

Anycast Polarization in The Wild



ASM Rizvi
USC CS and ISI
Akamai Technologies



Tingshan Huang
Akamai Technologies



Rasit Esrefoglu
Akamai Technologies



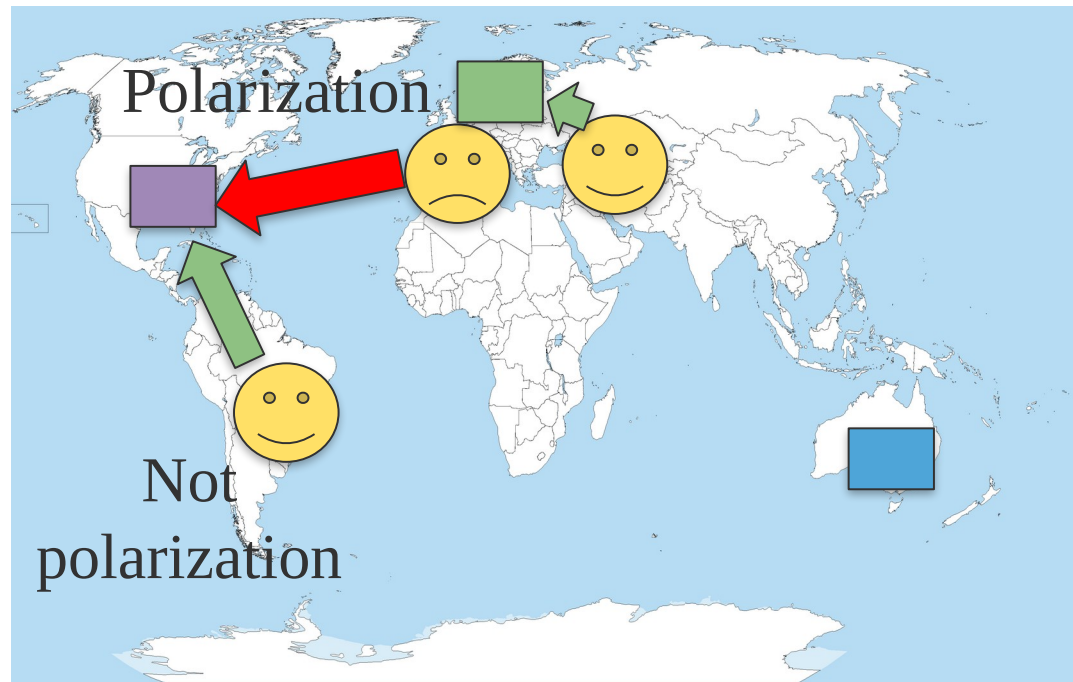
John Heidemann
USC and ISI

March 11th, 2024

Passive and Active Measurement Conference,
PAM-2024

Anycast Polarization Adds Latency

Client from Europe is going to North America




Cross-continent traffic
adds 100 ms of extra
latency

When there is an anycast site in Europe

Contributions



Reasons for
polarization and
algorithm to find them



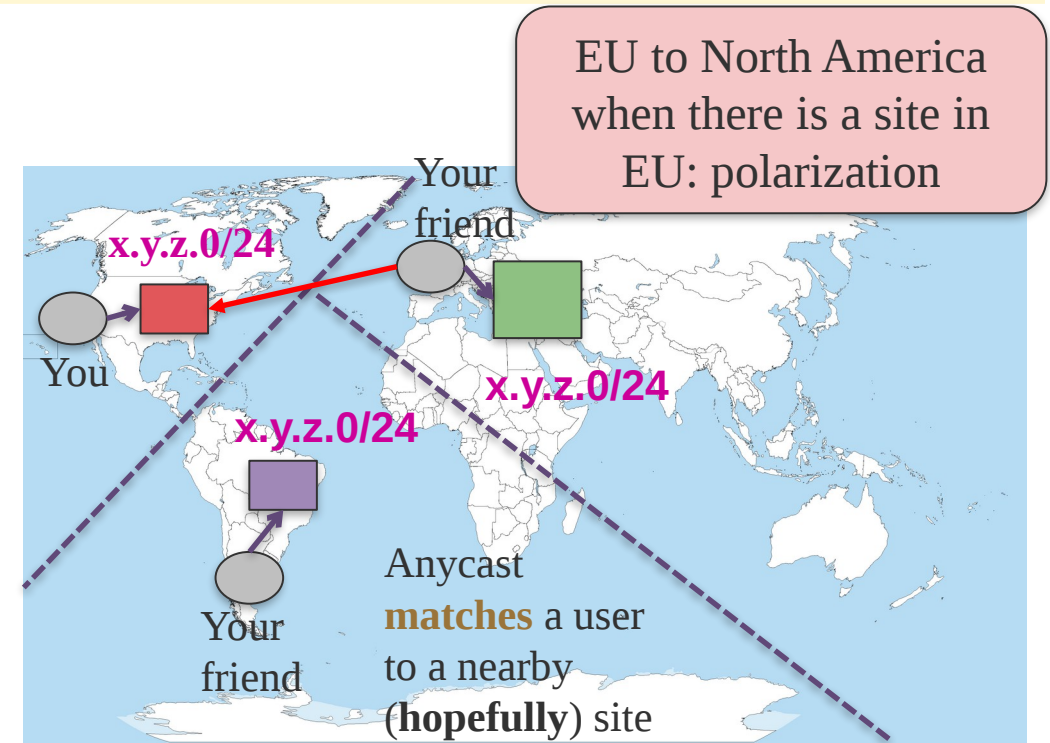
Scan 7986 anycast
prefixes to see
polarization in the wild



Two studies of
operator solutions
to polarization

Ideal Anycast World

- Clients go to the nearby location
 - Out of multiple anycast sites
- Improved latency
 - Unlike unicast
 - Since clients are going to the nearby location
- But not always in reality



Anycast **divides** the world into catchments

Progress

- Finding reasons behind polarization
- Ways to find polarization problem and their reasons
- How common is polarization?
- How can operators mitigate polarization?

