USC Viterbi
School of Engineering
Department of Computer Science

Impactful Routing Research with the PEERING Testbed

Combining intradomain emulation with *real* BGP connectivity

Ethan Katz-Bassett (University of Southern California)

with:
Brandon Schlinker and Kyriakos Zarifis (USC students)
Italo Cunha (UFMG Brazil)
Nick Feamster (Georgia Tech)

Supported By: Google NSF amsix amsterdam internet exchange

USC Viterbi
School of Engineering
Department of Computer Science

Impactful Routing Research with the PEERING Testbed

Combining intradomain emulation with *real* BGP connectivity

Ethan Katz-Bassett (University of Southern California)

with:
Brandon Schlinker and Kyriakos Zarifis (USC students)
Italo Cunha (UFMG Brazil)
Nick Feamster (Georgia Tech)

Supported By: Google NSF amsix amsterdam internet exchange

BGP limits today's networks

CDN Overlay Network

SDN Data Center

Advanced Cellular Network

BGP limits today's networks

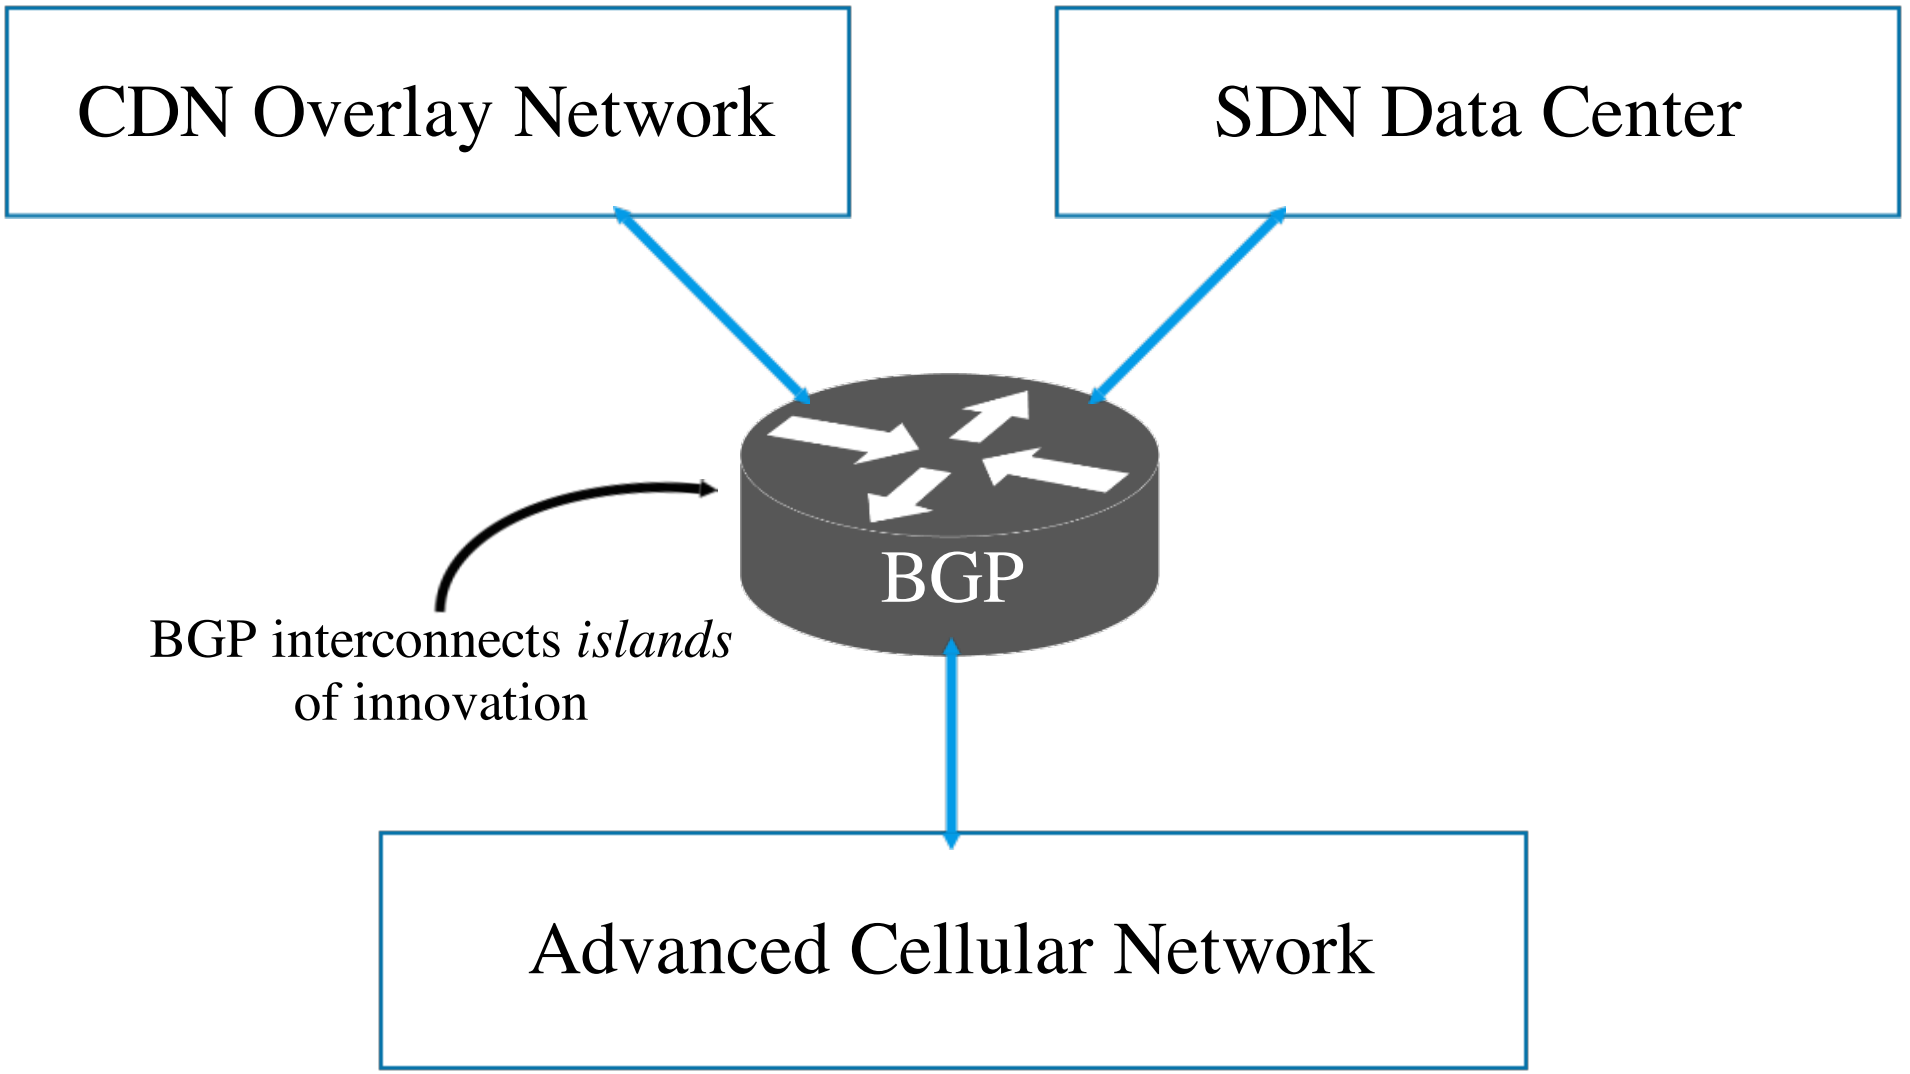

What's so bad about BGP?

BGP contributes to many of the Internet's
fundamental problems

Examples of problems created by BGP

BGP contributes to many of the Internet's **fundamental problems**

BGP design results in:

- Poor performance *(inflated routes)*
- Security vulnerabilities *(route hijacking)*
- Longer outages (*lengthy convergence times*)
- Routing failures (*route redistribution issues*)
- QoS problems in gaming, VoIP (*path oscillations*)

(the list goes on…)

How do we improve BGP?

Remainder of Talk:

- Why is ***impactful*** BGP research and innovation hard?

How do we improve BGP?

Remainder of Talk:

- Why is ***impactful*** BGP research and innovation hard?

