Measuring Anycast Routing:

Catchment, RTT, and Optimal routing of your DNS deployment

OARC-46, EDINBURGH, MAY’26

Remi Hendriks

UNIVERSITY OF TWENTE.

Introduction
Motivation

- Anycast: Making an IP address available in multiple Points of Presence (PoPs)
  - *How*? - Announcing an IP prefix at multiple locations using BGP

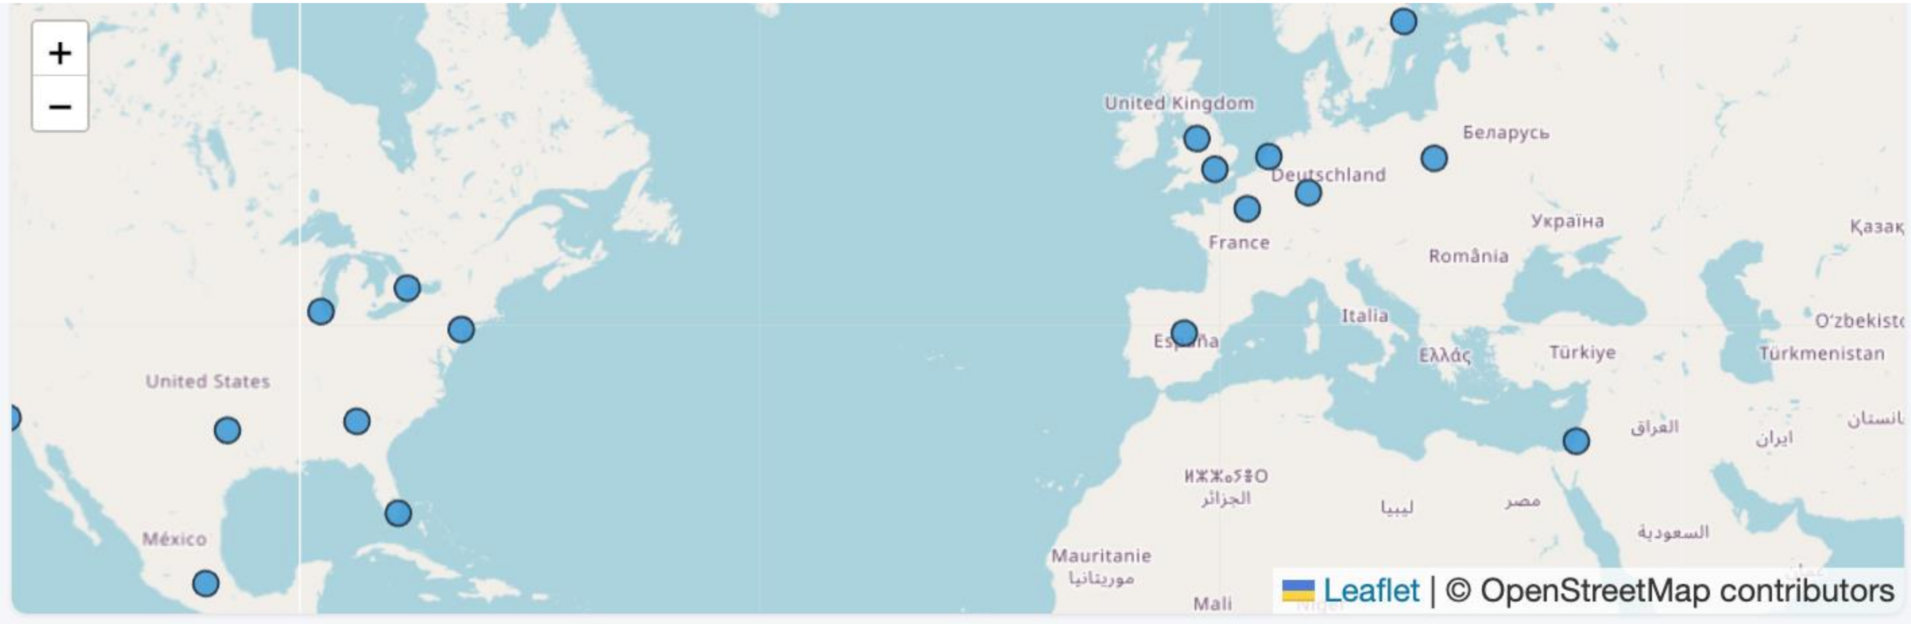

Introduction
Motivation

- Anycast: Making an IP address available in multiple Points of Presence (PoPs)
  - *How*? - Announcing an IP prefix at multiple locations using BGP
  - *Why?* - **Distribute load**

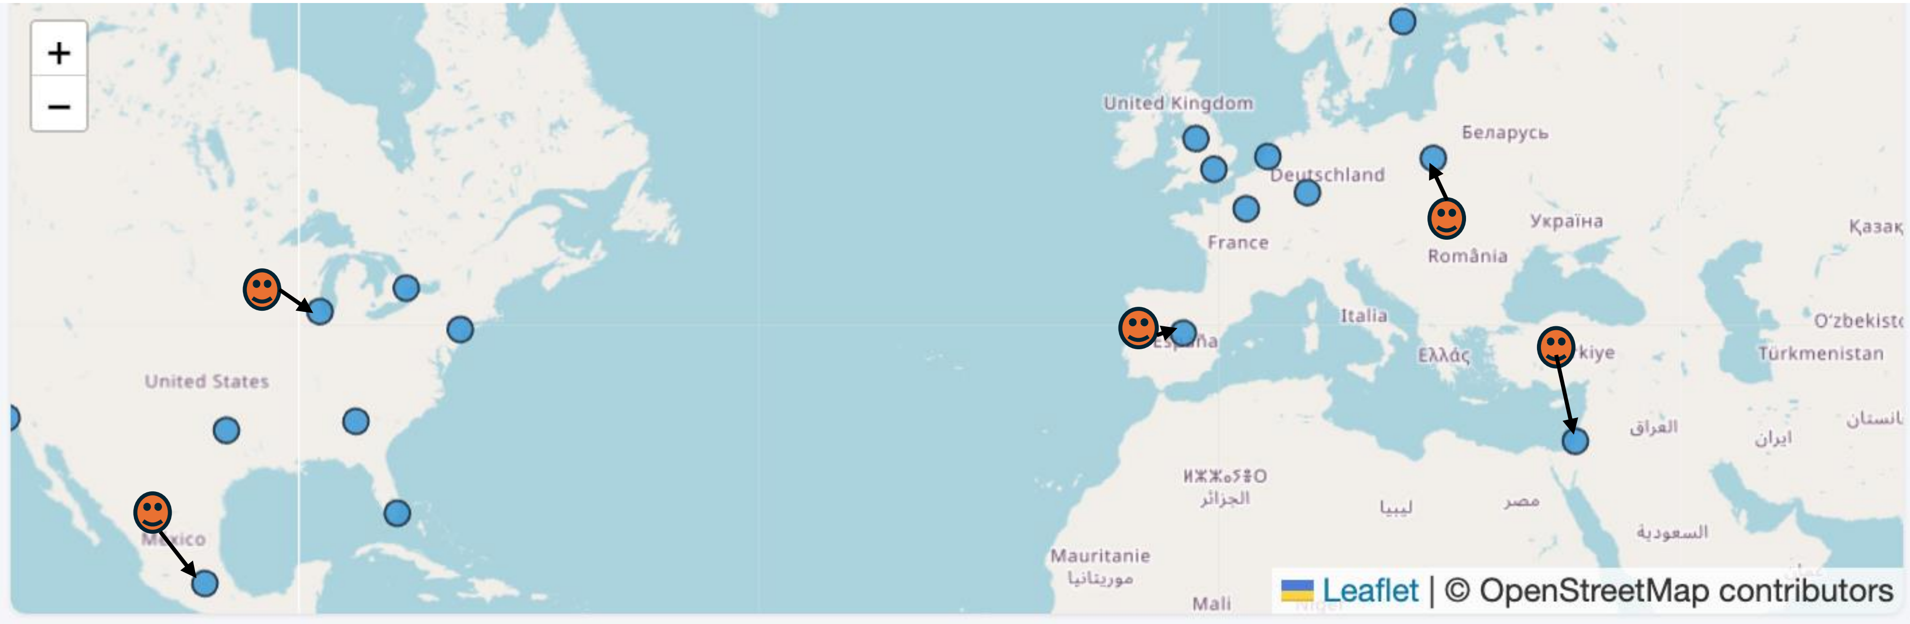

Introduction
Motivation

- Anycast: Making an IP address available in multiple Points of Presence (PoPs)
  - *How*? - Announcing an IP prefix at multiple locations using BGP
  - *Why?* - Distribute load, **reduce latency**

