Withdrawing the BGP Re-Routing Curtain

Understanding the Security Impact of BGP Poisoning via Real-World Measurements

Jared M. Smith, *Kyle Birkeland, Tyler McDaniel, Max Schuchard*
University of Tennessee, Knoxville
volsec.org

FYRE

Internet Routing: Theory into Practice

- Security systems assume how complex infrastructures like the Internet work
  - **Claim:** *"Protocol implies X works, so X must work in practice"*
  - **Methodology:** *"Inference and passive measurement are enough"*
  - **Assumption:** *"Common logic suggests X does not work, so X must not work"*

- **Our goal:** To understand how real-world Internet routing behavior impacts published security literature
  - Actively measure the ability conduct BGP poisoning
  - Re-evaluate systems measured only in simulation, passively, or with inferences
  - Examine if common logic about the Internet holds

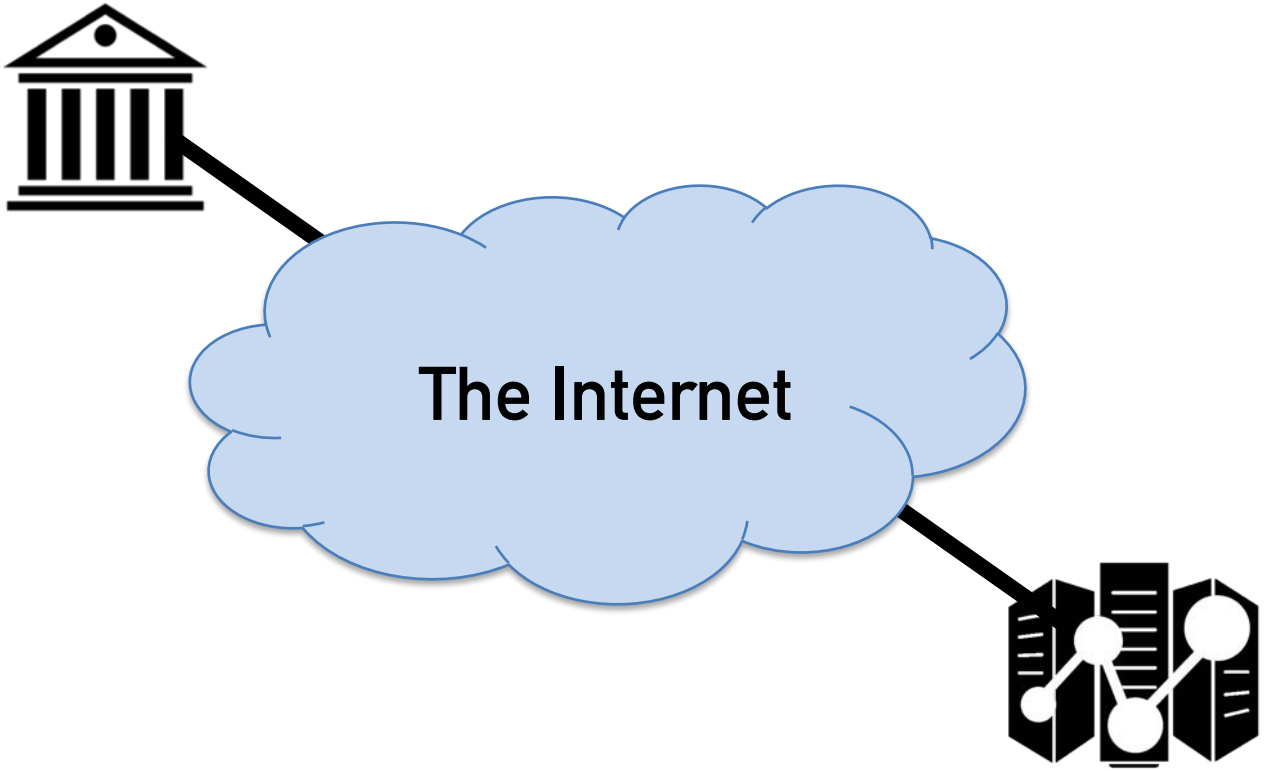
The Internet

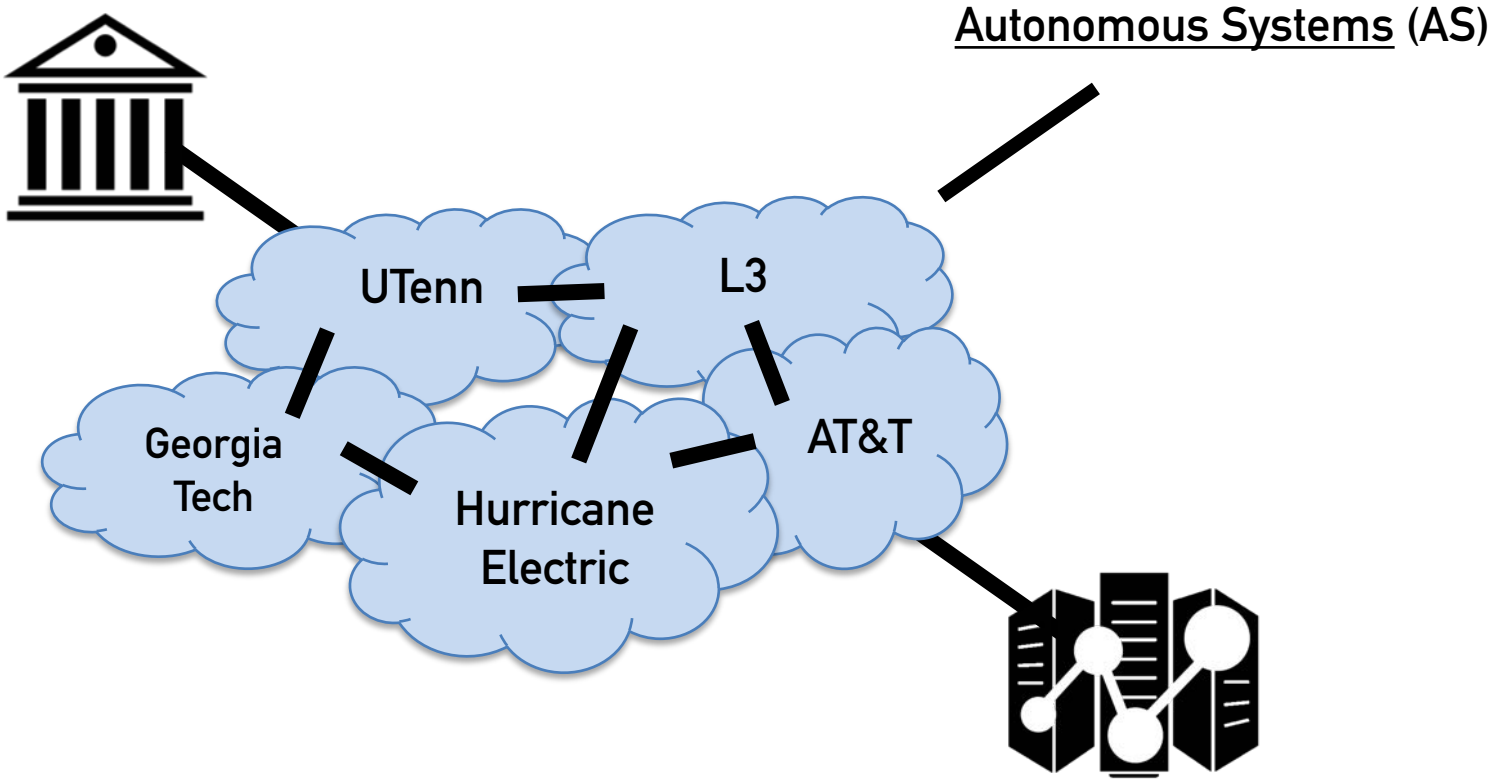
Autonomous Systems (AS)
UTenn
L3
Georgia Tech
Hurricane Electric
AT&T

BGP

Border Gateway Protocol

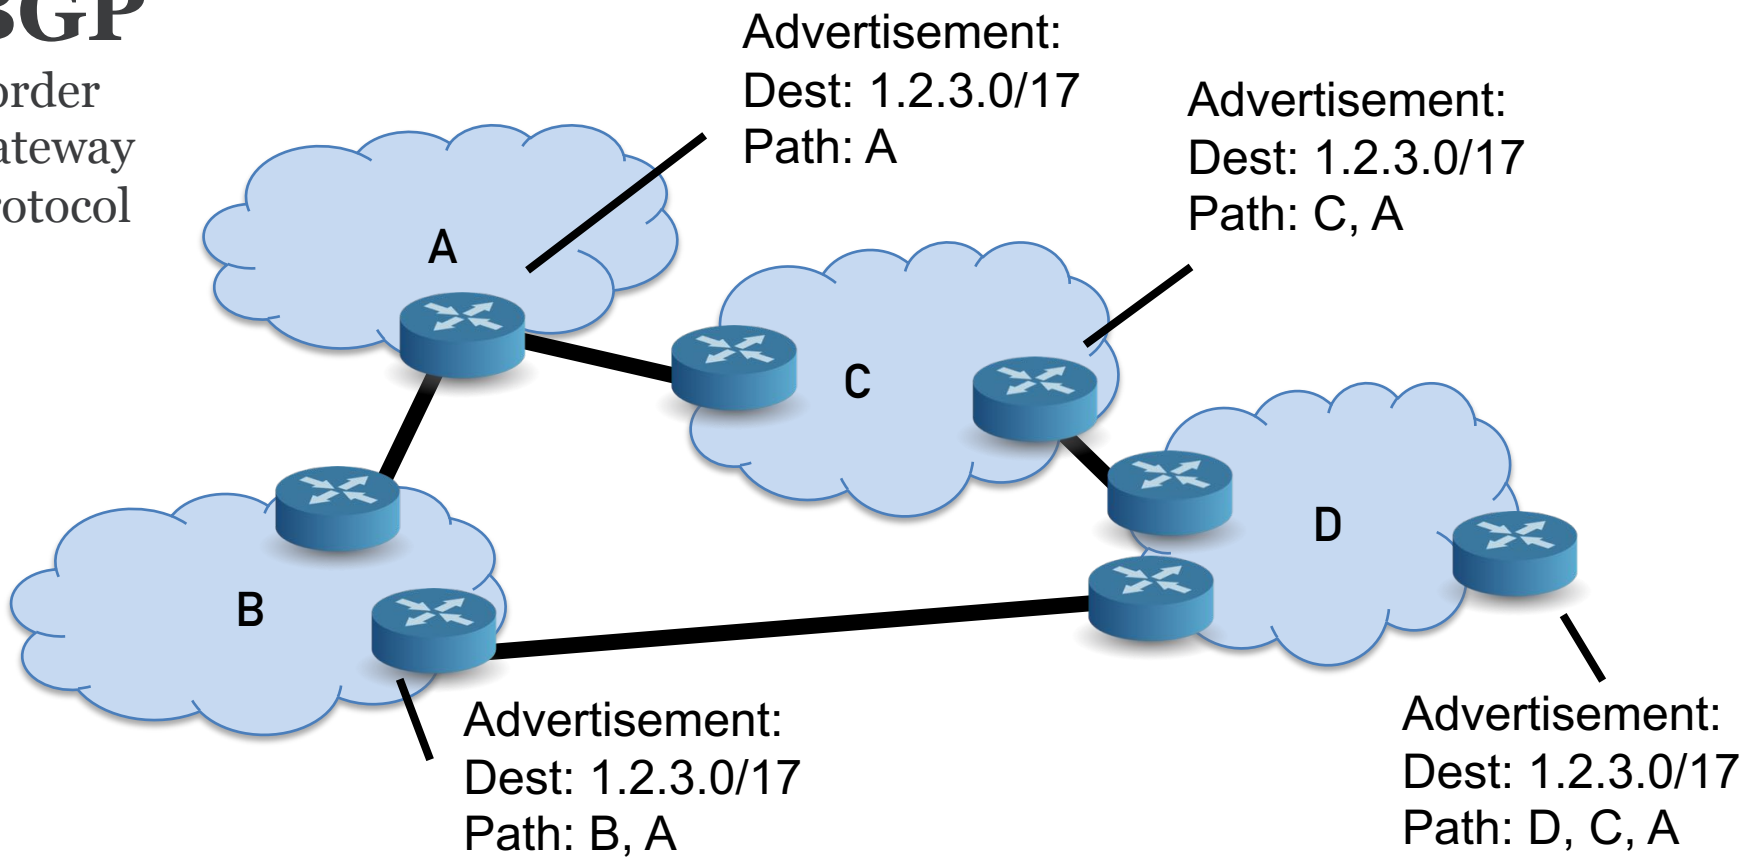

Inbound Path Manipulation

- Mechanisms give hints for which inbound path to take
  - Example: Multi-Exit Discriminator (MED)

- We can use side-effects of protocol-compliant behavior
  - Example: **BGP Poisoning**