Reasons Behind Polarization



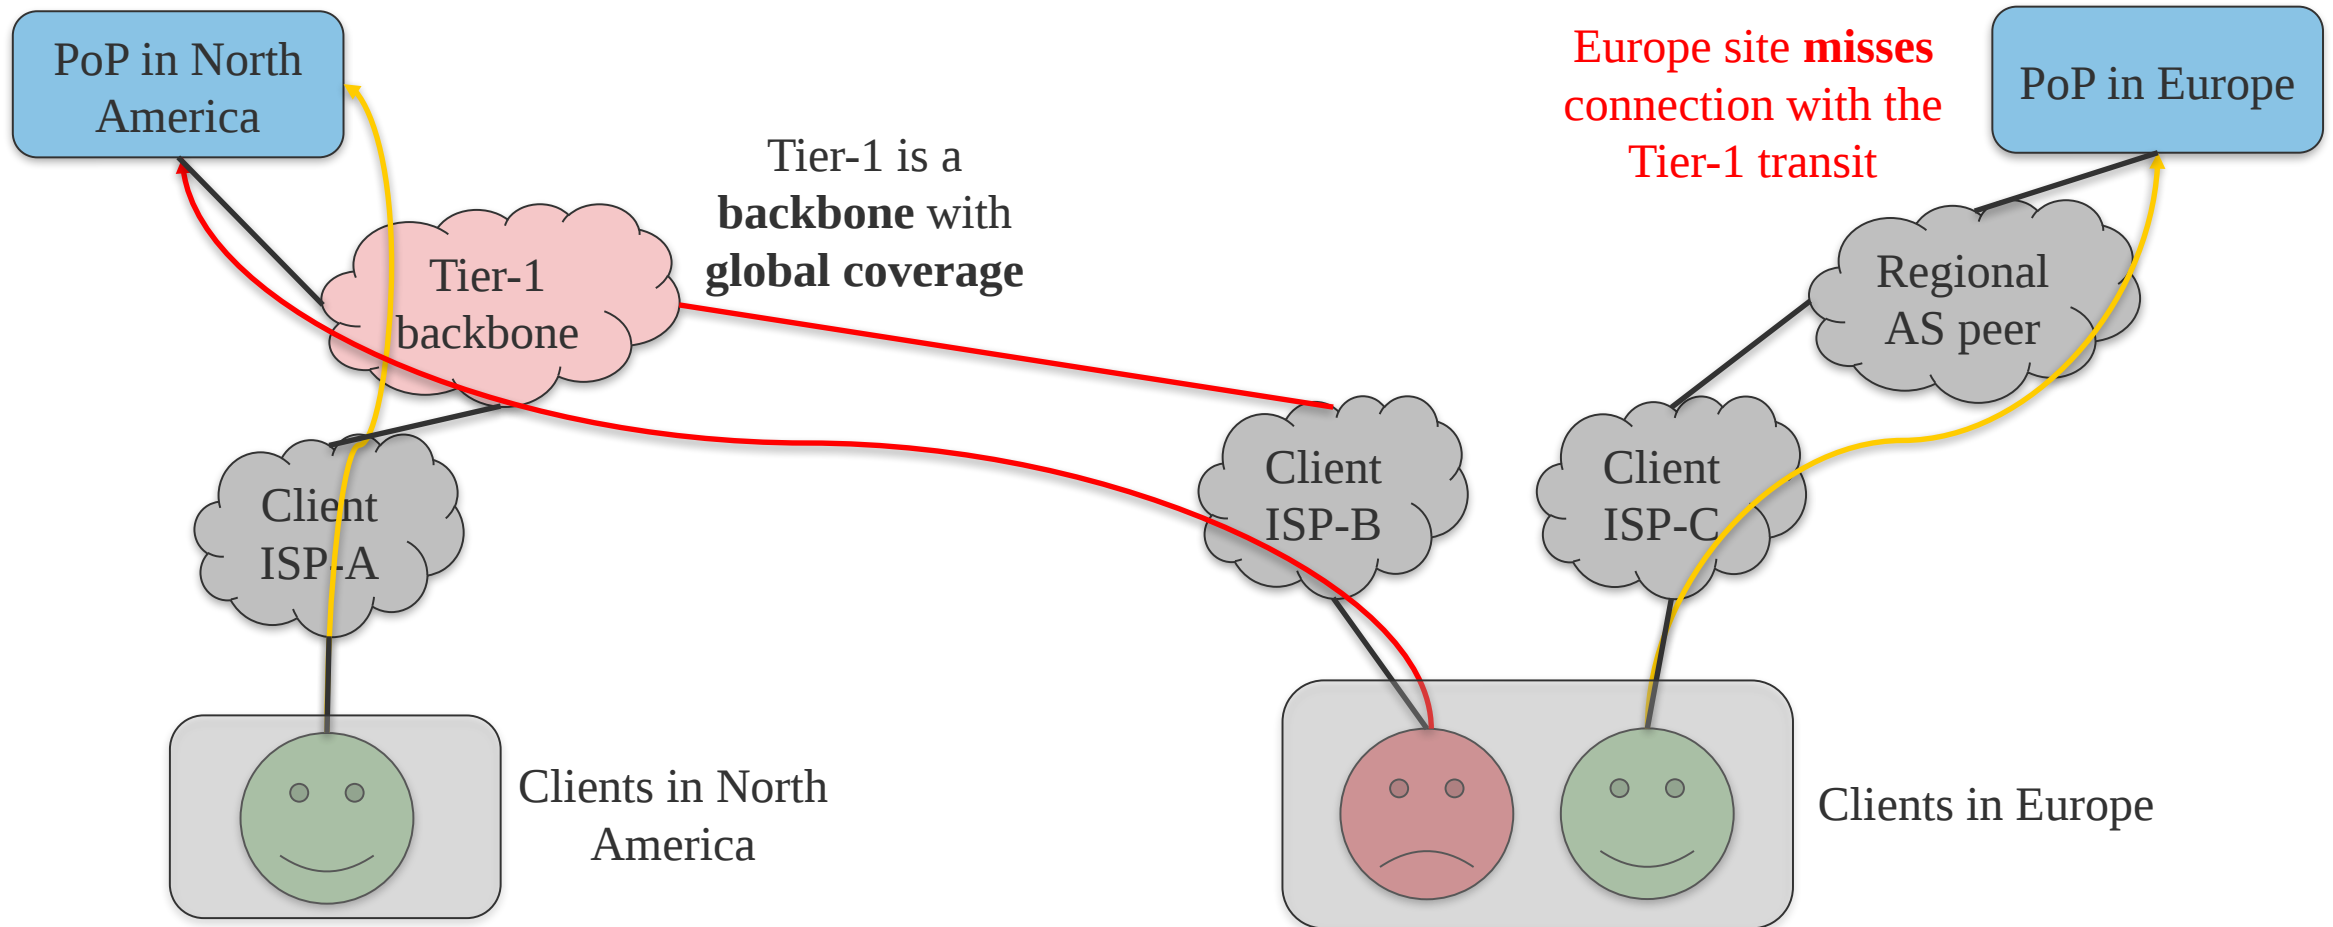
Two key reasons:

Multi-PoP backbone problem

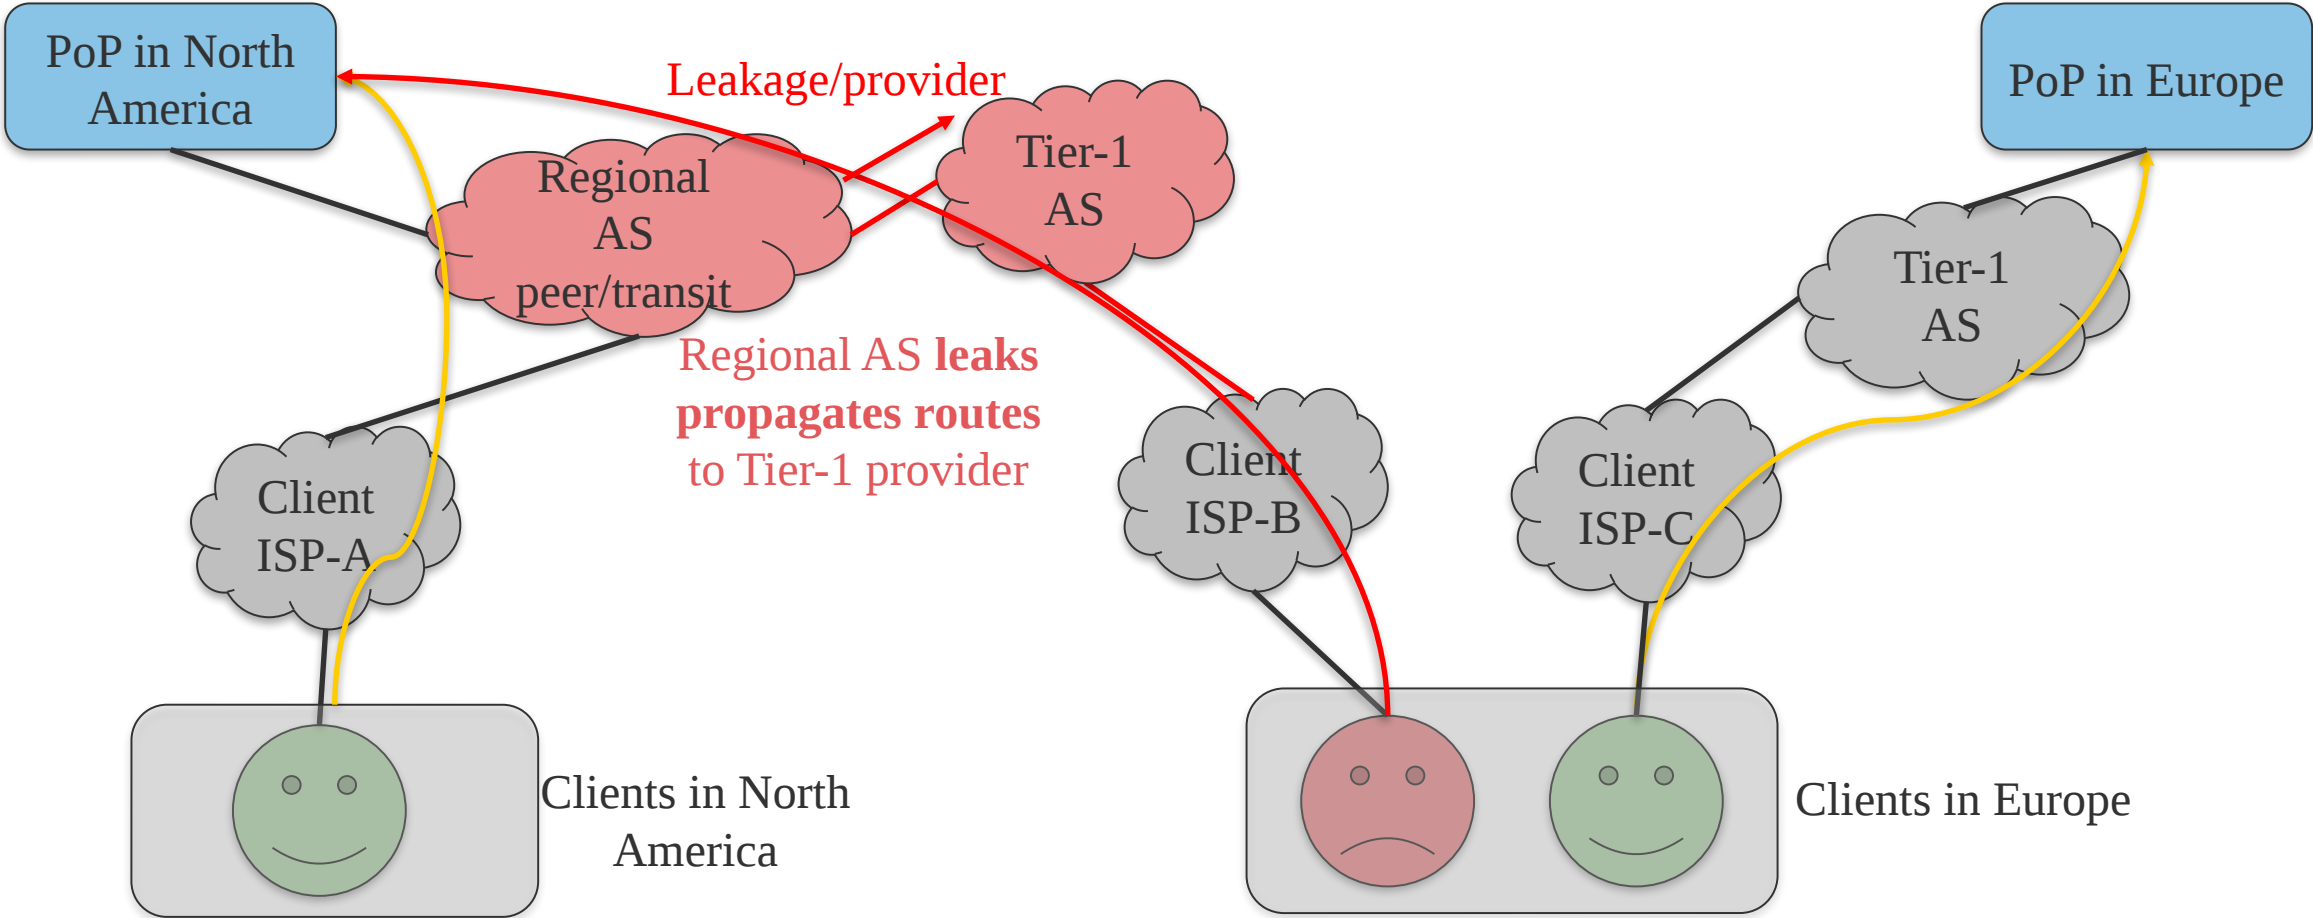
Regional leakage problem

We find both of them in the wild

Why Polarization?: Multi-PoP Backbone Problem



Why Polarization?: Regional Leakage Problem

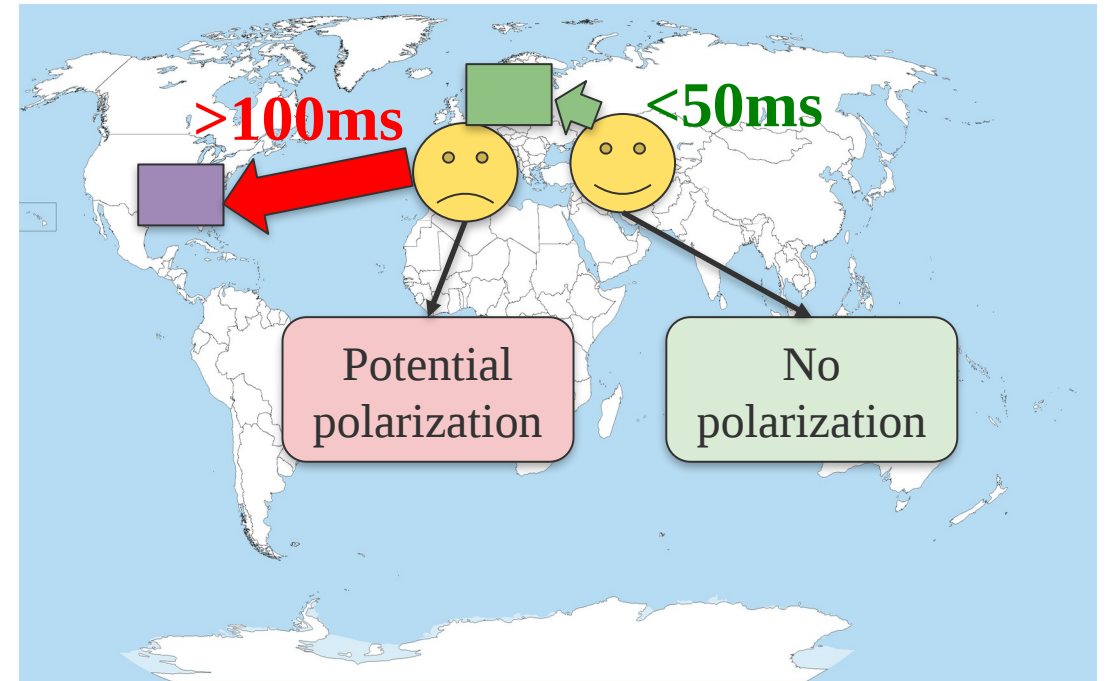


Progress

- Key reasons behind polarization
- Ways to find polarization problem and their reasons
- How common is polarization?
- How can operators mitigate polarization?

Finding Polarization

- List all the anycast prefixes
 - Used the [Manycast2] dataset of known anycast prefixes
- Ping to the anycast pingable IPs
 - From global RIPE Vantage Points (VPs)
- Potential polarization if
 - Some VPs show **<50ms** AND
 - Some VPs show **>100ms**
 - Latencies that are clearly between continents



Finding Root Causes Using Traceroutes

Input: Traceroutes from RIPE VPs

For each path in the traceroute:

find penultimate_as_hop with polarization

if penultimate_as_hop with polarization in Tier1 ASes:

if no_other_paths contain the same Tier1 AS to the nearby site:
multi-pop backbone problem++

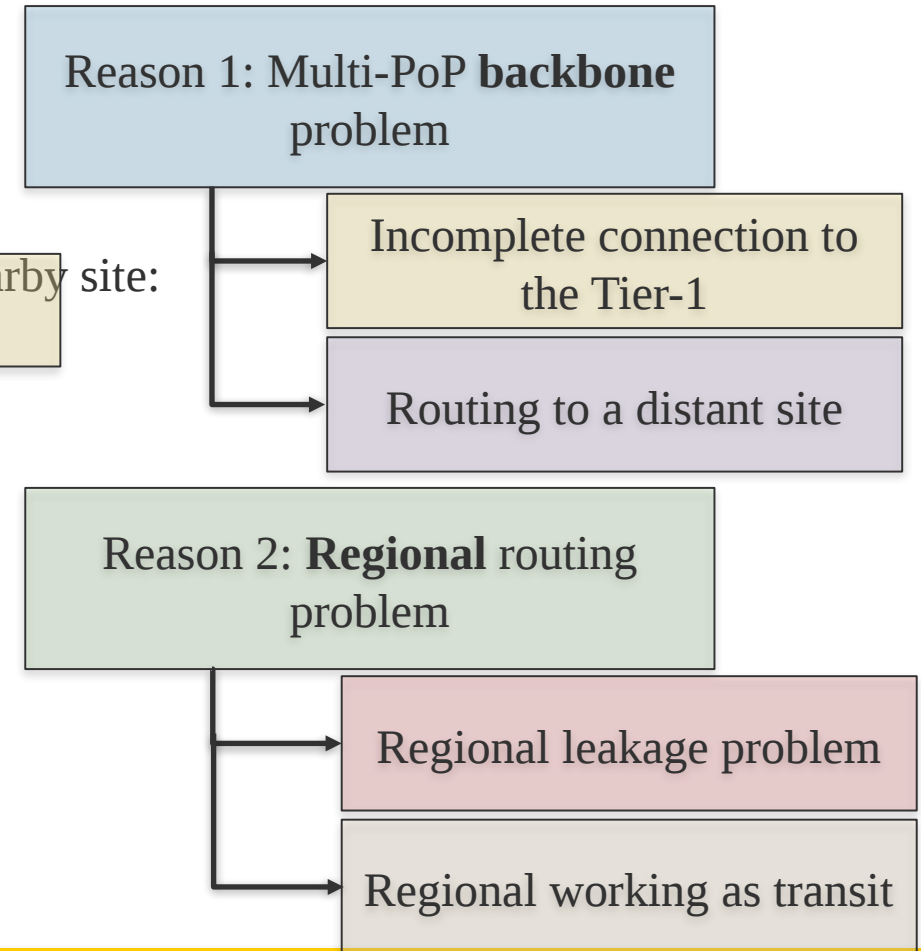
else:
routing problem++

else:

if other_paths contain Tier1 AS:
regional leaking problem++

else:
regional working as transit++

Output: Root causes of the polarization



Progress

- Finding reasons behind polarization
- Ways to find polarization problem and their reasons
- **How common is polarization?**
- How can operators mitigate polarization?

Polarization in Numbers

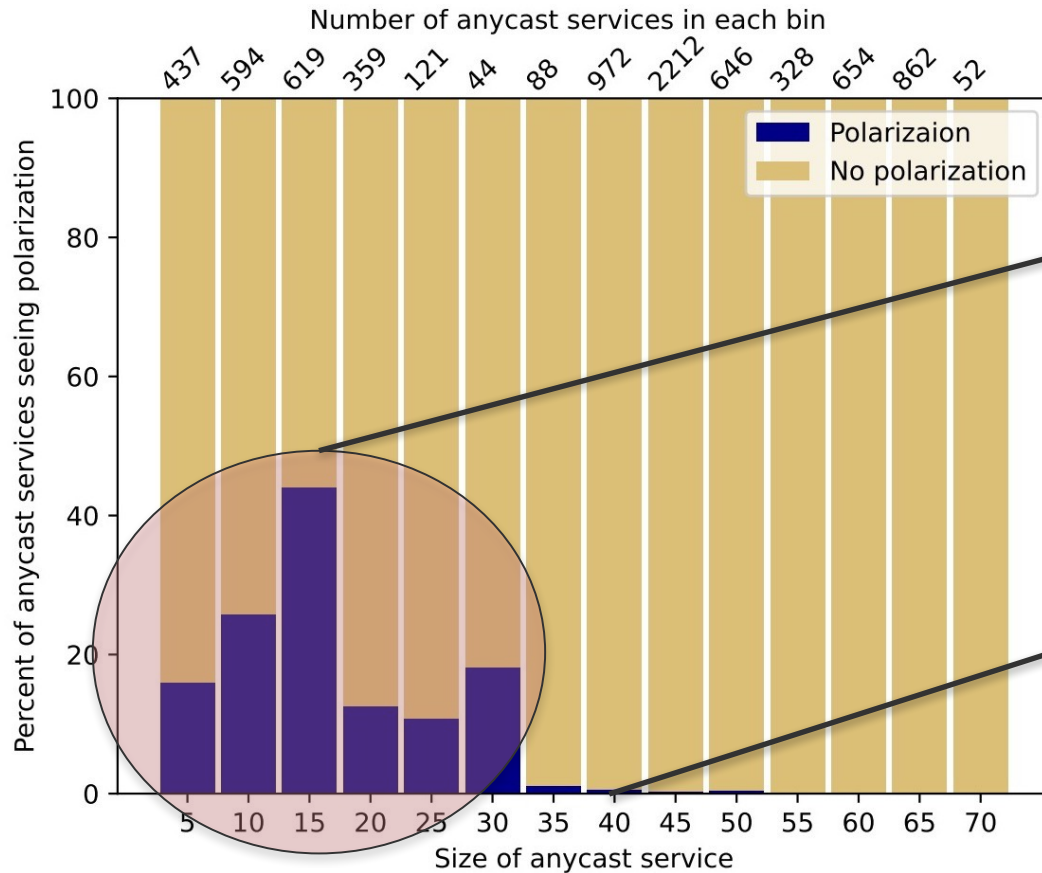
Category	Count	% Confirmed
Known anycast prefixes	7986	100
No observed potential polarization	5713	72
Potential polarization	2273	28
In continents outside EU and NA	1647	20
In EU and NA	626	8
No class found (a)	161	2
Only multi-pop backbone problem (b)	232	3
Only regional leakage (c)	89	1
Both classes (d)	144	2
Multi-pop backbone problem (b+d)	376	5
Incomplete Tier-1 connections	218	3
Routing to a distant site	158	2
Leaking regional problem (c+d)	233	3
Leakage by regional	177	2
Leakage by regional transits	56	1

2273 known anycast prefixes have polarization

We confirmed 18 polarization problems with the operators

18
9
9
0
9
6
3

Impacts of Number of Sites over Polarization



Most common with less than 35 anycast sites

Not completely uncommon with more than 35 sites

Fig: Impacts of number of sites (for EU and NA polarizations)

Impacts of Polarization

Provider	No. of /24 prefixes	Reason of the problem	Common AS to the distant site	Source cont.	Med.	95th	Example
Anon-CDN-2	14	Incomplete Tier-1	AS1299	EU NA	12 25	173 88	Poland to USA
Anon-Cloud-1	1	Incomplete Tier-1	AS3356, AS6939	EU NA	51 55	158 162	Canada to Germany
Anon-CDN-6	67	Incomplete Tier-1 (peers as transits)	AS6762	EU NA	16 39	122 141	Greece to USA
Anon-CDN-3	9	Exceptional incomplete Tier-1	AS6453	EU NA	32 25	113 39	Germany to USA
Anon-CDN-1	67	Leaking regional transit	AS1273	EU NA	21 88	53 150	USA to UK
Anon-DNS-1	18	Leaking regional peer	AS4826	EU NA	174 37	314 214	USA to Australia
Anon-CDN-5	67	Leaking regional peer	AS7473	EU NA	39 24	248 252	USA to Singapore
Anon-CDN-7	67	Leaking regional peer (merging org)	AS209	EU NA	29 26	84 65	Finland to USA
Anon-DNS-2	216	Leaking regional Incomplete Tier-1	AS4637 AS1299	EU NA	165 81	301 254	Canada to Tokyo

Big difference between median and 95th percentile


Cross-continent traffic

Polarization in different CDN, Clouds, and DNS

Progress

- Finding reasons behind polarization
- Ways to find polarization problem and their reasons
- How common is polarization?
- How can operators mitigate polarization?


