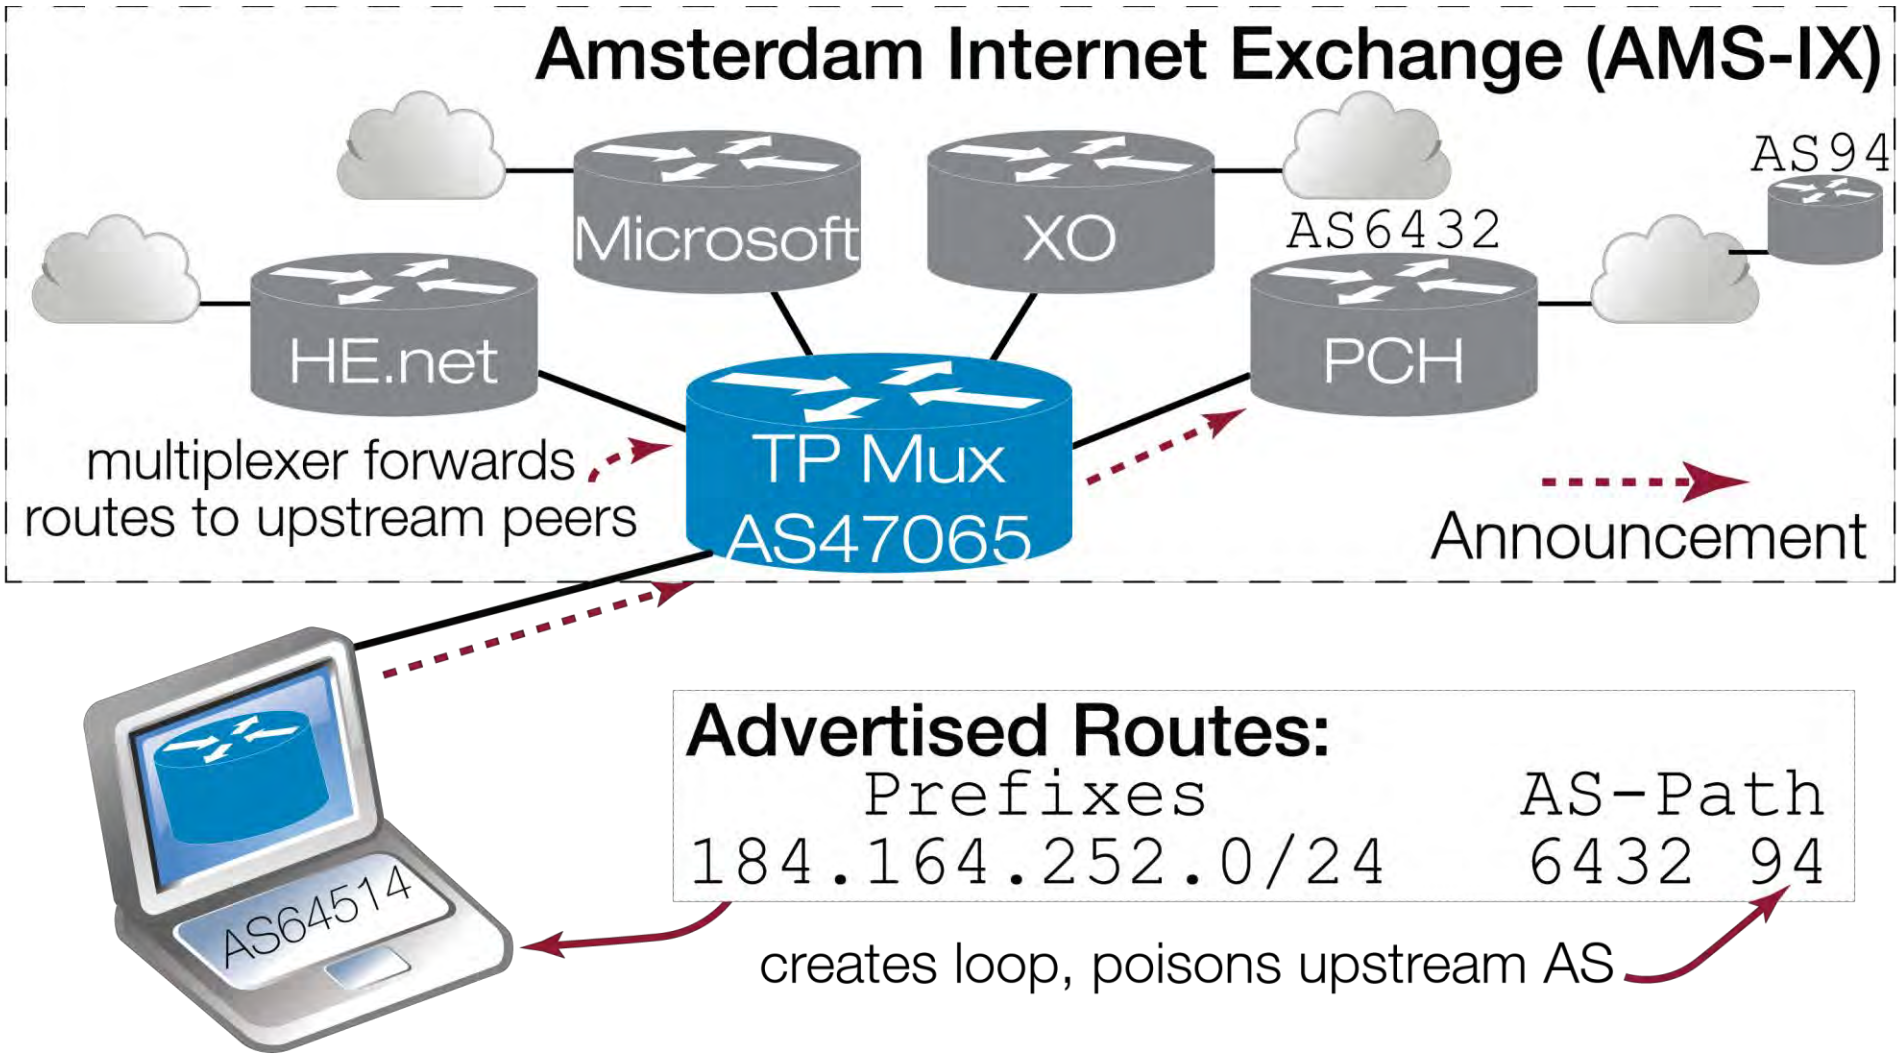

Transit Portal

A Closer Look at the AMS-IX Multiplexer

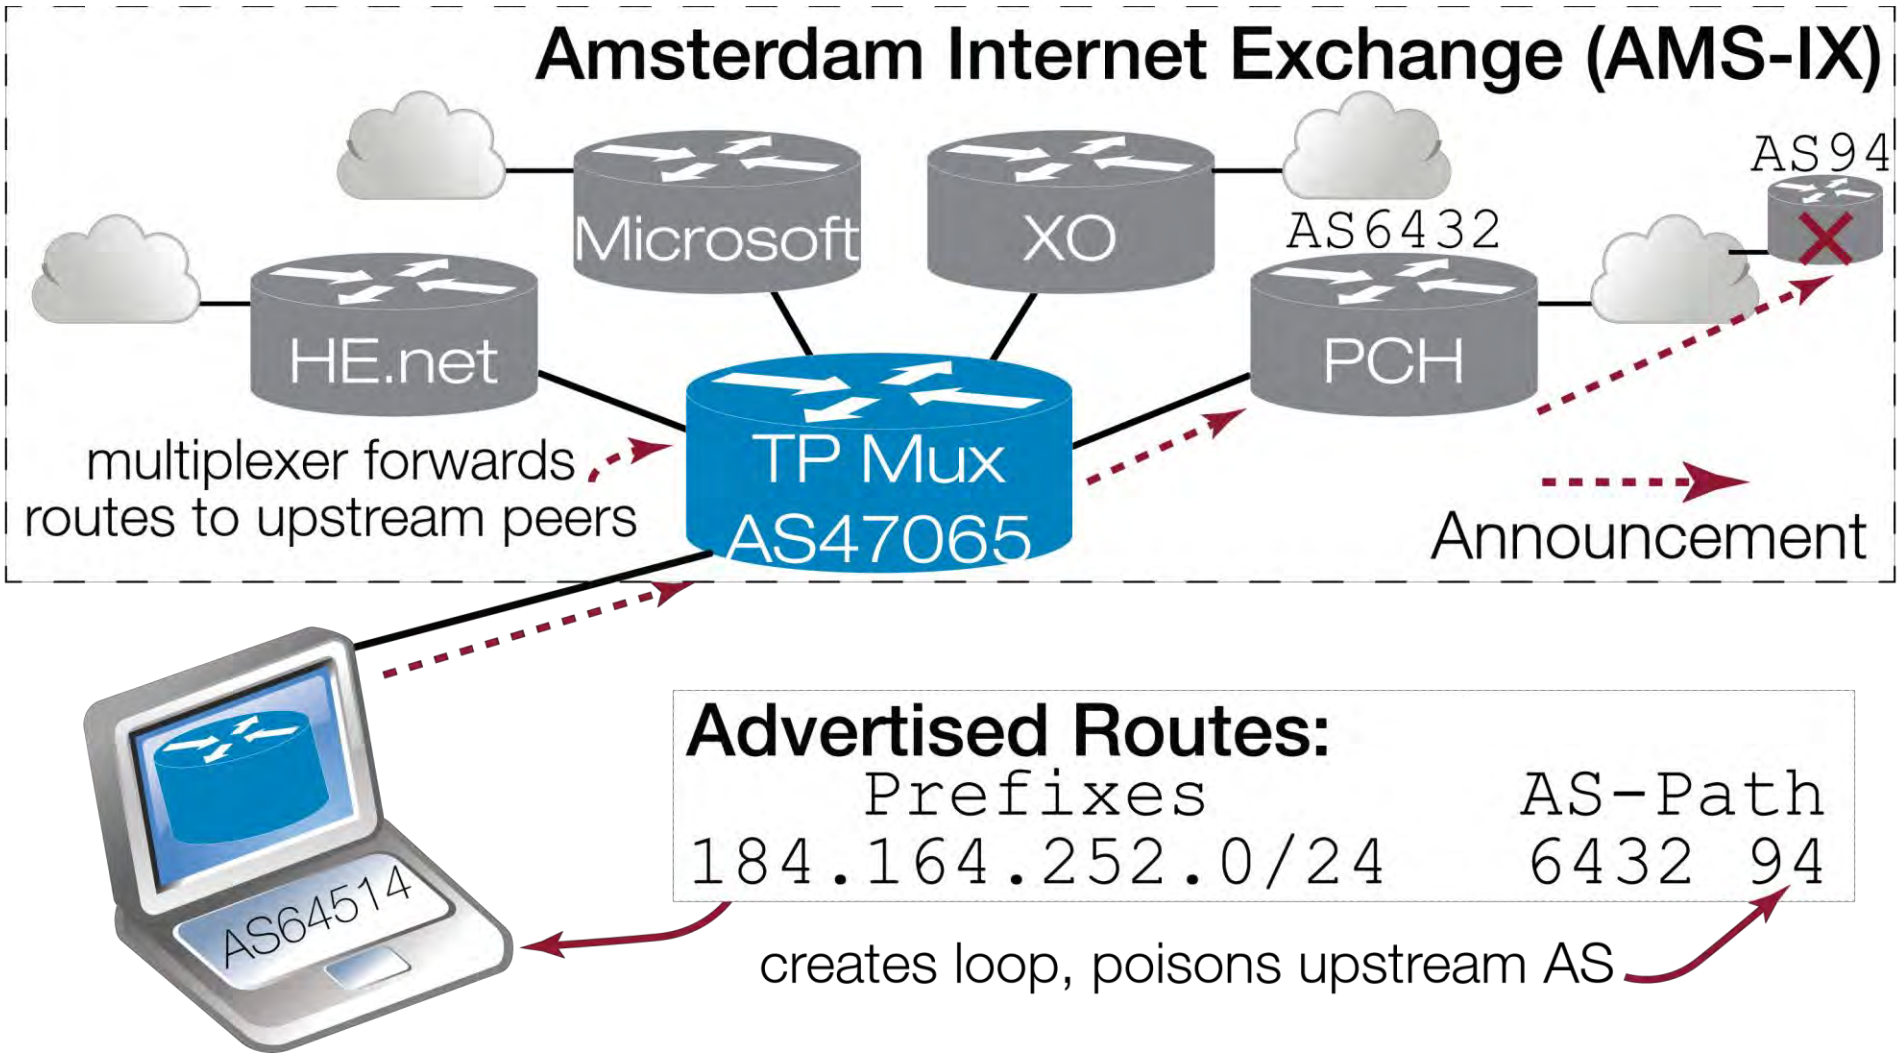

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

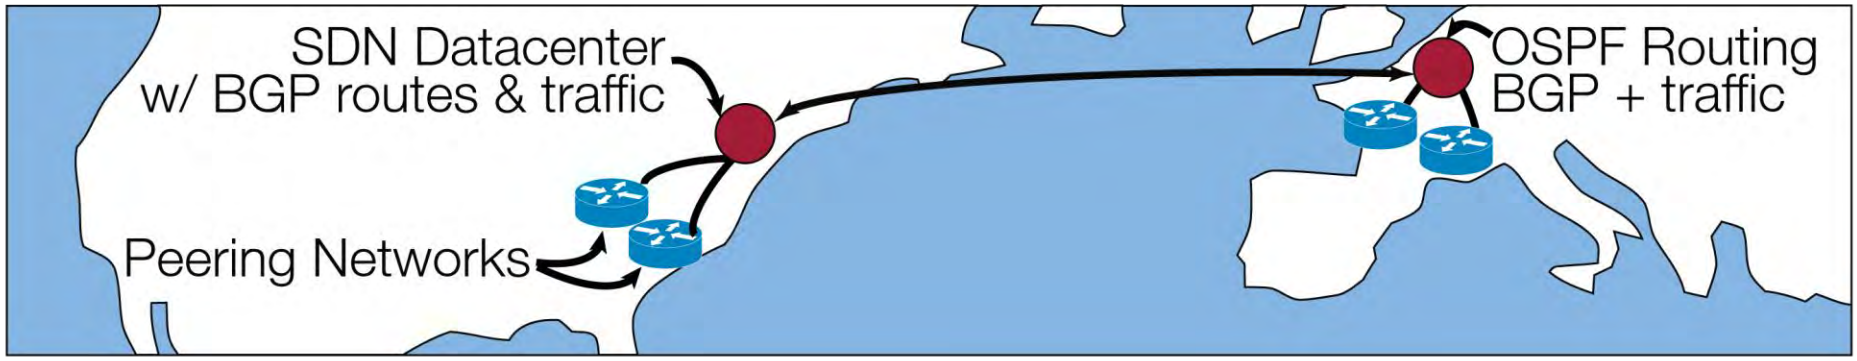

Wish List for Anycast Experiment:

1. Intercontinental private wide-area network (WAN) between US and EU
2. Different routing domains in each datacenter (SDN & OSPF)
3. Ability to exchange *routes and traffic* via BGP with other peers
4. Rich peering connectivity with *REAL* providers in US and EU