- Our PEERING testbed enables ***impactful*** BGP research
  Provides ***control*** and ***realism*** needed to tackle BGP problems

Limited existing tools for BGP research

Simulation

- Provides complete control and visibility
- But only as accurate as the inputs, and we do not know how to accurately model Internet peering topology or policies

Limited existing tools for BGP research

Simulation

- Provides complete control and visibility
- But only as accurate as the inputs, and we do not know how to accurately model Internet peering topology or policies

Measurements of real routes

- Traceroutes, route collectors, BGP beacons, looking glasses
- Realistic, but passive observation only

How does PEERING help?

Remainder of Talk:

- How PEERING testbed enables ***impactful*** BGP research
  Provides ***control*** and ***realism*** needed to tackle BGP problems

PEERING Testbed Enables BGP Research

PEERING:

Pairing **E**mulated **E**xperiments with **R**eal **I**nterdomain **N**etwork **G**ateways

With PEERING, a researcher or network operator:

- ***Emulates*** an AS, including its topology and routing policies
- ***Connects*** the designed AS to *real* ASes on the Internet via BGP
- ***Controls*** the AS, including its exchange of traffic and routes

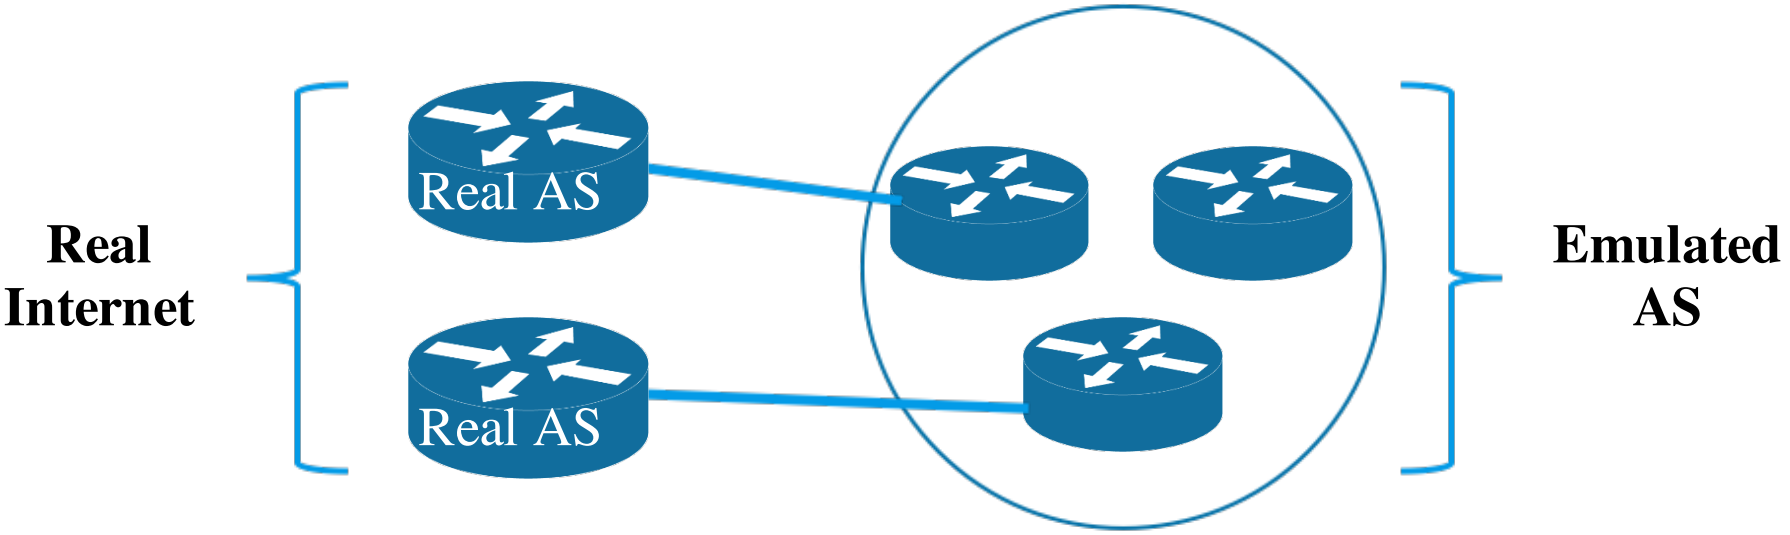

PEERING at AMS-IX (who kindly host us)

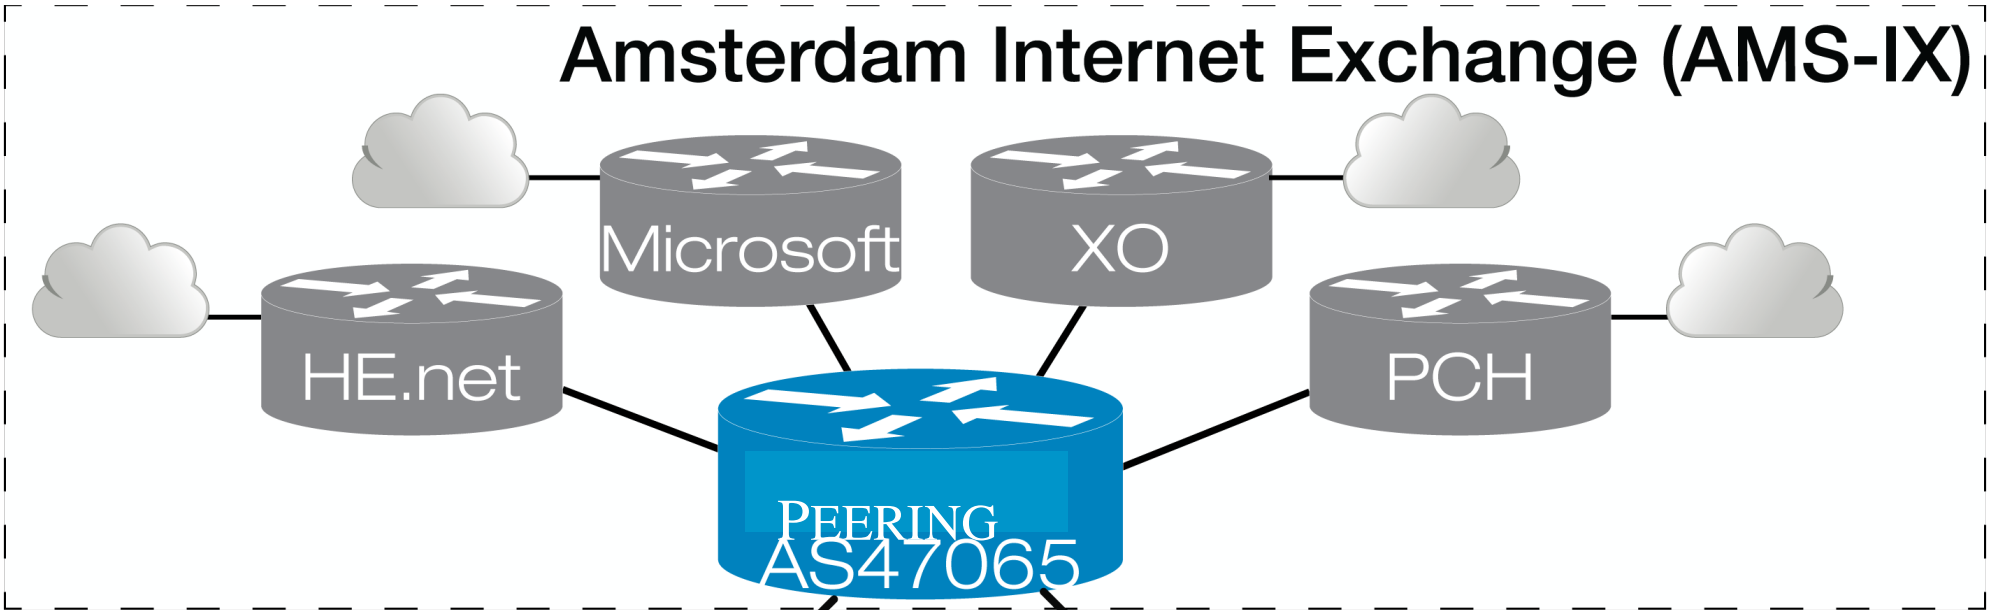

PEERING at AMS-IX (who kindly host us)