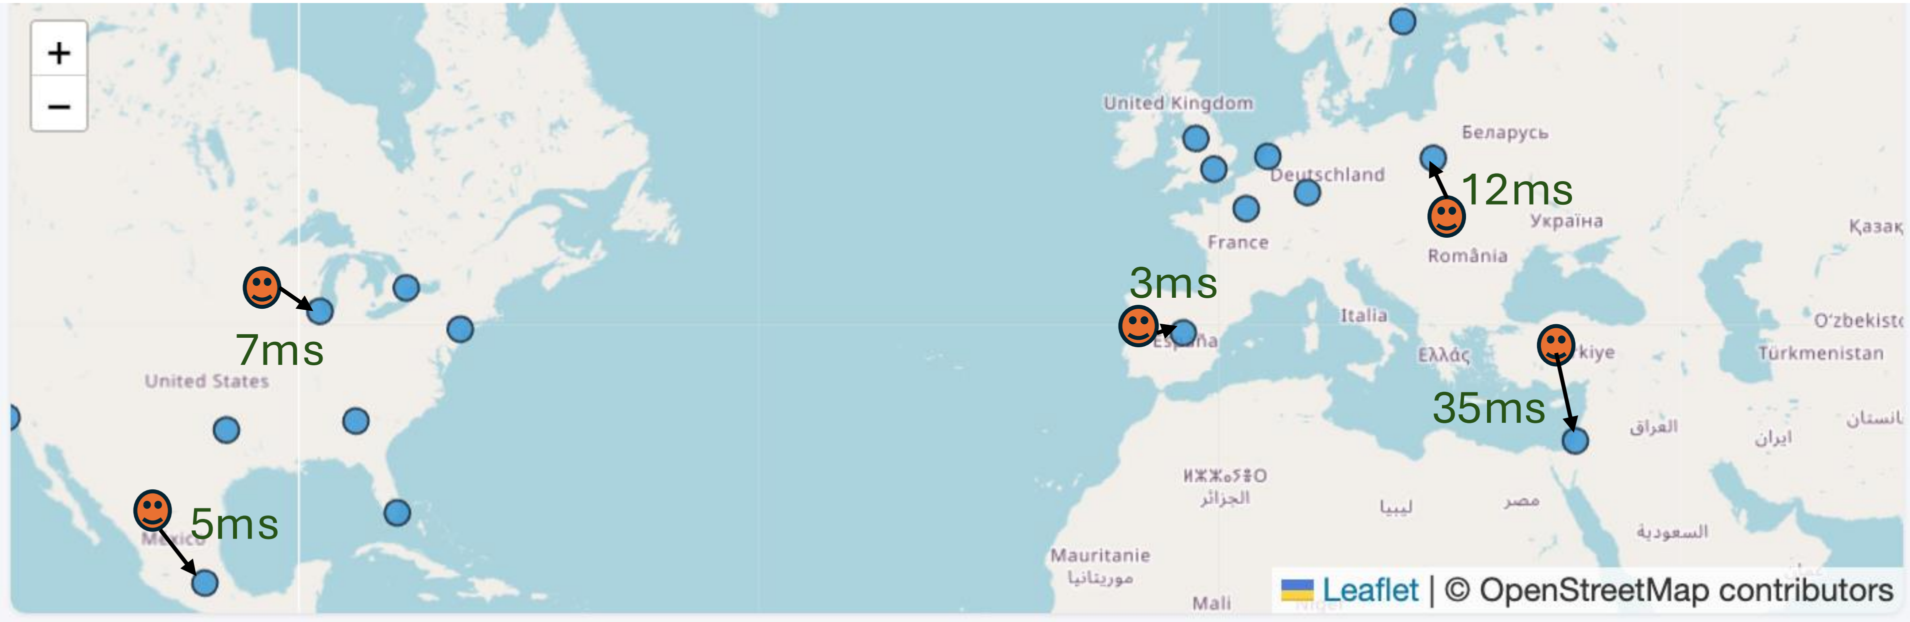

Introduction
Motivation

- Anycast: Making an IP address available in multiple Points of Presence (PoPs)
  - *How*? - Announcing an IP prefix at multiple locations using BGP
  - *Why?* - Distribute load, reduce latency, **resiliency**

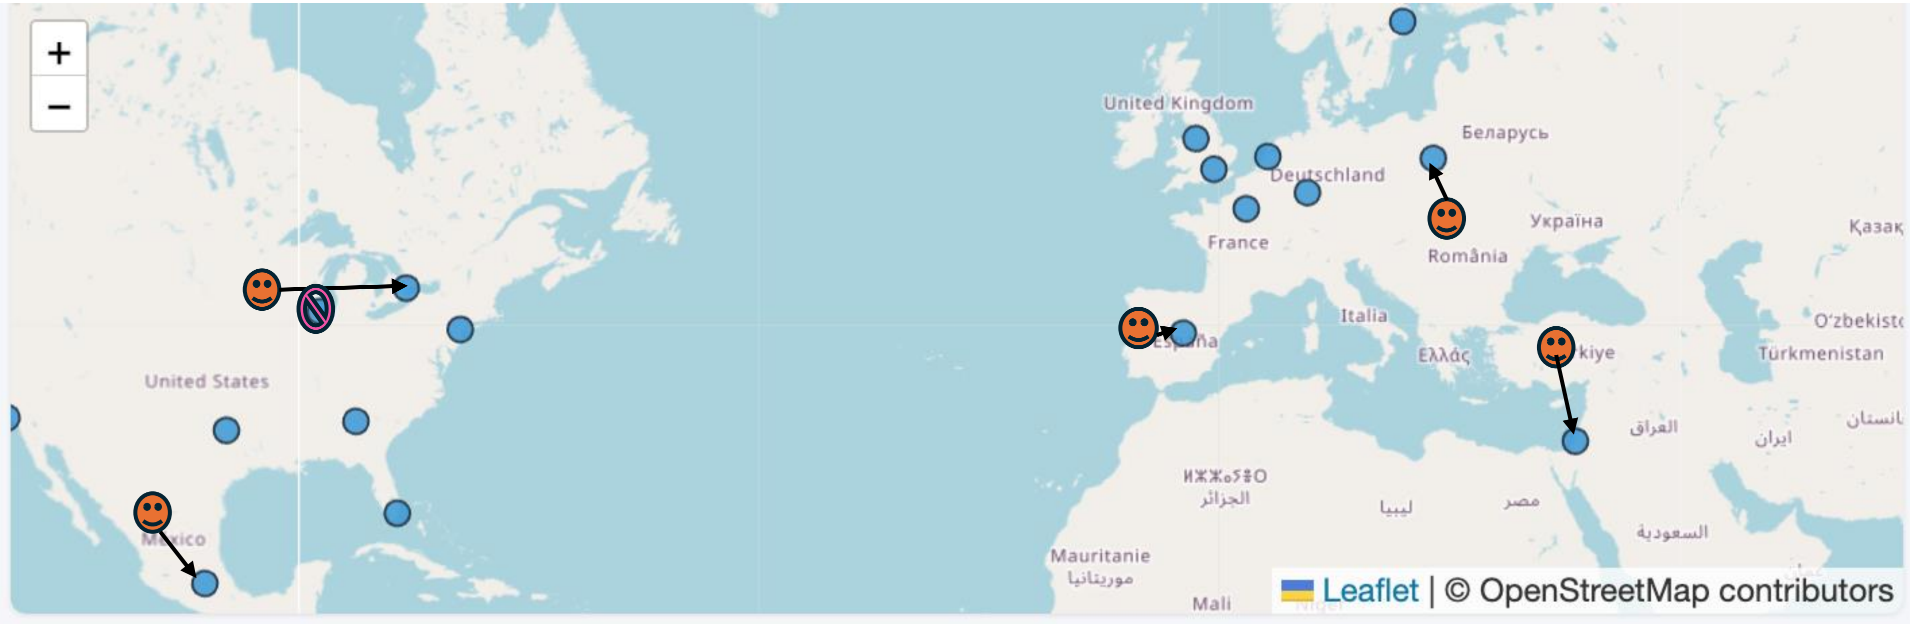

Introduction
Anycast and the DNS

- Anycast is widely used
  - Over 1,000 ASes deploy anycast

NETWORK STATISTICS (HIGH CONFIDENCE) – 2026-04-16

	IPv4	IPv6	Union
Anycast ASes ⓘ	1,082	580	1,235
Unique BGP prefixes (backing_prefix) ⓘ	5,385	2,889	8,273
MOAS prefixes ⓘ	220	313	533

Statistics from manycast.net

Introduction
Anycast and the DNS

- Anycast is widely used
  - Over 1,000 ASes deploy anycast
  - 63% of these run anycast DNS services

NETWORK STATISTICS (HIGH CONFIDENCE) – 2026-04-16

	IPv4	IPv6	Union
Anycast ASes ⓘ	1,082	580	1,235
Unique BGP prefixes (backing_prefix) ⓘ	5,385	2,889	8,273
MOAS prefixes ⓘ	220	313	533

Statistics from manycast.net

Introduction
Anycast and the DNS

- Anycast is widely used
  - Over 1,000 ASes deploy anycast
  - 63% of these run anycast DNS services
- DNS anycast deployments
  - All root letters

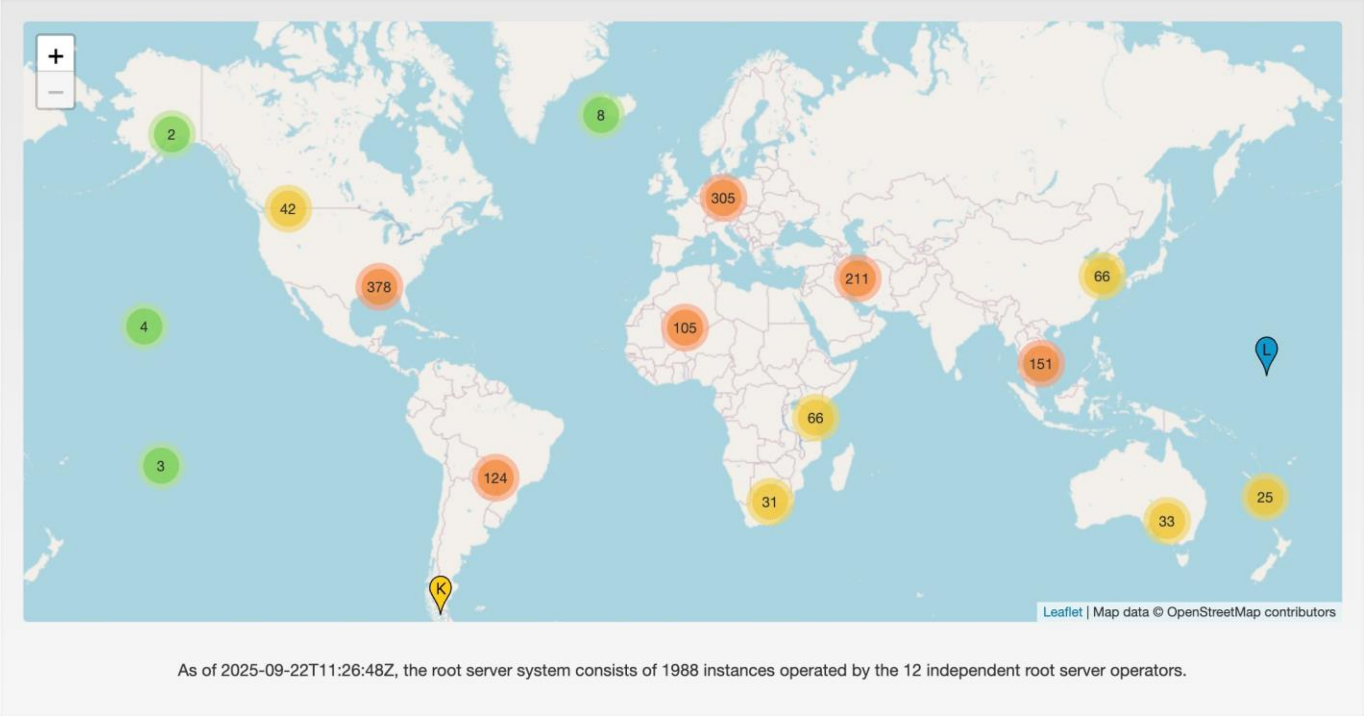

Source: root-servers.org
"All 12 DNS root-letters are deployed using anycast; combined totaling 1,988 sites"