BGP Poisoning

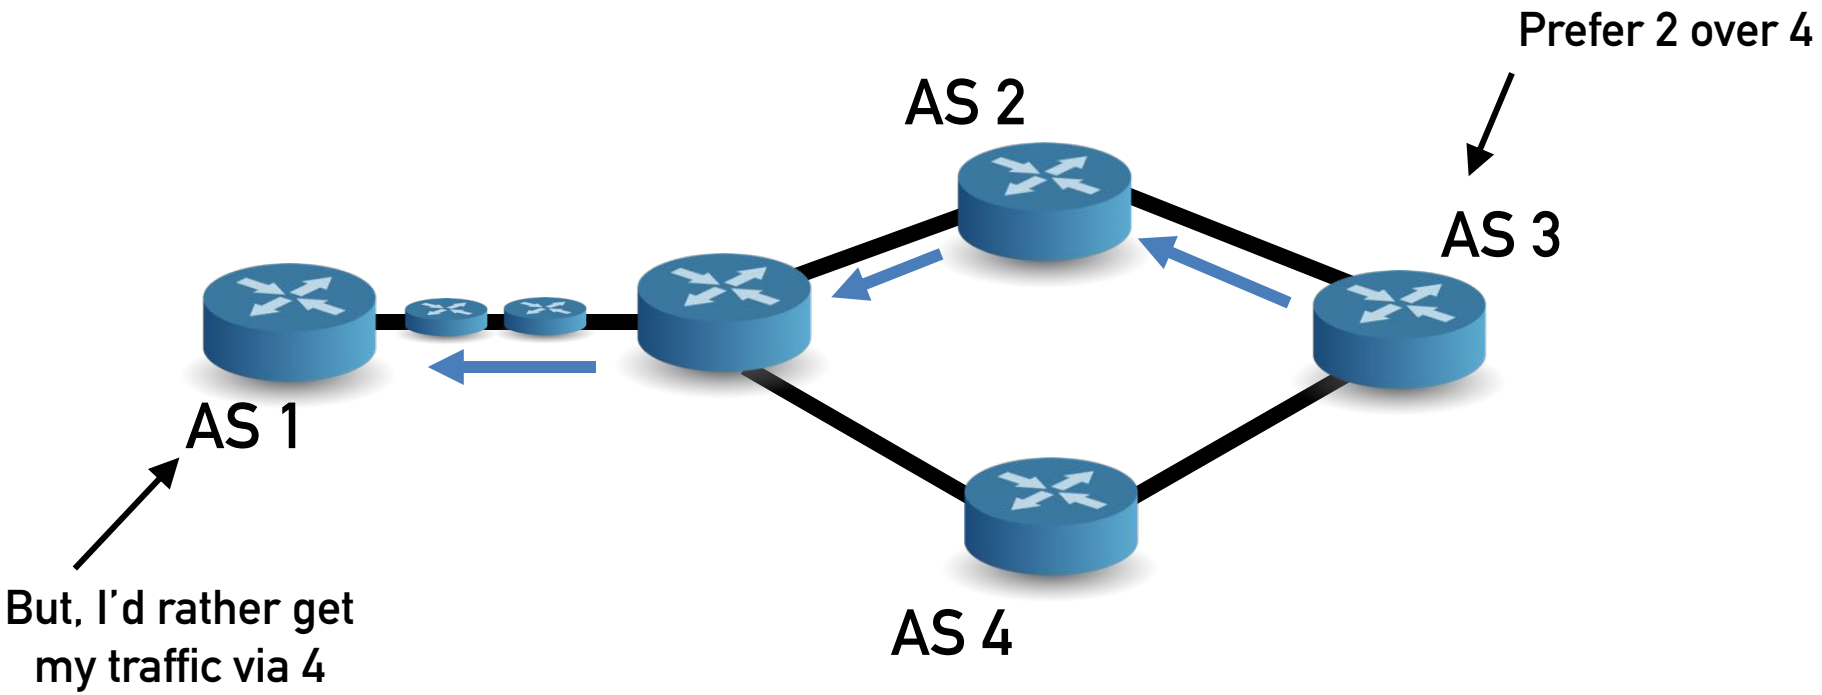

BGP Poisoning

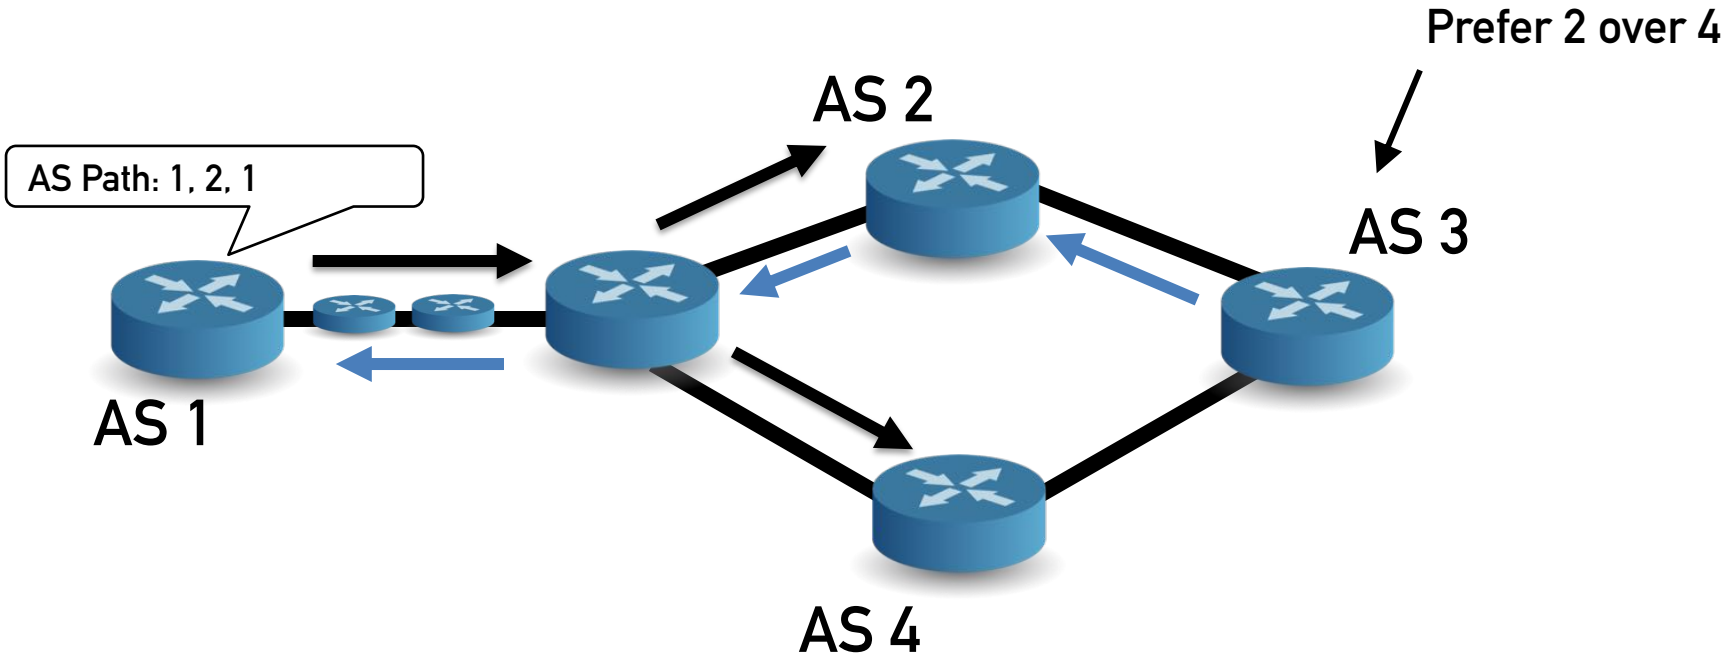

BGP Poisoning

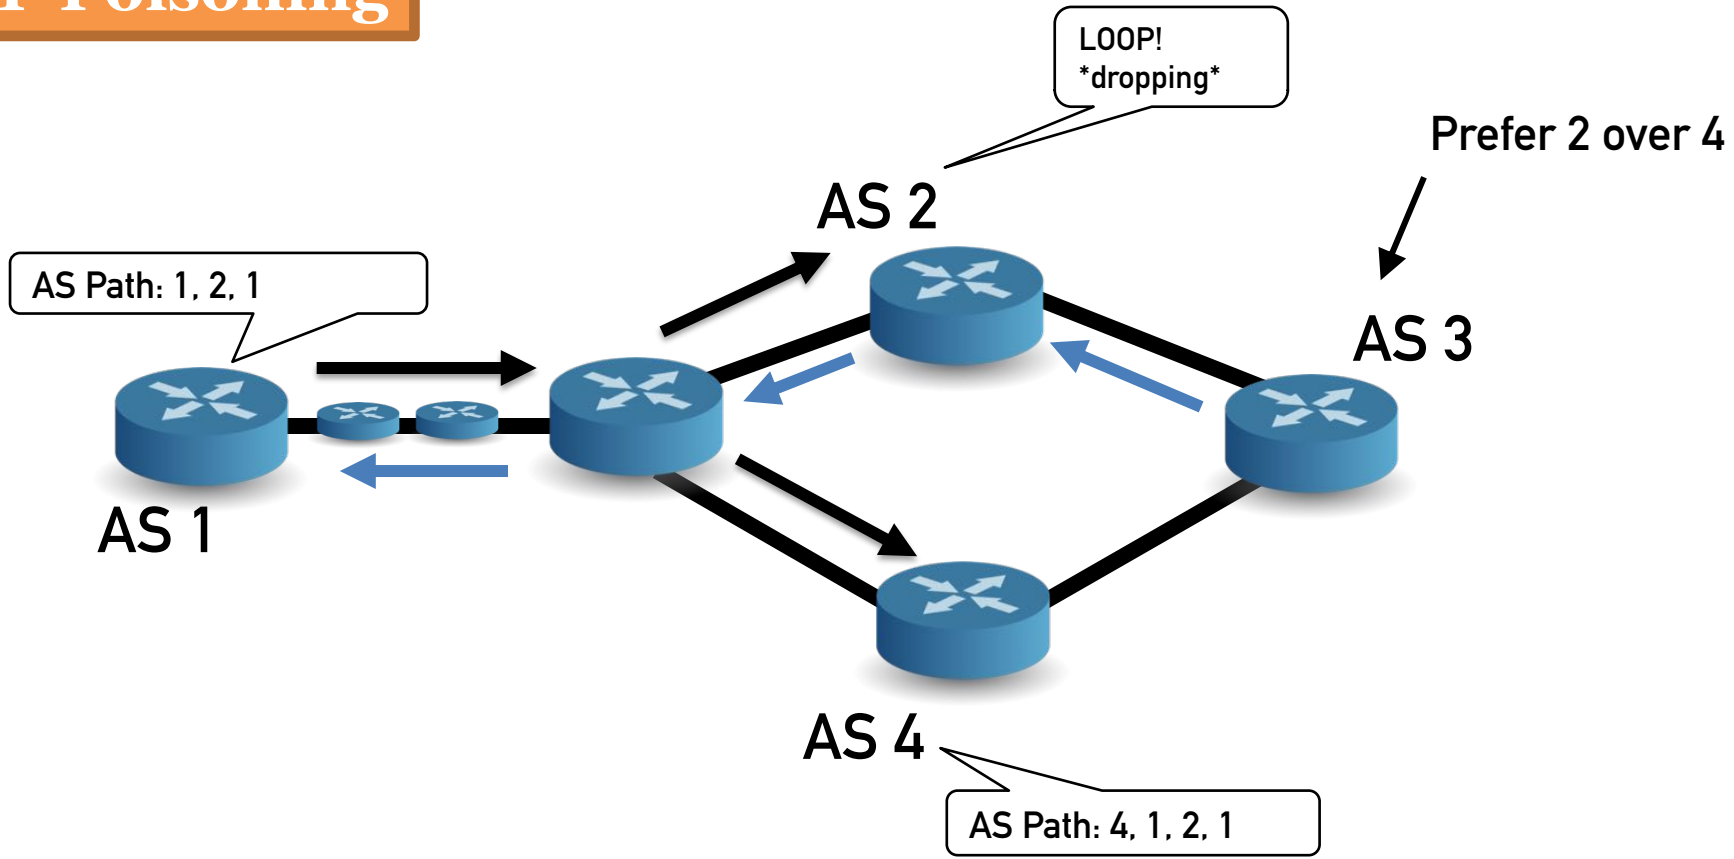

BGP Poisoning

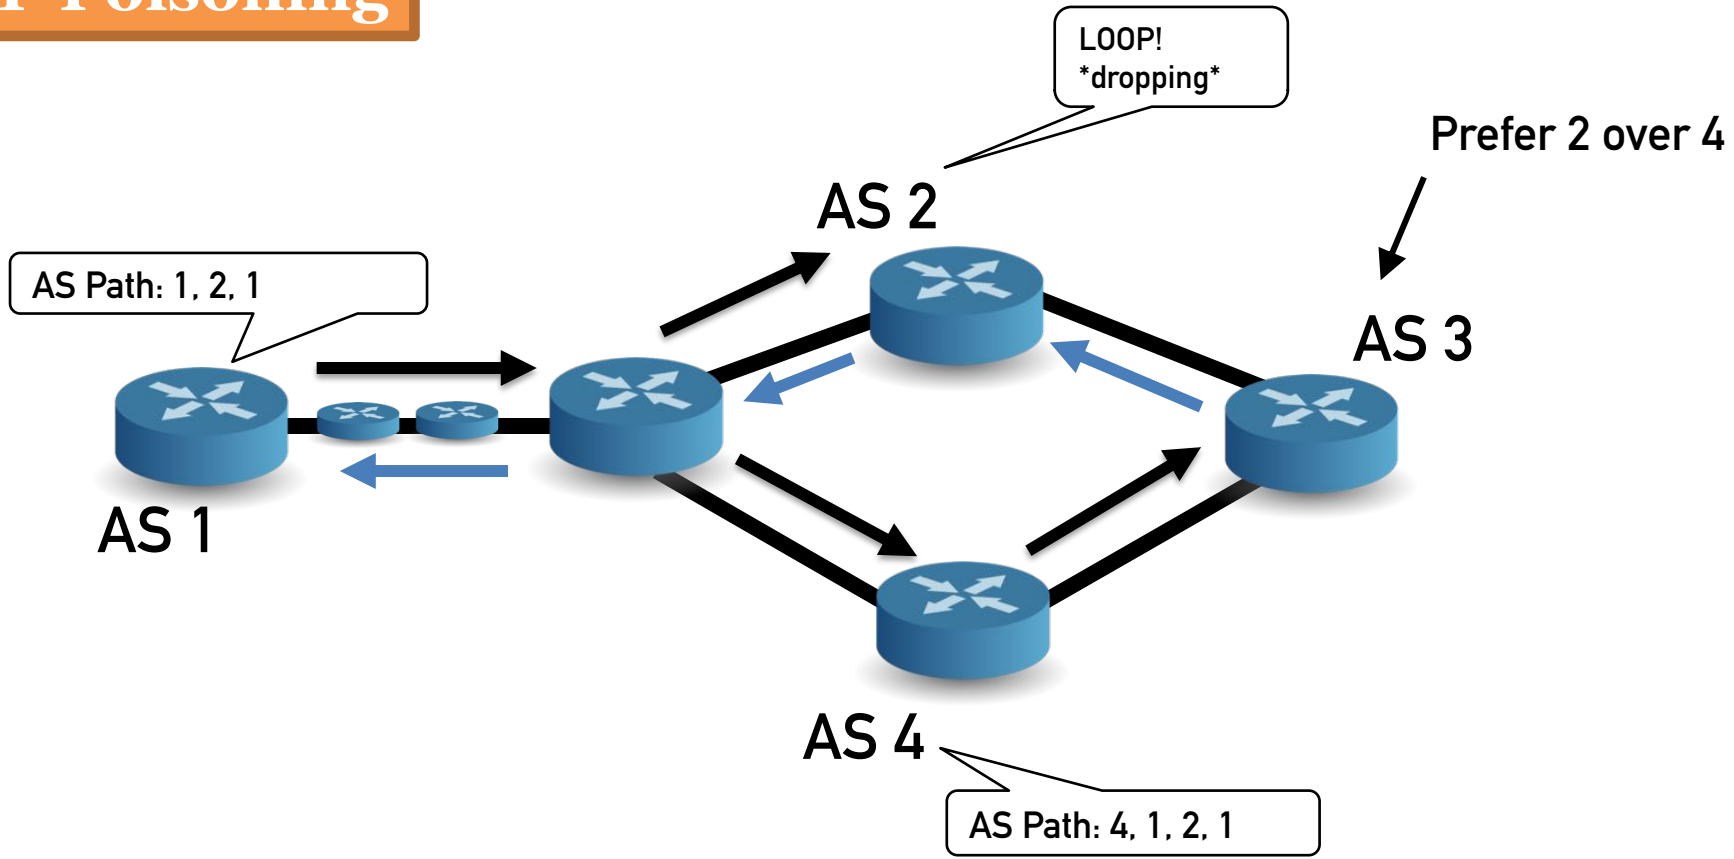

BGP Poisoning

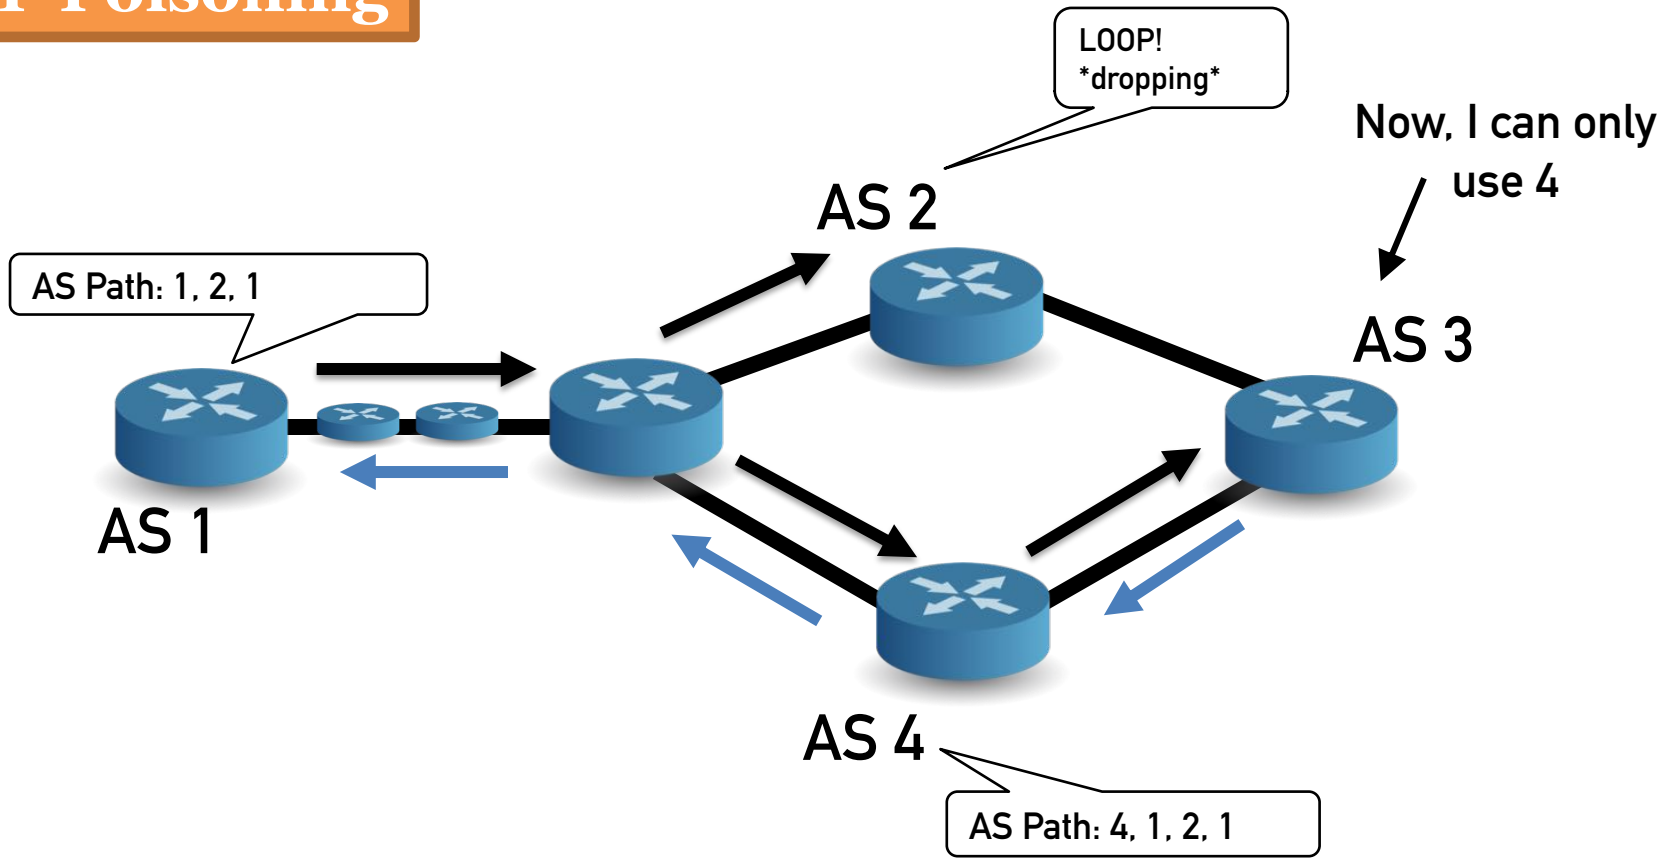

Nyx: Routing Around Congestion

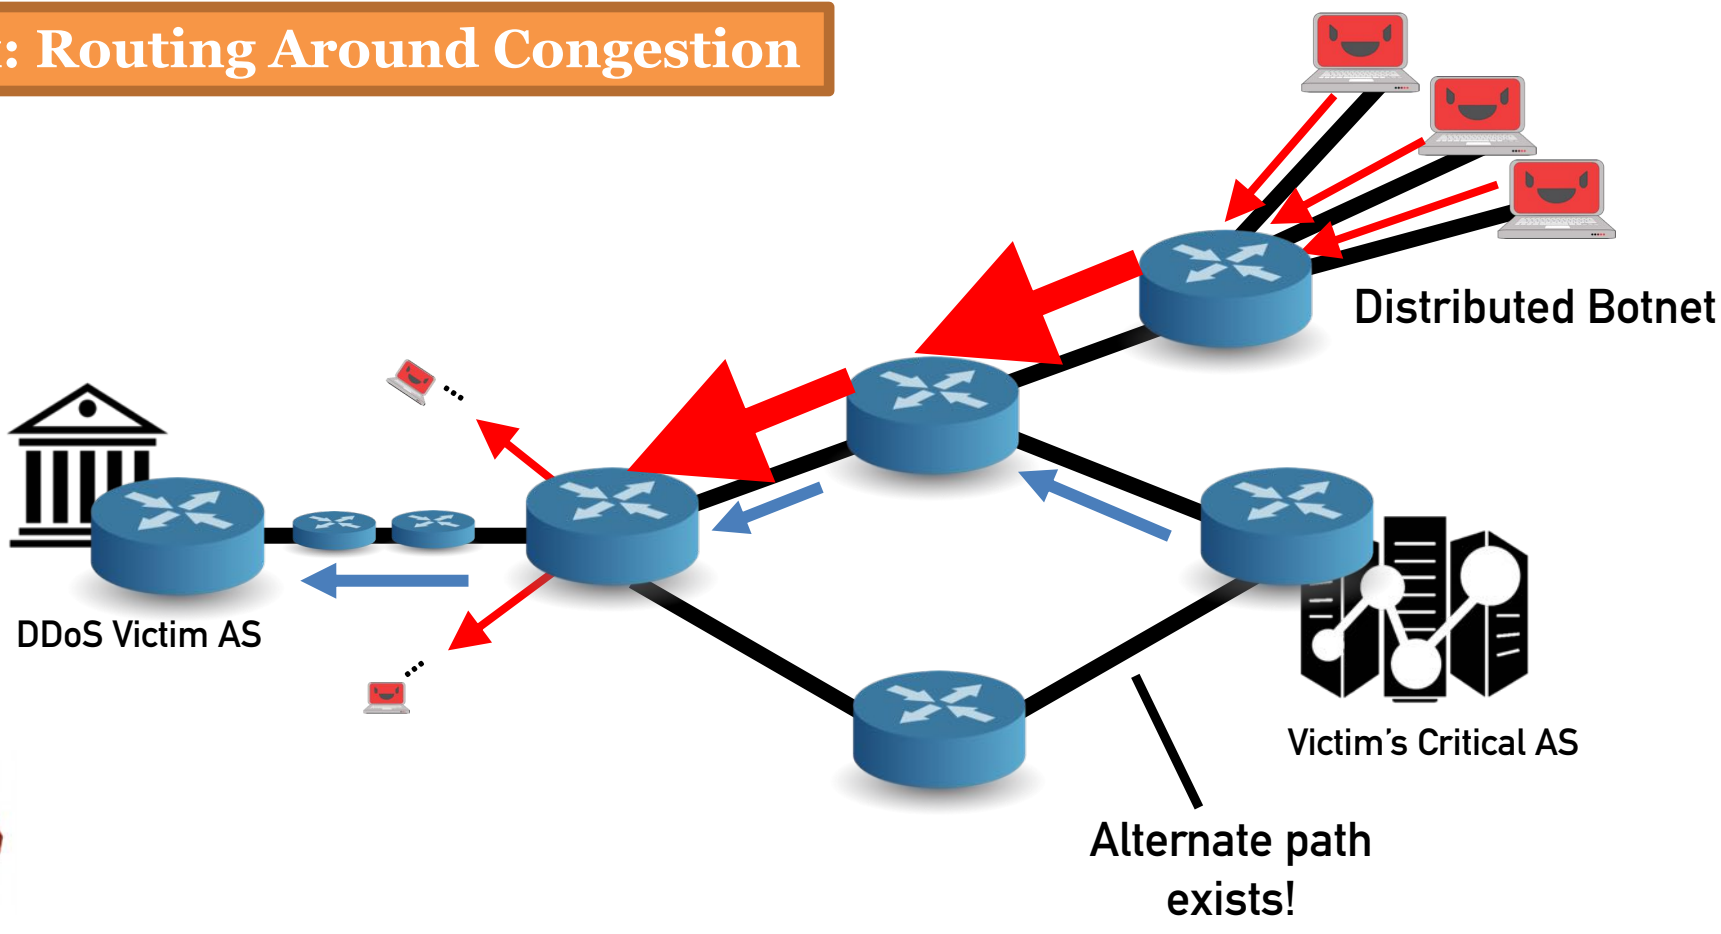

IEEE S&P

Nyx: Routing Around Congestion

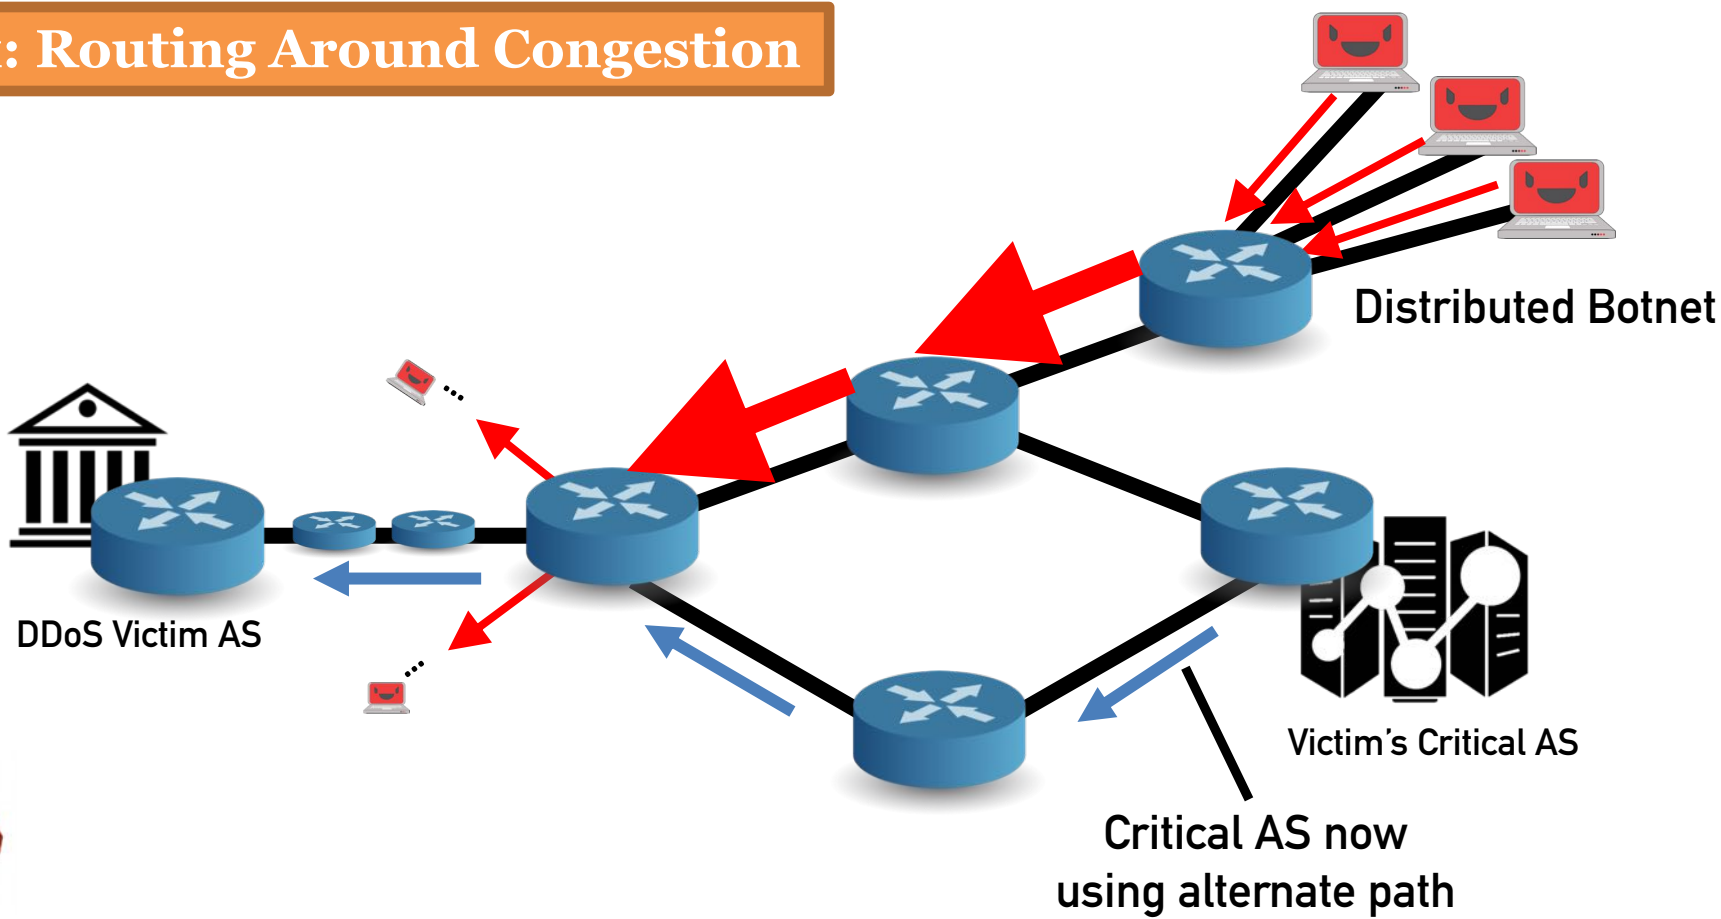

IEEE S&P

Relevant Security Literature

- **Nyx (DDoS Defense – S&P 2018)**
- RAD (Censorship Circ. – CCS 2012)
- Waterfall of Liberty (Censorship Circ. – CCS 2017)