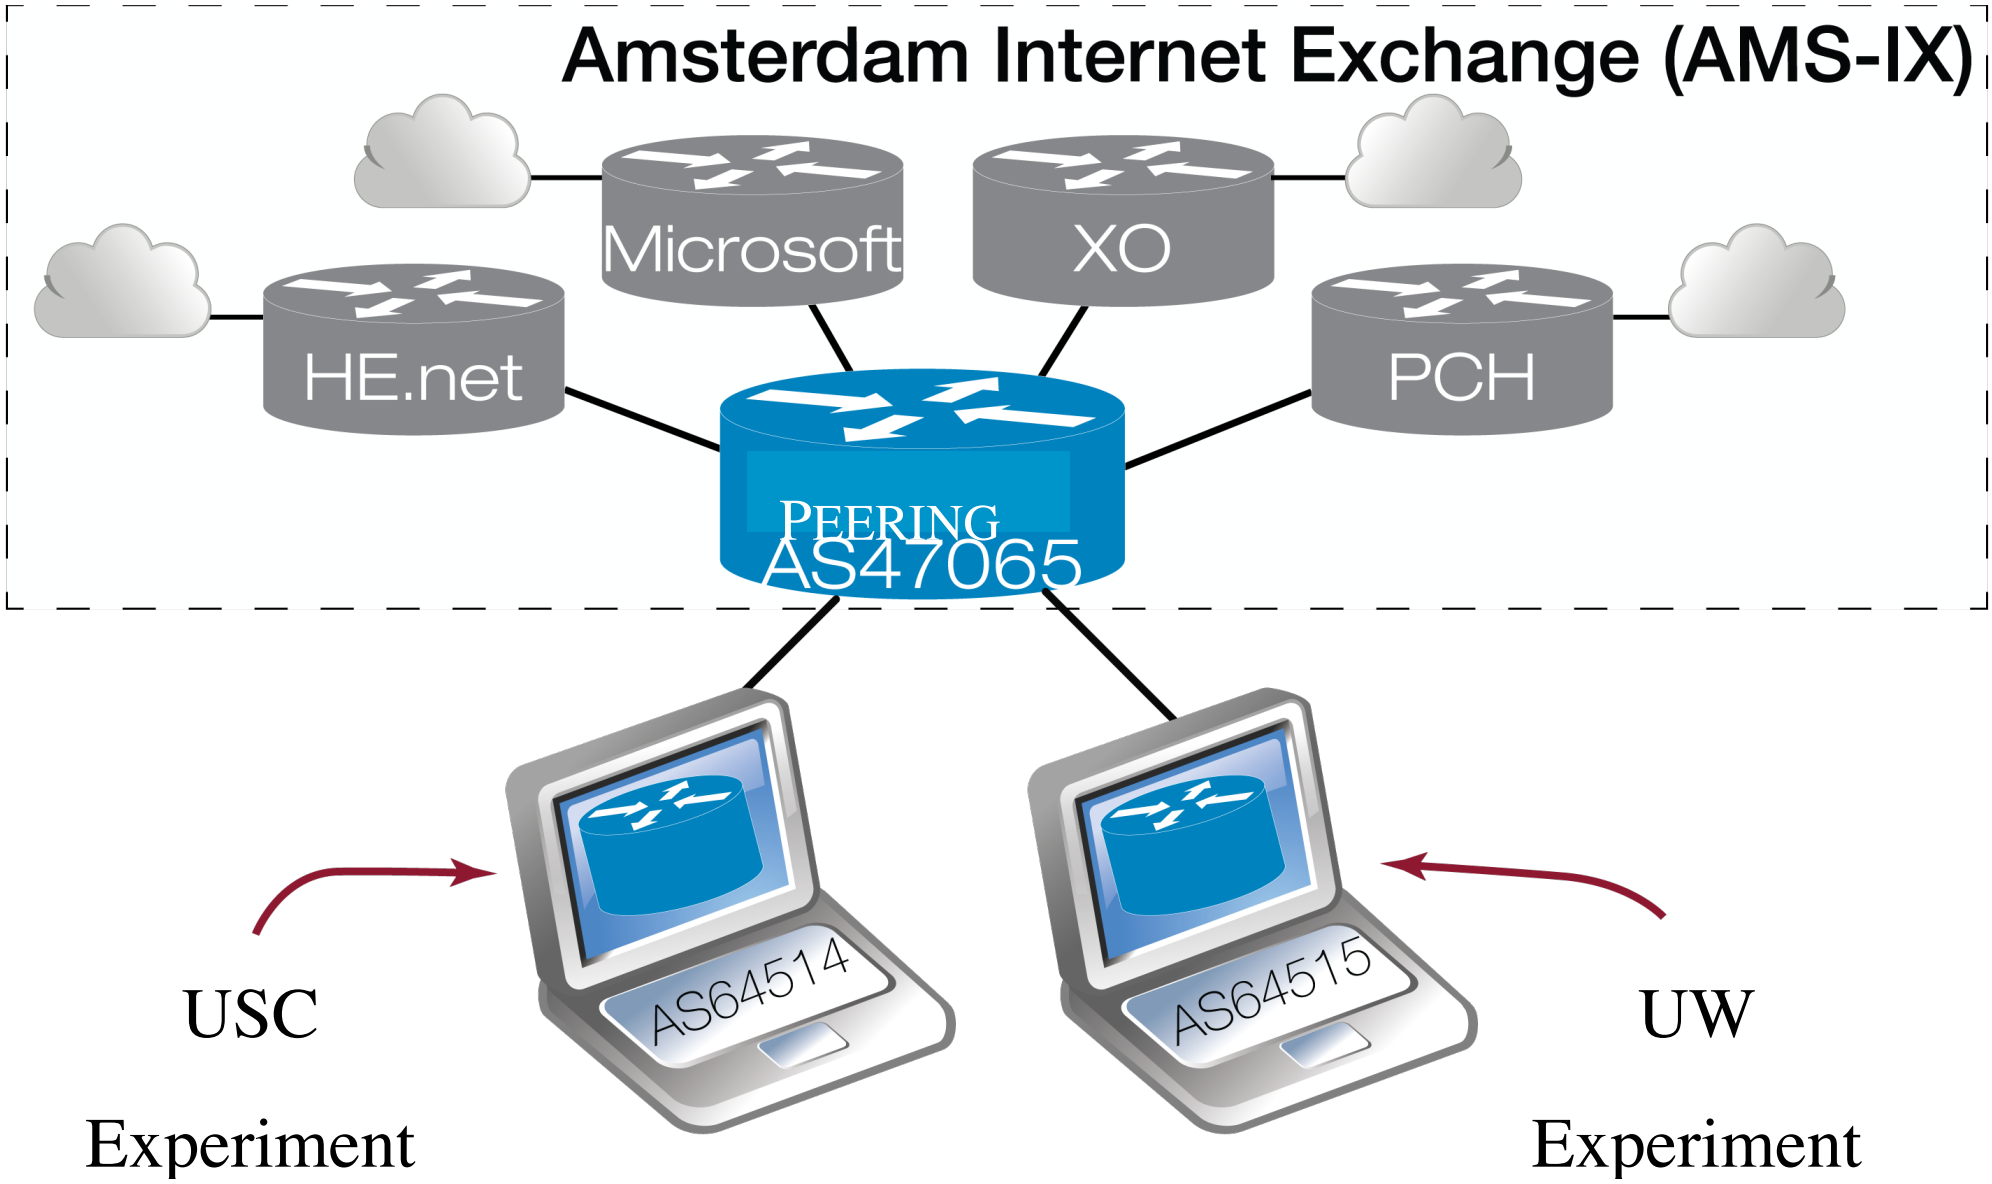

PEERING at AMS-IX (who kindly host us)