Let's build using Transit Portal + Enhanced Mininet

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

① Mininet container initialized with no topology

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

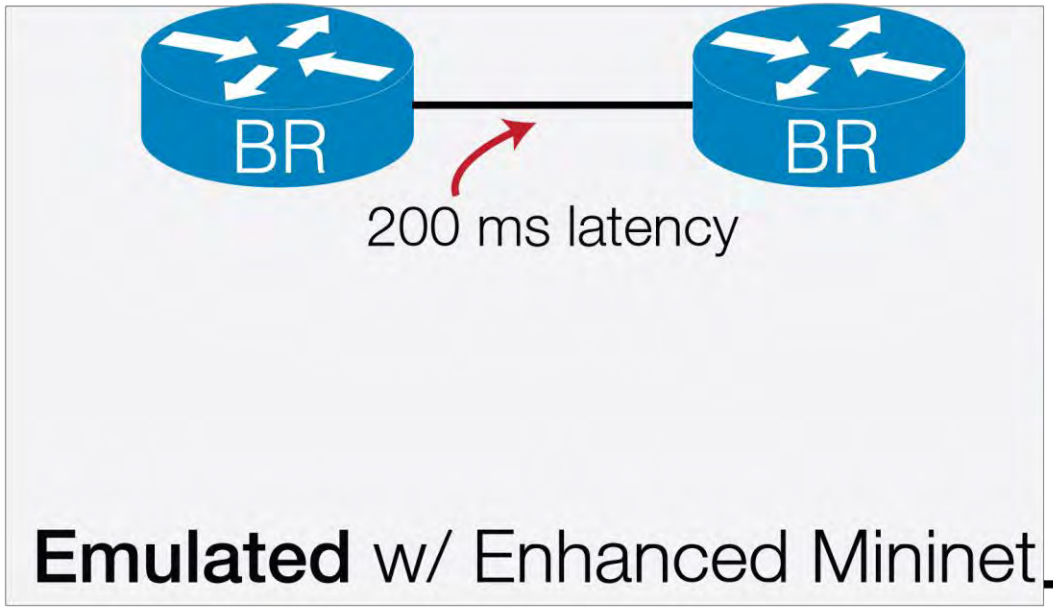

Executed on **single laptop**

② Border routers & virtual WAN link with latency

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

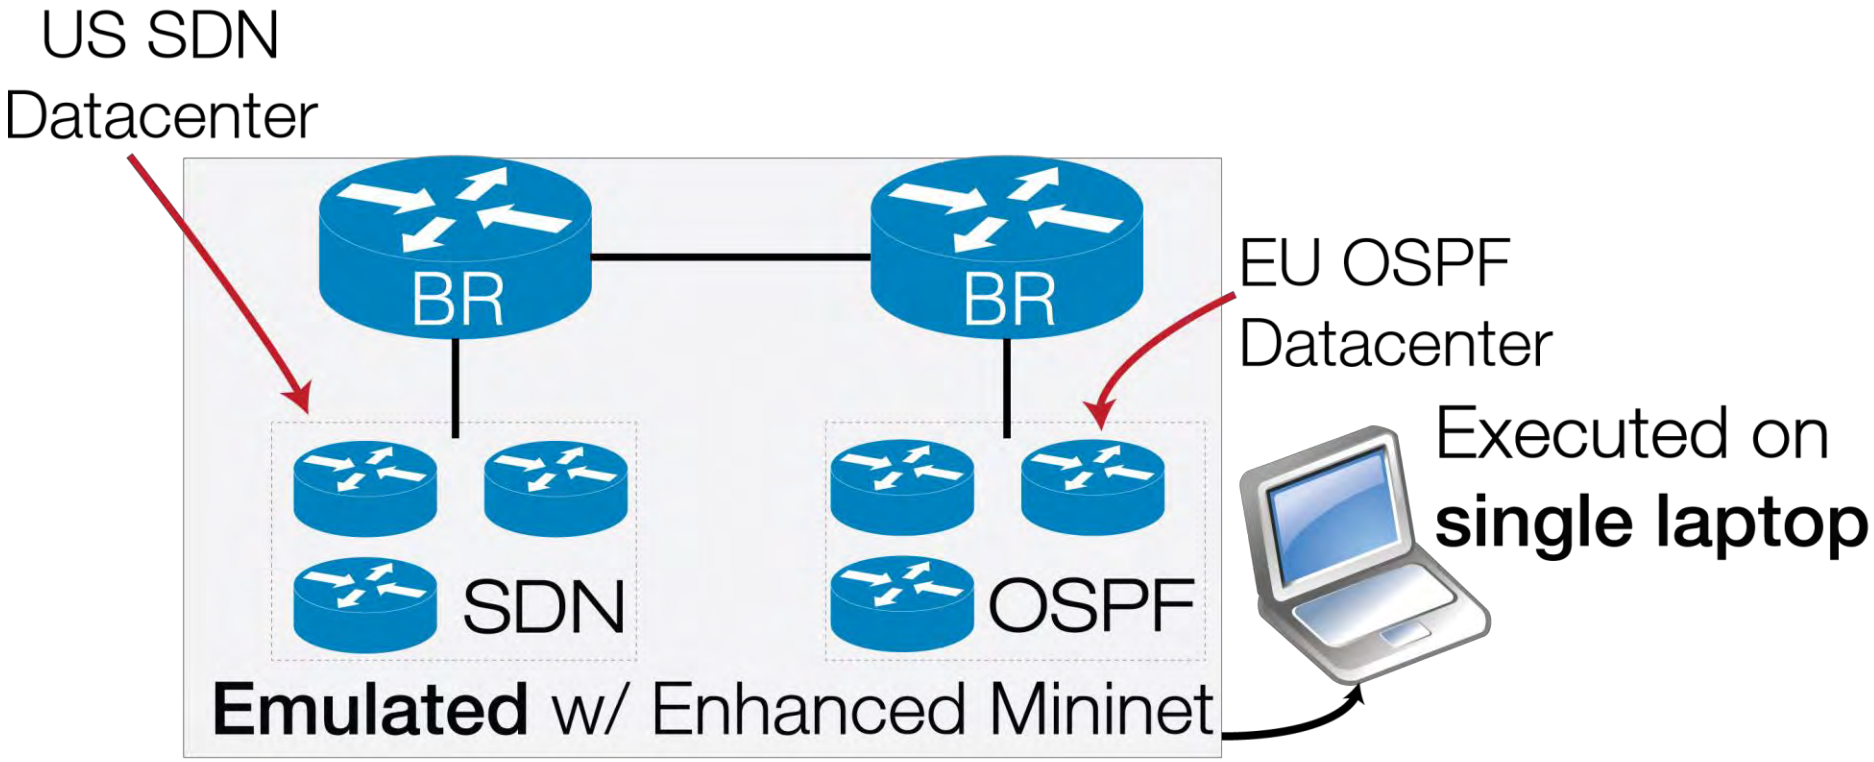