- On Feasibility of Re-Routing (Examination of Nyx - S&P 2019)
- …

Diverging Claims

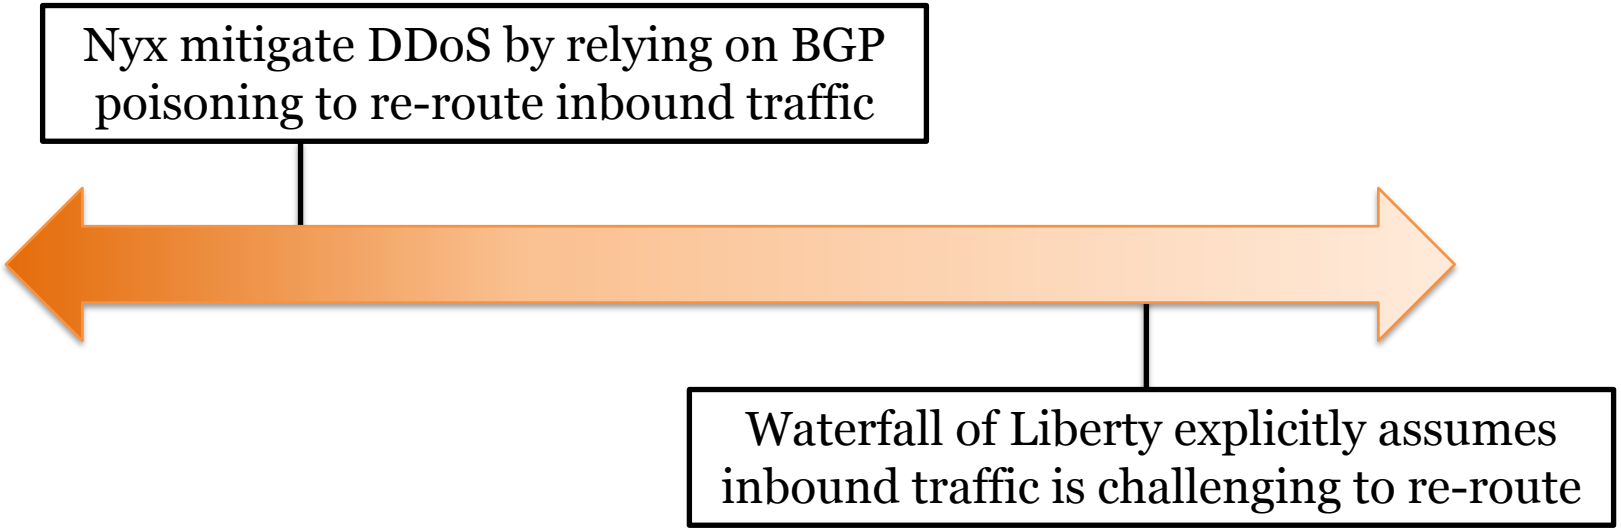

Diverging Claims

Ny
po

Nyx and *Waterfall of Liberty* are **built on polar opposite assumptions**, but **not tested on the live Internet**

Waterfall of Liberty explicitly assumes inbound traffic is challenging to re-route

All of this literature makes assumptions about how BGP poisoning works…

In reality, problems may occur…

- An AS might realize its not actually on the path
- An AS might realize we're lying about the path
- An AS might think the path looks anomalous

“Here be dragons”