Introduction
Anycast and the DNS

- Anycast is widely used
  - Over 1,000 ASes deploy anycast
  - 63% of these run anycast DNS services
- DNS anycast deployments
  - All root letters
  - Auth NS (*e.g.*, ccTLDs)

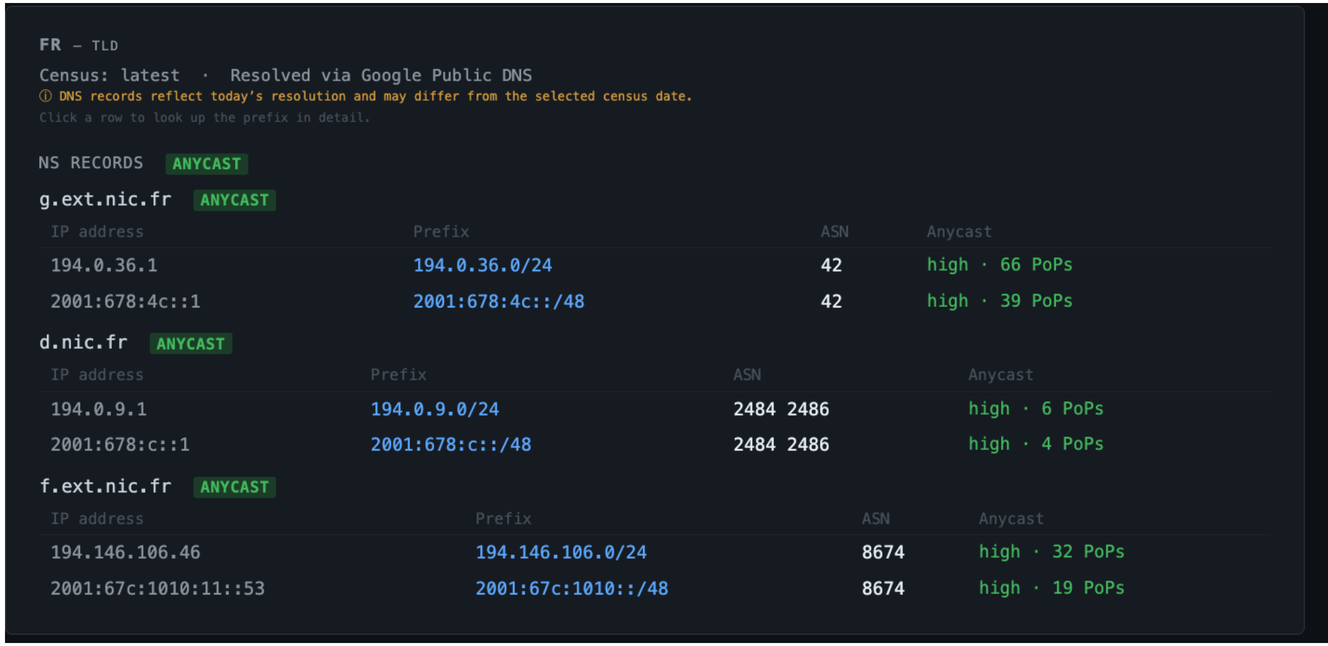

Source: manycast.net/?q=.fr
"Example: NS anycast deployments for .fr (France)"

Introduction

Anycast and the DNS

- Anycast is widely used
  - Over 1,000 ASes deploy anycast
  - 63% of these run anycast DNS services
- DNS anycast deployments
  - All root letters
  - Auth NS (*e.g.,* ccTLDs)
  - Second-level NS

NS RECORDS ANYCAST

ns1.google.com ANYCAST

IP address	Prefix	ASN	Anycast
216.239.32.10	216.239.32.0/24	15169	high · 9 PoPs
2001:4860:4802:32::a	2001:4860:4802::/48	15169	high · 34 PoPs

ns4.google.com ANYCAST

IP address	Prefix	ASN	Anycast
216.239.38.10	216.239.38.0/24	15169	high · 14 PoPs
2001:4860:4802:38::a	2001:4860:4802::/48	15169	high · 34 PoPs

ns2.google.com ANYCAST

IP address	Prefix	ASN	Anycast
216.239.34.10	216.239.34.0/24	15169	high · 6 PoPs
2001:4860:4802:34::a	2001:4860:4802::/48	15169	high · 34 PoPs

ns3.google.com ANYCAST

IP address	Prefix	ASN	Anycast
216.239.36.10	216.239.36.0/24	15169	high · 9 PoPs
2001:4860:4802:36::a	2001:4860:4802::/48	15169	high · 34 PoPs

Source: https://manycast.net/?q=google.com
"Example: NS anycast deployments for google.com"

Introduction
Anycast and the DNS

- Anycast is widely used
  - Over 1,000 ASes deploy anycast
  - 63% of these run anycast DNS services
- DNS anycast deployments
  - All root letters
  - Auth NS (*e.g.*, ccTLDs)
  - Second-level NS
  - Resolvers

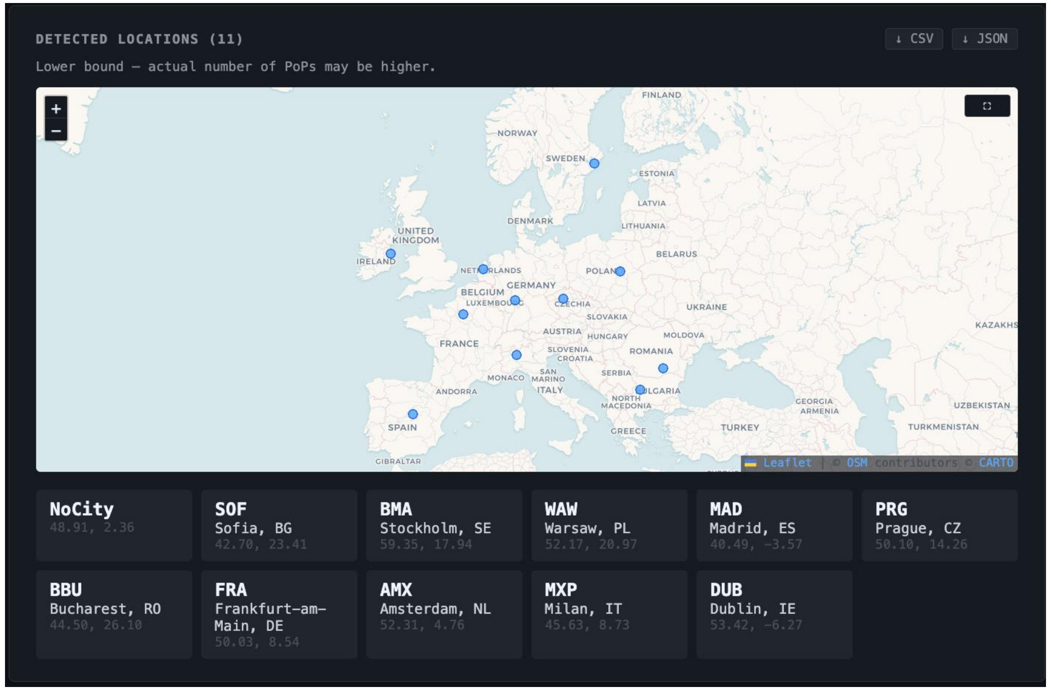

Source: https://manycast.net/?q=86.54.11.1
"Example: DNS4EU public DNS resolver locations"

Introduction
Problem Statement

- Anycast relies on BGP to route clients
- But.. BGP routes based on shortest path -> not performance
  - Transits may perform cold-potato routing toward distant PoPs
  - Remote-peering at IXPs can route client sub-optimally

Introduction
Problem Statement

- Anycast relies on BGP to route clients
- But.. BGP routes based on shortest path -> not performance
- Case A) Client routes to a distant PoP