Ways to Improve




Transit connectivity
covering different
continents



More sites with
proper connectivity

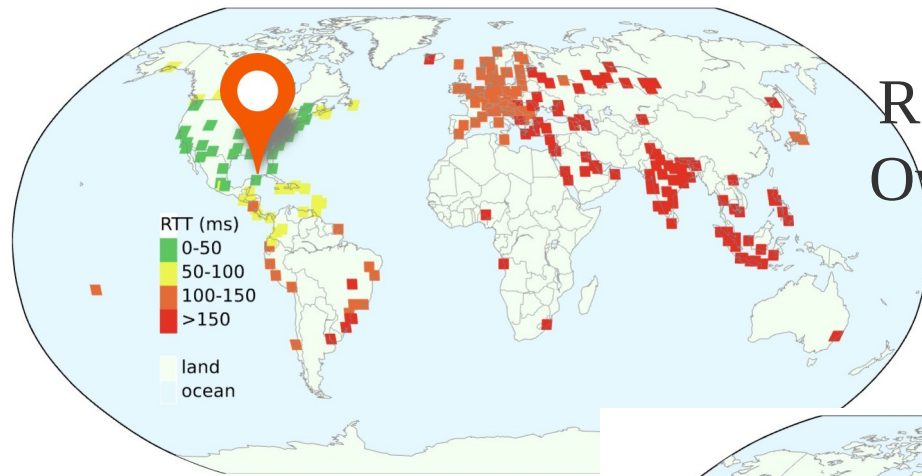


Vigilant about route
leakage



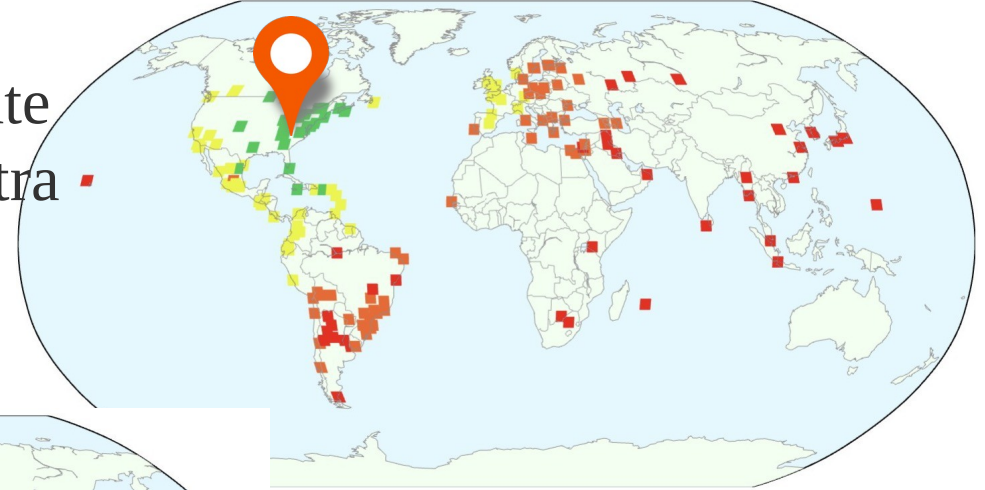
Proper routing
nudges

Case Study: Polarization Adds Latency



Catchment: Dallas,
USA

Red dots indicate
Over **100ms** extra
latency

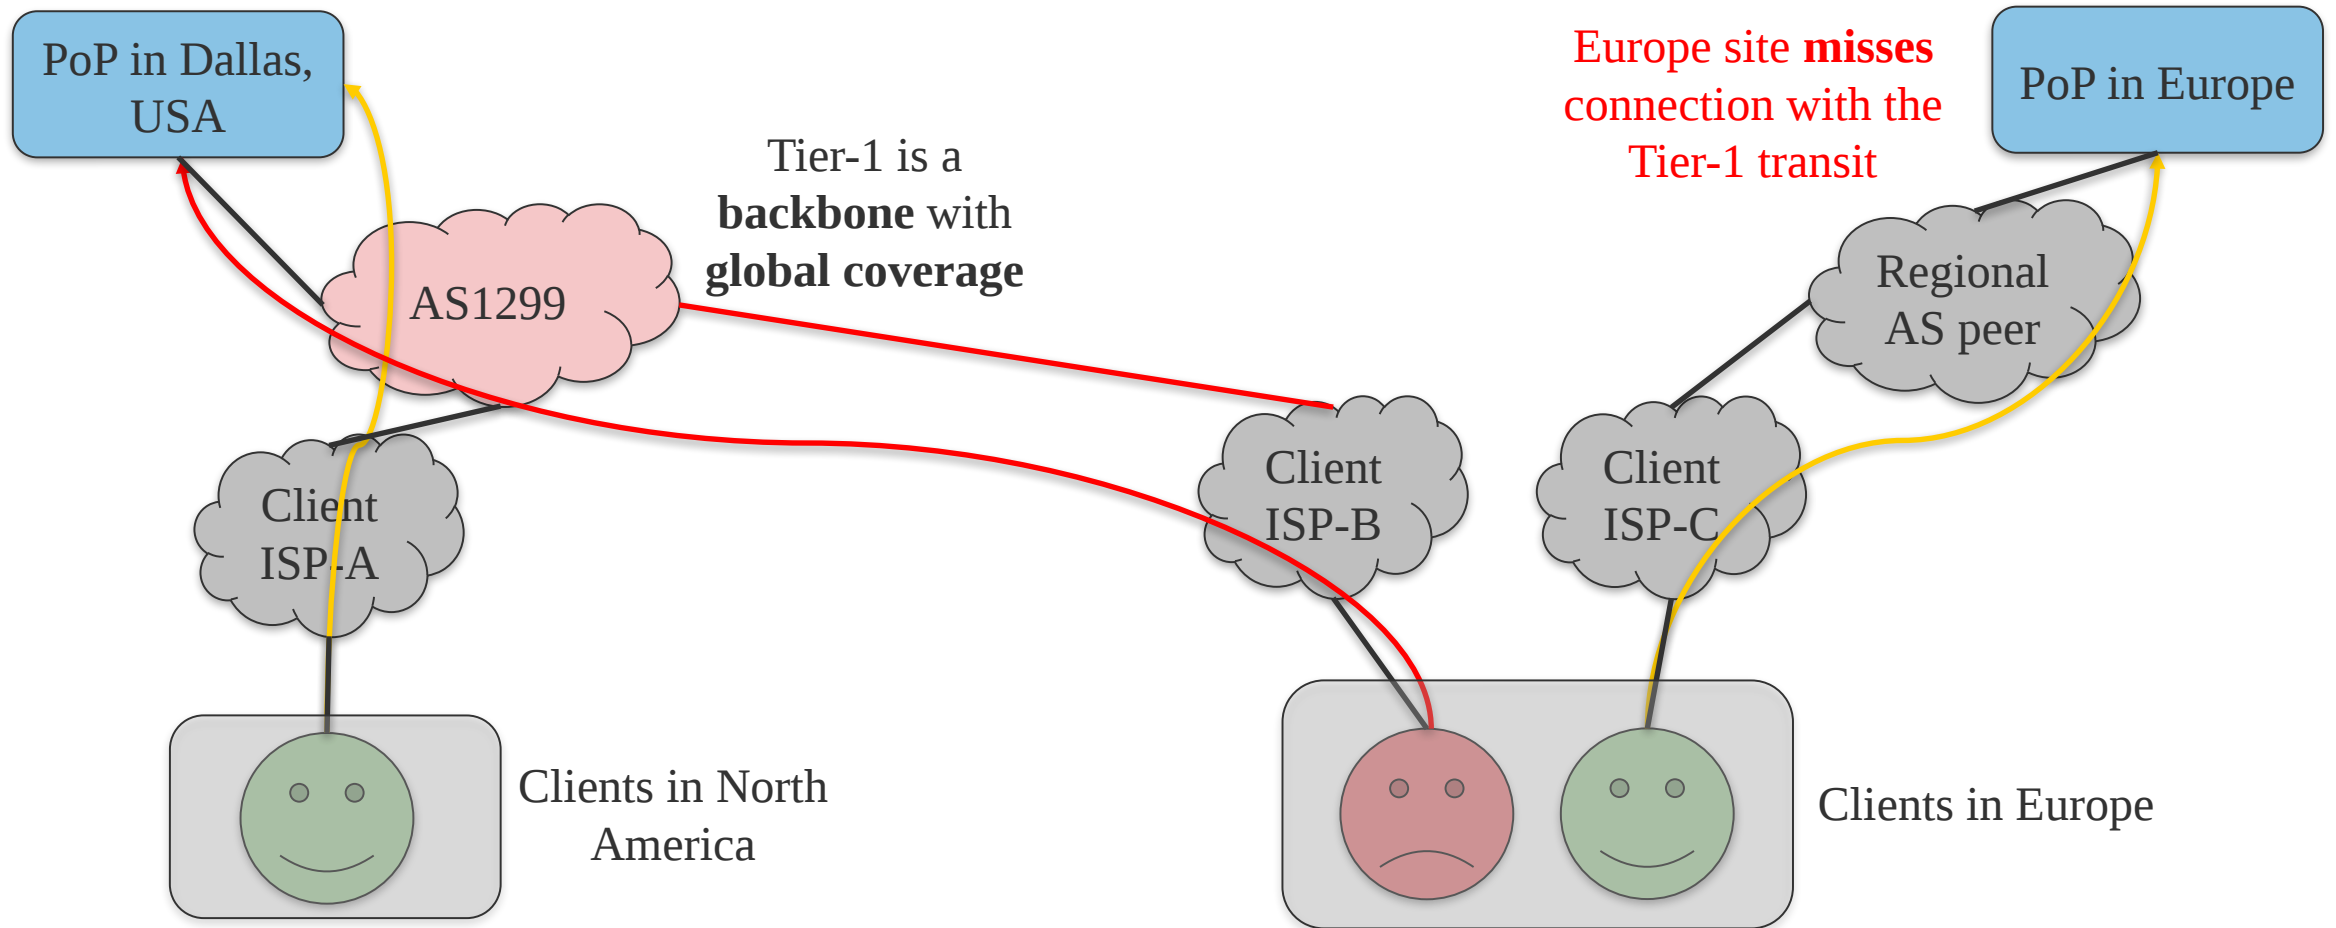


Catchment: Virginia,
USA

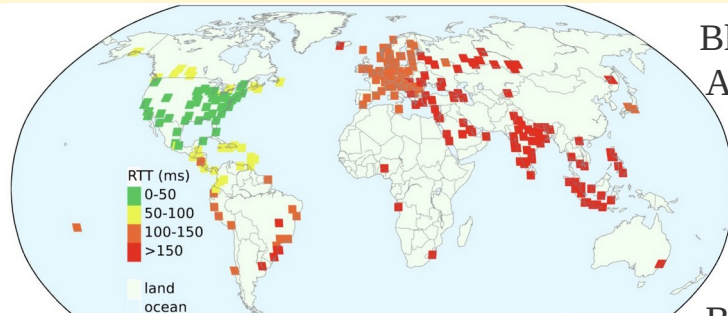


Catchment: Milan,
Italy

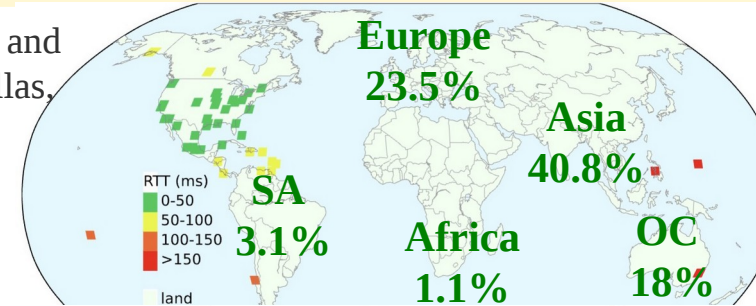

Why Polarization?: Multi-PoP Backbone Problem in Dallas Site