Internet Topology

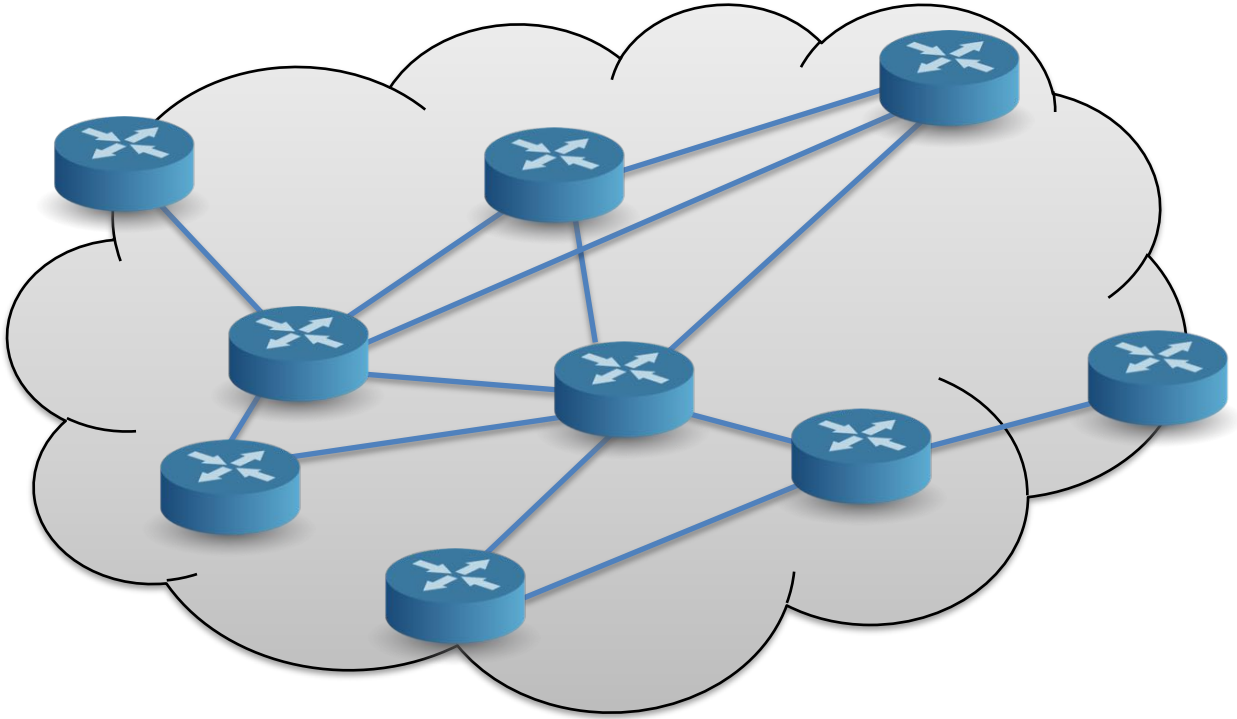

Sending BGP Advertisements

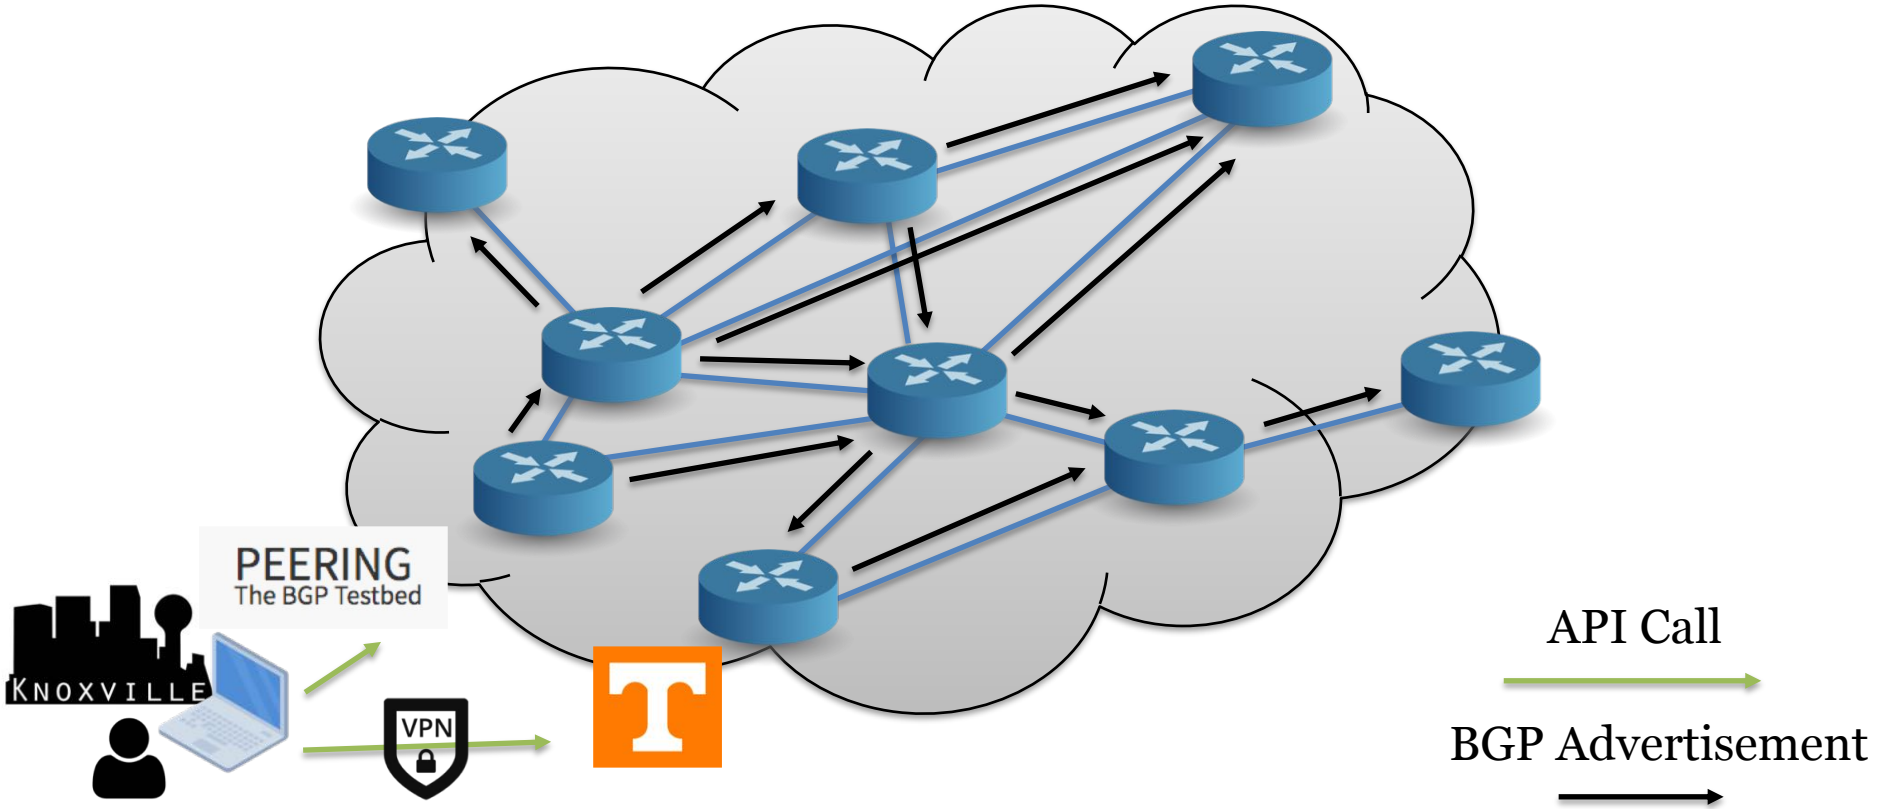

Collecting BGP Updates

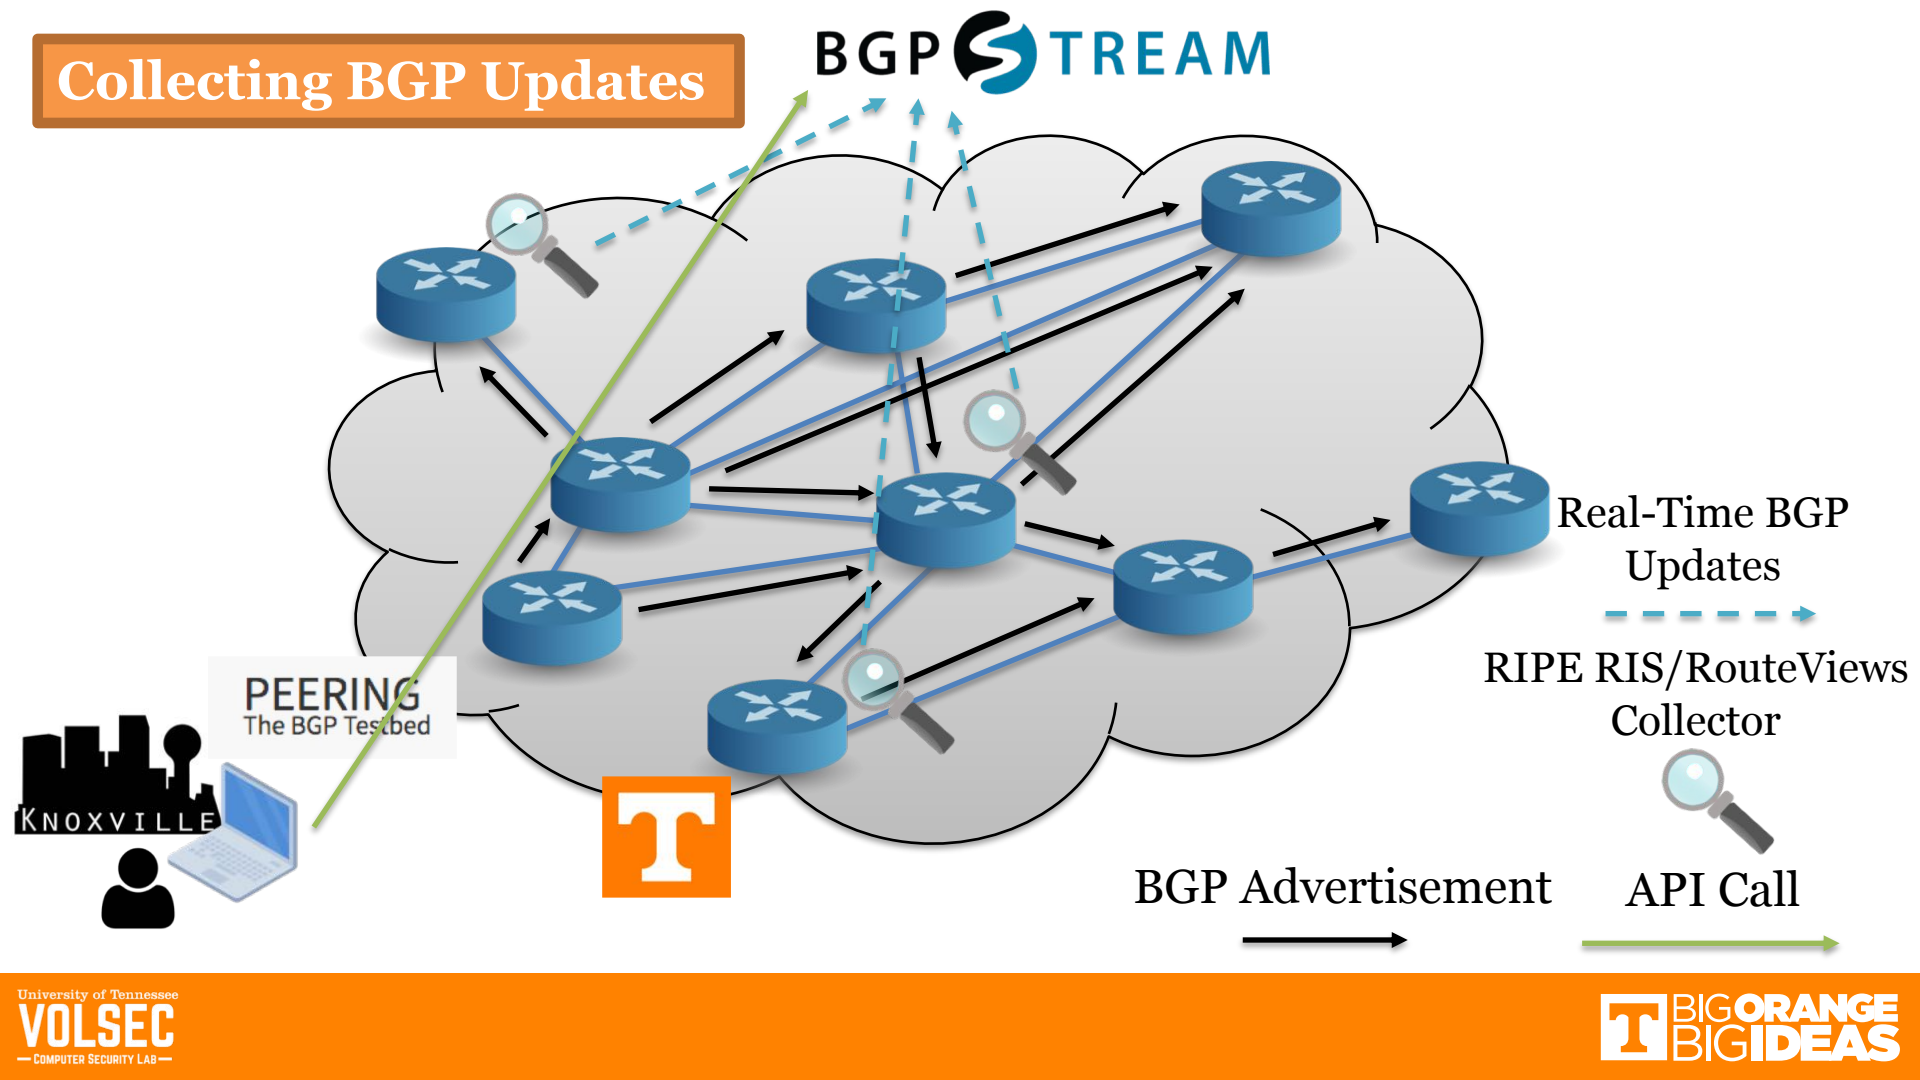

Sending Traceroutes

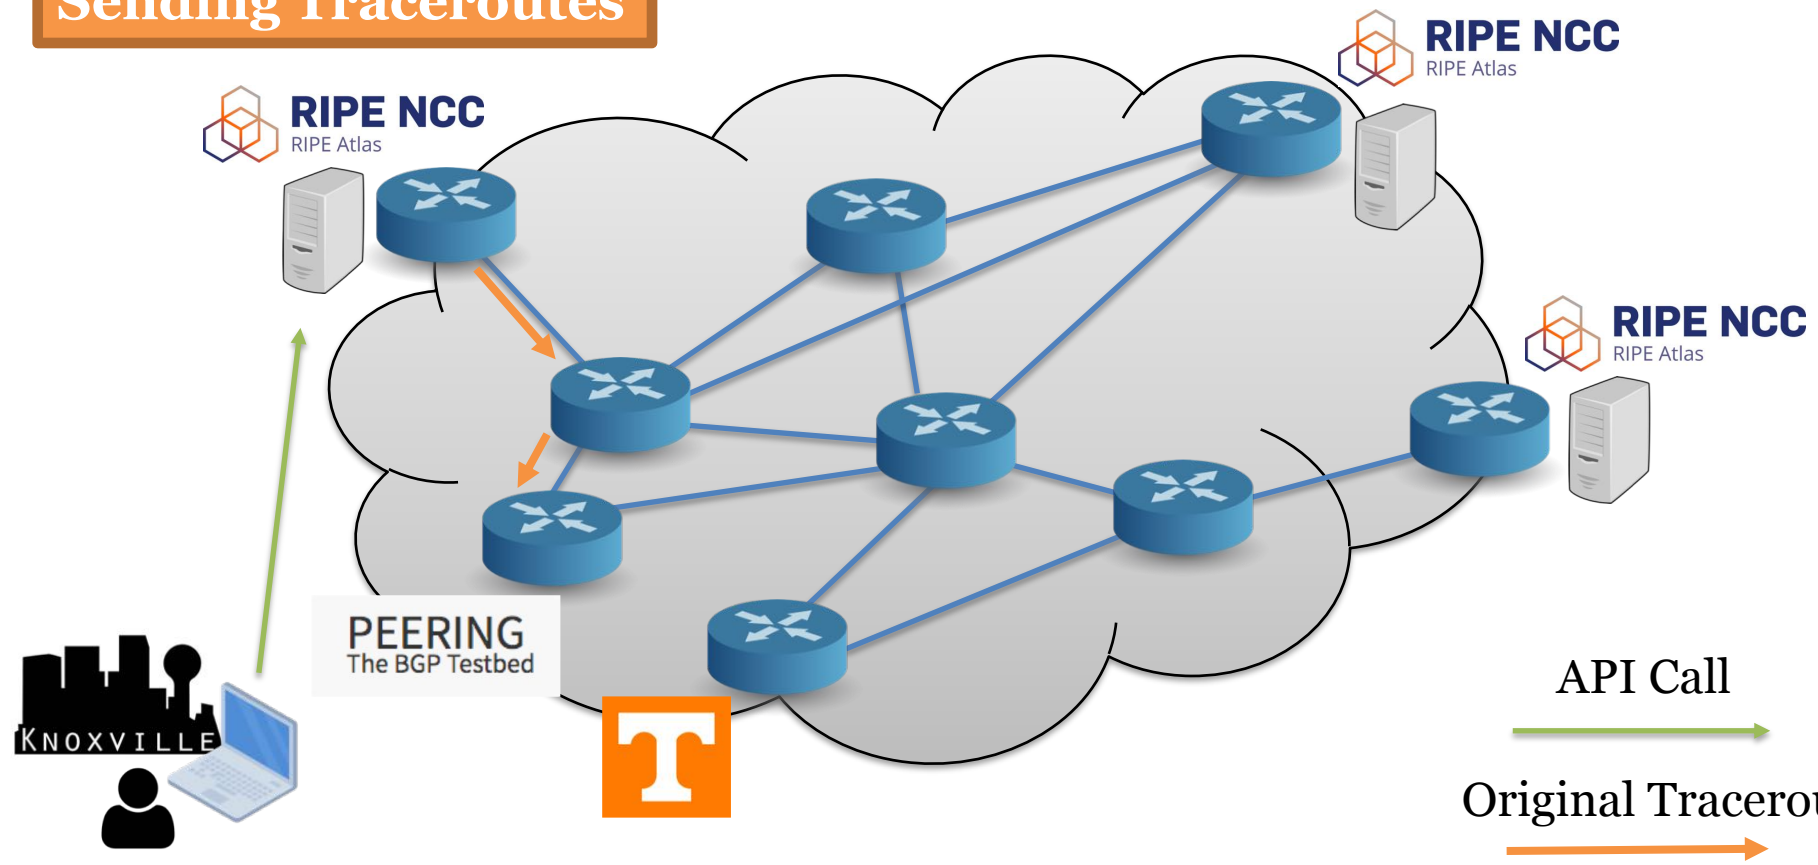

Infrastructure Details

BGP Advertisements

14 PoPs, 3 countries

Traceroutes

5,000 vantage points

BGP Updates

32 collectors

Automated experiment software:
https://github.com/volsec/active-bgp-measurement

It’s free! You can use this infrastructure!

Infrastructure Details

BGP Advertisements

free, application

Traceroutes

free

BGP Updates

free

Open Source:

https://github.com/volsec/active-bgp-measurement

Experimental Ethics

- Announced to and engaged with network operators
- No production traffic affected
- Minimal traffic sent along re-routed paths (< 1 Kbps)
- Normal BGP announcements (no malformed)
- Conformed to ISP filtering policies

All Experiments

1. **Ability to re-route across entire original AS-path**
2. **Performance of original versus new paths**
3. **Real-world comparison with prior simulations**
4. Predicting who can re-route w/ BGP poisoning

---

5. **Propagating long poisoned paths**
6. **Filtering of certain poisoned ASes**
7. Filtering of long poisoned paths
8. Routing Working Groups behavior

---

9. Default route prevalence
10. Reachability of /25's

How well can an AS re-route with poisoning?

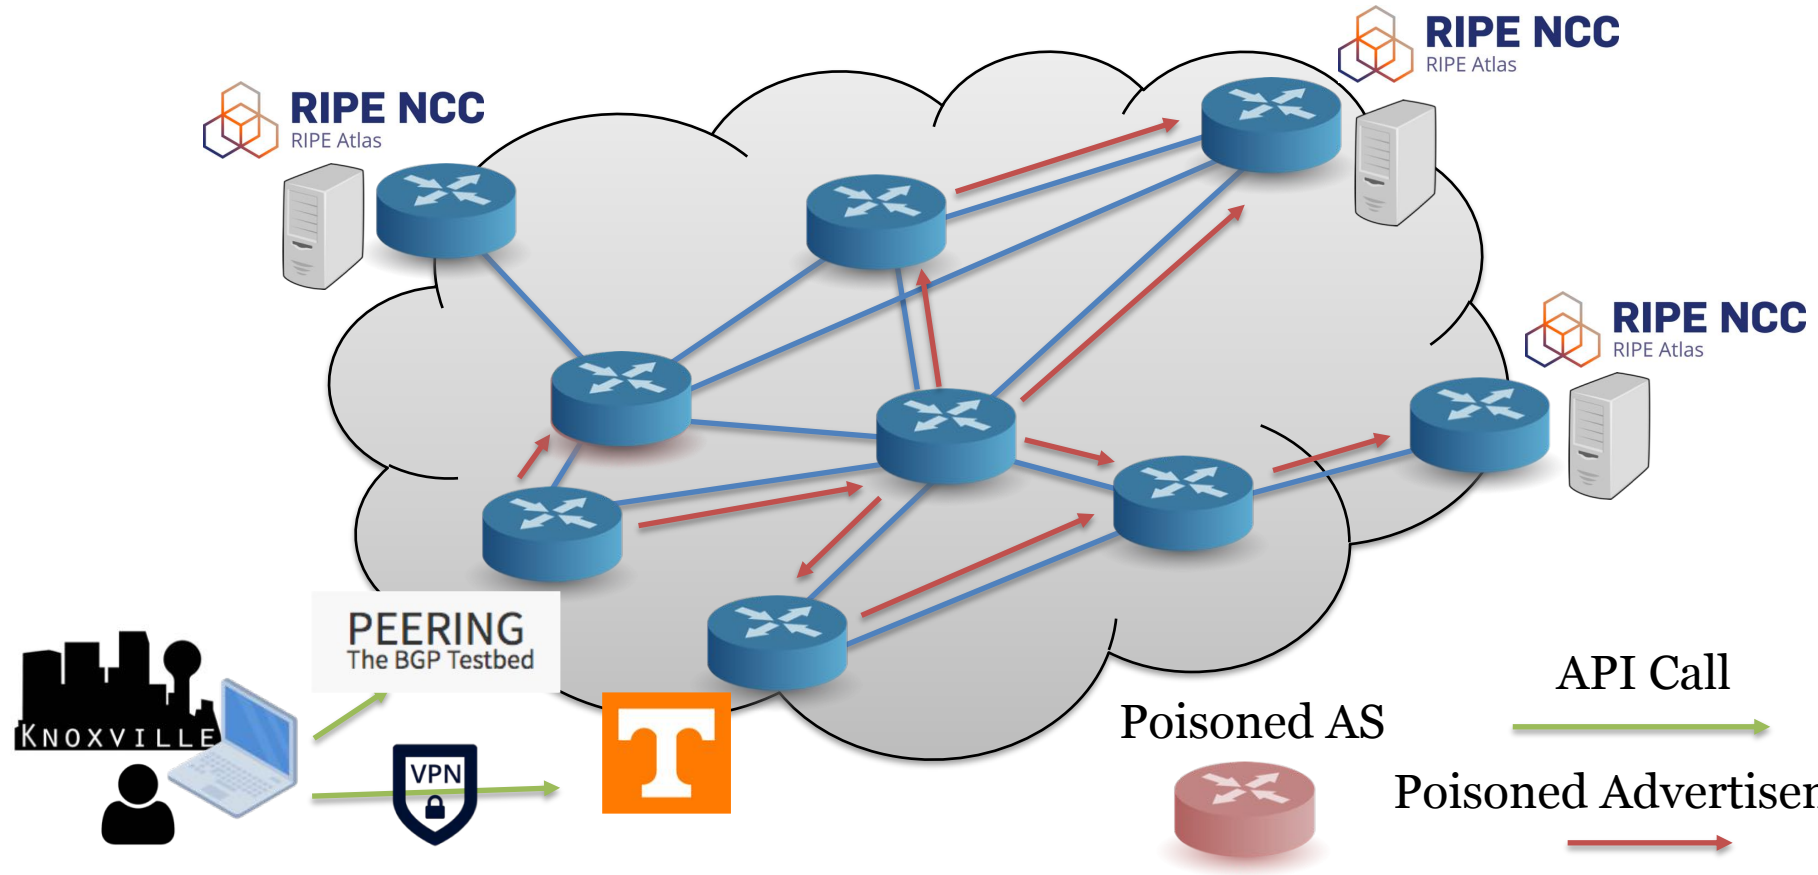

How well can an AS re-route with poisoning?

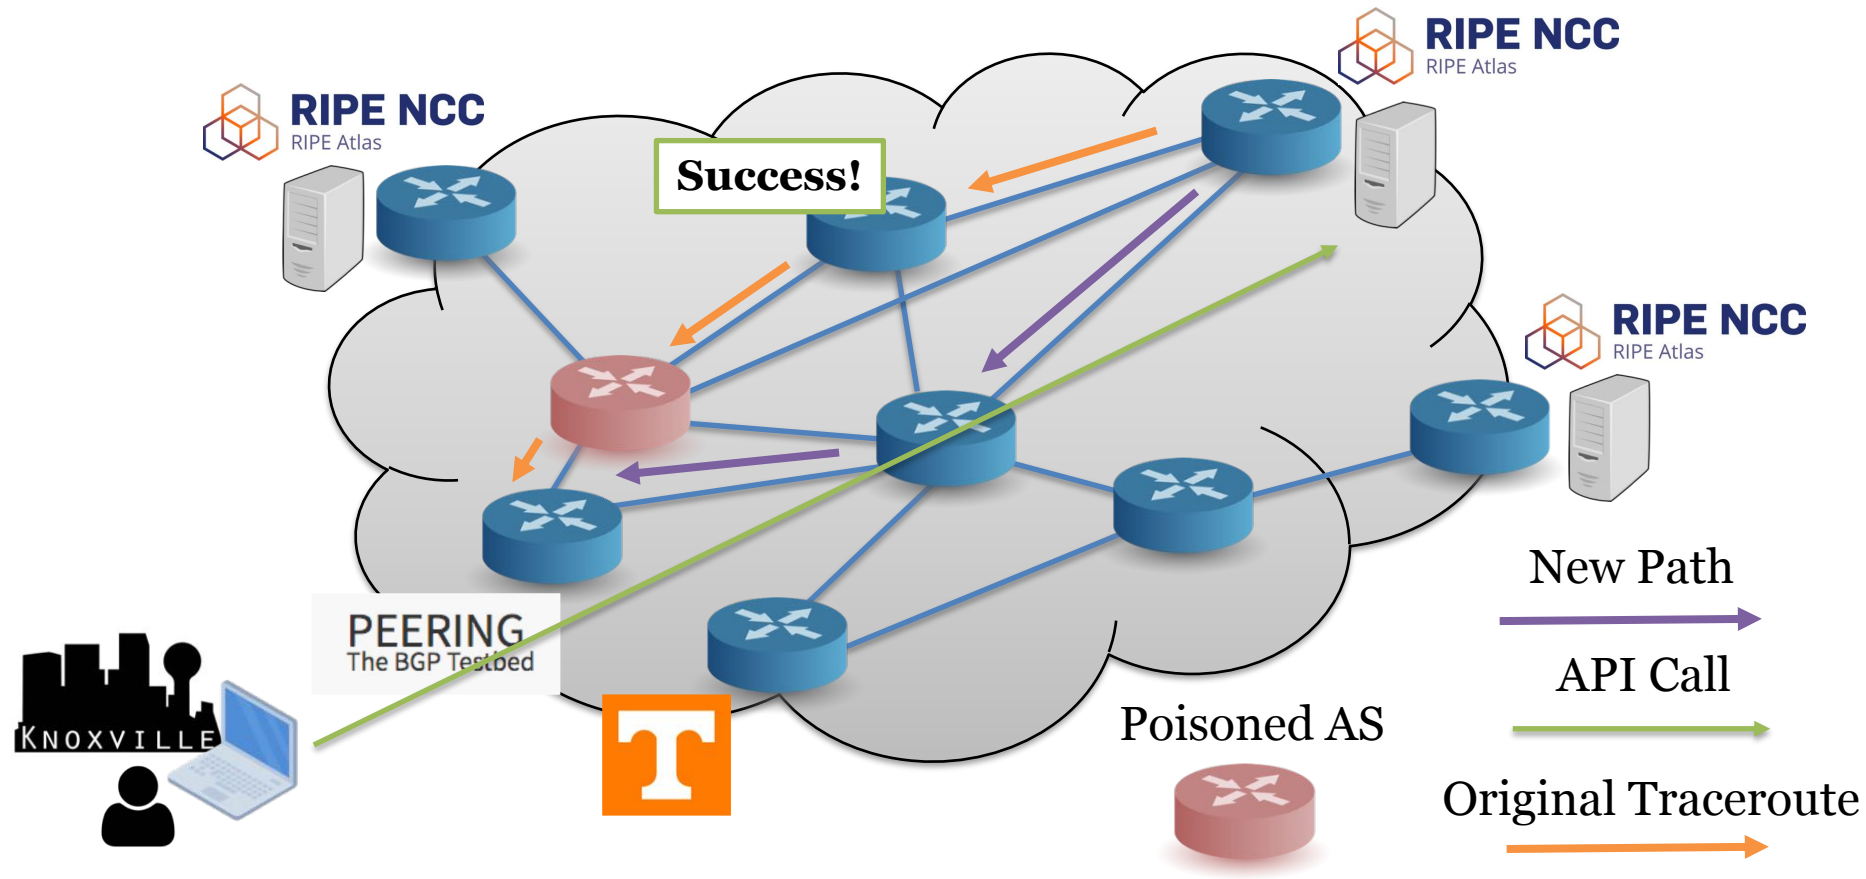

How well can an AS re-route with poisoning?

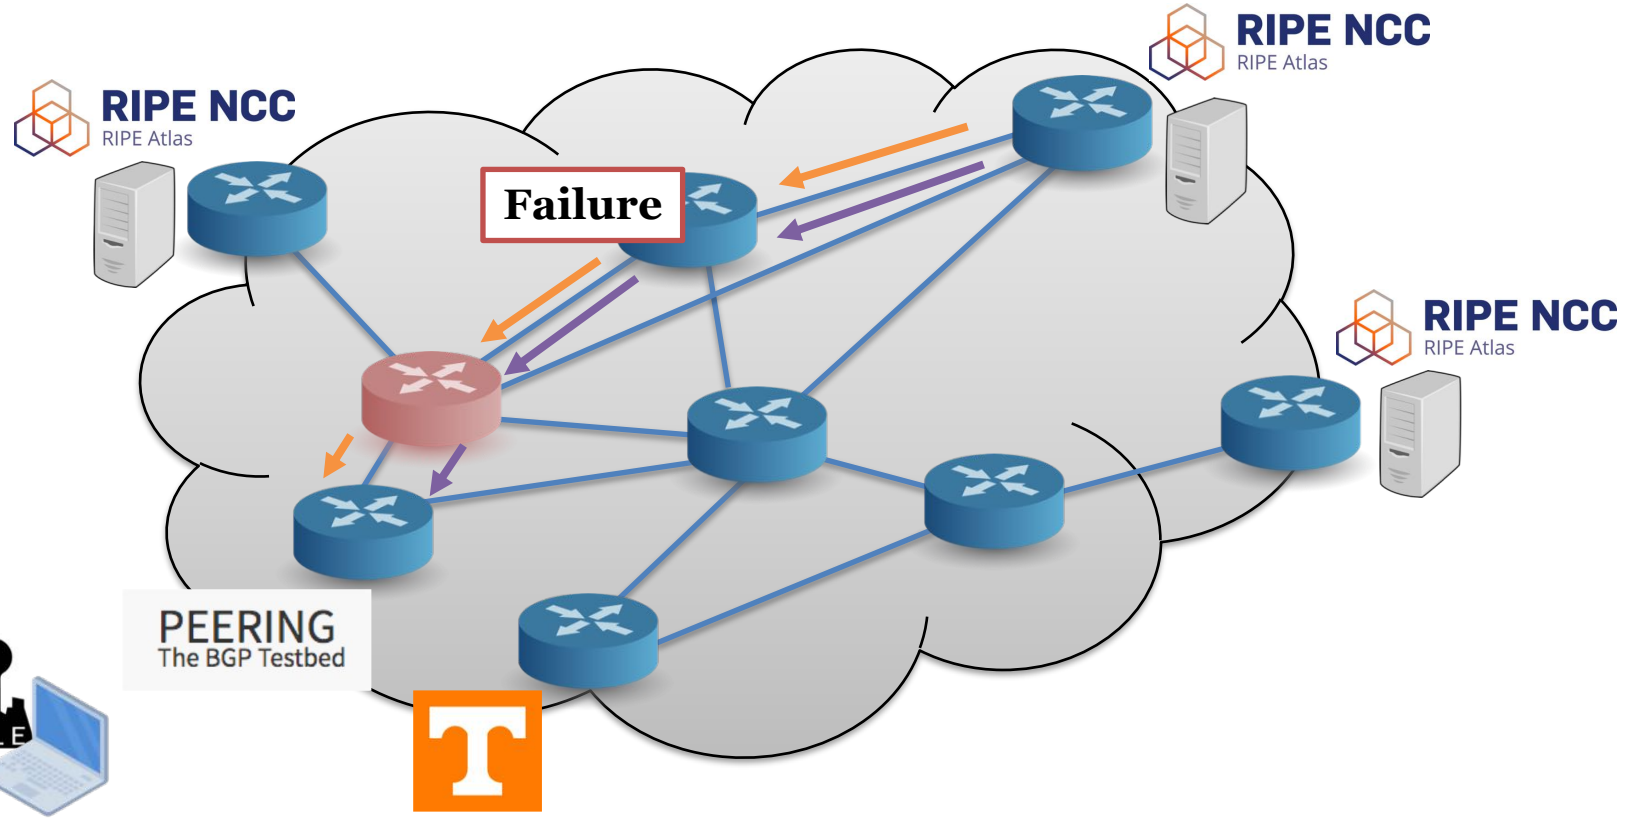

How well can an AS re-route with poisoning?