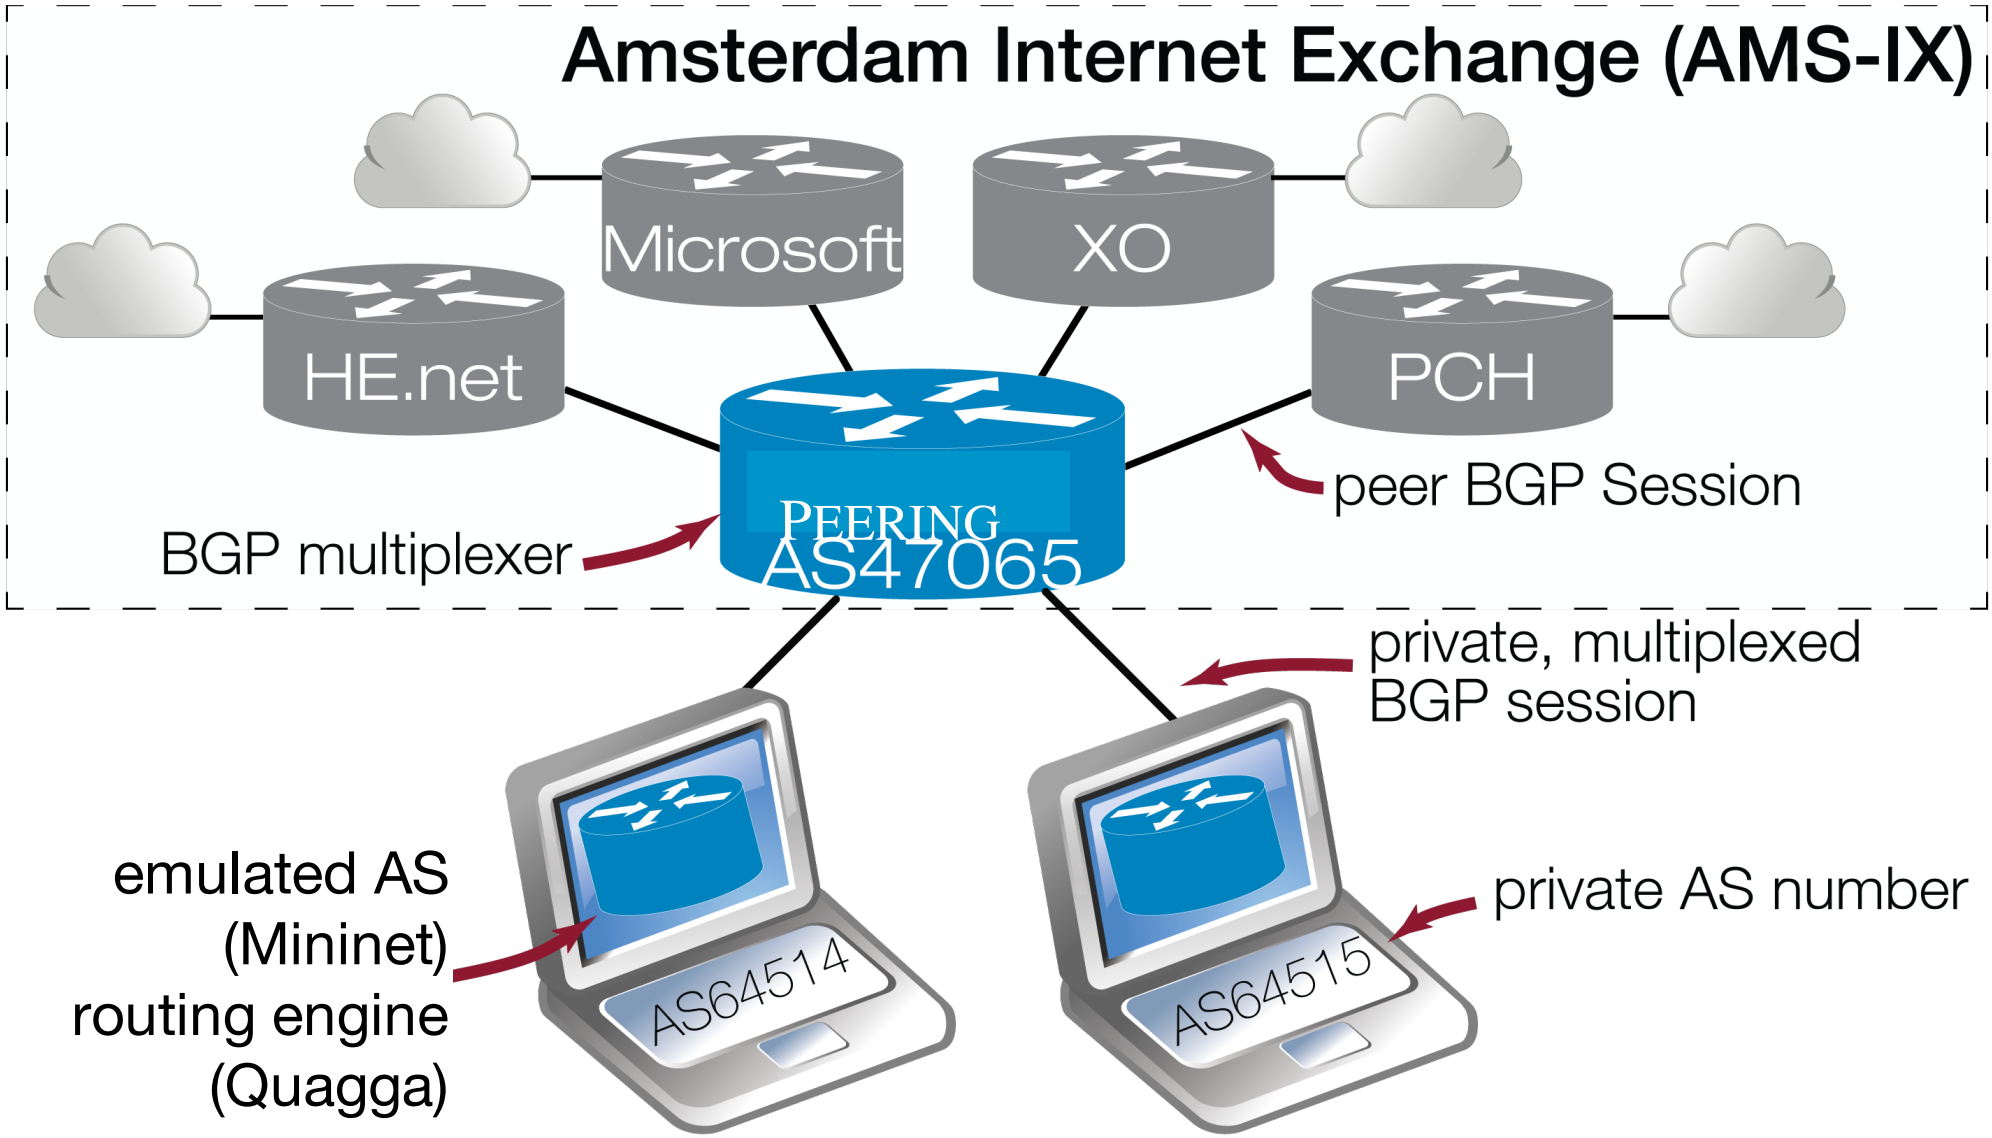

Making Announcements via AMS-IX

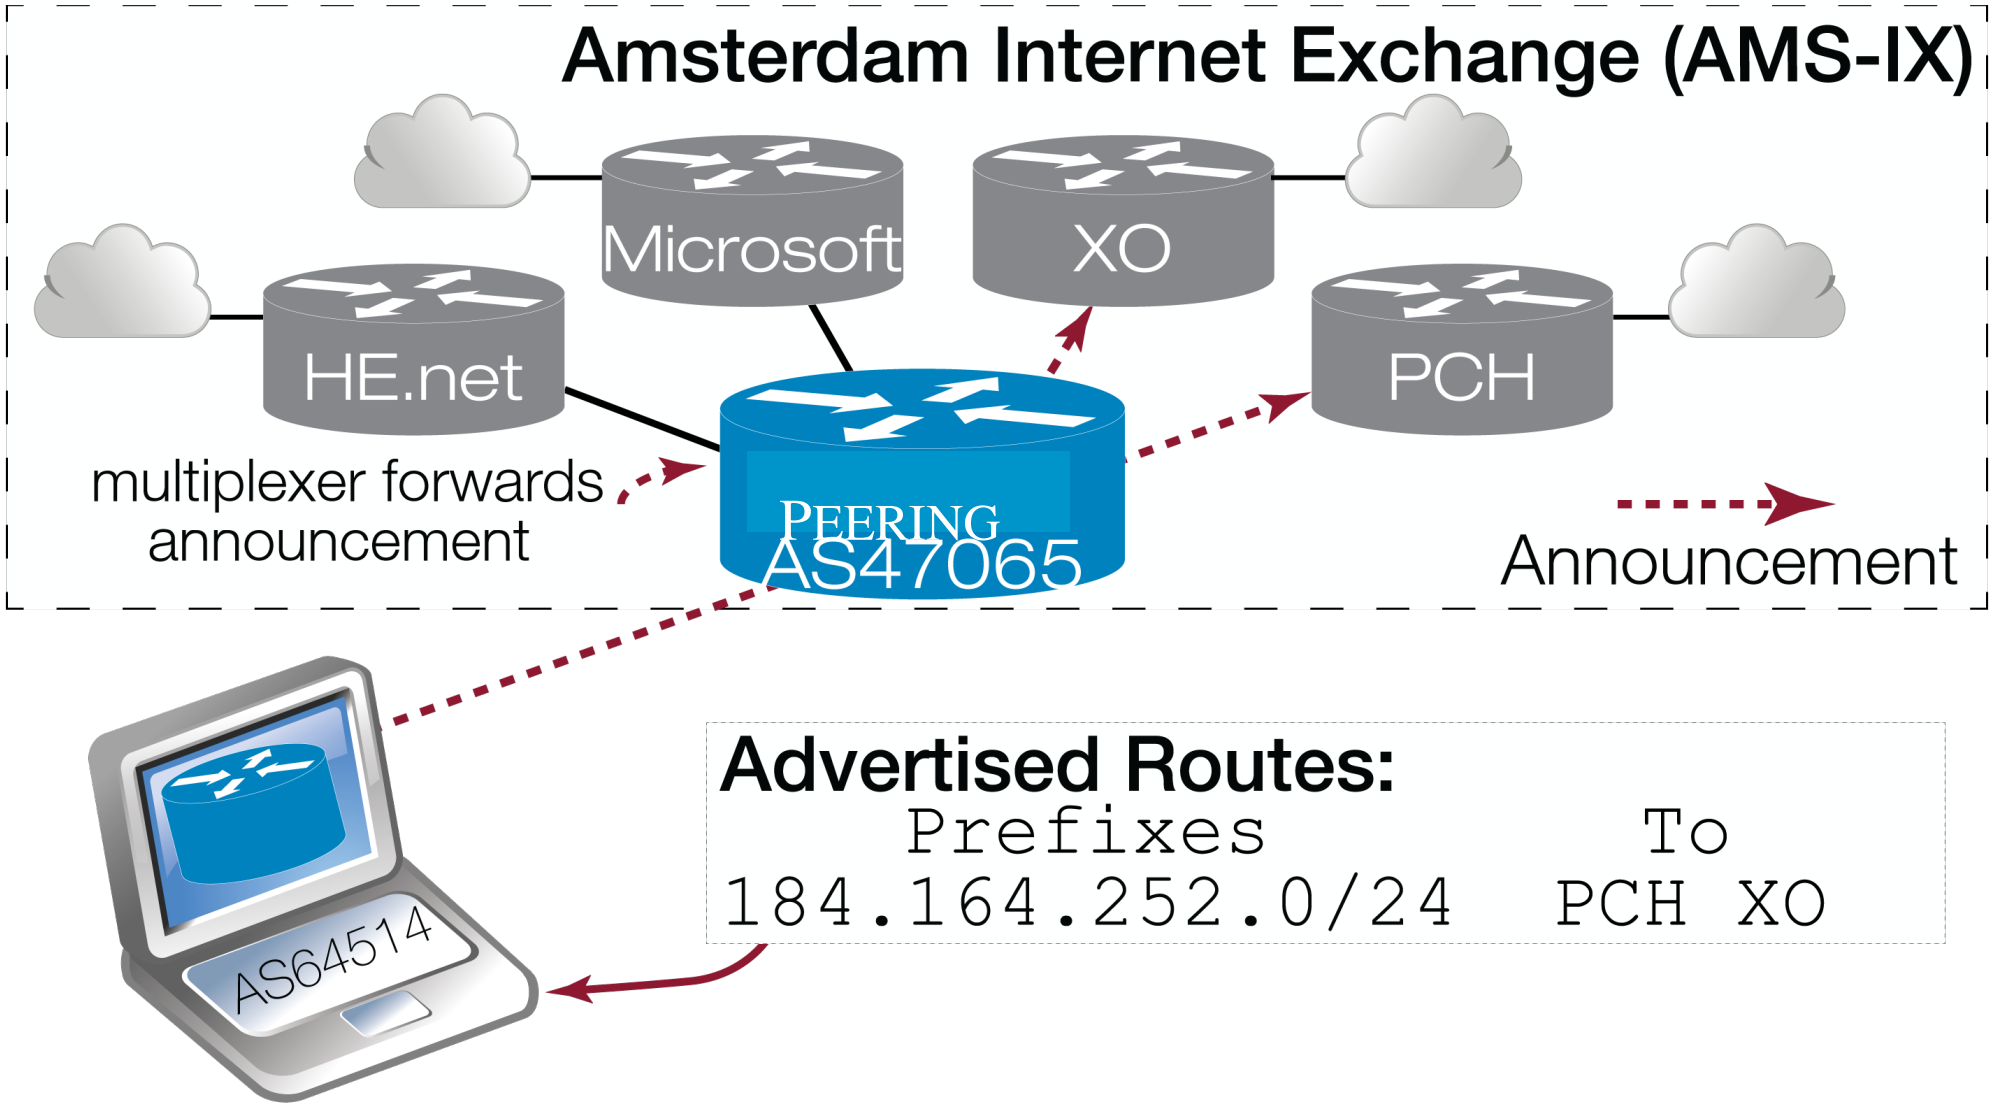

Receiving Routes via AMS-IX

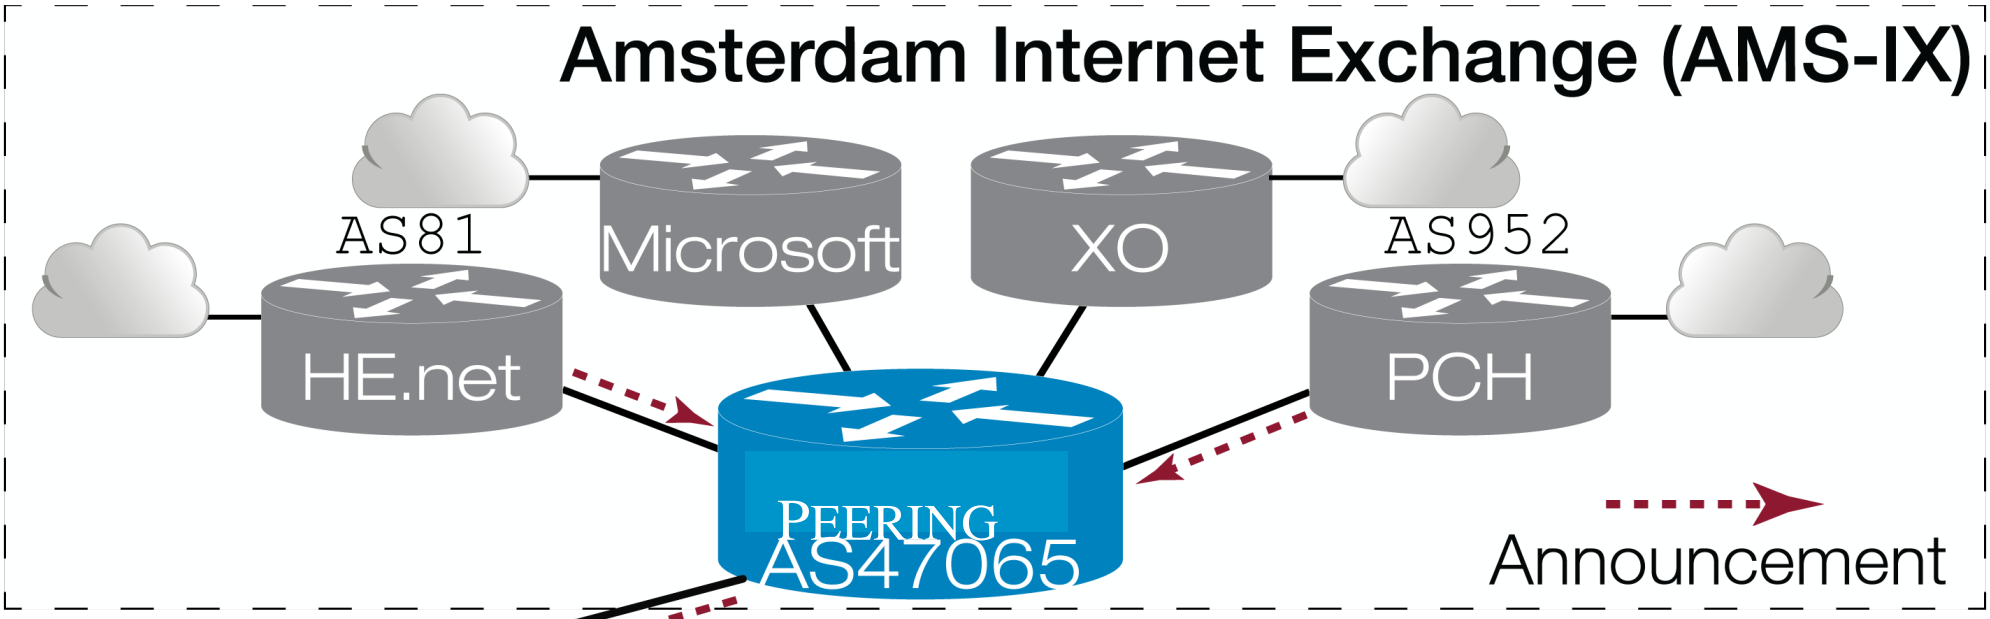

announced routes received by multiplexer
transparently forwarded to private AS

Received Routes:

```
    Prefixes          AS-Path
>77.78.39.0/24        952
 77.78.39.0/24        81   952
```