③ Internal routing domains created with Quagga

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

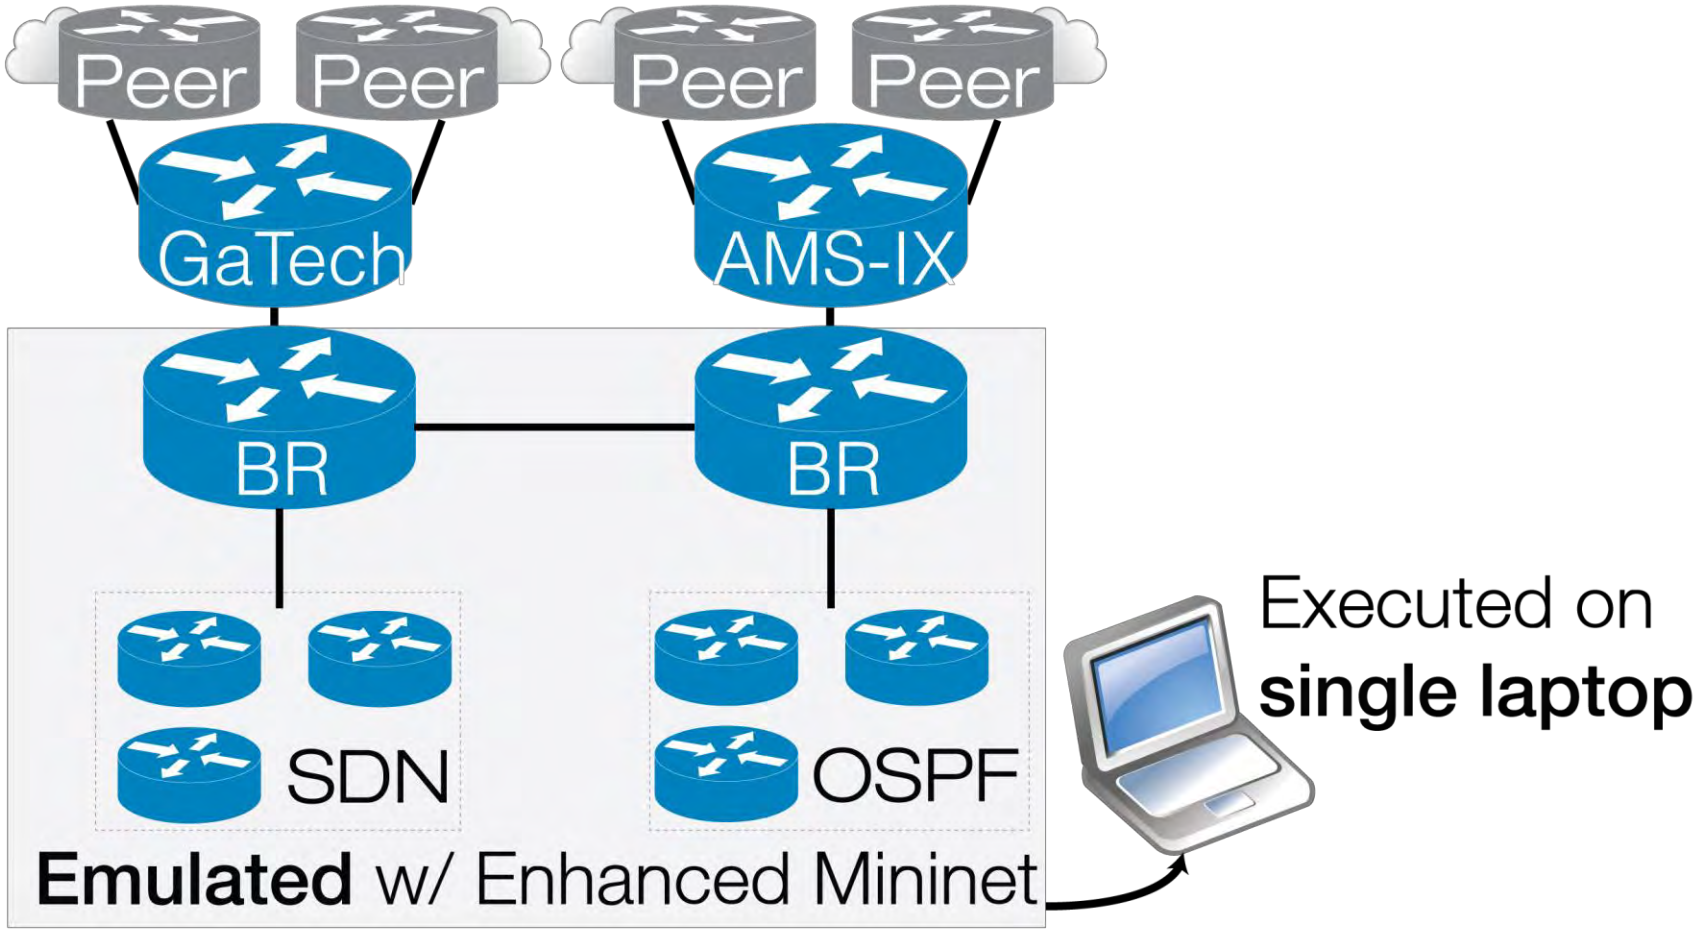

④ Transit Portal Mux connectivity established

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

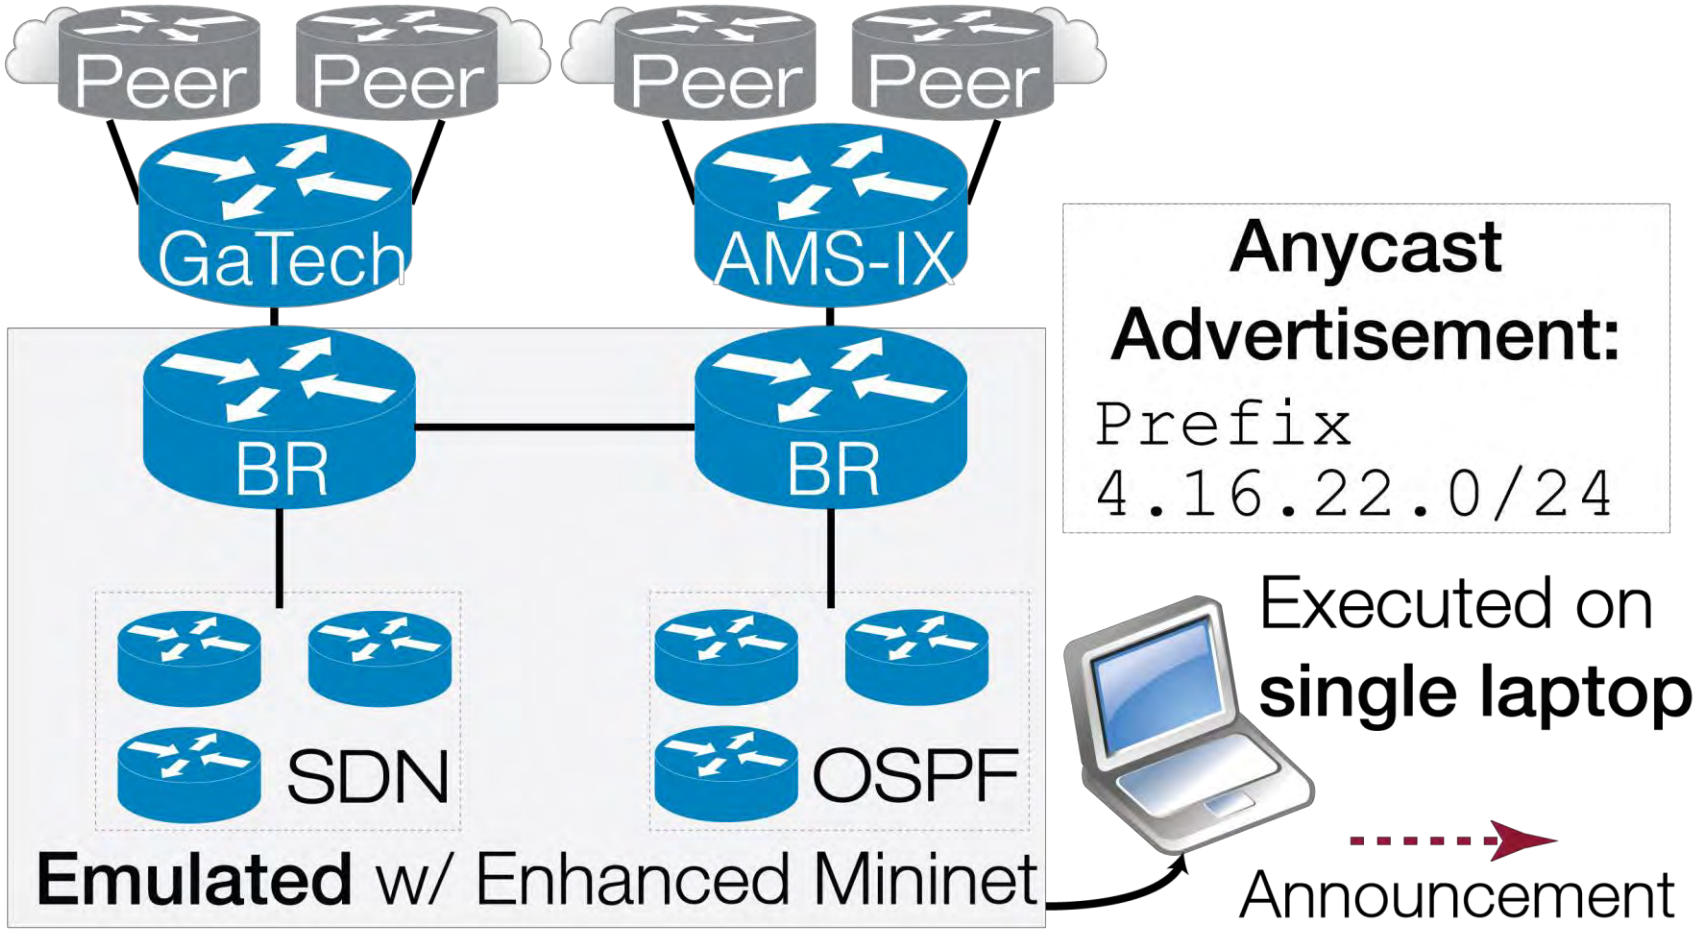

⑤ Anycast announcements propogate

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

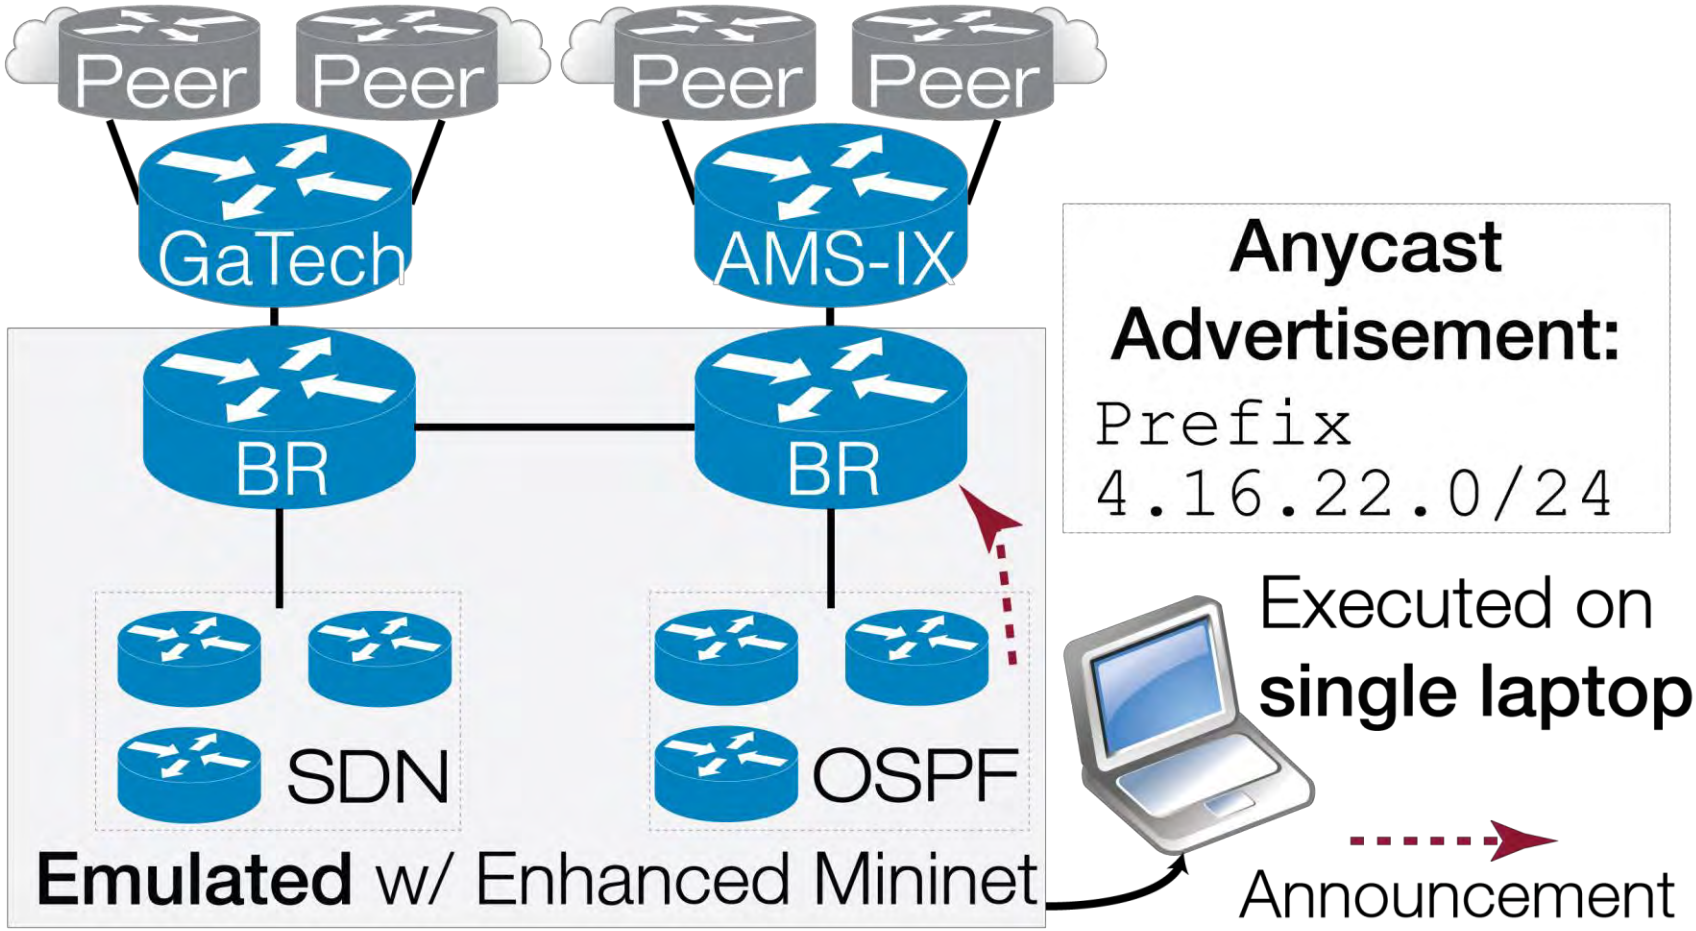