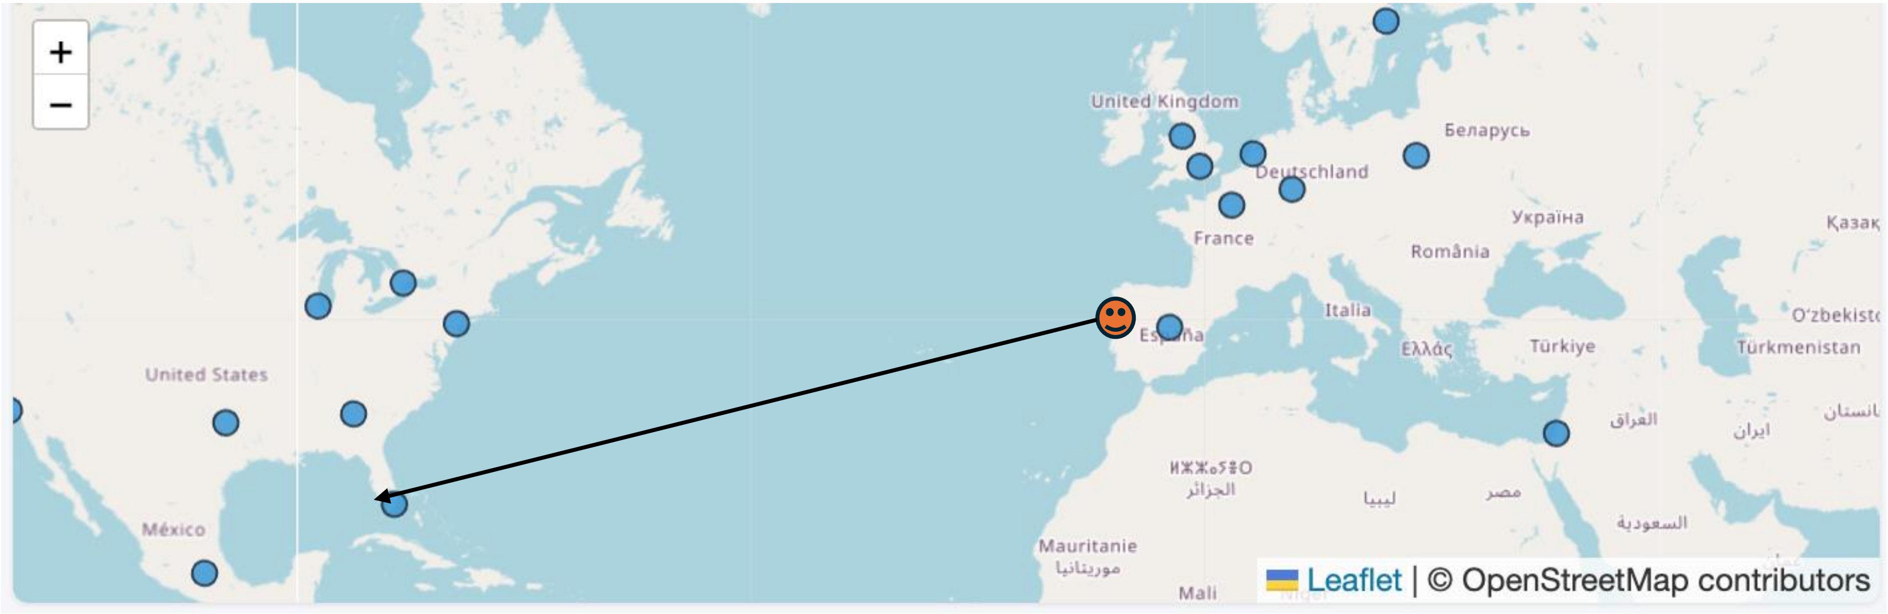

Introduction
Problem Statement

- Anycast relies on BGP to route clients
- But.. BGP routes based on shortest path -> not performance
- Case B) Client routes to a nearby PoP, with a long/slow path
  - *E.g.,* boomerang routing

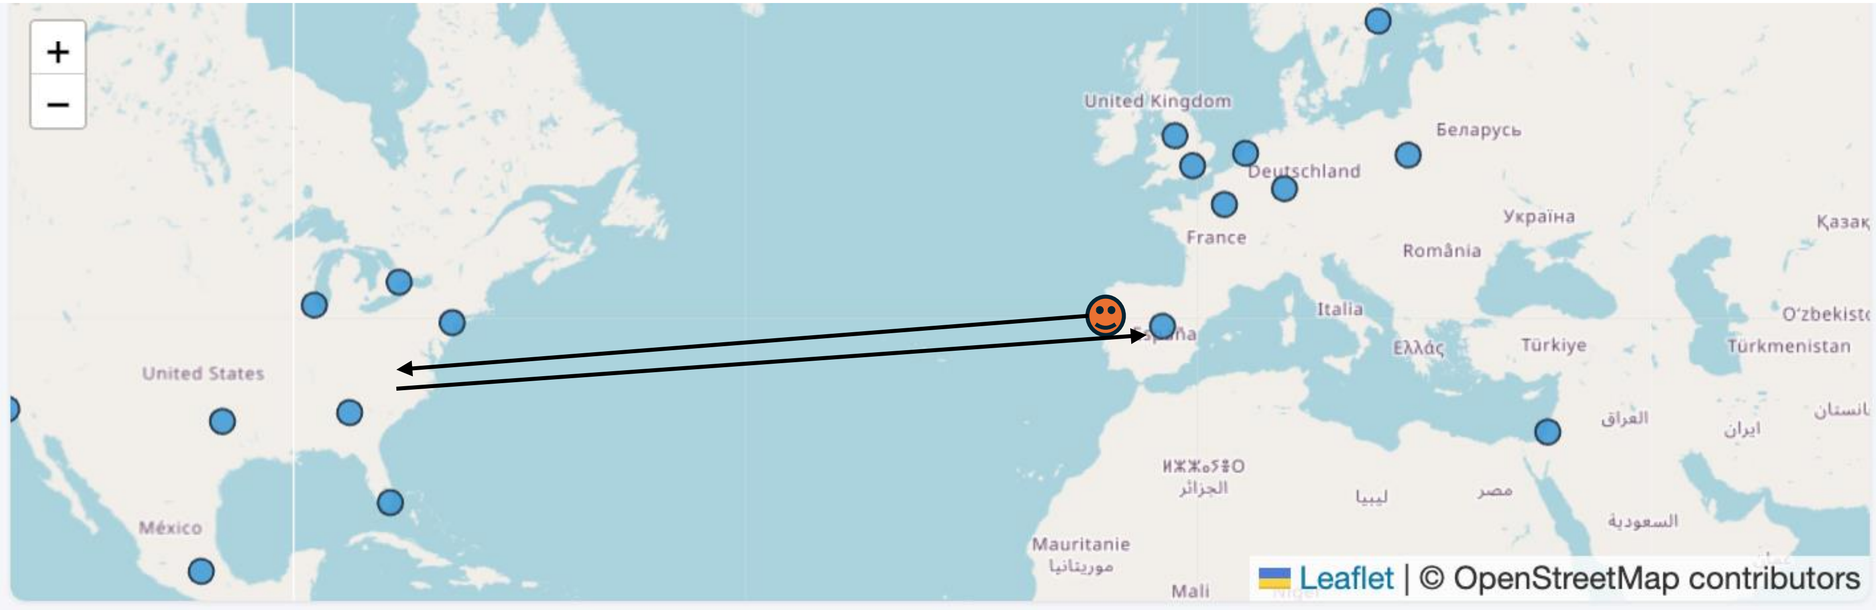

Introduction
Motivation

- Anycast routing can be sub-optimal
  - Clients reaching distant PoPs
    - High latencies, long paths
  - Poor load distribution amongst PoPs
- Performance metrics help with
  - Infrastructure (expansion) planning
  - Making traffic engineering decisions
- But how to get these metrics?

Monitoring DNS routing

Option 1) Passive DNS traffic analysis

Pros:

- Data already available
- No burden of active measurements
- High coverage (all existing clients)

Cons:

- Cannot prevent issues pro-actively
- RTT unknown
  - *Can be inferred from DNS over TCP traffic (requires random truncation)
- Not available when strategizing new PoP placements

Monitoring DNS routing

Option 2) Probing platforms

Pros:

- Can measure on-demand
  - *E.g.*, to measure effects of traffic changes
- RTTs available
- Does not require user traffic
  - Experimental deployments/prefixes

Cons:

- Rate-limiting restrictions
- Often credit costs involved
- Limited coverage

RIPE Atlas probe locations

Monitoring DNS routing

Option 3) Probing responsive hosts with anycast source IP

Previous work
Verfploeter

- Verfploeter maps anycast catchments *[1]*
  - I.e., which anycast site ‘catches’ which traffic
- How?
  - Ping host with anycast source IP
  - See where the ping reply comes back