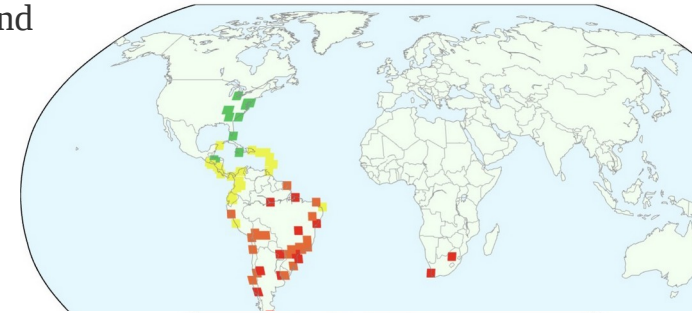
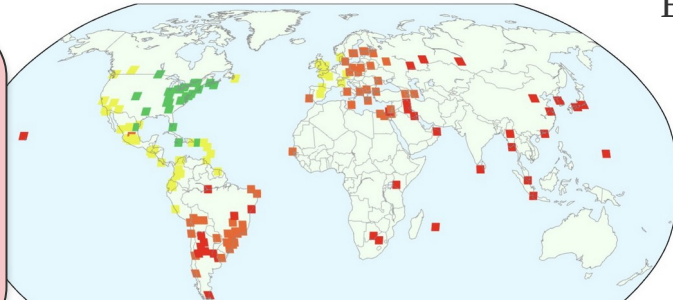
Improvement after Routing Changes



Blocking AS-3356 and AS-1299 from Dallas, USA



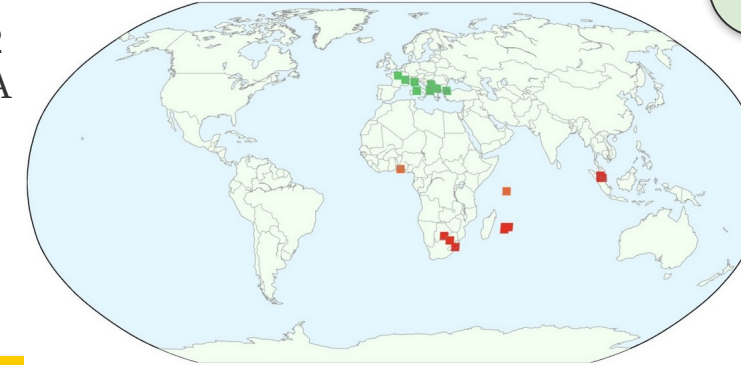
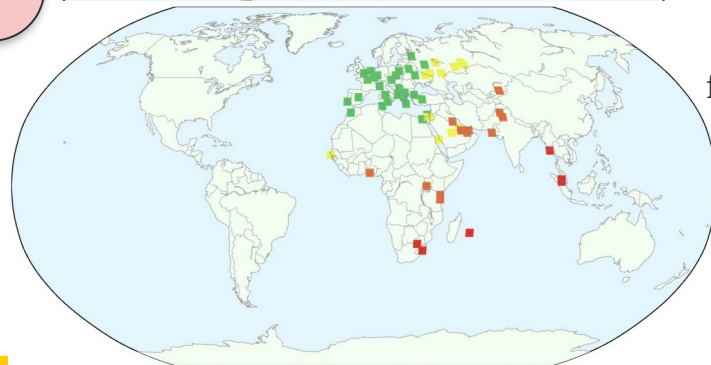

Blocking AS-209 and AS-6762 from Virginia, USA



More red dots (over 100ms latency)

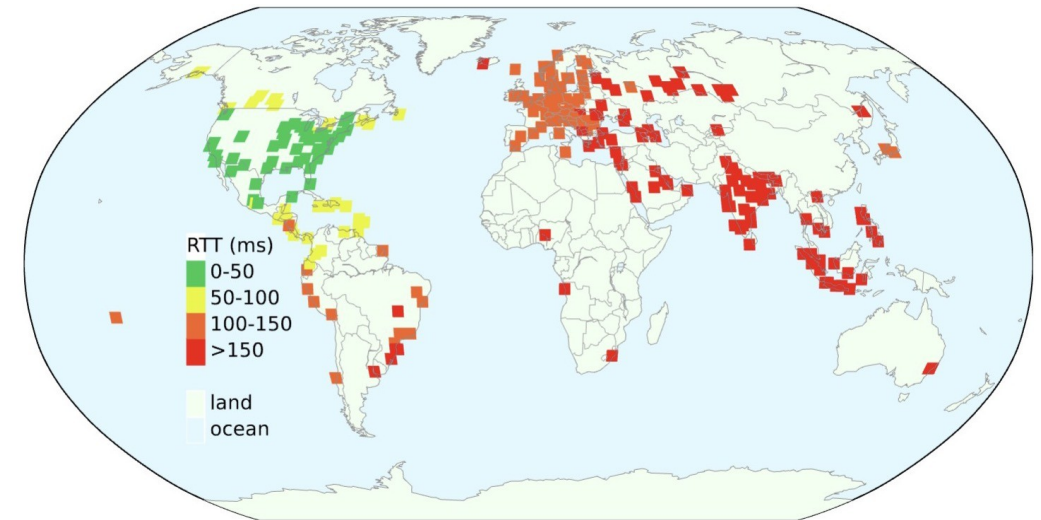
Less number of red dots (over 100ms latency)

Blocking AS-6762 from Virginia, USA



Conclusion

- Polarization is not uncommon
- Two key reasons
 - Multi-PoP backbone problem
 - Regional leakage problem
- Routing changes can improve