⑥ Anycast route originates from EU OSPF domain

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

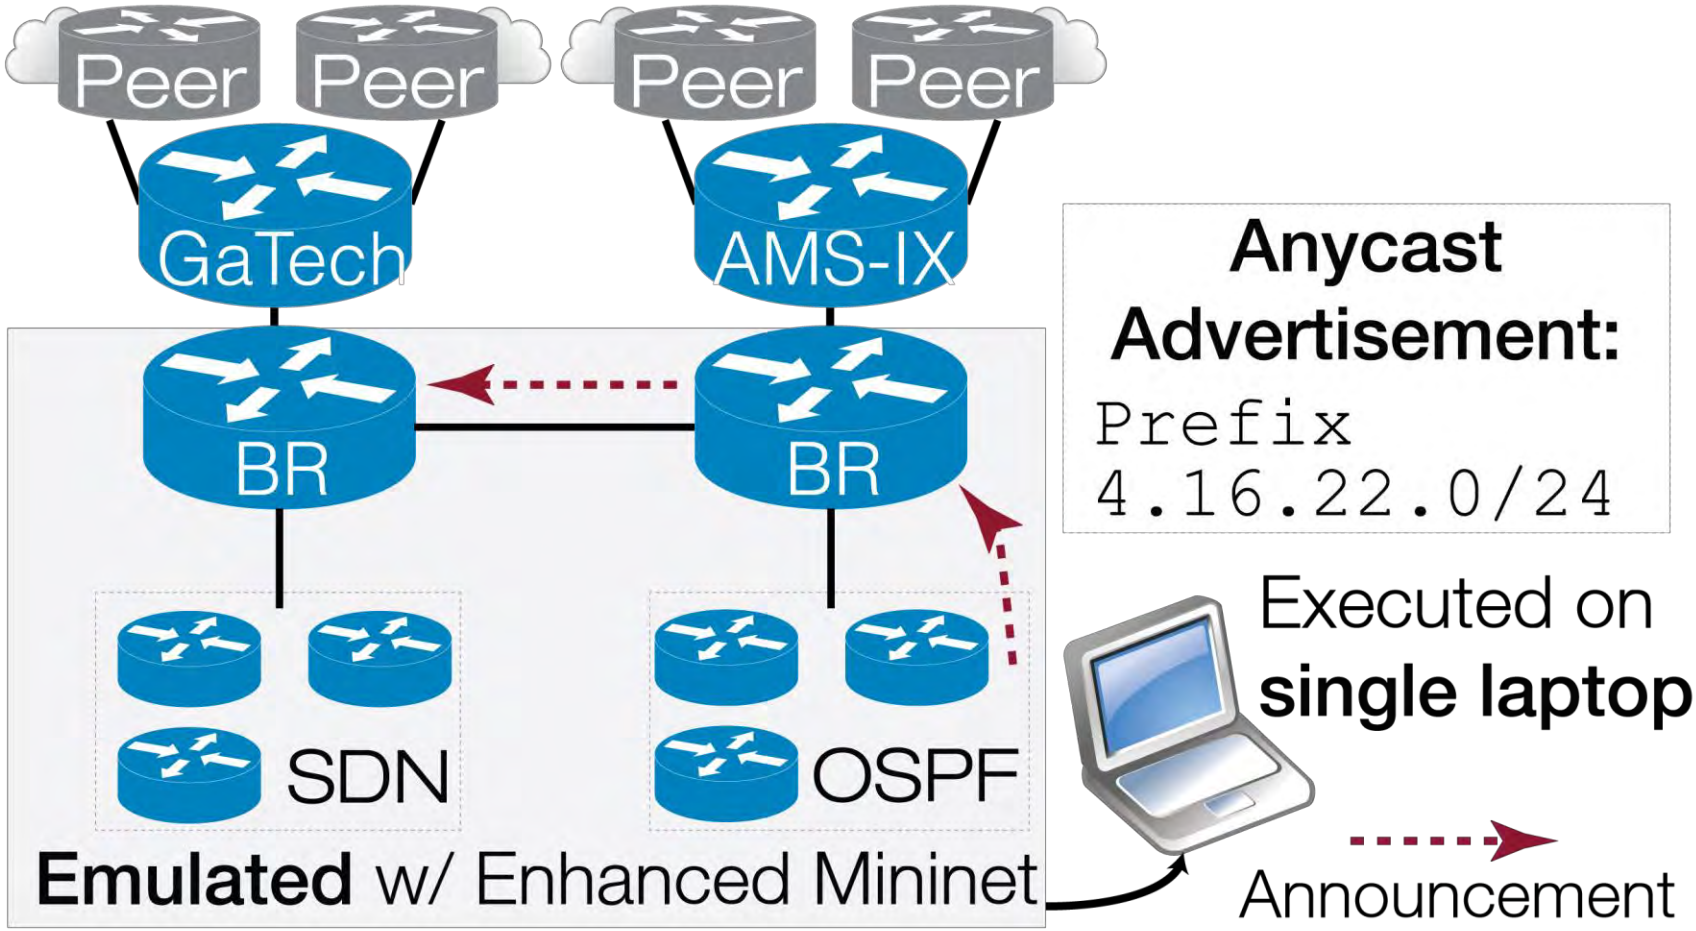

⑦ Propogates to US border router

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

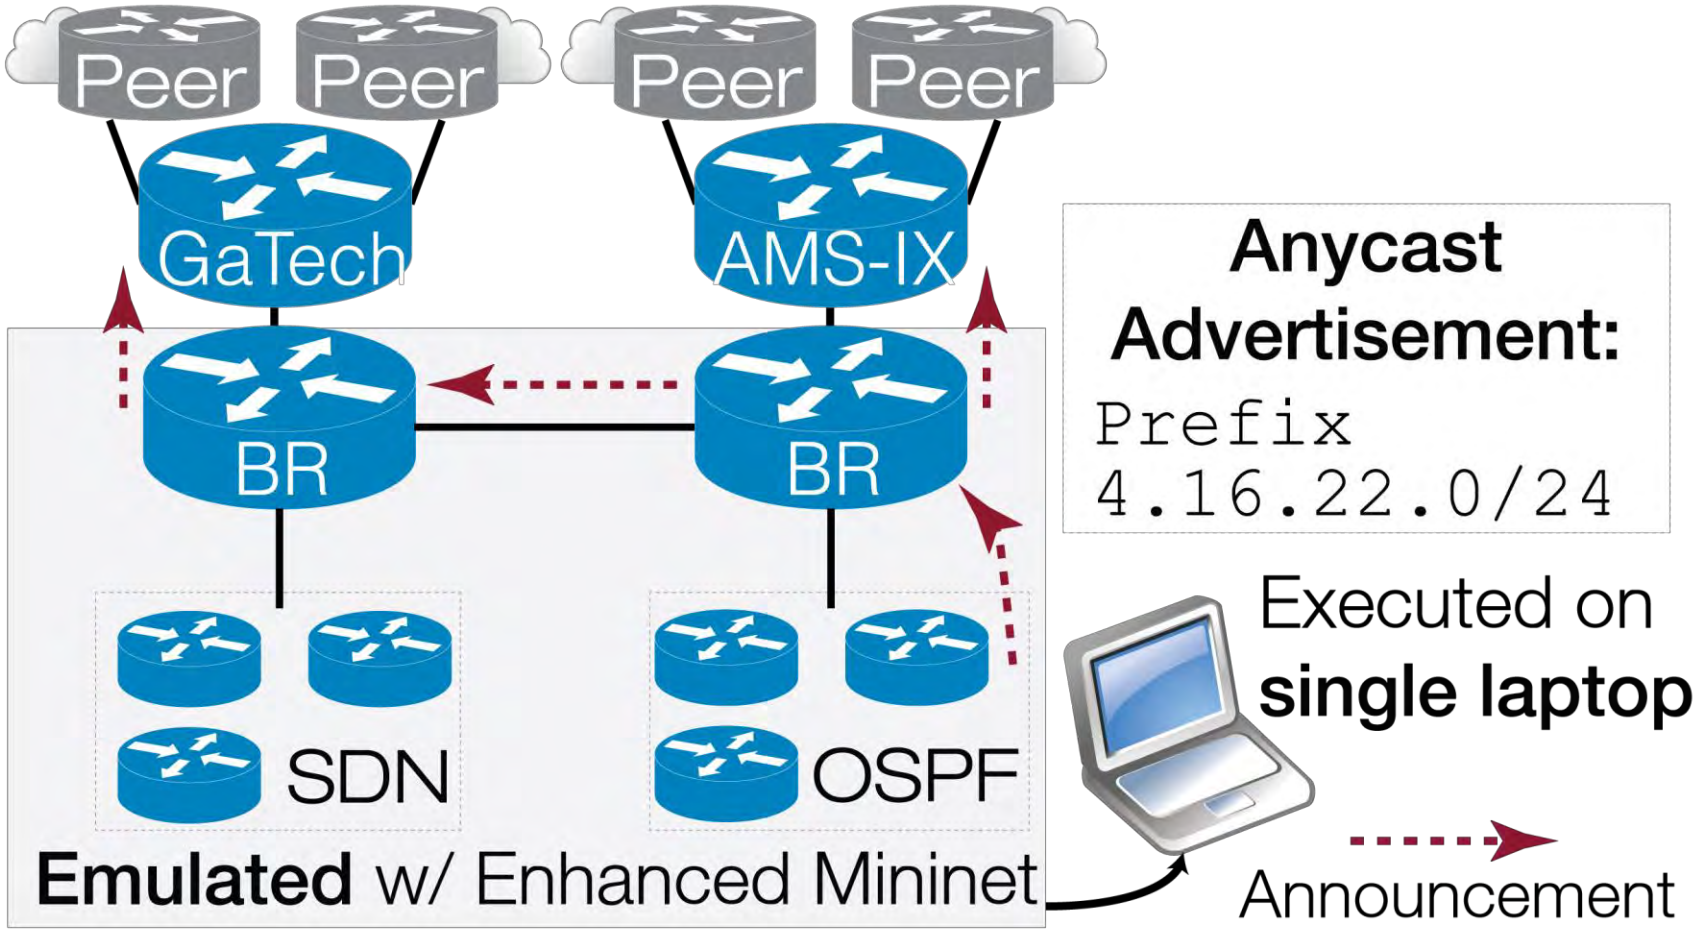

⑧ Propogates through Transit Portal Muxes

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

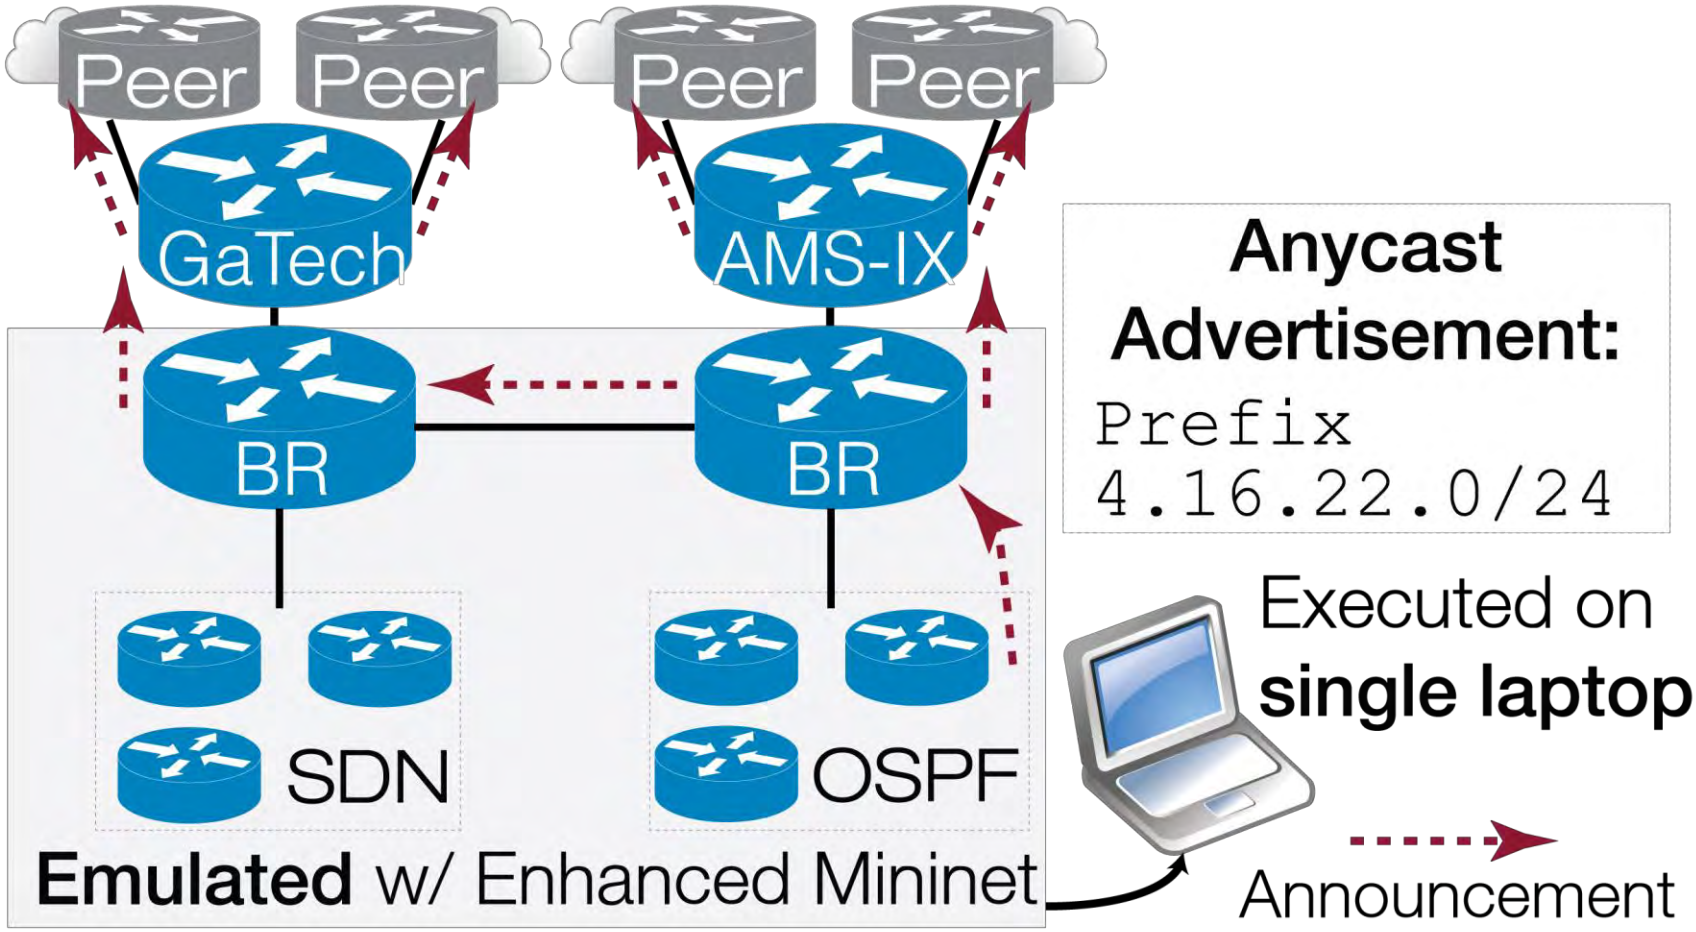