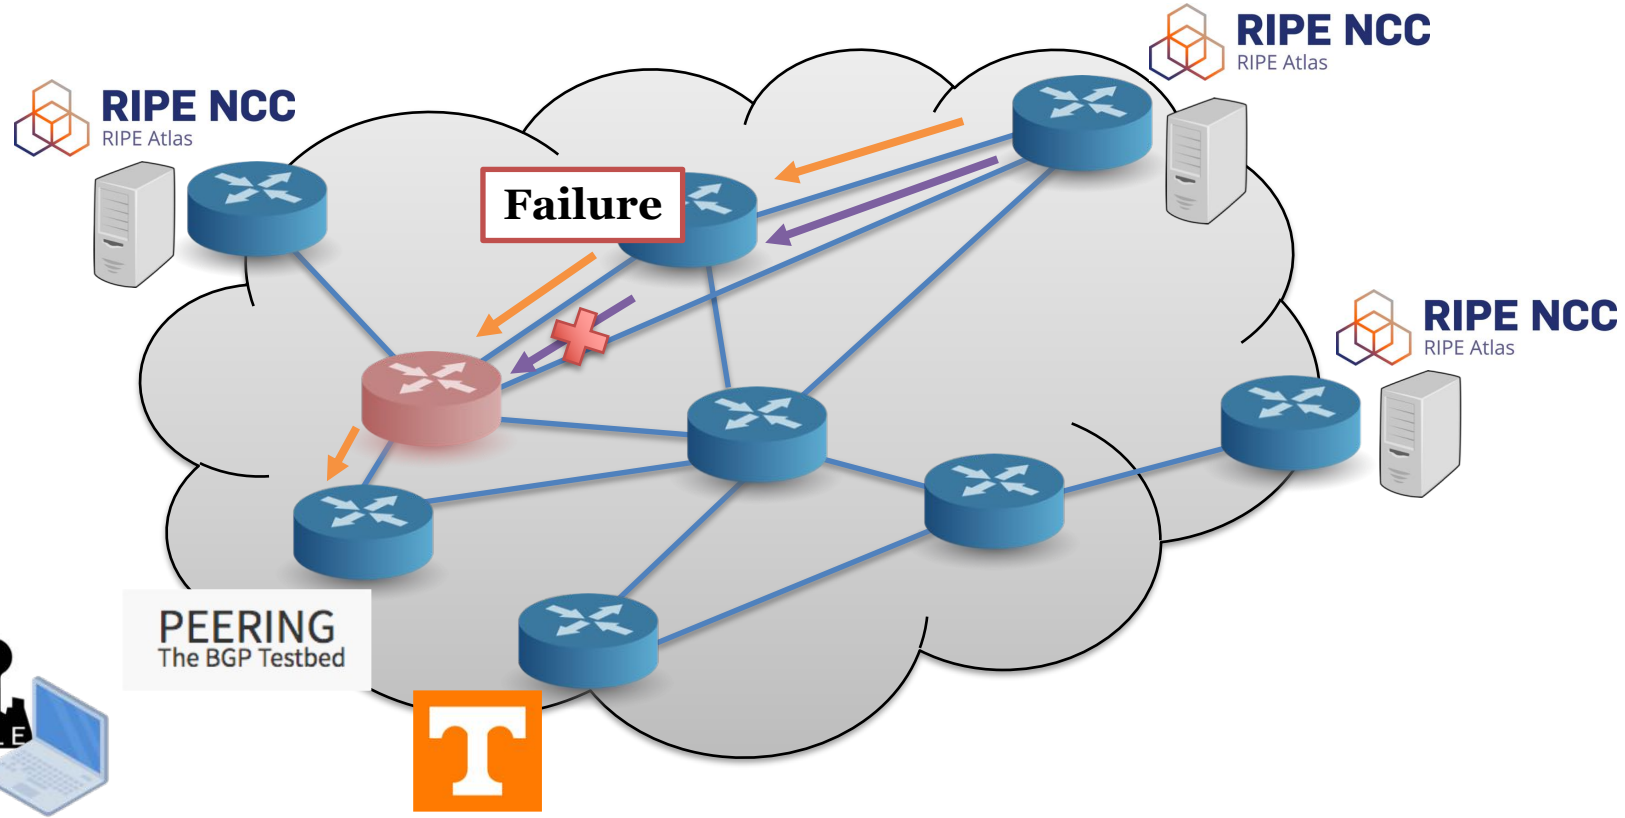

High-Level Findings

1,460/1,888 (77%)
successful cases of poisoning

6.45
avg. new ASes discovered

2.03 for 6.45
avg. poisons needed/avg. new ASes

2.25
avg. new paths discovered

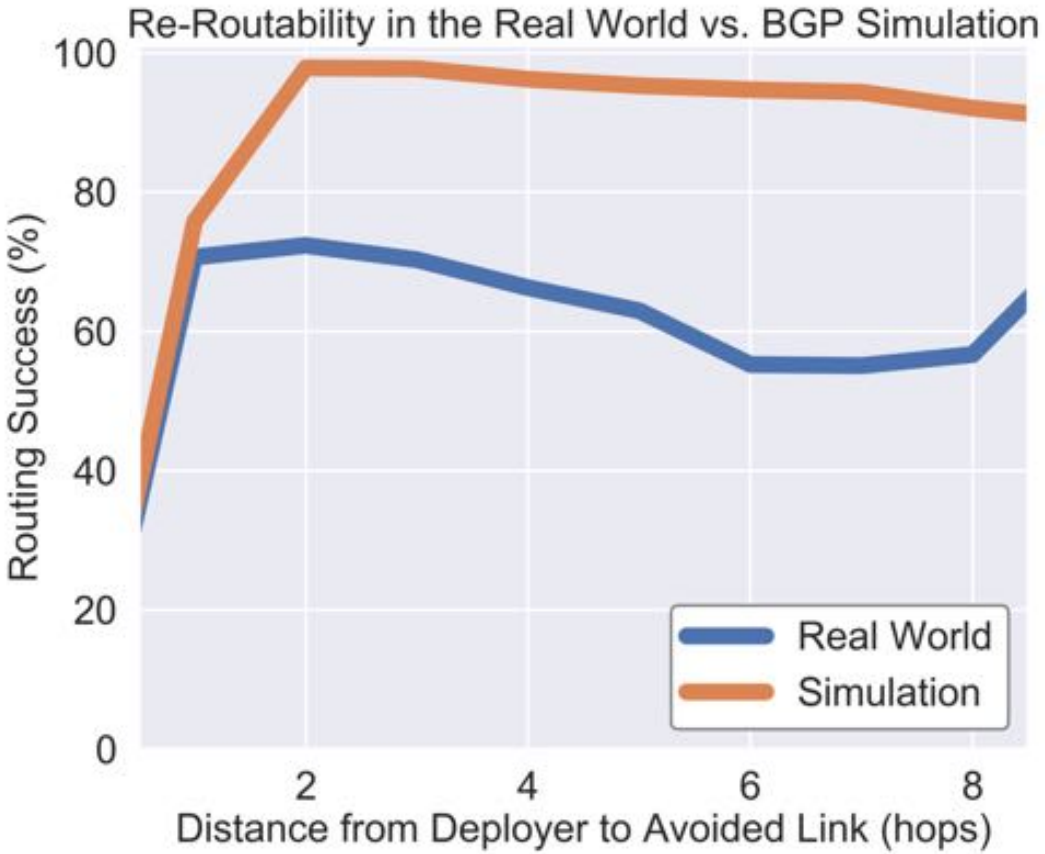
Re-Routability in the Real World vs. BGP Simulation
Routing Success (%)
100
80
60
40
20
0
2
4
6
8
Distance from Deployer to Avoided Link (hops)
Real World
Simulation

Security Implications

- Real-world evidence supports poisoning-enabled systems
- Security systems need to account for poisoning
- Success in simulation ***does not guarantee*** success in the real-world

Are alternate routes slower?

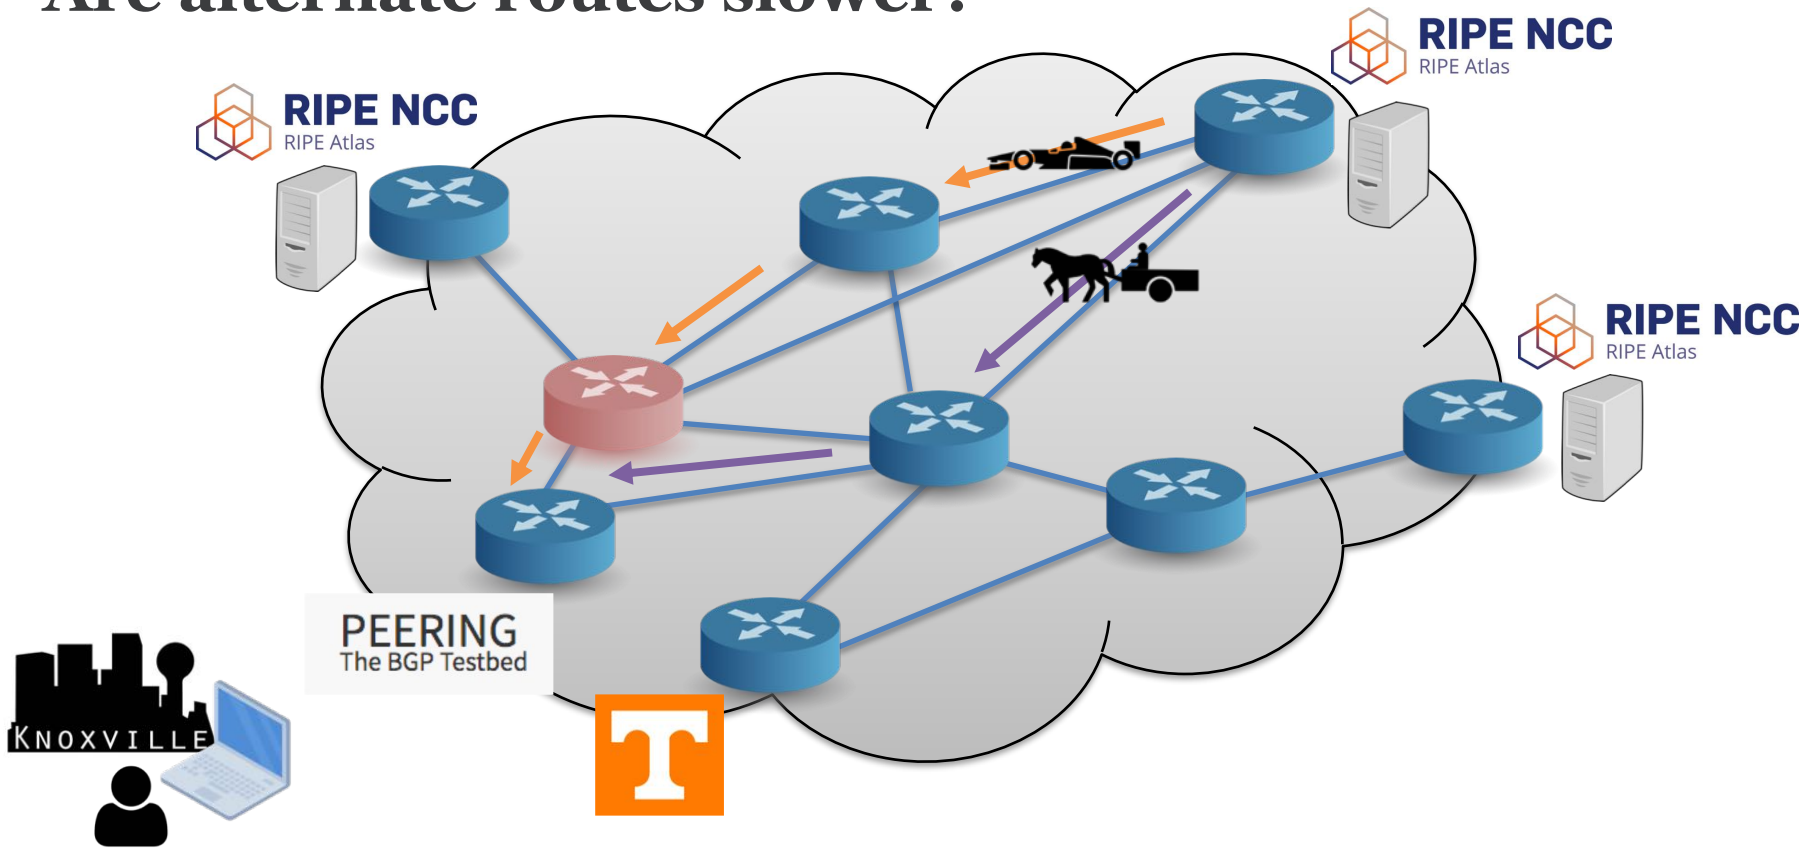

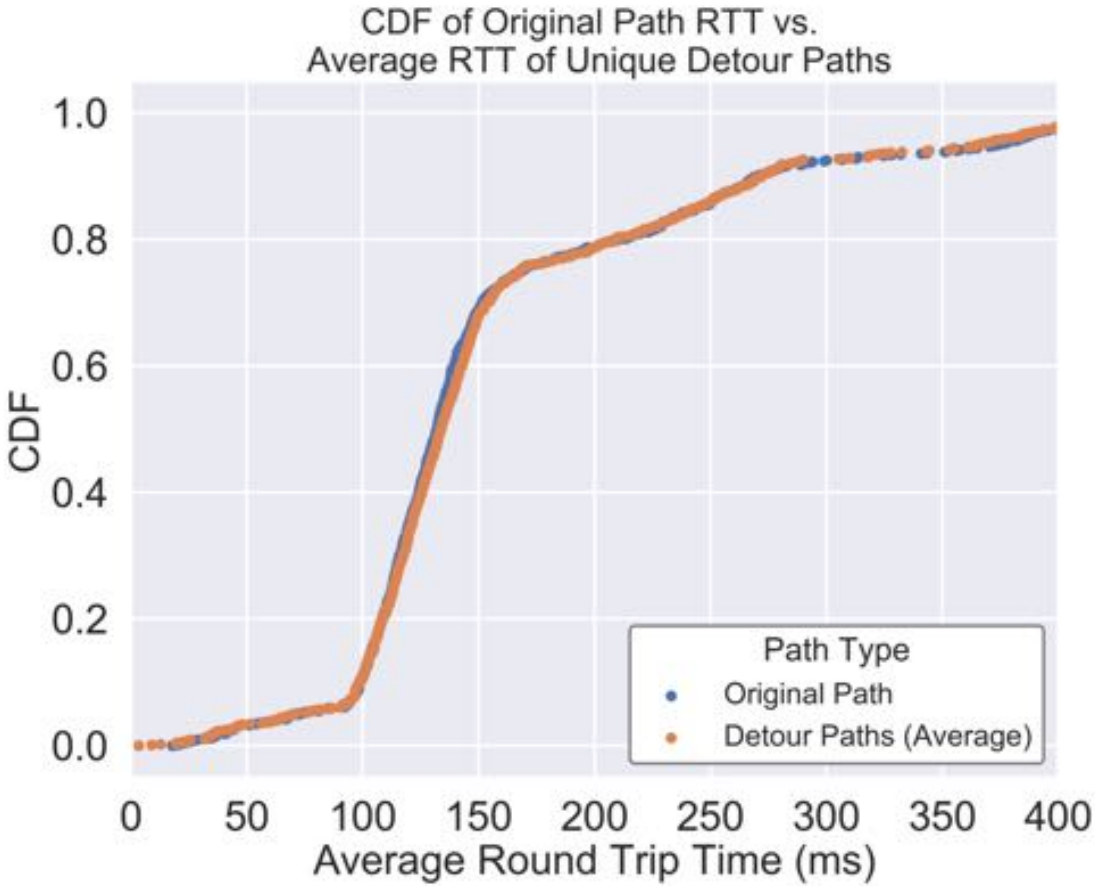
CDF of Original Path RTT vs.
Average RTT of Unique Detour Paths
CDF
1.0
0.8
0.6
0.4
0.2
0.0
0
50
100
150
200
250
300
350
400
Average Round Trip Time (ms)
Path Type
Original Path
Detour Paths (Average)

Security Implications

- Common logic suggests Internet paths not used by default would be less favorable
- Impacts the likelihood of operators deploying systems like Nyx

Are long paths filtered?

Baseline: 2 collectors saw path

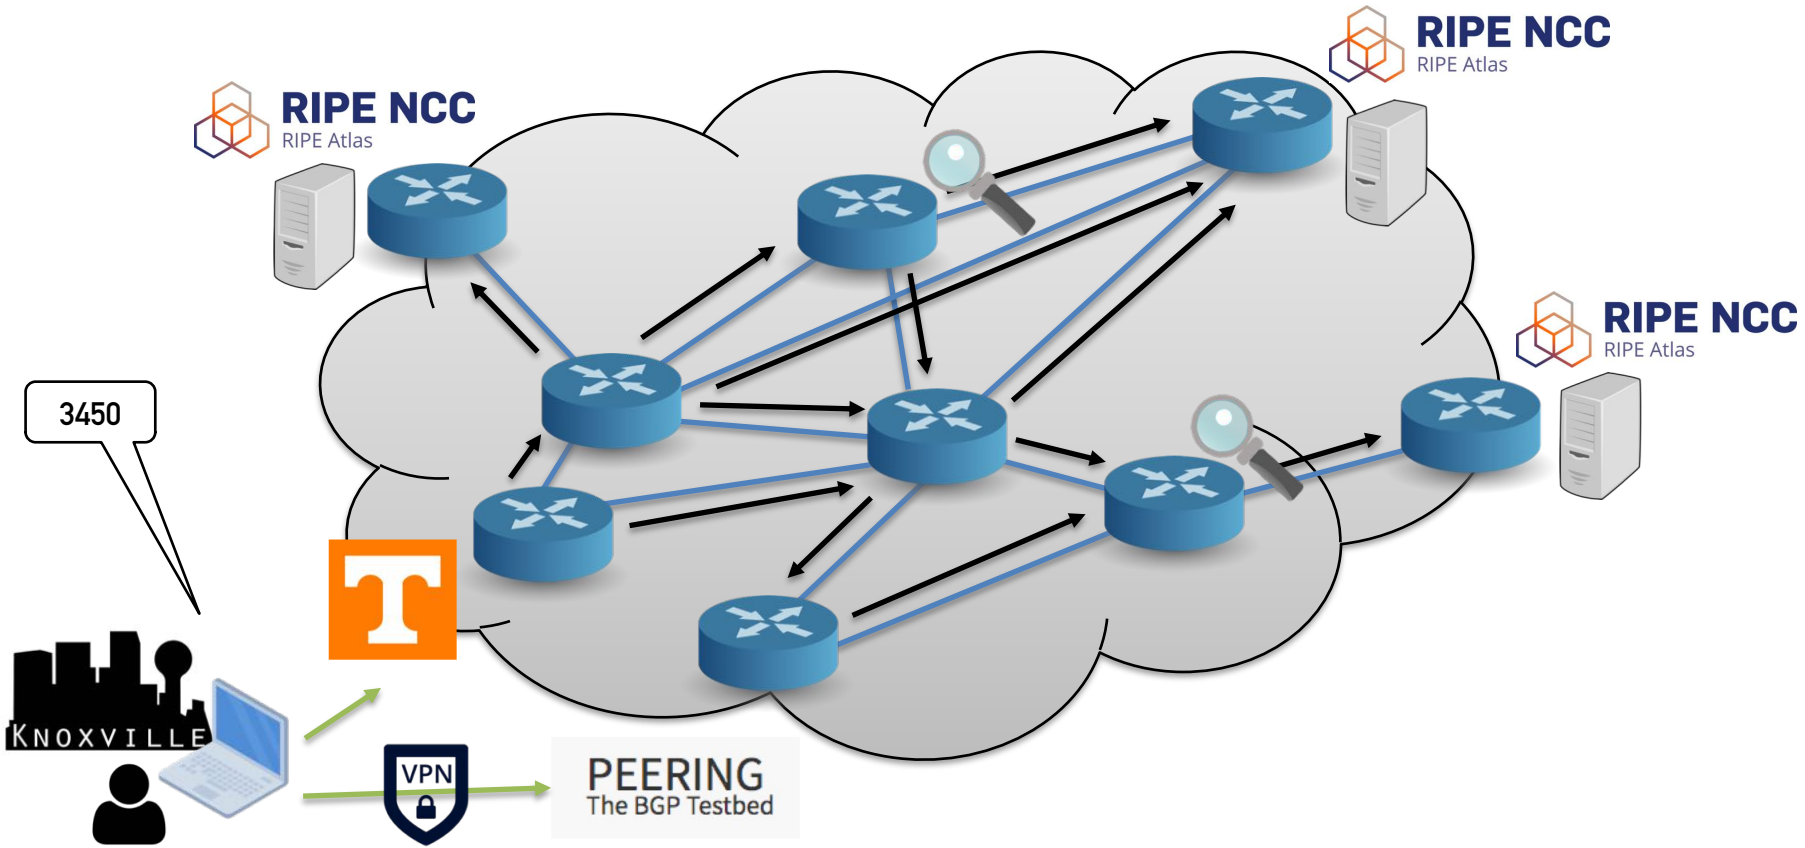

Are long paths filtered?