Backup Slides

Community Strings Are Important

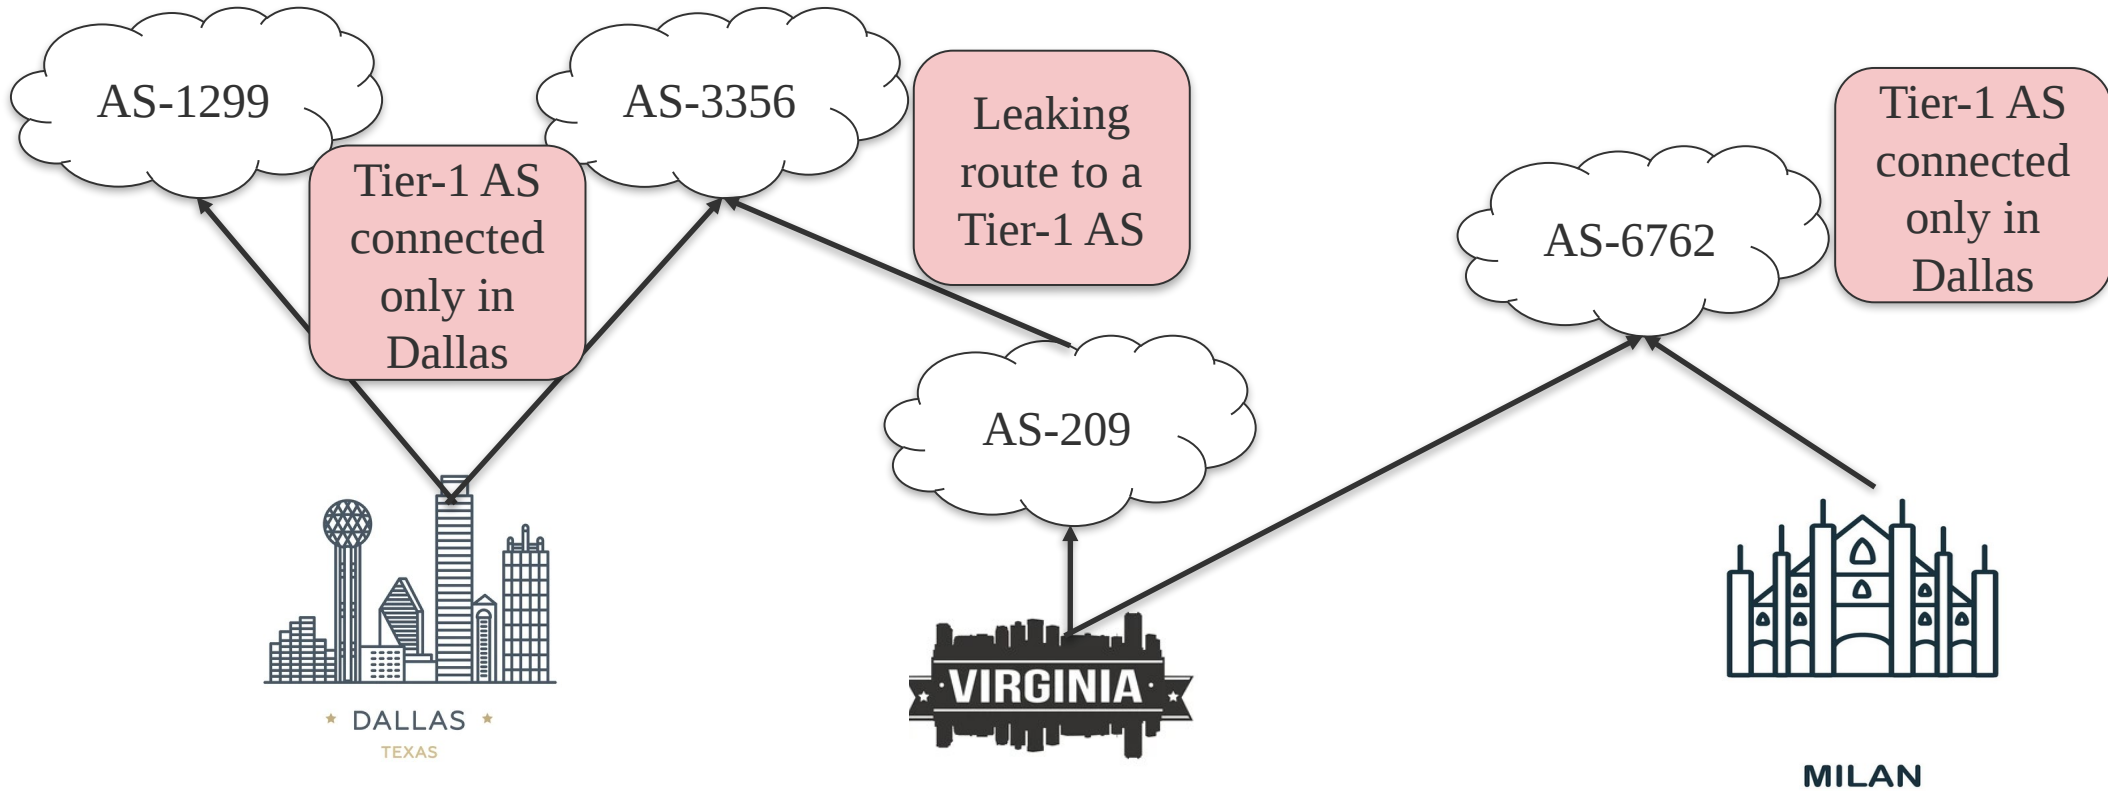
Routing changes	Improvement (%)					
	Africa	Asia	Europe	North America	Oceania	South America
Announcement to AS-7473 only within Asia Pacific and block others (using community strings)	10.5	12.3	34.8	19.7	2.3	1.9
Announcement to AS-7473 with two prepending	3.8	-9.6	6.5	6.4	6.3	3.8

Community string better improves

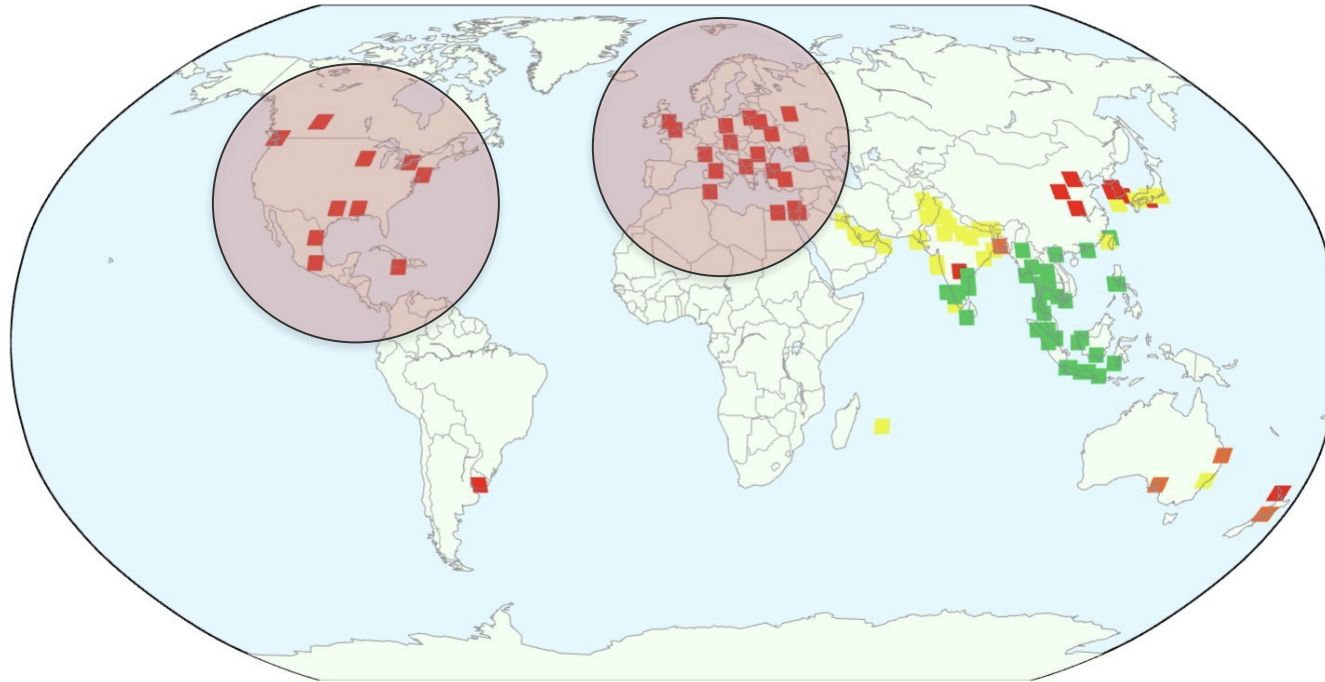
AS-7473 is important for Asia. So, we should announce it within Asia, but block it in other continents