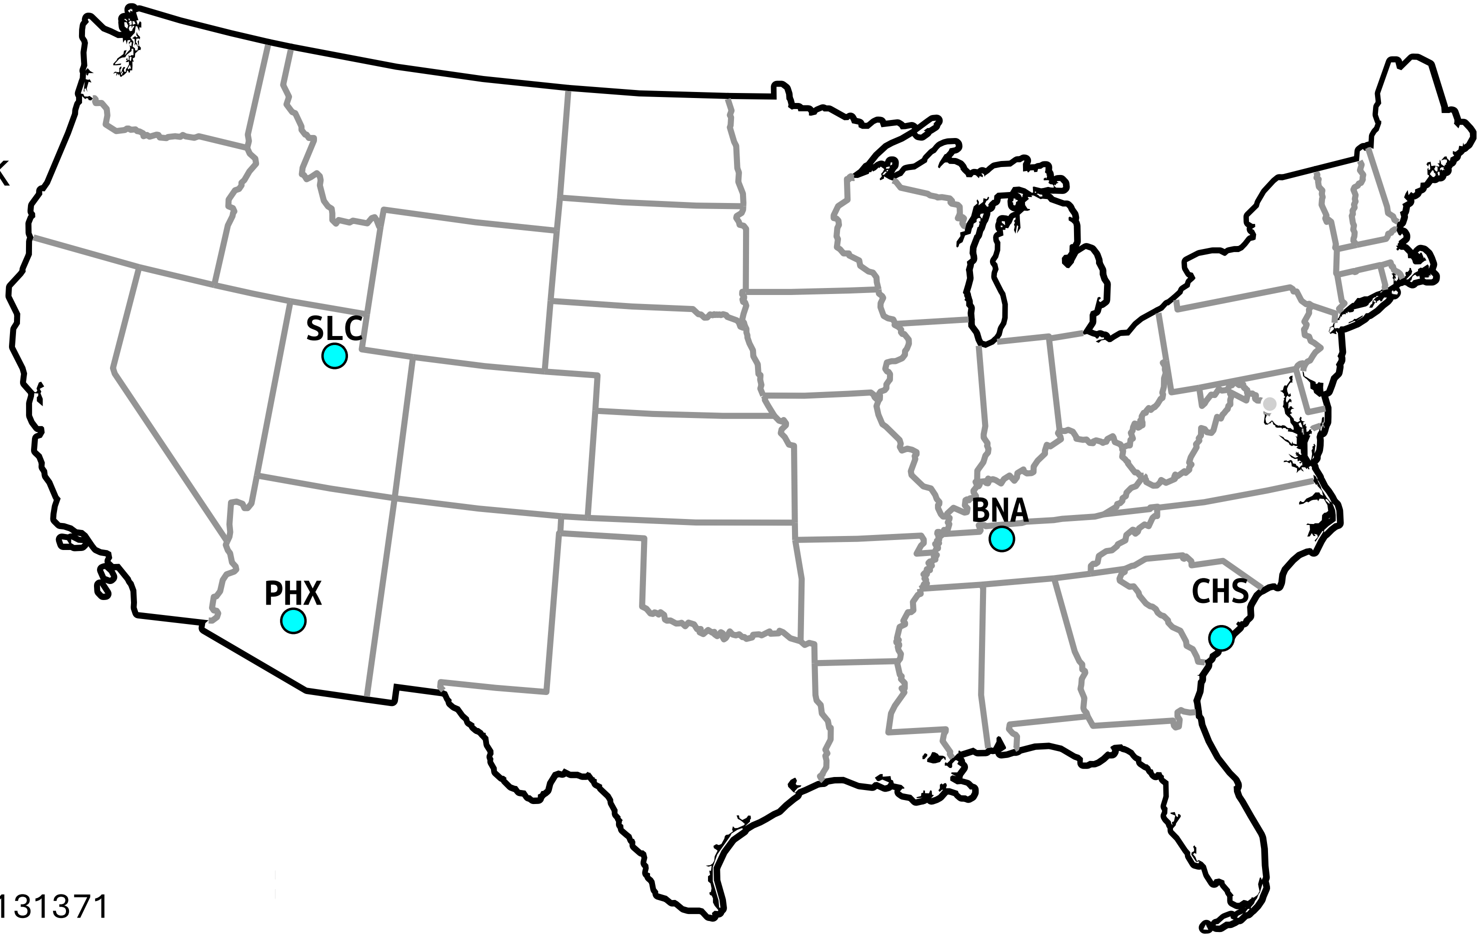

[1] https://dl.acm.org/doi/10.1145/3131365.3131371

Previous work
Verfploeter

- Verfploeter maps anycast catchments *[1]*
  - I.e., which anycast site ‘catches’ which traffic
- How?
  - Ping host with anycast source IP
  - See where the ping reply comes back

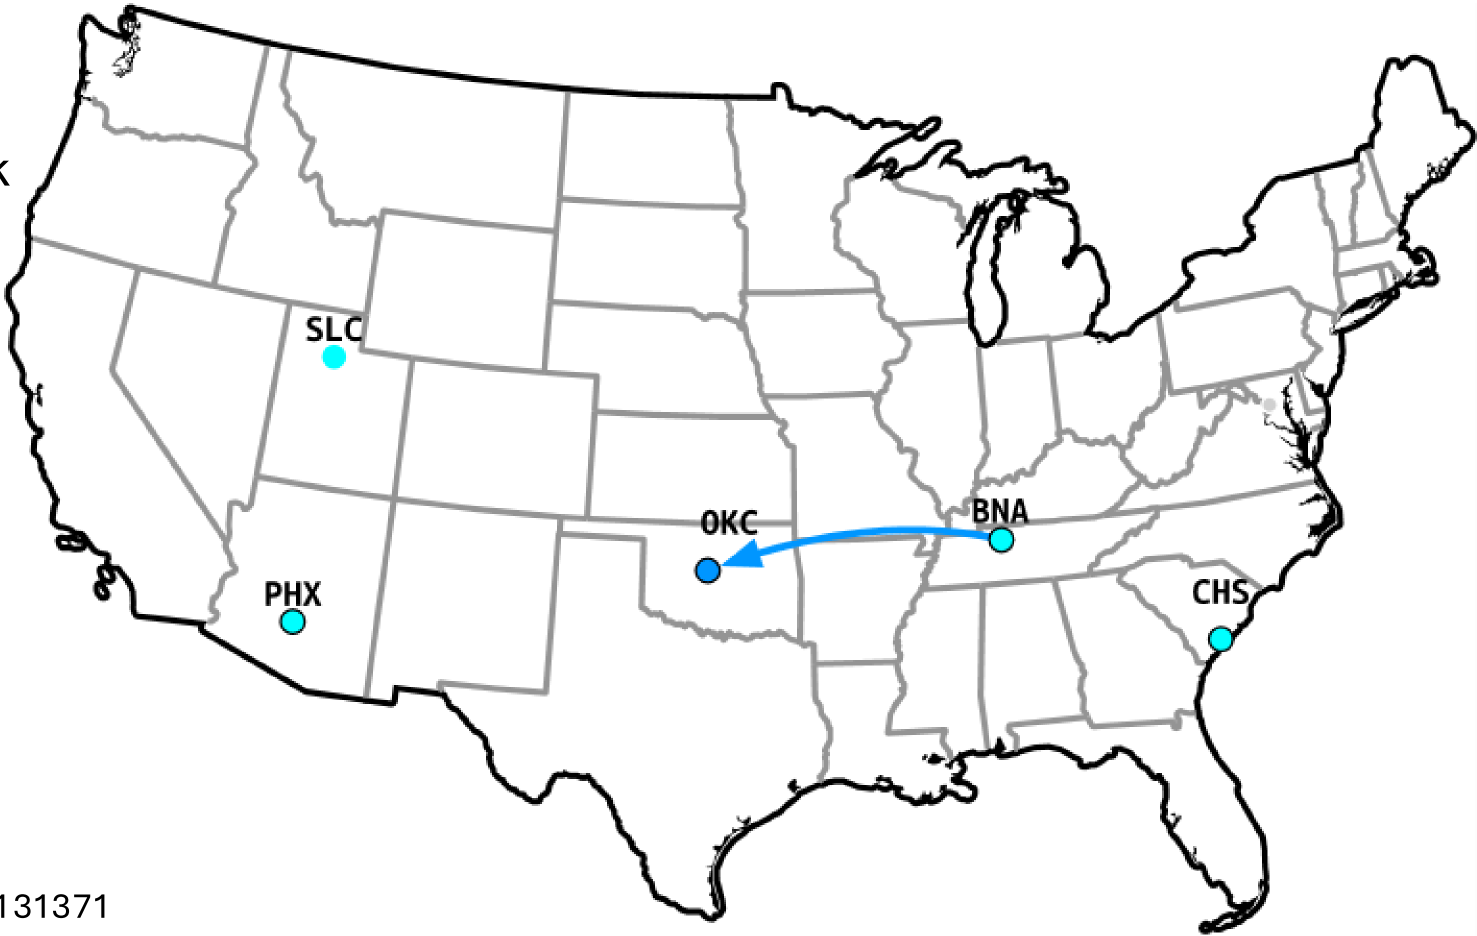

[1] https://dl.acm.org/doi/10.1145/3131365.3131371

Previous work
Verfploeter

- Verfploeter maps anycast catchments *[1]*
  - I.e., which anycast site ‘catches’ which traffic
- How?
  - Ping host with anycast source IP
  - See where the ping reply comes back