Exchanging Traffic via AMS-IX

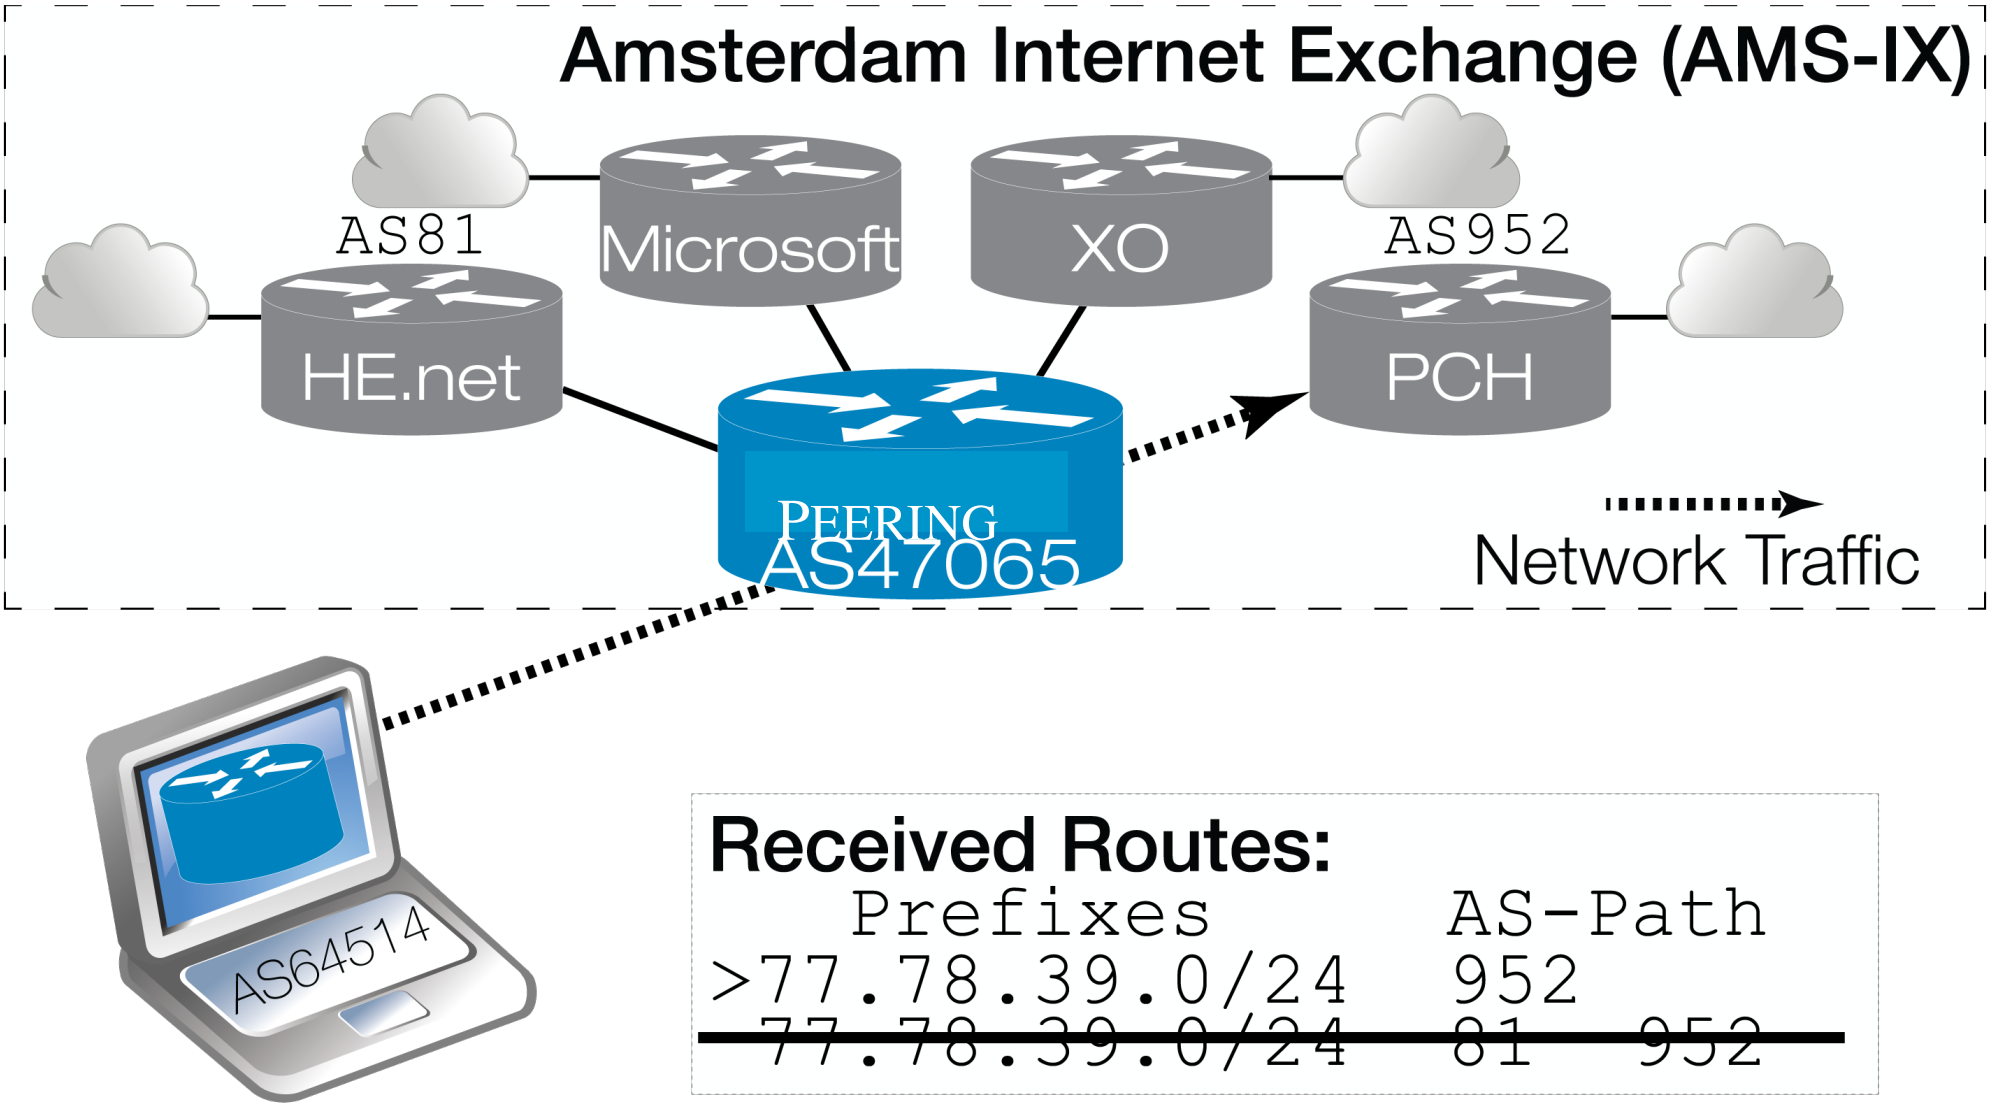

Isolation of Experiments at AMS-IX

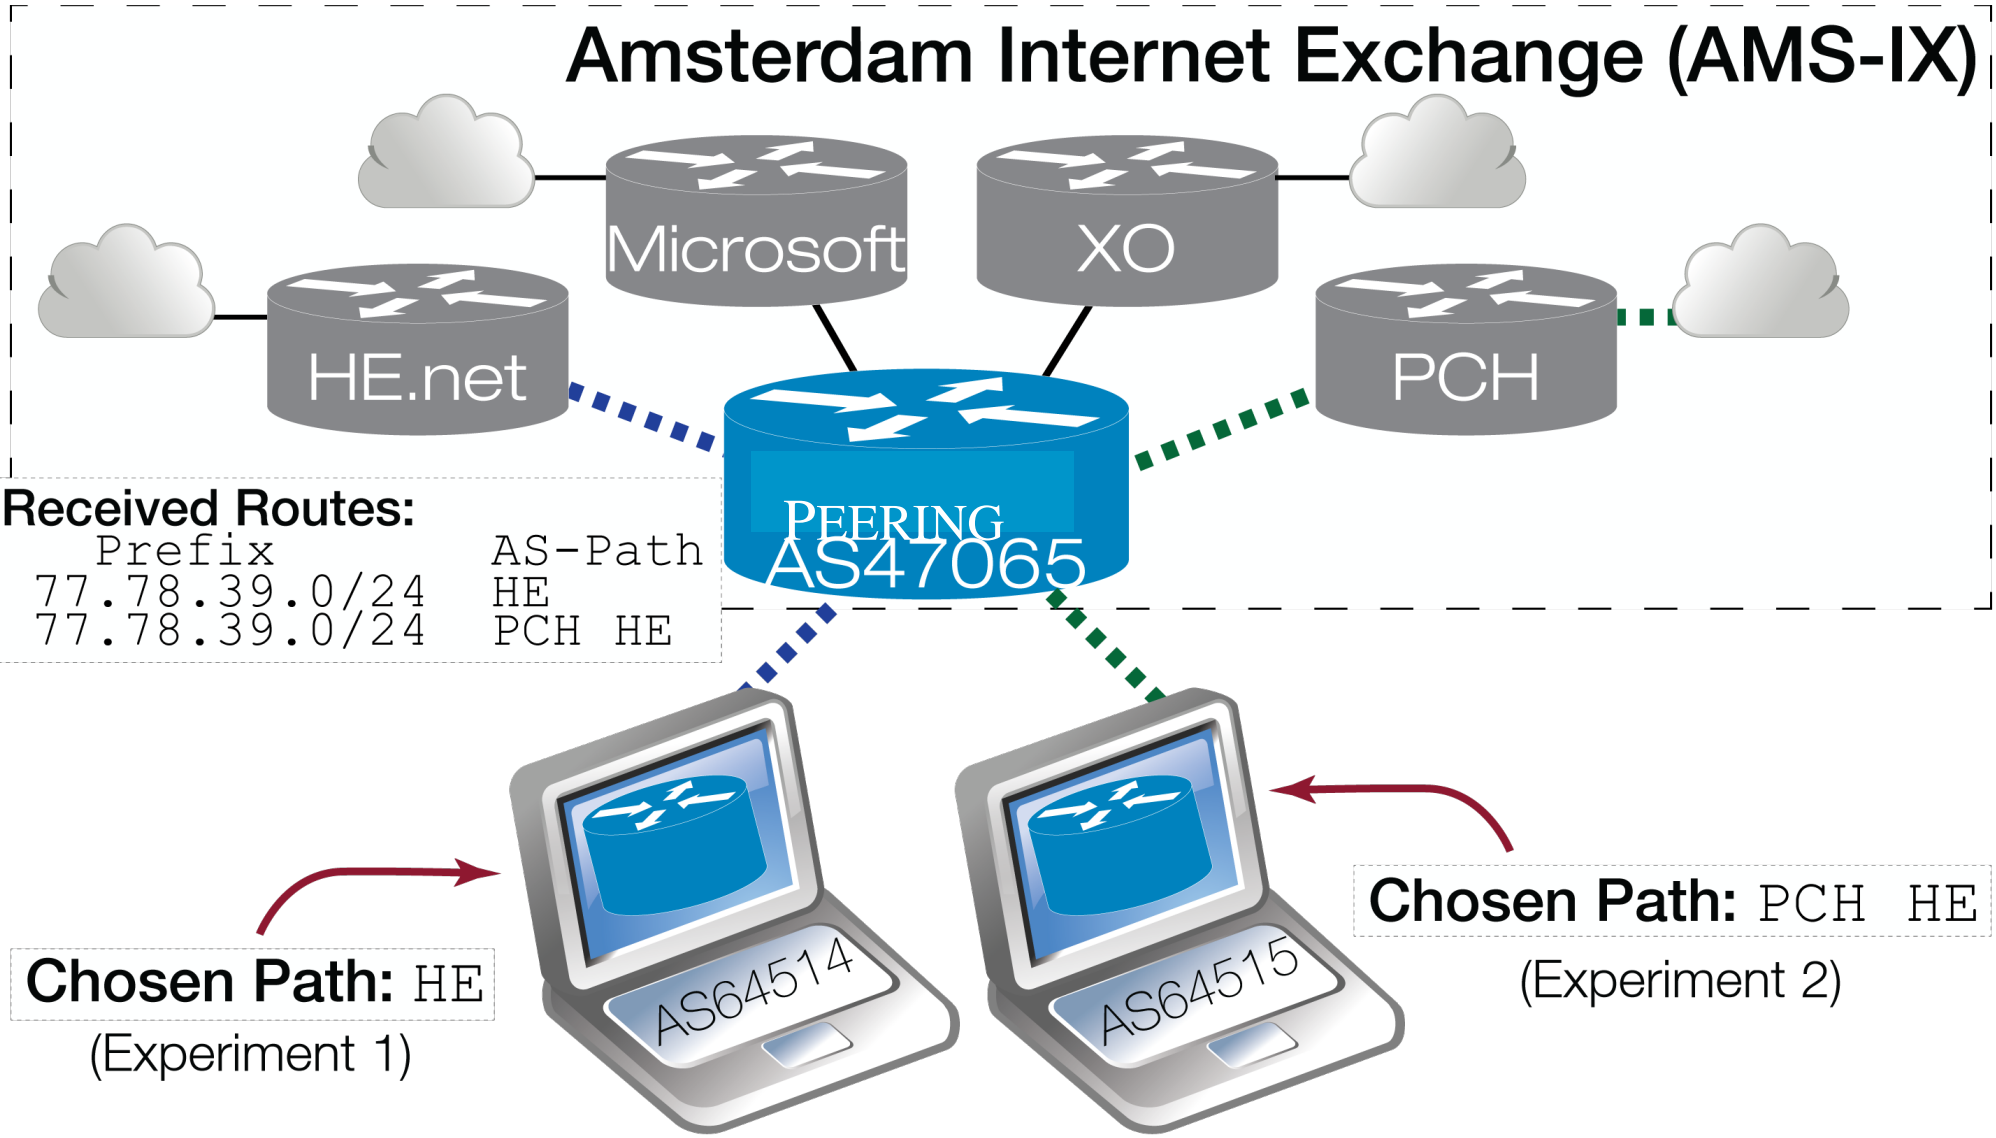

Why a community testbed?

Solve problems once

- Establish peerings and PoPs
- Safety
  - Filter spoofing, hijacks, leaks
  - Rate limit announcements and traffic

Centralize relationships

- ASes have a relationship with us, not each experiment
  - AS47065 is an academic research ASN
- Stable peerings
- Central point-of-contact for ops and support
- Centralize trust

PEERING's use in prior research

Software Defined Internet Exchange (SDX) [SIGCOMM14]

- SDX provides a peering fabric giving ASes fine-grained control
- **PEERING** enabled SDX experiments to exchange real Internet traffic

PEERING's use in prior research

Software Defined Internet Exchange (SDX) [SIGCOMM14]

- SDX provides a peering fabric giving ASes fine-grained control
- **PEERING** enabled SDX experiments to exchange real Internet traffic

PoiRoot: Investigating Root Causes of Path Changes [SIGCOMM '13]

- Made announcements via **PEERING** to partially infer AS's policies
- **PEERING** let us induce changes, providing ground truth for evaluation

PEERING's use in prior research

Software Defined Internet Exchange (SDX) [SIGCOMM14]

- SDX provides a peering fabric giving ASes fine-grained control
- **PEERING** enabled SDX experiments to exchange real Internet traffic

PoiRoot: Investigating Root Causes of Path Changes [SIGCOMM '13]

- Made announcements via **PEERING** to partially infer AS's policies
- **PEERING** let us induce changes, providing ground truth for evaluation

LIFEGUARD: Practical Repair of Persistent Route Failures [SIGCOMM '12]

- Used crafted BGP announcements to route around routing failures
- **PEERING** let us evaluate (path availability, convergence, filtering, etc.)

PEERING's use in prior research

Software Defined Internet Exchange (SDX) [SIGCOMM14]

- SDX provides a peering fabric giving ASes fine-grained control
- **PEERING** enabled SDX experiments to exchange real Internet traffic

PoiRoot: Investigating Root Causes of Path Changes [SIGCOMM '13]

- Made announcements via **PEERING** to partially infer AS's policies
- **PEERING** let us induce changes, providing ground truth for evaluation

LIFEGUARD: Practical Repair of Persistent Route Failures [SIGCOMM '12]

- Used crafted BGP announcements to route around routing failures
- **PEERING** let us evaluate (path availability, convergence, filtering, etc.)