Long Path: 1/2 collectors saw path (50%)

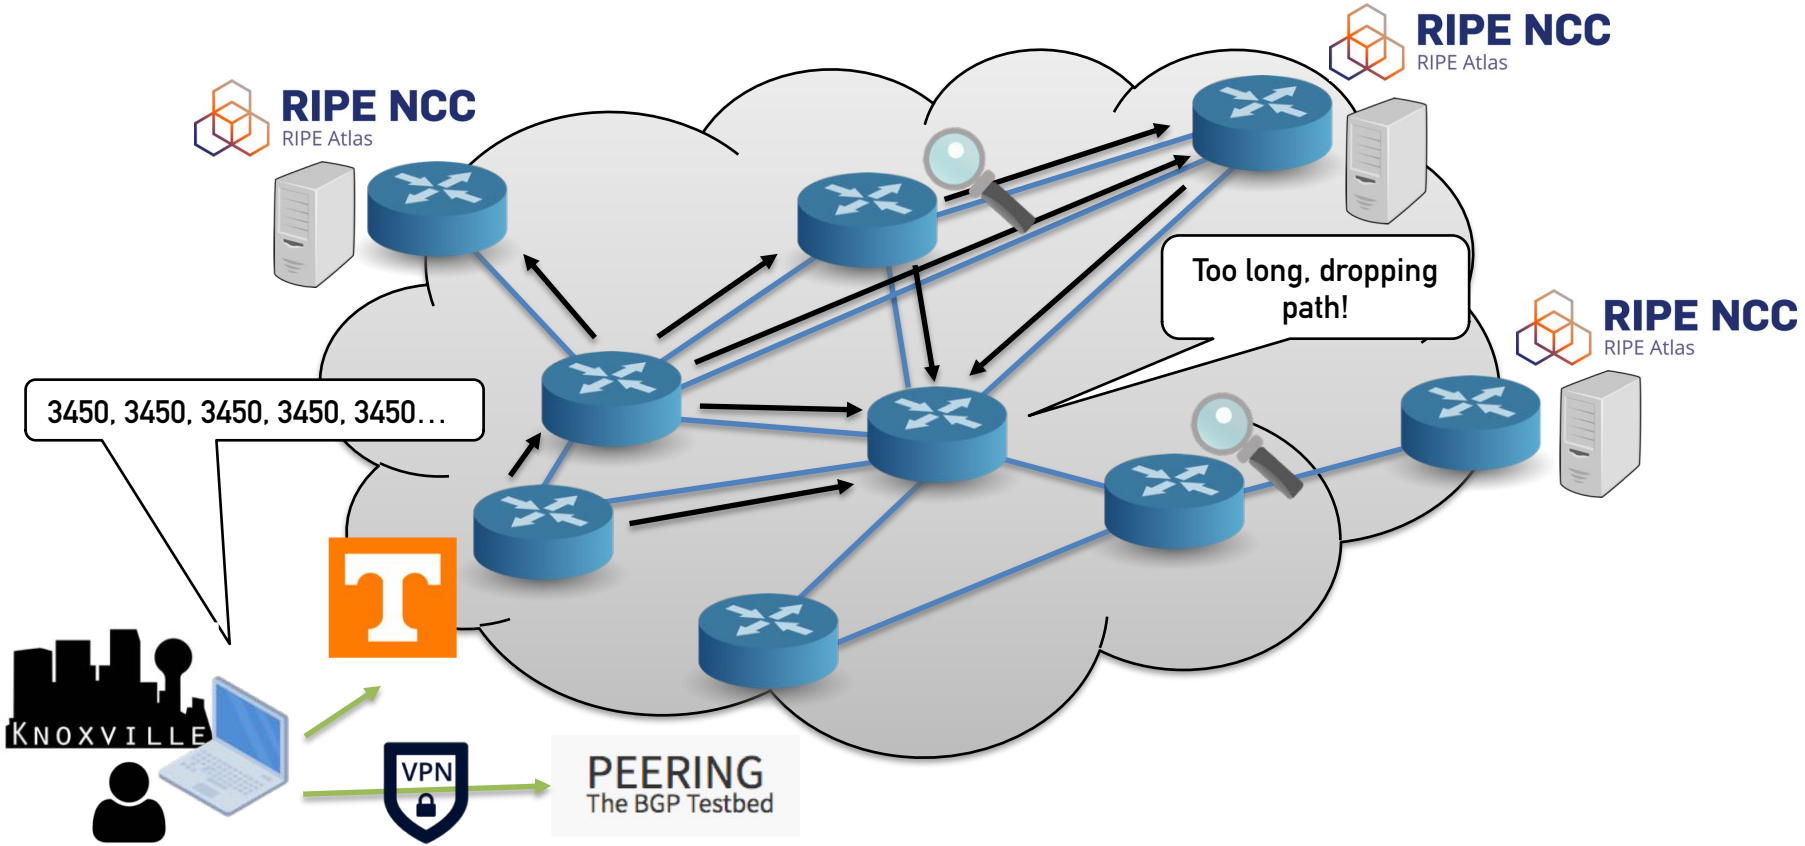

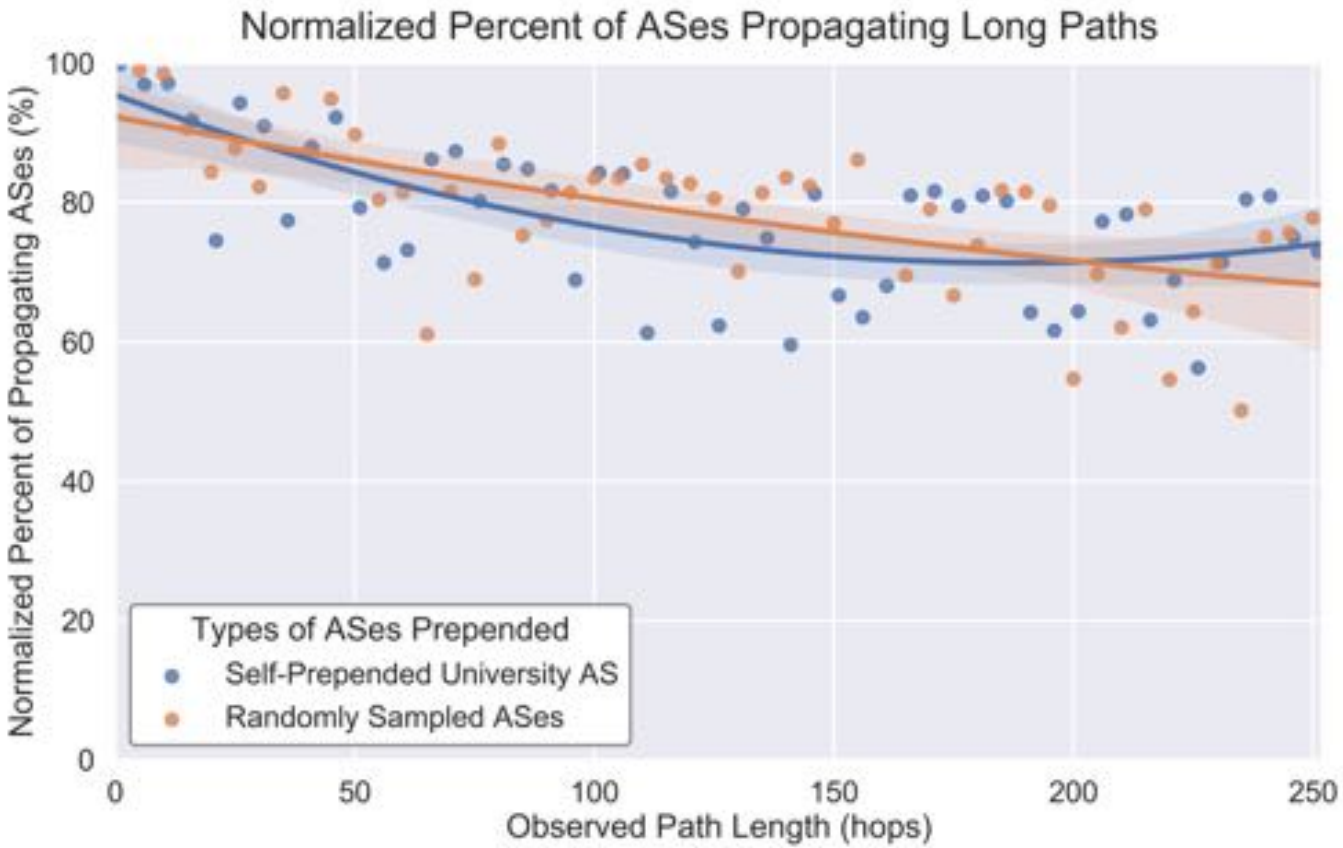
Normalized Percent of ASes Propagating Long Paths
Normalized Percent of Propagating ASes (%)
100
80
60
40
20
0
0
50
100
150
200
250
Observed Path Length (hops)
Types of ASes Prepended
Self-Prepended University AS
Randomly Sampled ASes

Security Implications

- Maximum AS path length of 255 needs to be accounted for in poisoning-enabled systems
- Network operator groups also claim they filter anomalous paths

Does the size of the poisoned AS affect filtering?

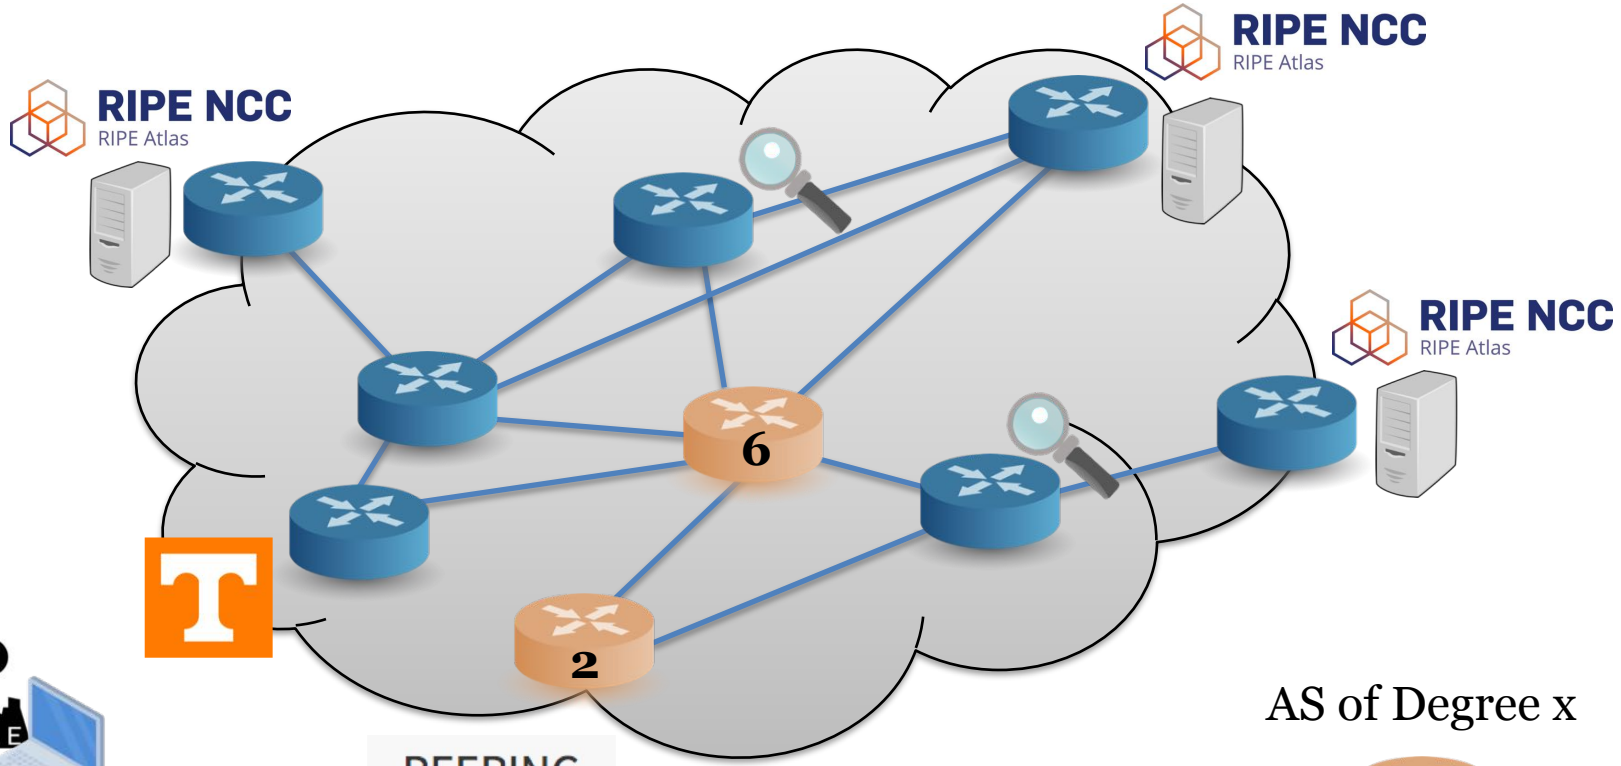

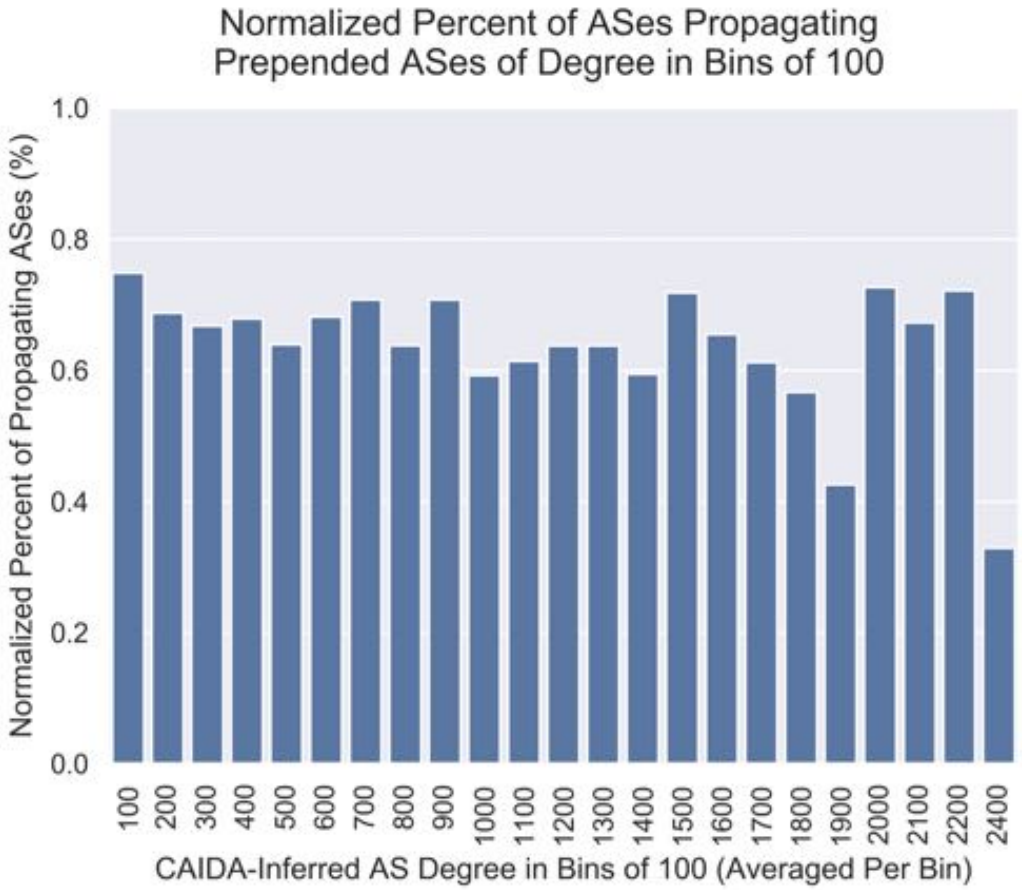
Normalized Percent of ASes Propagating
Prepended ASes of Degree in Bins of 100
Normalized Percent of Propagating ASes (%)
1.0
0.8
0.6
0.4
0.2
0.0
100
200
300
400
500
600
700
800
900
1000
1100
1200
1300
1400
1500
1600
1700
1800
1900
2000
2100
2200
2400
CAIDA-Inferred AS Degree in Bins of 100 (Averaged Per Bin)

Security Implications

- Common logic suggests operators may filter weird behavior
  - Filtering poisoned ASes that run the Internet → seems intuitive
  - Not filtering poisoned ASes that you do not often see in advertisements → also seems intuitive

Diverging Claims

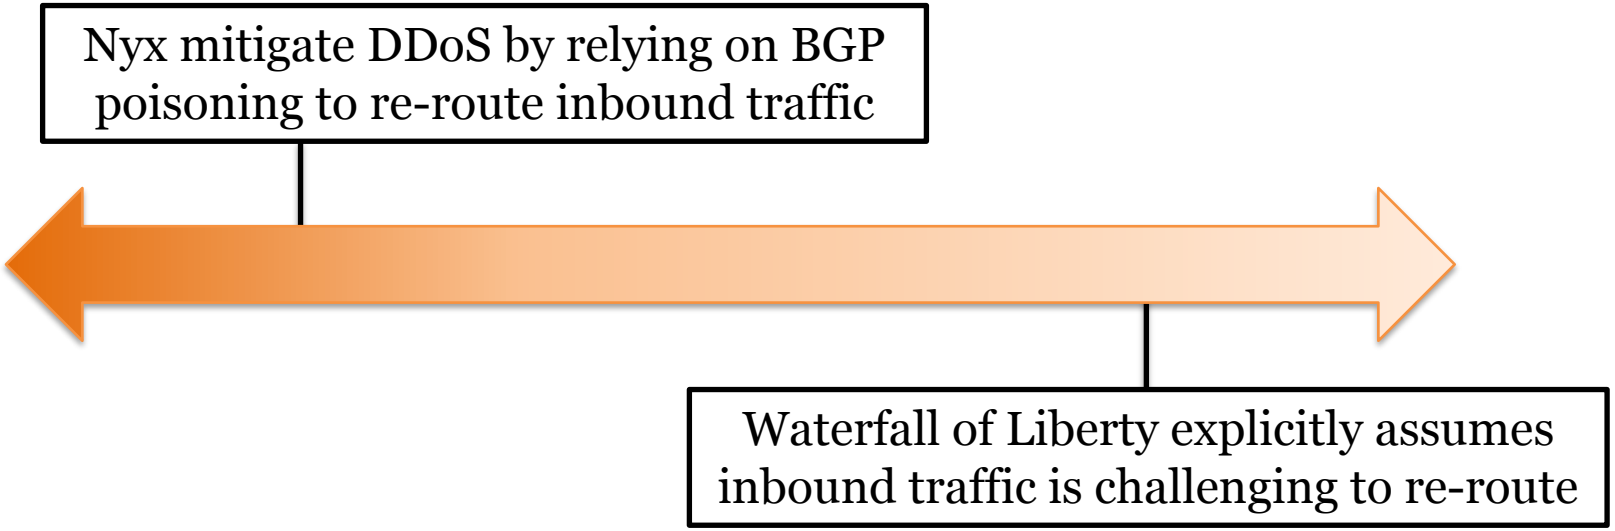

Diverging Claims

Nyx mitigate DDoS by relying on BGP poisoning to re-route inbound traffic

Yet, *Nyx* and *Waterfall of Liberty* **can both work in practice.**

Waterfall of Liberty explicitly assumes inbound traffic is challenging to re-route

We should publish and disseminate our work **after we have tested** our assumptions **in the same environment** where we intend to deploy our work.