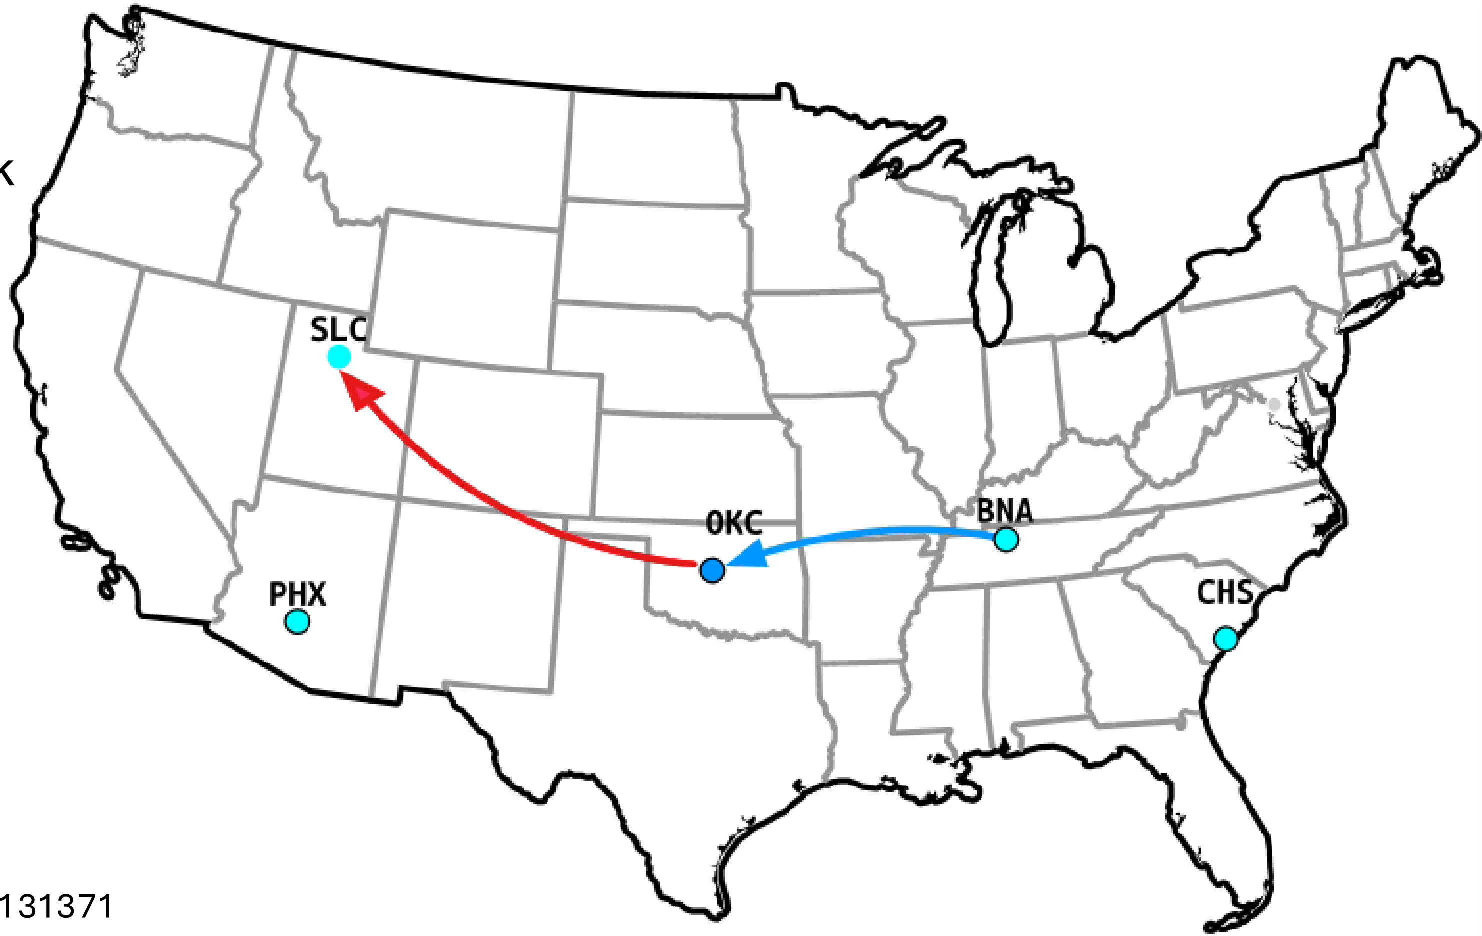

[1] https://dl.acm.org/doi/10.1145/3131365.3131371

Previous work
Verfploeter

- Verfploeter maps anycast catchments *[1]*
  - I.e., which anycast site ‘catches’ which traffic
- How?
  - Ping host with anycast source IP
  - See where the ping reply comes back
- Advantages
  - Covers 4 million /24s with ICMP responsive hosts
    - 70k+ ASes
  - Can measure pro-actively / on-demand
    - Experimental prefixes, traffic changes, …
- Disadvantages
  - Cannot measure RTTs
  - Not always clear what is the ‘optimal’ site
  - Nearby catchment does not guarantee a good path
    - Boomerang routing

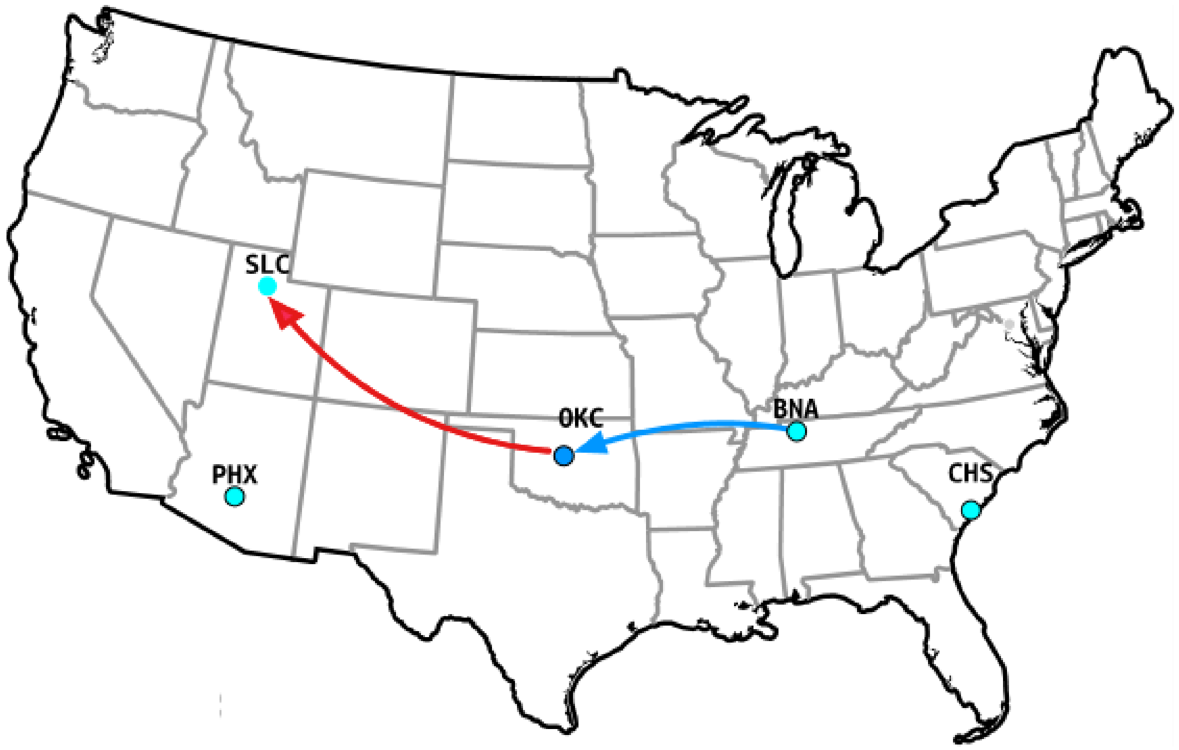

[1] https://dl.acm.org/doi/10.1145/3131365.3131371

Our open-source tooling

- Builds on the anycast probing method
- Designed as a ‘Swiss knife’
  - Large variety of supported measurements
- Low resource usage
- 8 MB static binary or docker image
- Supports TCP (SYNACK) and UDP
  - Increases coverage by 10%+
  - Measure catchments using multiple protocols
- Supports IPv6

Use cases
Anycast RTT

- **Probe 1: assess catchment (discovery probe)**
- Probe 2: probe from catcher (measurement probe)
    - RTT = receive time – transmit time

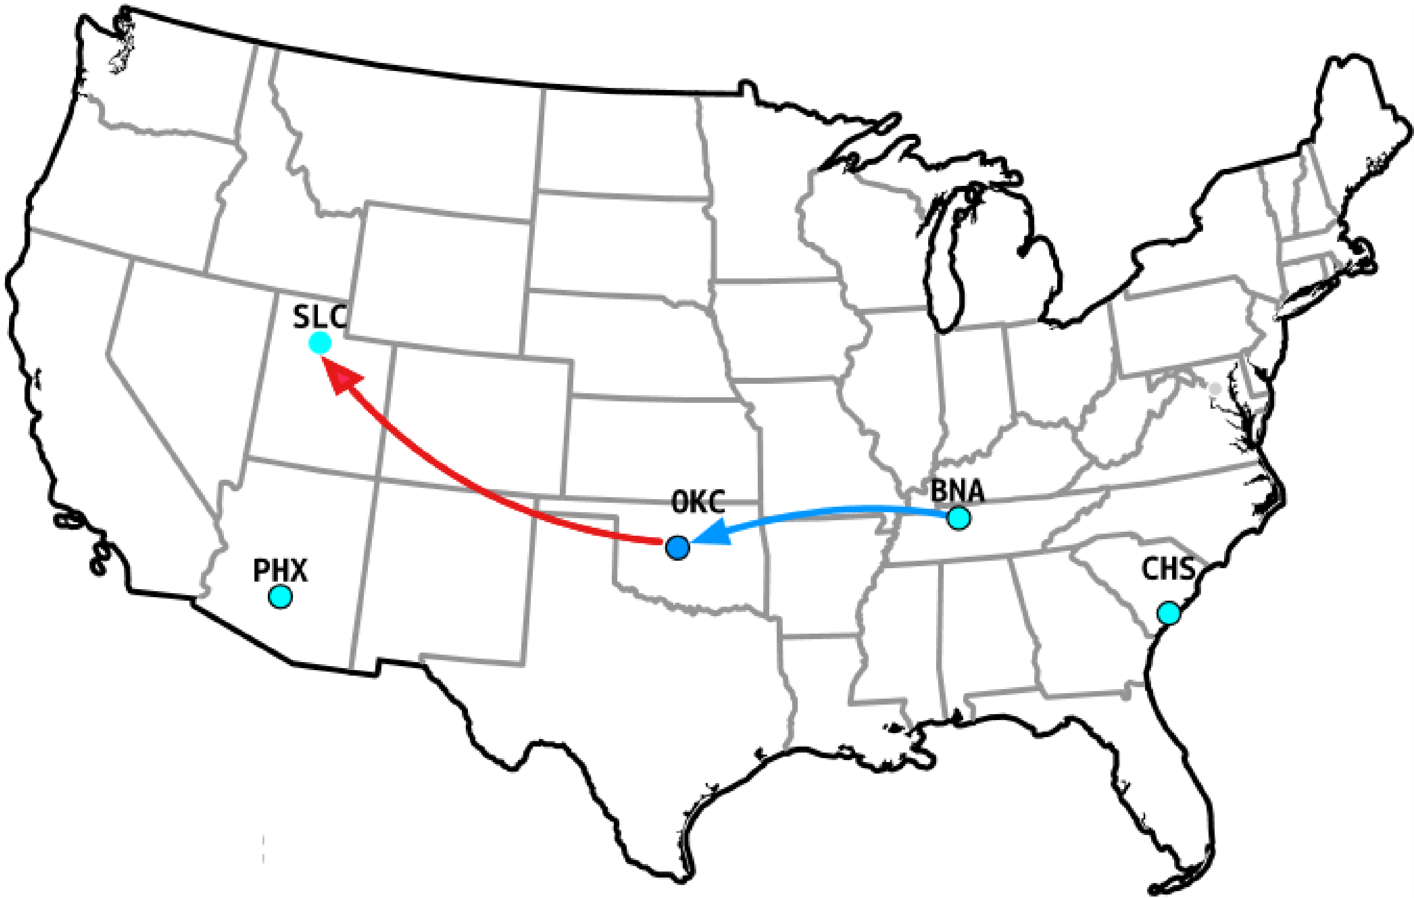

Use cases
Anycast RTT

- Probe 1: assess catchment (discovery probe)
- **Probe 2: probe from catcher (measurement probe)**
  - RTT = receive time – transmit time