⑨ Propogates to peers at Georiga Tech and AMS-IX

Anycast Experiment

Building with Transit Portal and Enhanced Mininet

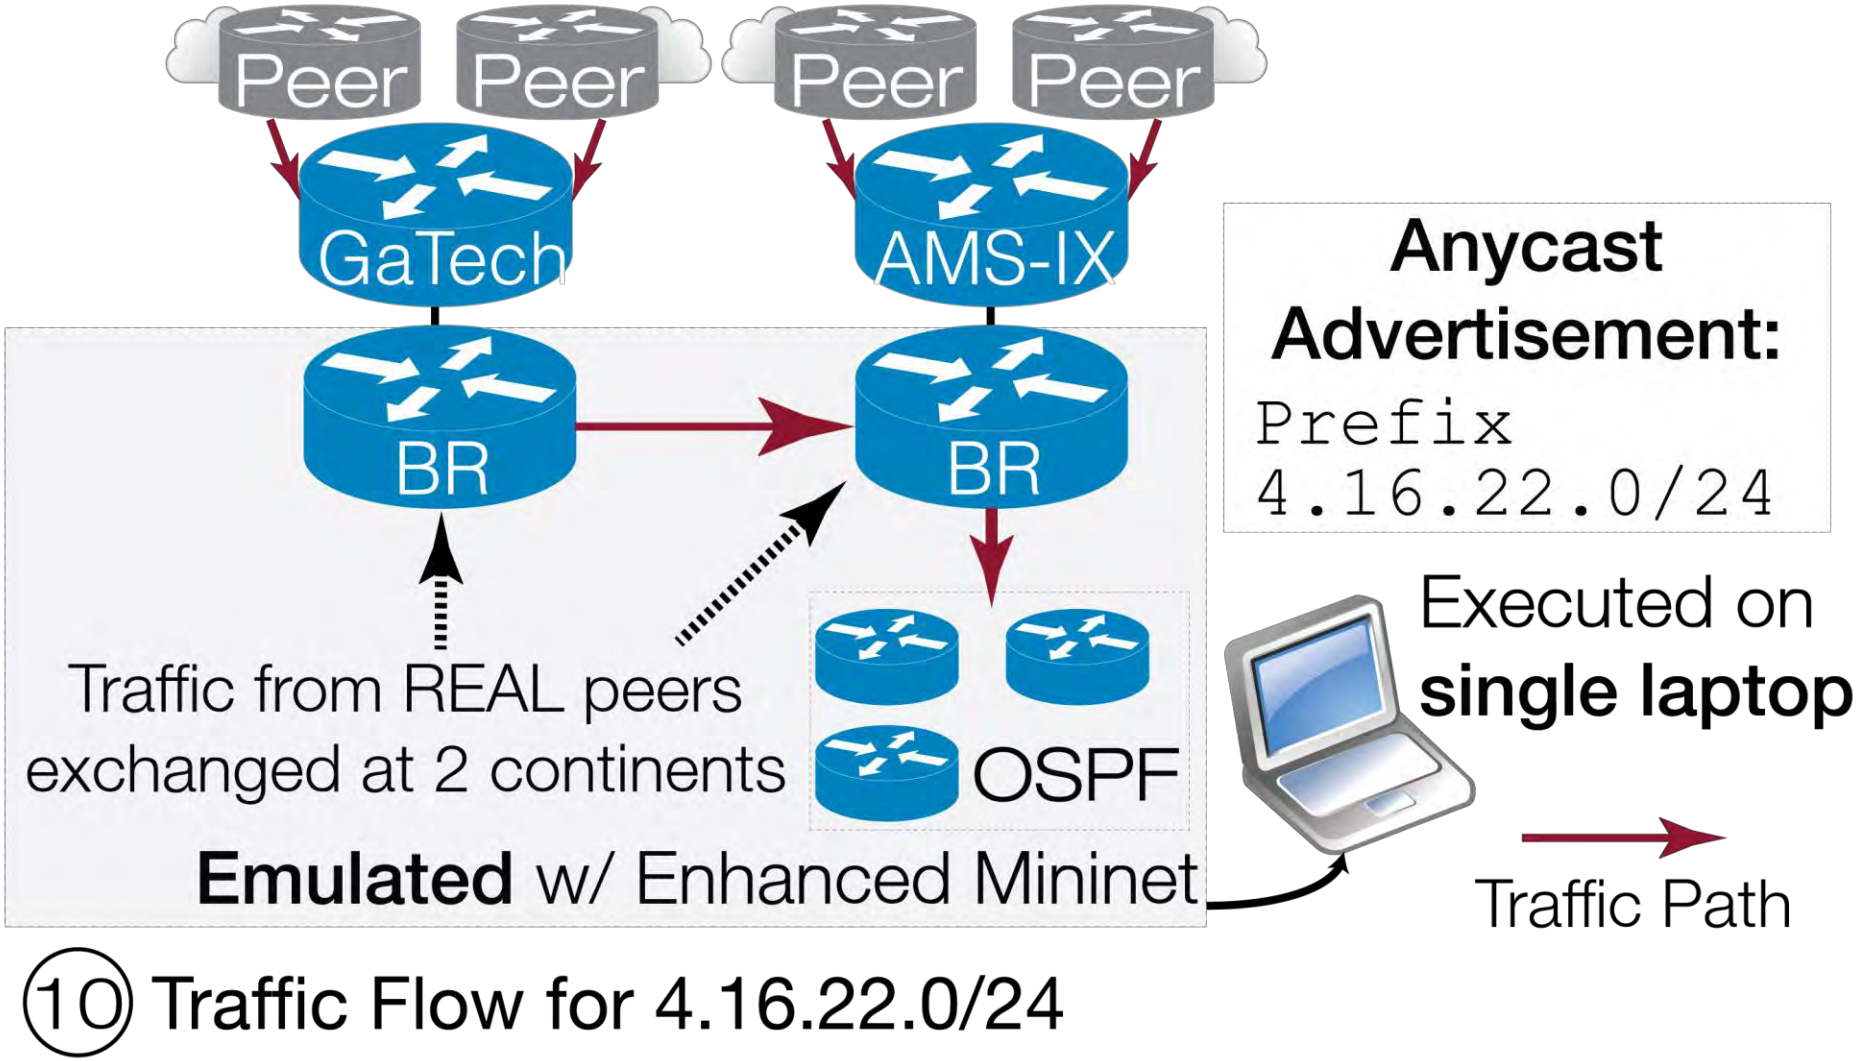

10 Traffic Flow for 4.16.22.0/24

Transit Portal and Enhanced Mininet

Capabilities


```
> sudo mn
```

+

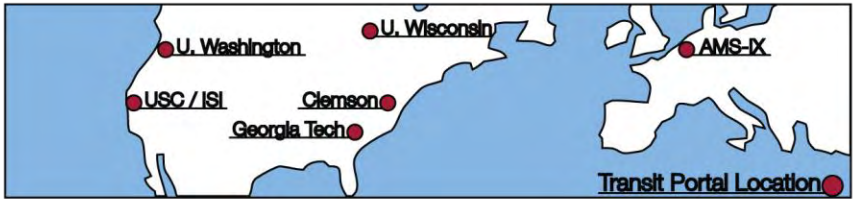

① Users can emulate complex ISP / AS topologies

- ⓐ Comprehensive emulation of intra & inter-domain