Conclusion

- BGP poisoning **works in most cases**
- Systems which assume the opposite **can still deploy** in areas where poisoning is harder
- Common logic of Internet behavior is **not always accurate**
- All Internet security research should be **actively tested on the Internet** if the research targets the Internet for deployment

University of Tennessee
VOLSEC
COMPUTER SECURITY LAB

Jared M. Smith

Twitter: *jaredthecoder*

Email: *jms@vols.utk.edu*

Web: *volsec.org*

BACKUP

RPKI During Poisoning

$$\{ AS_{orig},\ AS_{BL_1},\ AS_{BL_2}, \ldots,\ AS_{BL_N}, \underbrace{AS_{orig}}_{\text{For RPKI}} \} \quad (1)$$

$$\{ \underbrace{AS_3,\ AS_2,\ AS_1}_{\text{Actual Path}},\ \overbrace{AS_{orig}}^{\text{Packet at Dest}}\ , \underbrace{AS_{BL_1}, \ldots,\ AS_{BL_N}}_{\text{Irrelevant for Forwarding}} \} \quad (2)$$

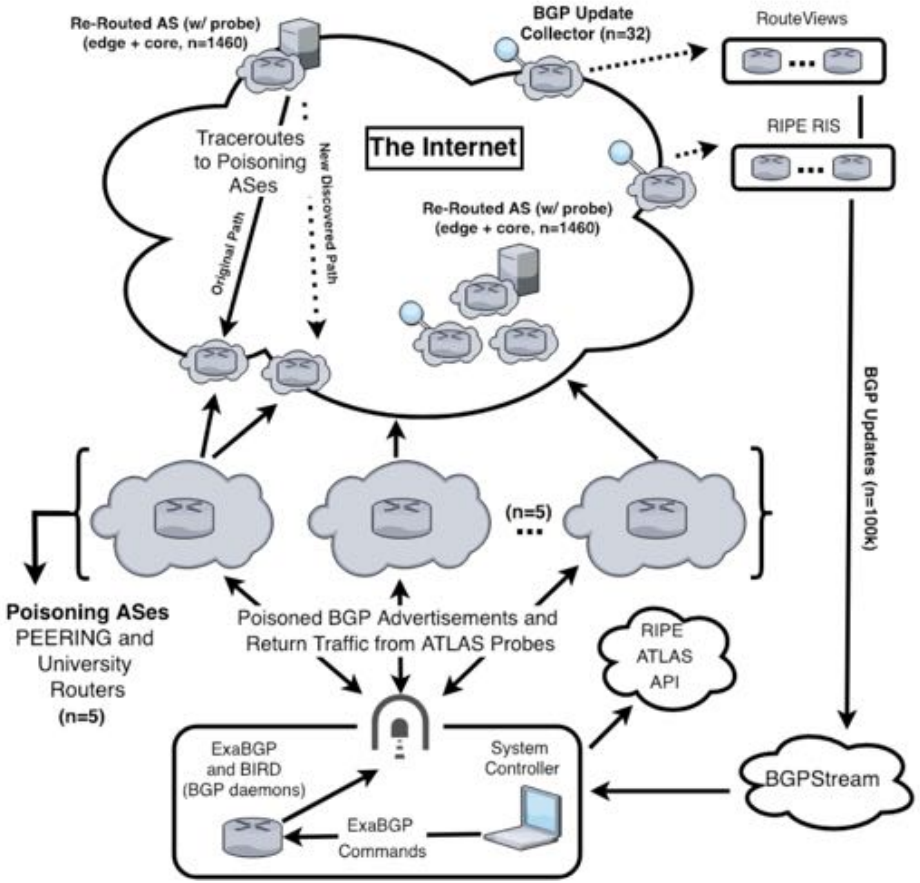
Re-Routed AS (w/ probe)
(edge + core, n=1460)
BGP Update
Collector (n=32)
RouteViews
RIPE RIS
The Internet
Traceroutes
to Poisoning
ASes
Original Path
New Discovered Path
Re-Routed AS (w/ probe)
(edge + core, n=1460)
BGP Updates (n=100k)
(n=5)
Poisoning ASes
PEERING and
University
Routers
(n=5)
Poisoned BGP Advertisements and
Return Traffic from ATLAS Probes
RIPE
ATLAS
API
ExaBGP
and BIRD
(BGP daemons)
System
Controller
ExaBGP
Commands
BGPStream

Infrastructure Numbers

Infrastructure	Source
5 BGP routers	PEERING and UT
8 IP prefixes	PEERING and UT
5,000+ distinct vantage points	RIPE ATLAS
3 countries	US, Amsterdam, Brazil
32 BGP collectors	CAIDA BGPStream*

*Collects BGP Updates from RouteViews and RIPE RIS

How feasible is re-routing with BGP poisoning?

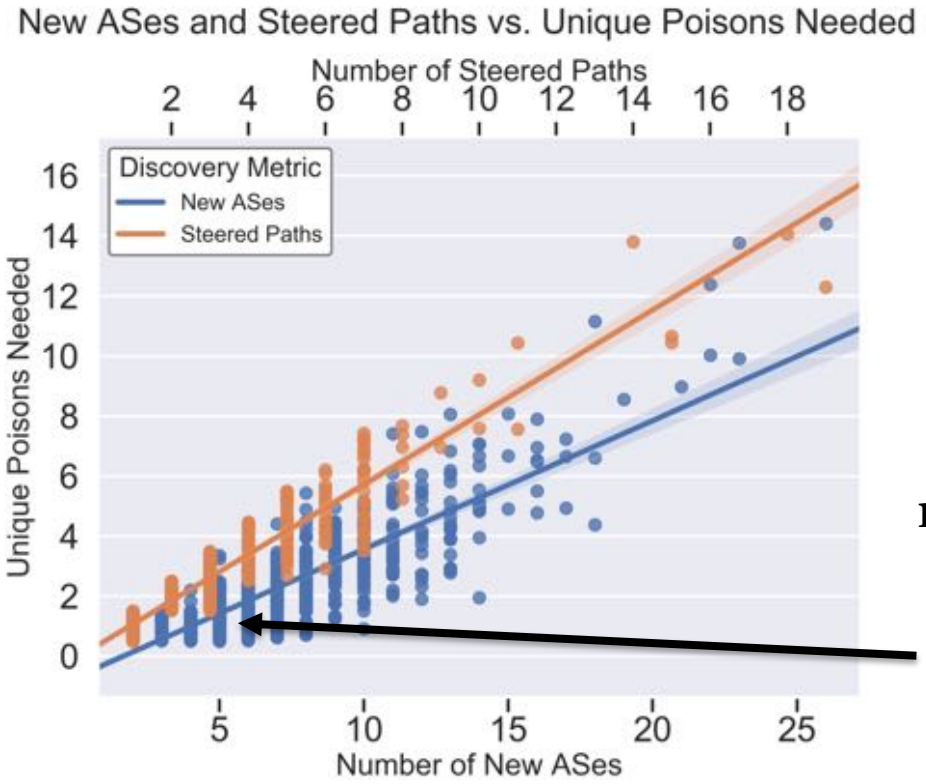

In practice, possible to re-route onto ~2.5 new alternate paths on average

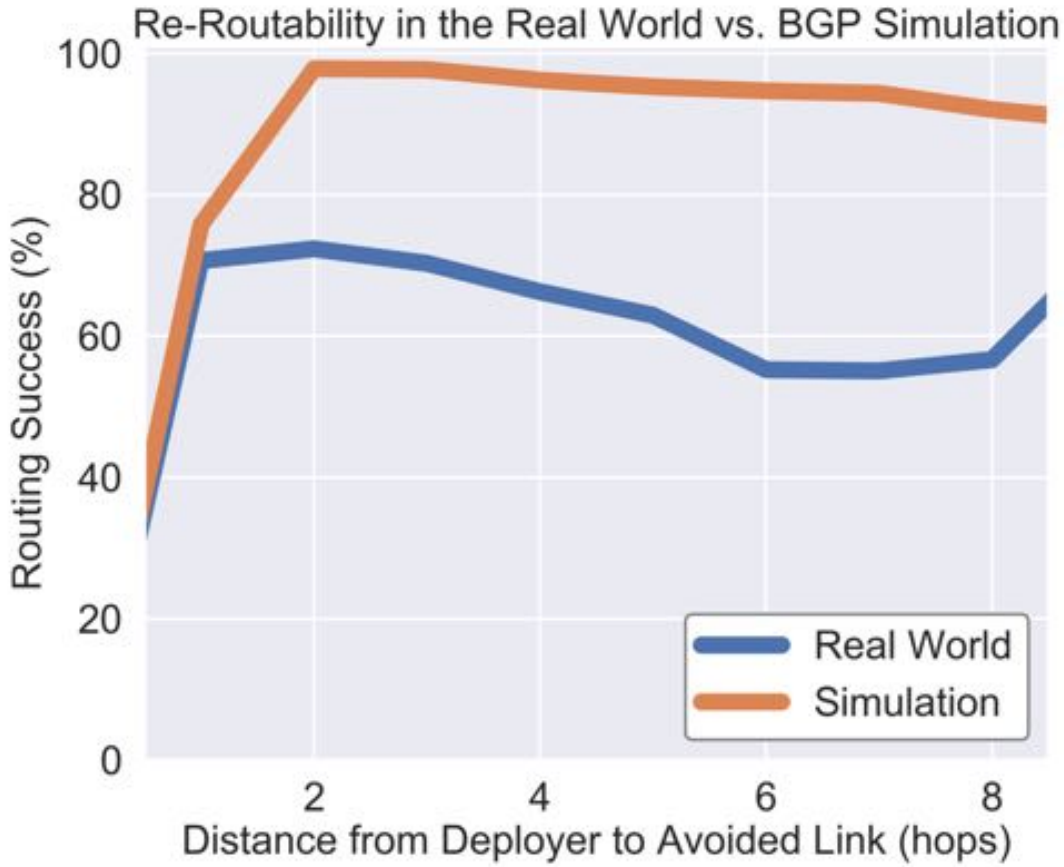
Re-Routability in the Real World vs. BGP Simulation
Routing Success (%)
100
80
60
40
20
0
2
4
6
8
Distance from Deployer to Avoided Link (hops)
Real World
Simulation

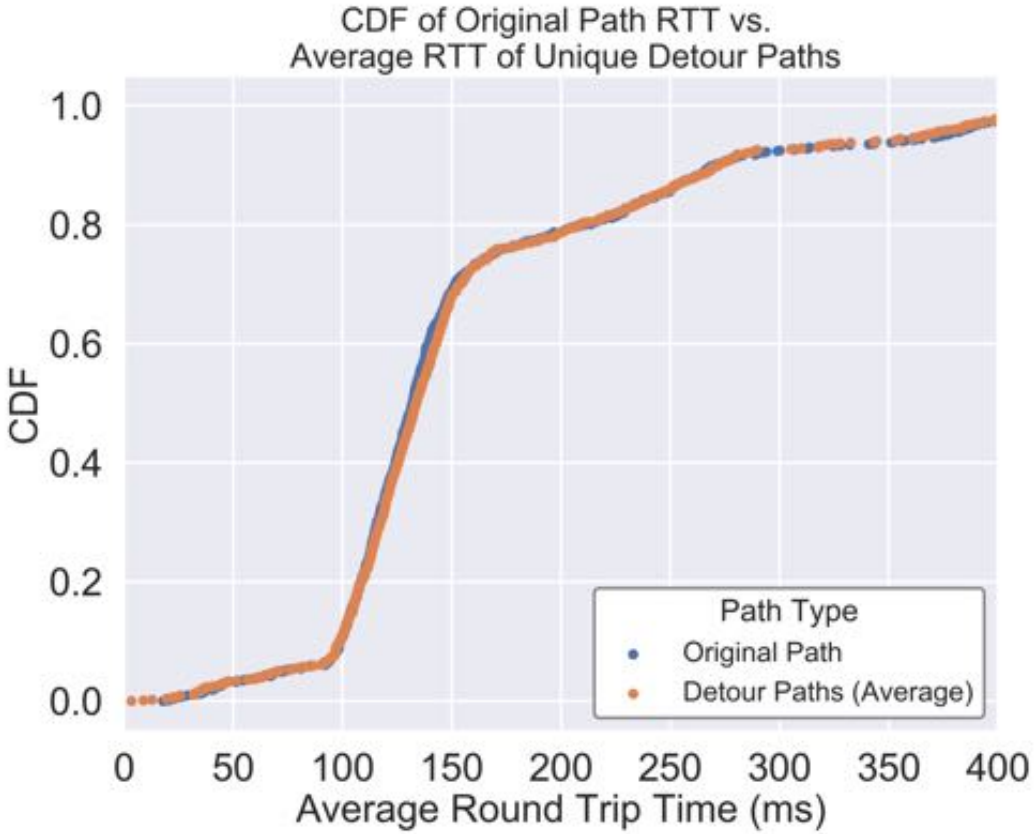
CDF of Original Path RTT vs.
Average RTT of Unique Detour Paths
1.0
0.8
0.6
0.4
0.2
0.0
CDF
0
50
100
150
200
250
300
350
400
Average Round Trip Time (ms)
Path Type
Original Path
Detour Paths (Average)

Graph-Theoretic Analysis of Return Paths

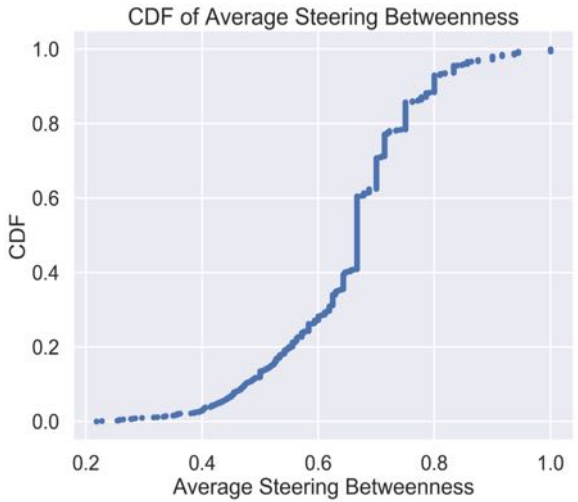

- Avg. Betweenness of 0.667
- Paths are not completely identical
- There is *some* diversity, but bottlenecks exist

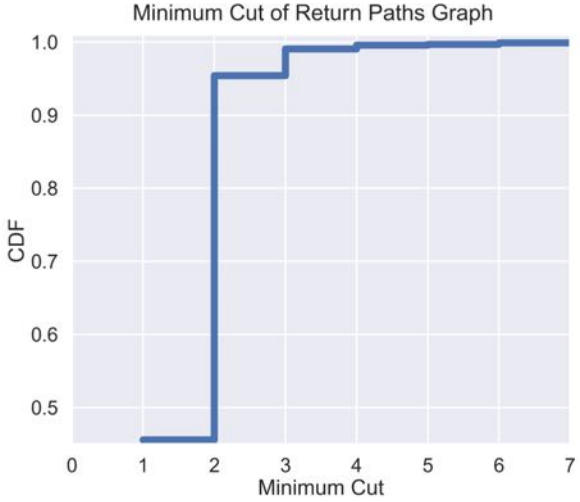

- Low min. cut means bottlenecks that Nyx/RAD cannot avoid
- For 90% of links, a bottleneck of at most 2 links occurs

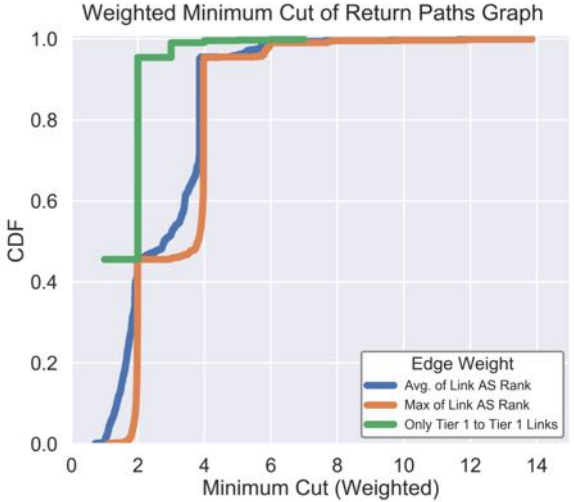

- Tier 1 ASes with inf. weight → bottlenecks **not** result of single unavoidable provider
- Within unweighted min cut → widely differing barriers to cut based on bandwidth

How well can we predict success with FRRP?

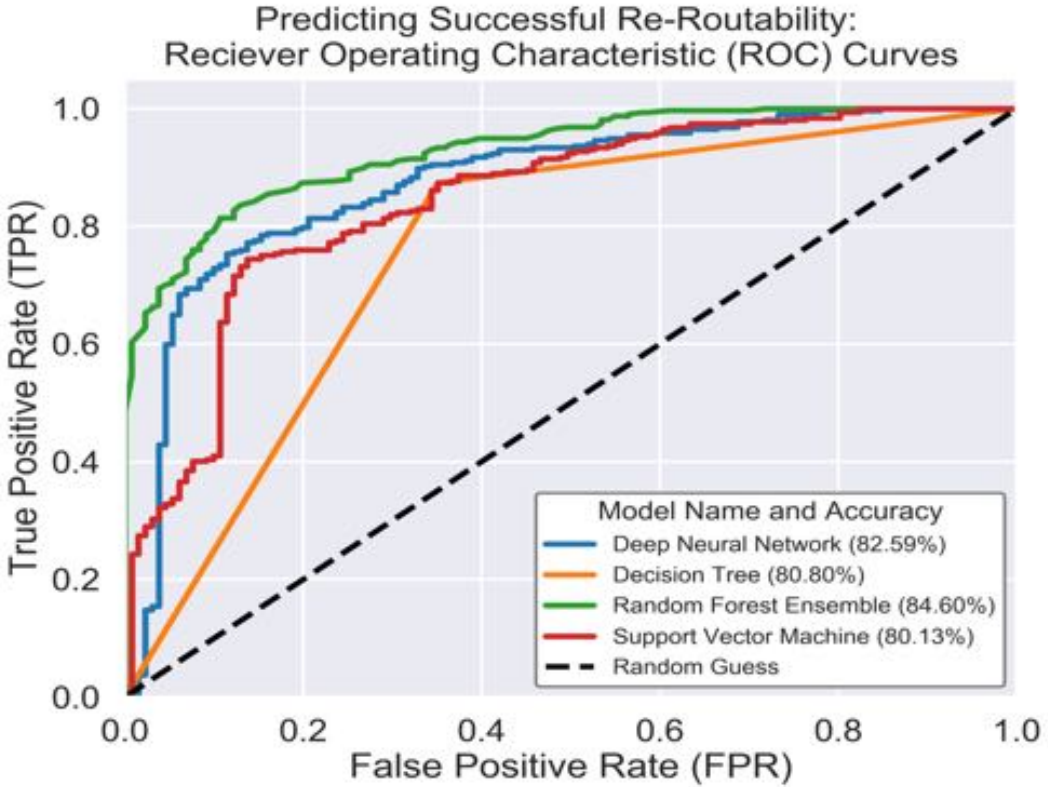

What link and AS properties are important for FRRP?