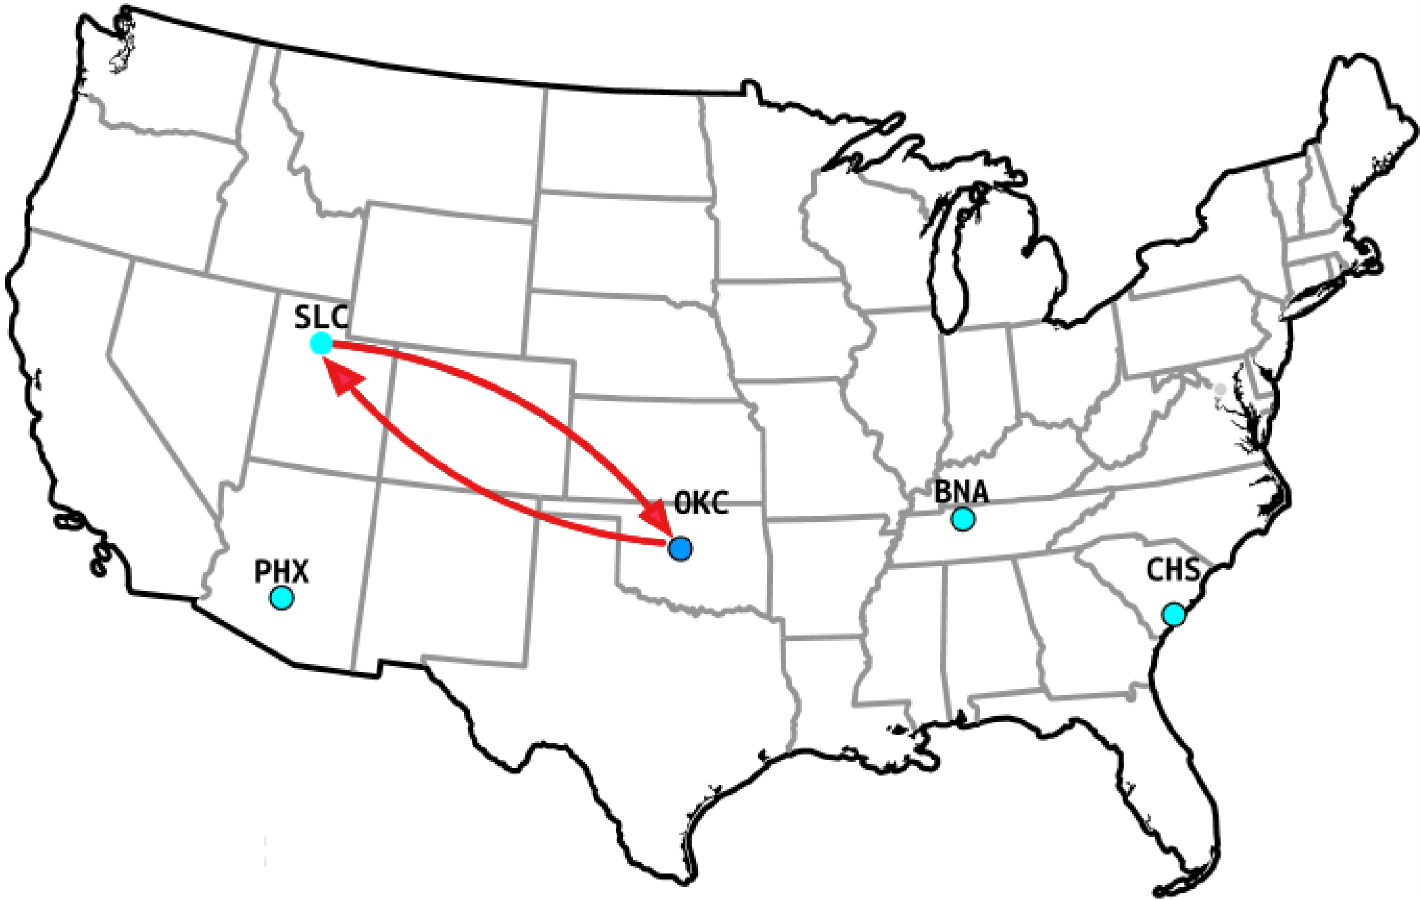

Validating Anycast RTTs

- Data from SIDN
- Tooling measures anycast RTTs (x-axis)
- TCP handshake latencies (y-axis)
  - From DNS over TCP queries

Takeaways

- Strong correlation
  - Measured RTTs are good indicators of performance
- RTTs mostly protocol agnostic
  - Similar for both TCP and ICMP

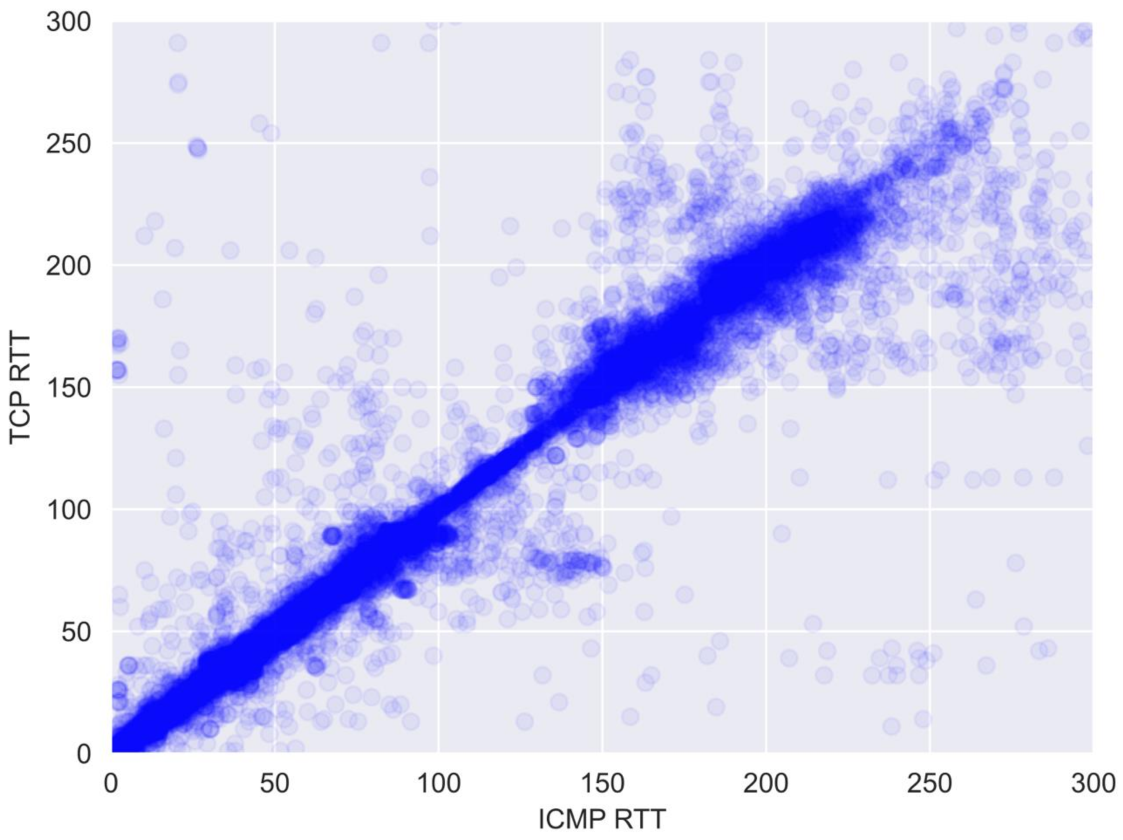

Figure 5: Scatterplot comparing our active ICMP latency measurement with passive TCP latency data.

Use cases
Inferring optimal Routing

- Probe target from all VPs using unicast IPs
  - Measures RTT with all possible sites
  - Site with lowest RTT -> ‘optimal’ site

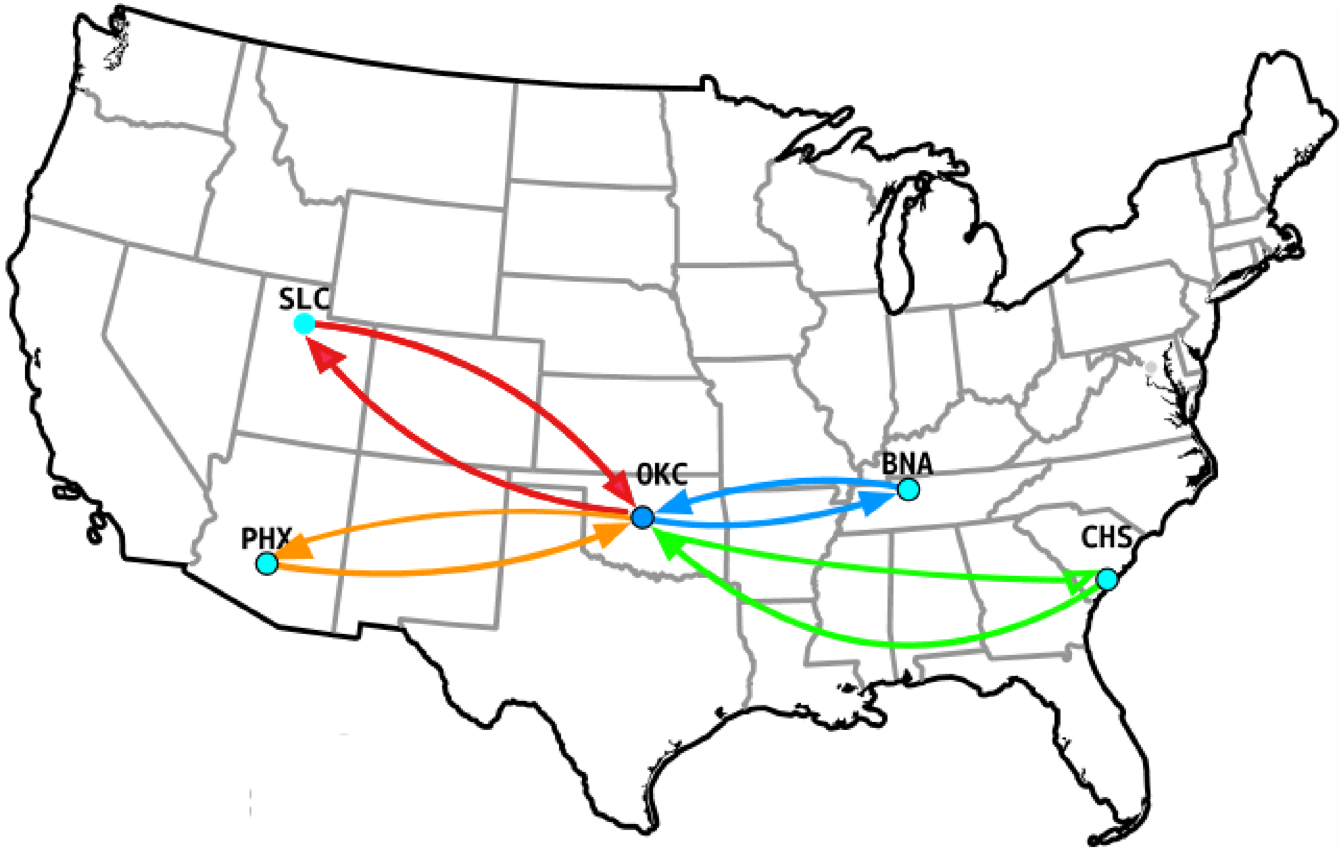

Use cases
Inferring optimal Routing

- Probe target from all VPs using unicast IPs
  - Measures RTT with all possible sites
  - Site with lowest RTT -> ‘optimal’ site
- Allows for:
  - Identifying best site
  - Quantifying possible RTT improvement