- ⓑ Isolated instances of routing engines (Quagga, BIRD)

② BGP route and traffic exchange via Transit Portal

- ⓐ Rich peering & BGP connectivity at multiple locations
- ⓑ Exposure to real (unknown) internet policies and failures

③ Entire ISP can be emulated from a laptop

- ⓐ Significantly reduced overhead, no VMs, easy config

Conclusion

Enabling Community's Future Experiments

> sudo mn + 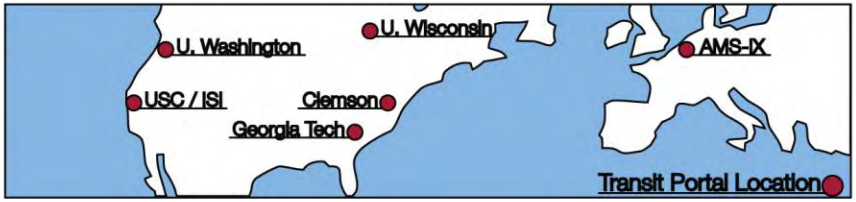

Intended to be a *Community Testbed*

- Wide-range of applicable experiments
- Enhancements to Mininet available on GitHub
- Access to Transit Portal available to researchers / operators
- **Contact:** bschlink@usc.edu | http://nsl.cs.usc.edu

Looking to expand Transit Portal further

- More IXP locations (AMS-IX first Transit Portal IXP)
- Collaboration opportunities with other univerisities
- Vantage points within Tier-1 ISPs