Community strings are important for precise control

Root Causes (Anycast Prefix-1)



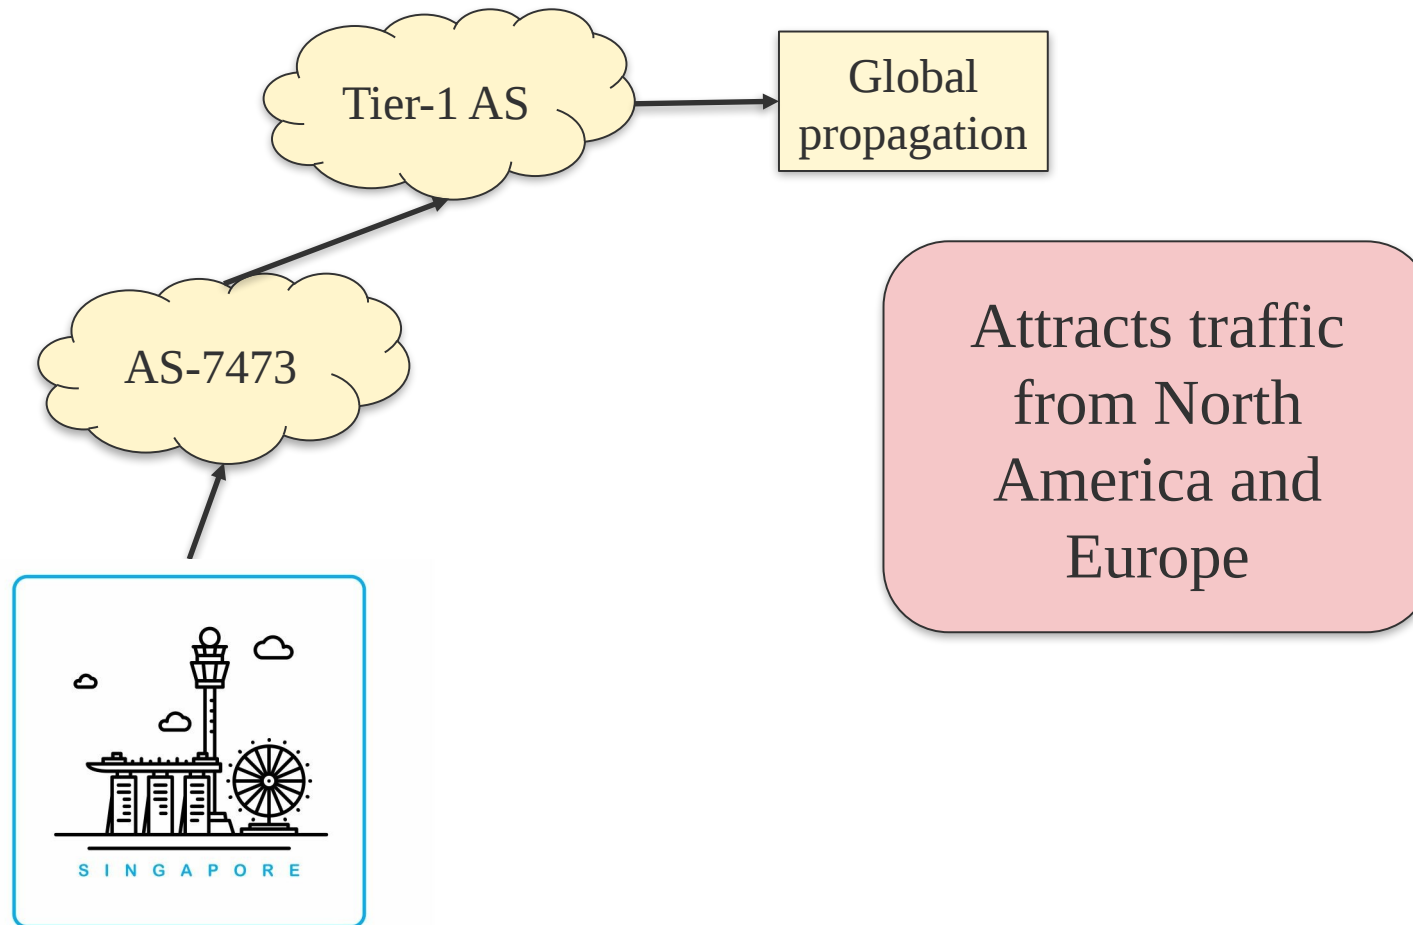
Another Case Study



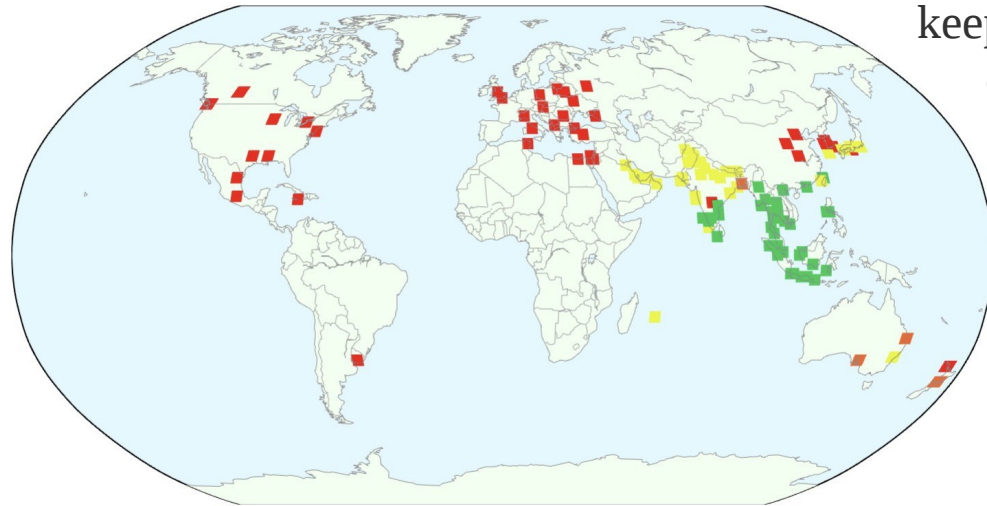
Catchment: Singapore

Traffic from the EU and North America going to Singapore (over 100ms latency)

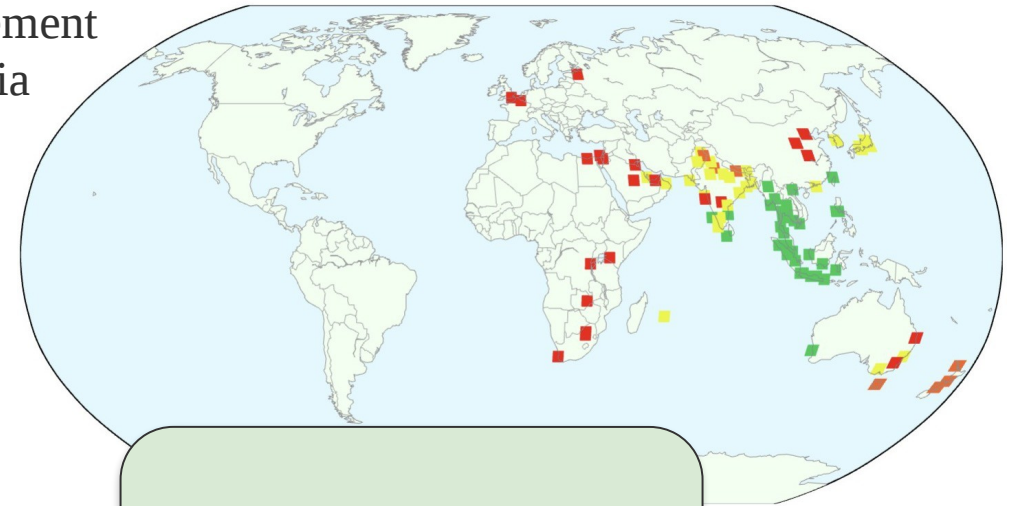

Root Cause (Anycast Prefix-2)



Improvement Using Community Strings



Community string to
keep the announcement
only within Asia
Pacific



No more North American
traffic going to Singapore