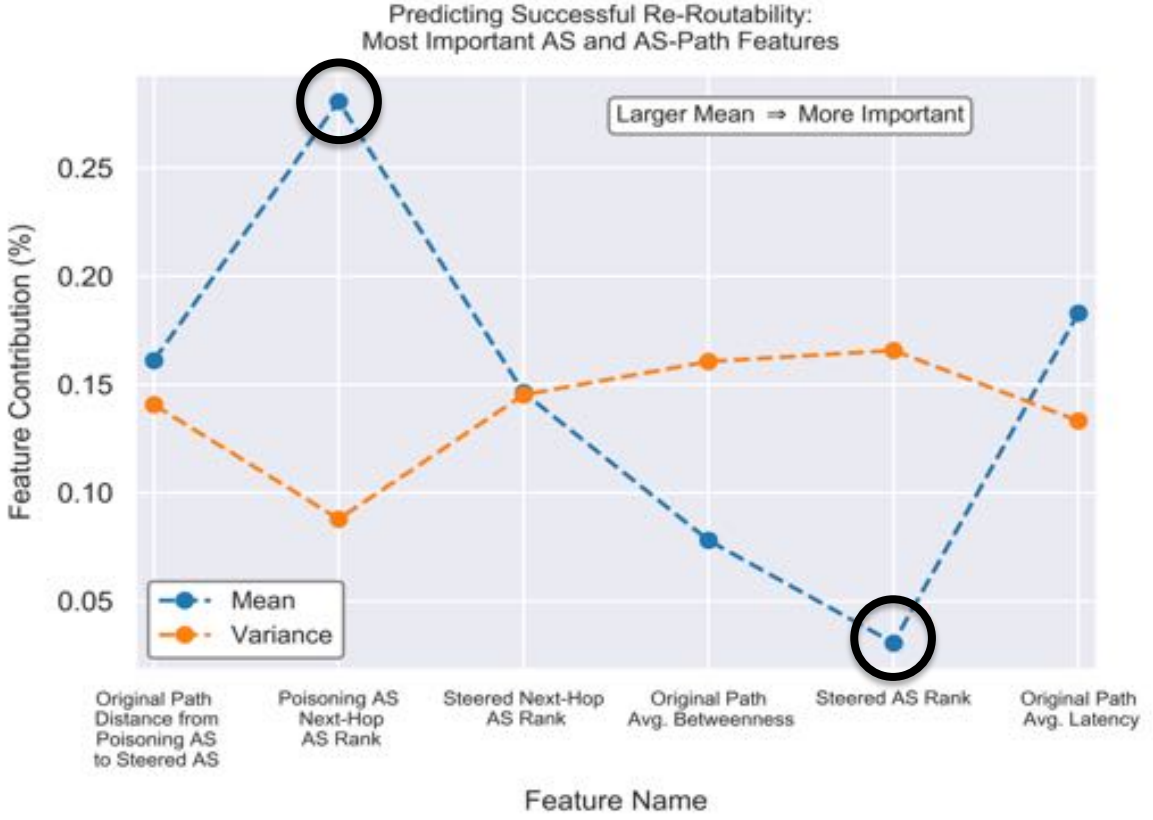

A Deeper Look at the Most Important Feature

Poisoning AS Next-Hop AS Rank

High Rank Matters

How long can (poisoned) paths be?

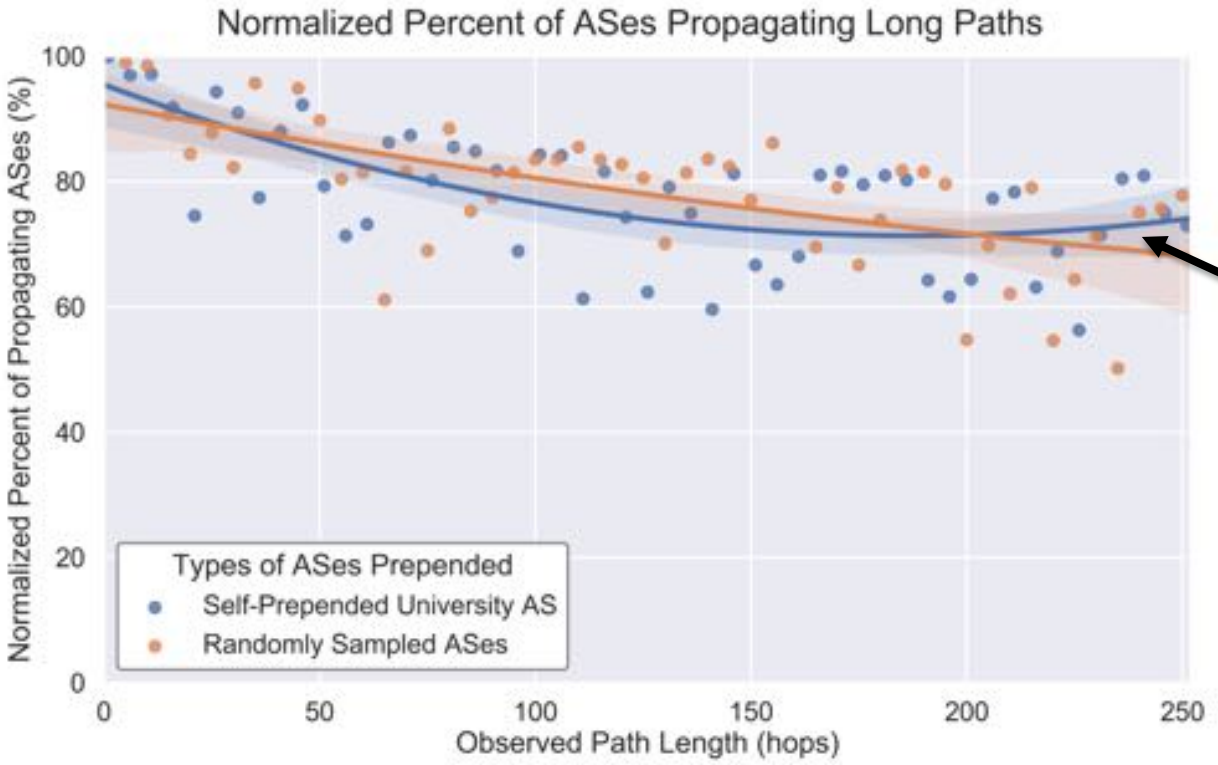

Propagation to 99% of the Internet at 250 AS-path length

How much do large ASes filter poisoned paths?

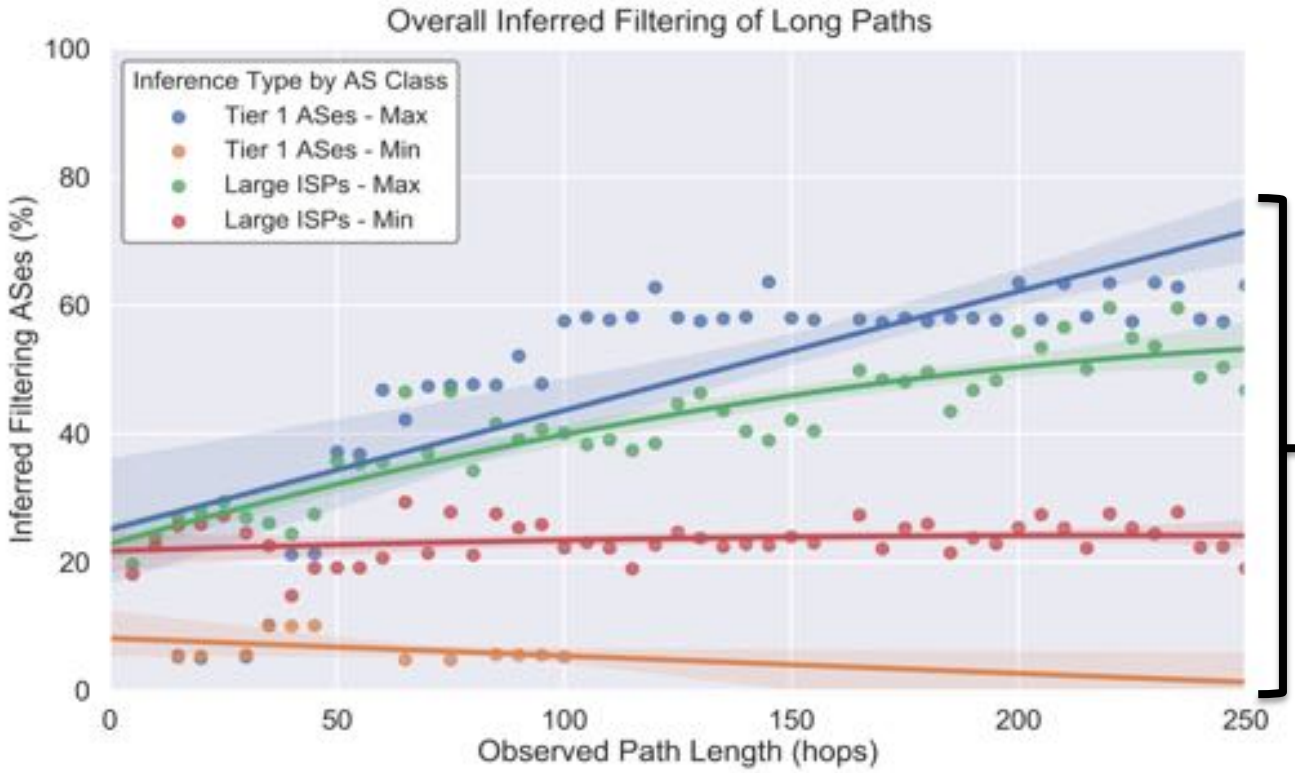

How much do small ASes filter poisoned paths?

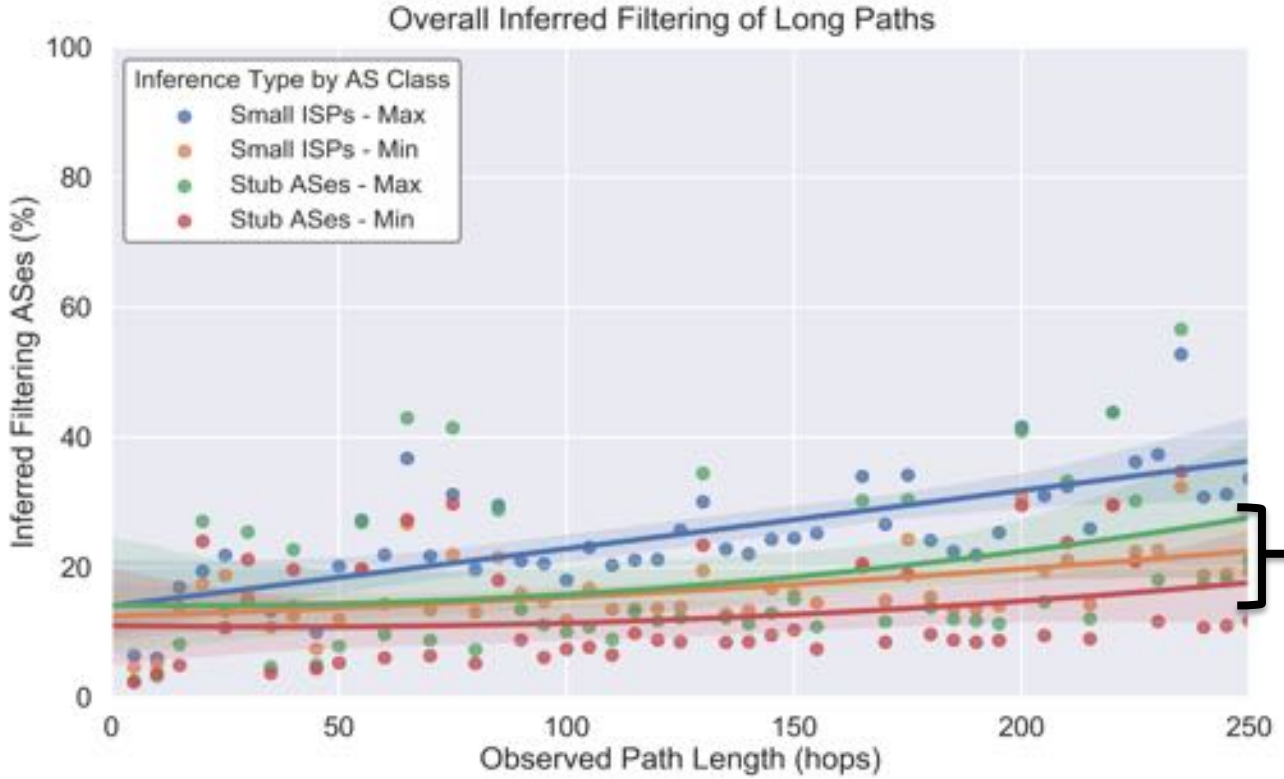

Do the Policy Leaders "*Walk the Walk*"?

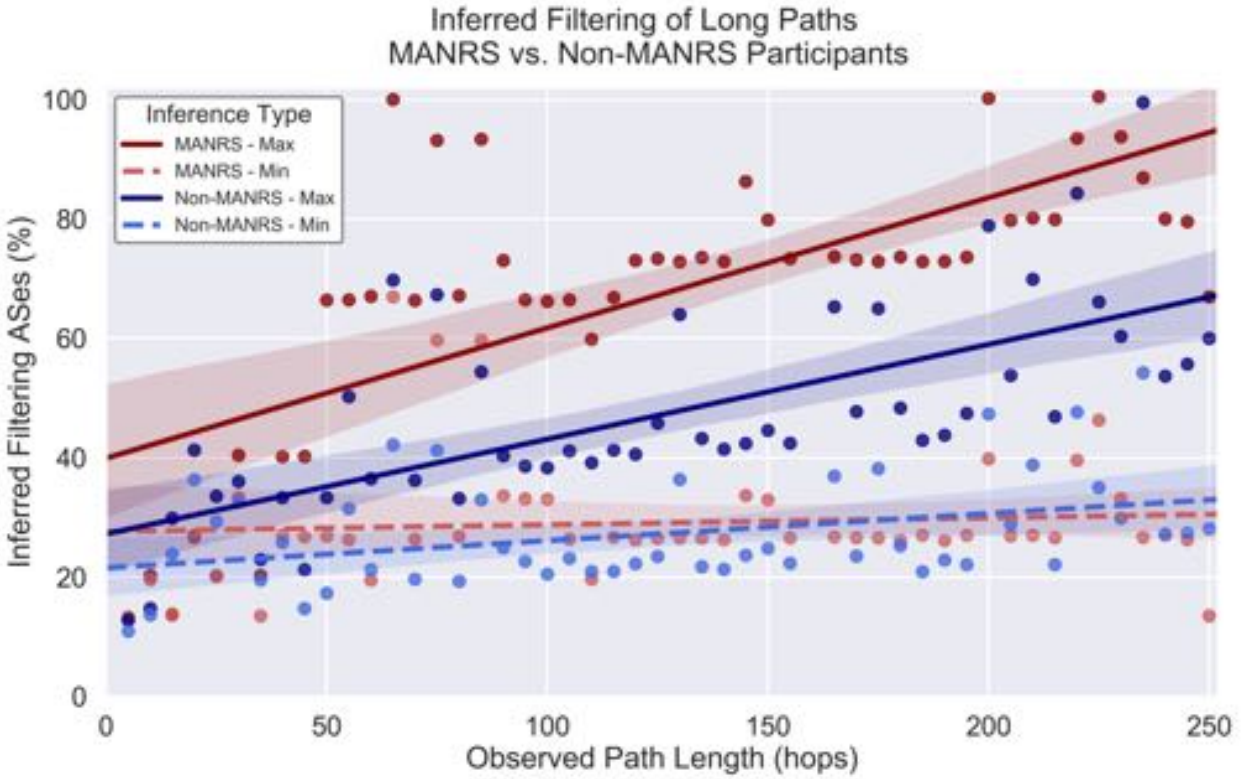

MANRS

"Mutually Agreed Norms for Routing Security"

Selected Participants (total=146):

- CenturyLink
- Charter
- Cogent
- Google
- Indiana U.
- …

Does AS-Degree of the Poisoned AS affect Filtering?

$Origin_{AS}$ $HighDegree_{AS}$ $Origin_{AS}$

...(in increments of 5)...

$Origin_{AS}$ $SmallDegree_{AS}$ $Origin_{AS}$

Rank by Degree	ASN and Name	Degree	Number of Customers	Registered Country by ASN	Normalized Propagation Percentage
1	6939 - Hurricane Electric	7064	1202	United States	11.9%
2	174 - Cogent	5352	5272	United States	11.6%
3	3356 - Level 3	4980	4898	United States	11.6%
4	24482 - SG.GS	3382	24	Singapore	96.1%
5	3549 - Level 3 GBLX	2538	2446	Unites States	11.6%
6	7018 - AT&T	2373	2330	United States	0.05%
7	58511 - Anycast	2351	13	Australia	60.1%
8	49605 - IVO	2193	11	Italy	66.7%
9	8492 - OBIT Ltd.	2153	46	Russia	71.4%
10	8220 - COLT Tech. Grp.	2143	716	United Kingdom	78.2%

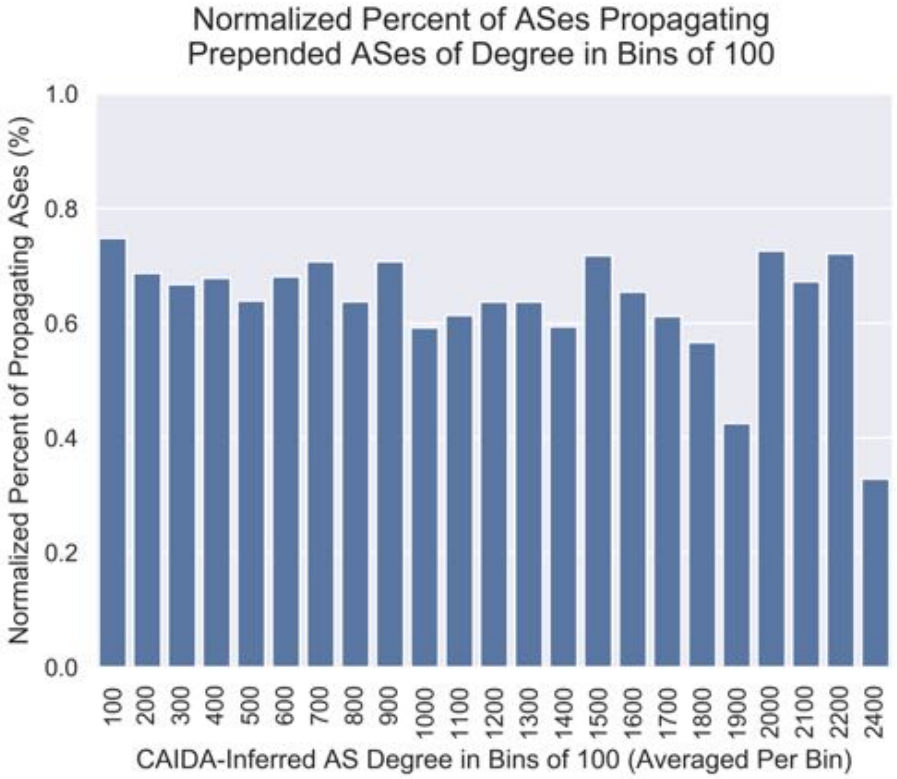

How has reachability changed since 2009?

Percent of ASes with default routes →

2009*: 77%
2018: 36.7%

Reachability of /25 prefixes →

	Control-Plane	Data-Plane
2009*:	1%	5%
2018:	50%	31%

*Bush et al. Internet Optometry, IMC 2009

Default Route Metrics

Measurement	Number of Instances
Fraction of Total Samples with Only 1 Provider (not multi-homed)	28.7% (419 / 1,460 total samples)
Fraction of Total Multi-Homed Samples with Default Routes	48.6% (506 / 1,041 multi-homed samples)
Fraction of Transit ASes with Default Routes	26.8% (196 / 731 total Transit ASes)
Fraction of Stub/Edge/Fringe ASes with Default Routes	36.7% (310 / 845 total Fringe ASes)

Comparison

2009*: 77% of Stubs had default routes (out of 24,224 **with ping**)
2018: 36.7% of Stubs had default routes (out of 845 **with traceroute**)

*Bush et al. Internet Optometry, IMC 2009

Reachability of /25 vs. /24

Prefix Length	Measurement	Findings	Timespan of Measurement
/25	BGP Observability	Seen at 21/37 (56.7%) collectors	96 hours of collection
/25	Traceroute Reachability	31% reached /25 prefix on average	7 hours; 5,000 distinct traceroutes every 1 hour
/24	BGP Observability	Seen at 34/37 (91.8%) collectors	96 hours of collection

Comparison

2009*: 1% of BGP Monitors Saw (11/615), 5% Data-Plane Reachability

2018: 50% of BGP Monitors Saw (21/37), 31% Data-Plane Reachability

*Bush et al. Internet Optometry, IMC 2009