ARROW: One Tunnel is Enough [SIGCOMM '14]

- Incrementally deployable architecture for tunneled transit as a service for remote customers
- Talk to Simon Peter (UW) this week at NANOG

PEERING's use in prior research

Software Defined Internet Exchange (SDX) [SIGCOMM14]

- SDX provides a peering fabric giving ASes fine-grained control
- **PEERING** enabled SDX experiments to exchange real Internet traffic

PoiRoot: Investigating Root Causes of Path Changes [SIGCOMM '13]

- Made announcements via **PEERING** to partially infer AS's policies
- **PEERING** let us induce changes, providing ground truth for evaluation

LIFEGUARD: Practical Repair of Persistent Route Failures [SIGCOMM '12]

- Used crafted BGP announcements to route around routing failures
- **PEERING** let us evaluate (path availability, convergence, filtering, etc.)

ARROW: One Tunnel is Enough [SIGCOMM '14]

- Incrementally deployable architecture for tunneled transit as a service for remote customers
- Talk to Simon Peter (UW) this week at NANOG

PECAN: Joint Content & Network Routing [SIGMETRICS '13, NANOG56]

- Measuring the benefits of joint traffic engineering and replica selection

Conclusion

Little innovation in BGP in past 20 years

- Yet BGP is at the root of fundamental Internet problems
- Researchers lacked tools to run meaningful BGP experiments

Conclusion

Little innovation in BGP in past 20 years

- Yet BGP is at the root of fundamental Internet problems
- Researchers lacked tools to run meaningful BGP experiments

PEERING provides a new approach to BGP research

- Connects emulated ASes with ***real*** ASes on the Internet
- Vision is to let researchers run the AS of their choice

Conclusion

Little innovation in BGP in past 20 years

- Yet BGP is at the root of fundamental Internet problems
- Researchers lacked tools to run meaningful BGP experiments

PEERING provides a new approach to BGP research

- Connects emulated ASes with ***real*** ASes on the Internet
- Vision is to let researchers run the AS of their choice

Looking for feedback from operator community

- What problems would you like to see PEERING tackle?
- Willing to peer with us or hosting us?
- **Contact:** bgpmux-noc@gtnoise.net

Conclusion

Little innovation in BGP in past 20 years

- Yet BGP is at the root of fundamental Internet problems
- Researchers lacked tools to run meaningful BGP experiments

PEERING provides a new approach to BGP research

- Connects emulated ASes with ***real*** ASes on the Internet
- Vision is to let researchers run the AS of their choice

Looking for feedback from operator community

- What problems would you like to see PEERING tackle?
- Willing to peer with us or hosting us?
- **Contact:** bgpmux-noc@gtnoise.net

Example PEERING Experiment

Testing Scalability of Software Defined Internet Exchange (SDX) [to appear, SIGCOMM '14]

- SDX provides a peering fabric giving ASes fine-grained control
  - Enables IXP peers to setup advanced policies (application specific peering)
  - BGP updates and controller policies define SDN fabric switch rules
- Built the SDX components – how do conduct ***realistic*** evaluation?

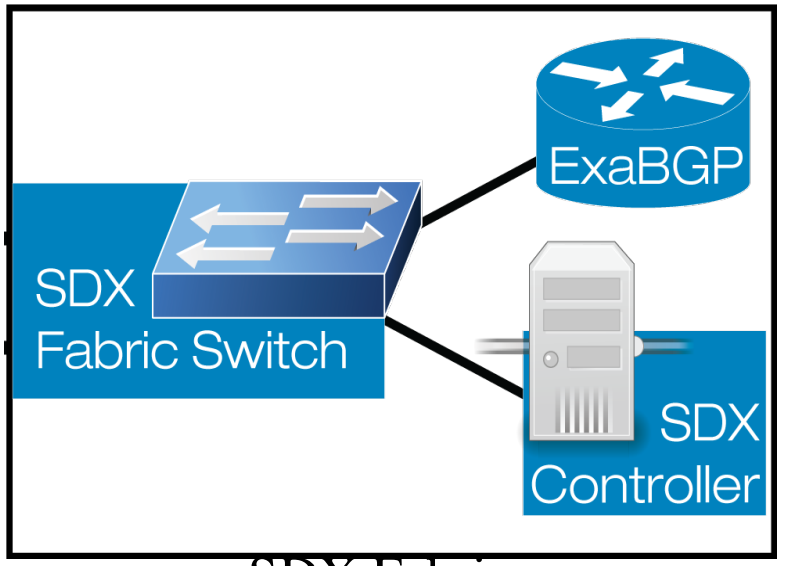

SDX Fabric

Example PEERING Experiment

Testing Scalability of Software Defined Internet Exchange (SDX) [to appear, SIGCOMM '14]

- SDX provides a peering fabric giving ASes fine-grained control
  - Enables IXP peers to setup advanced policies (application specific peering)
  - BGP updates and controller policies define SDN fabric switch rules
- Built the SDX components – how do conduct ***realistic*** evaluation?

Can we test how SDX would operate at *real* IXP?

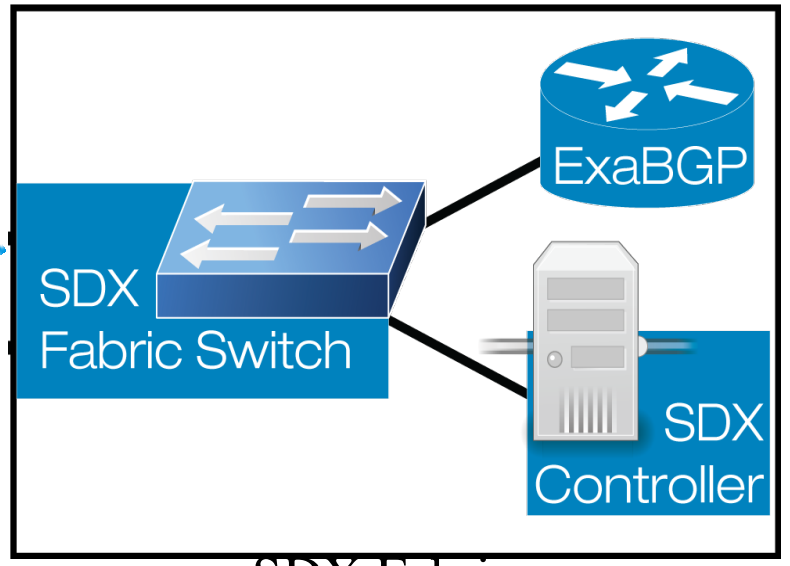

Example PEERING Experiment

Testing Scalability of Software Defined Internet Exchange (SDX) [to appear, SIGCOMM '14]

- SDX provides a peering fabric giving ASes fine-grained control
  - Enables IXP peers to setup advanced policies (application specific peering)
  - BGP updates and controller policies define SDN fabric switch rules
- Built the SDX components – how do conduct ***realistic*** evaluation?

Can we test how SDX would operate at *real* IXP?

Need to add real peers from *real* IXP!

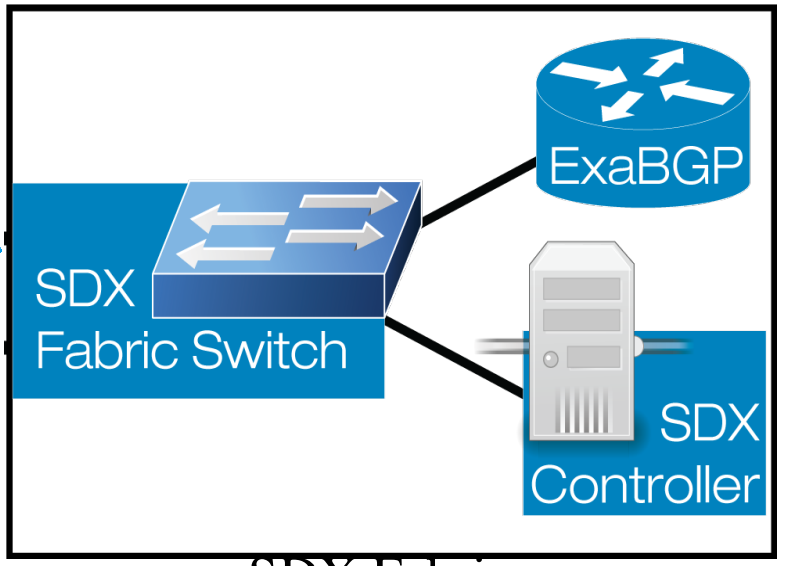

Example PEERING Experiment

Testing Scalability of Software Defined Internet Exchange (SDX) [to appear, SIGCOMM '14]

- Real AS located at AMS-IX ***virtually*** added to emulated SDX fabric
  - TP Mux transparently tunneled BGP sessions over VPN connection
  - ExaBGP maintains BGP sessions with peers, forwards to controller
- Controller updates SDX fabric in response to ***real*** BGP messages