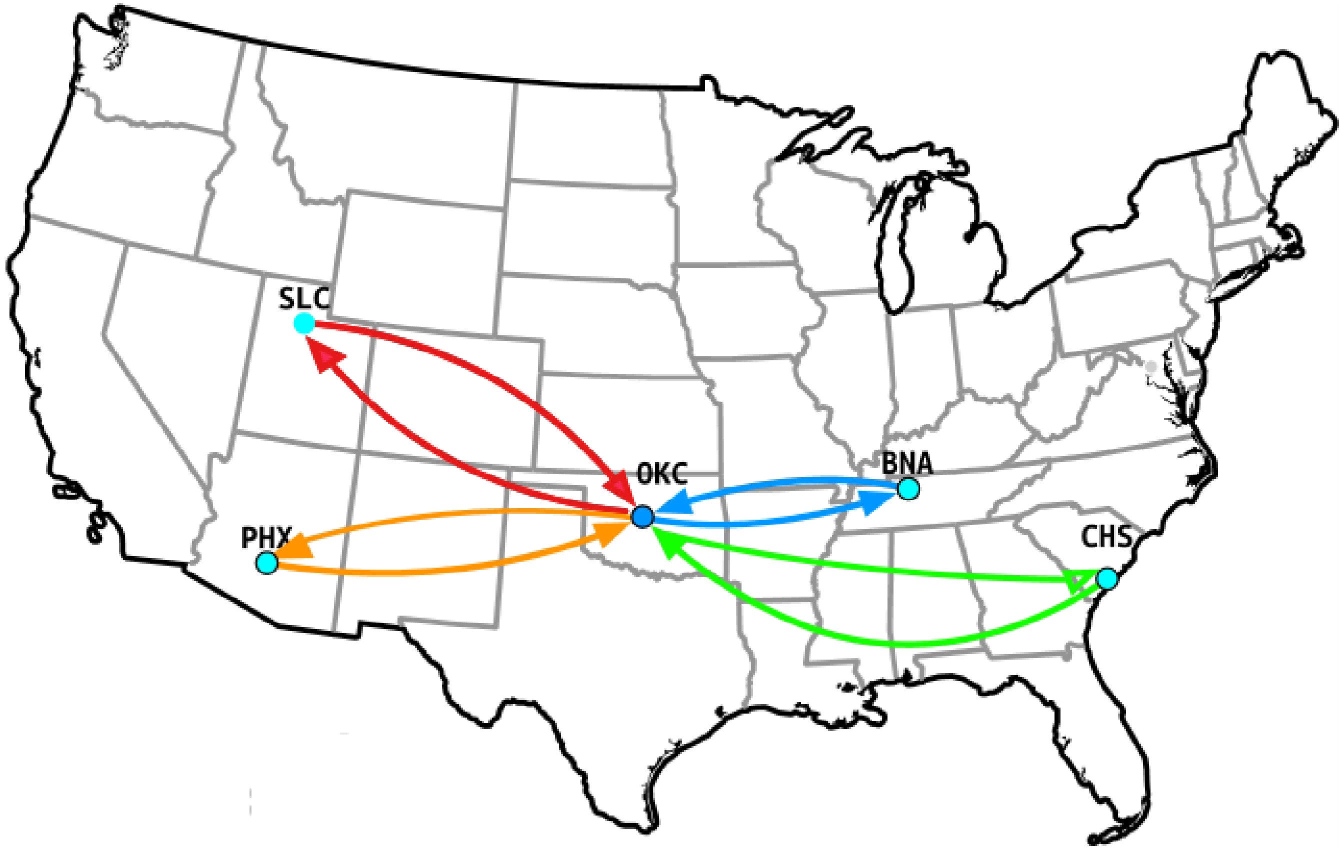

Results

- Measurement results for our anycast testbed
  - Deployed with Vultr (32 PoPs)
- Results meant to showcase the usefulness of the tooling
- Representative for those that anycast their DNS using Vultr

RTTs

- Latencies experienced by networks connecting to Vultr

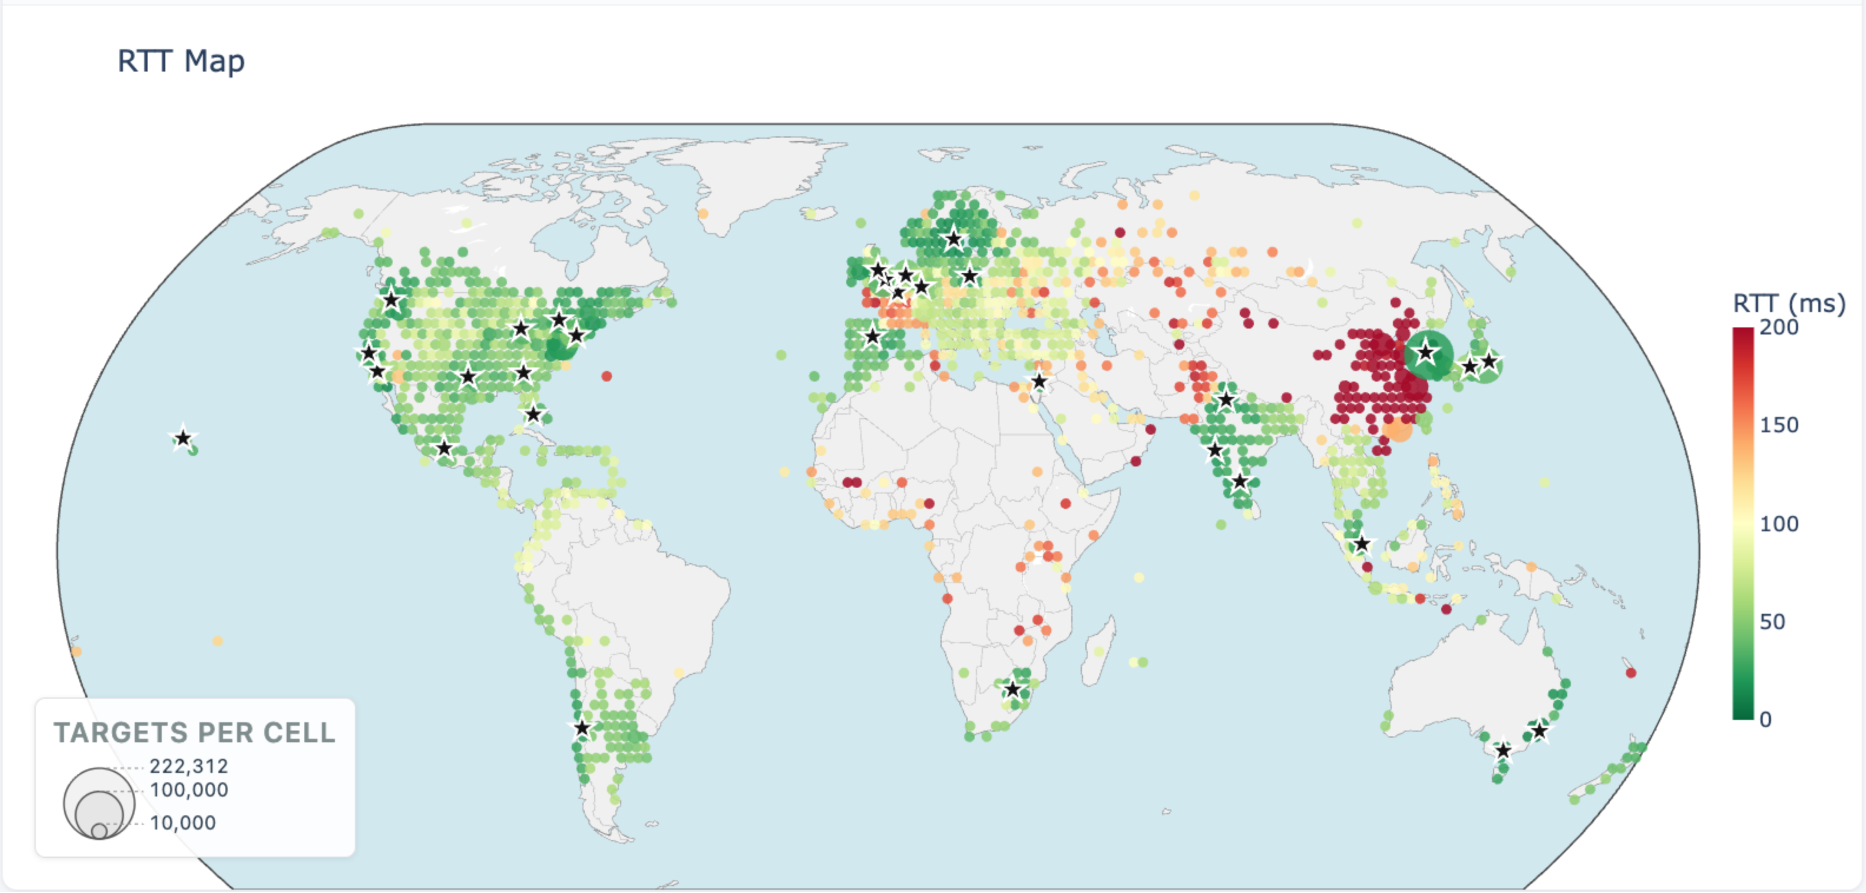

RTTs

- Optimal latencies for networks towards Vultr

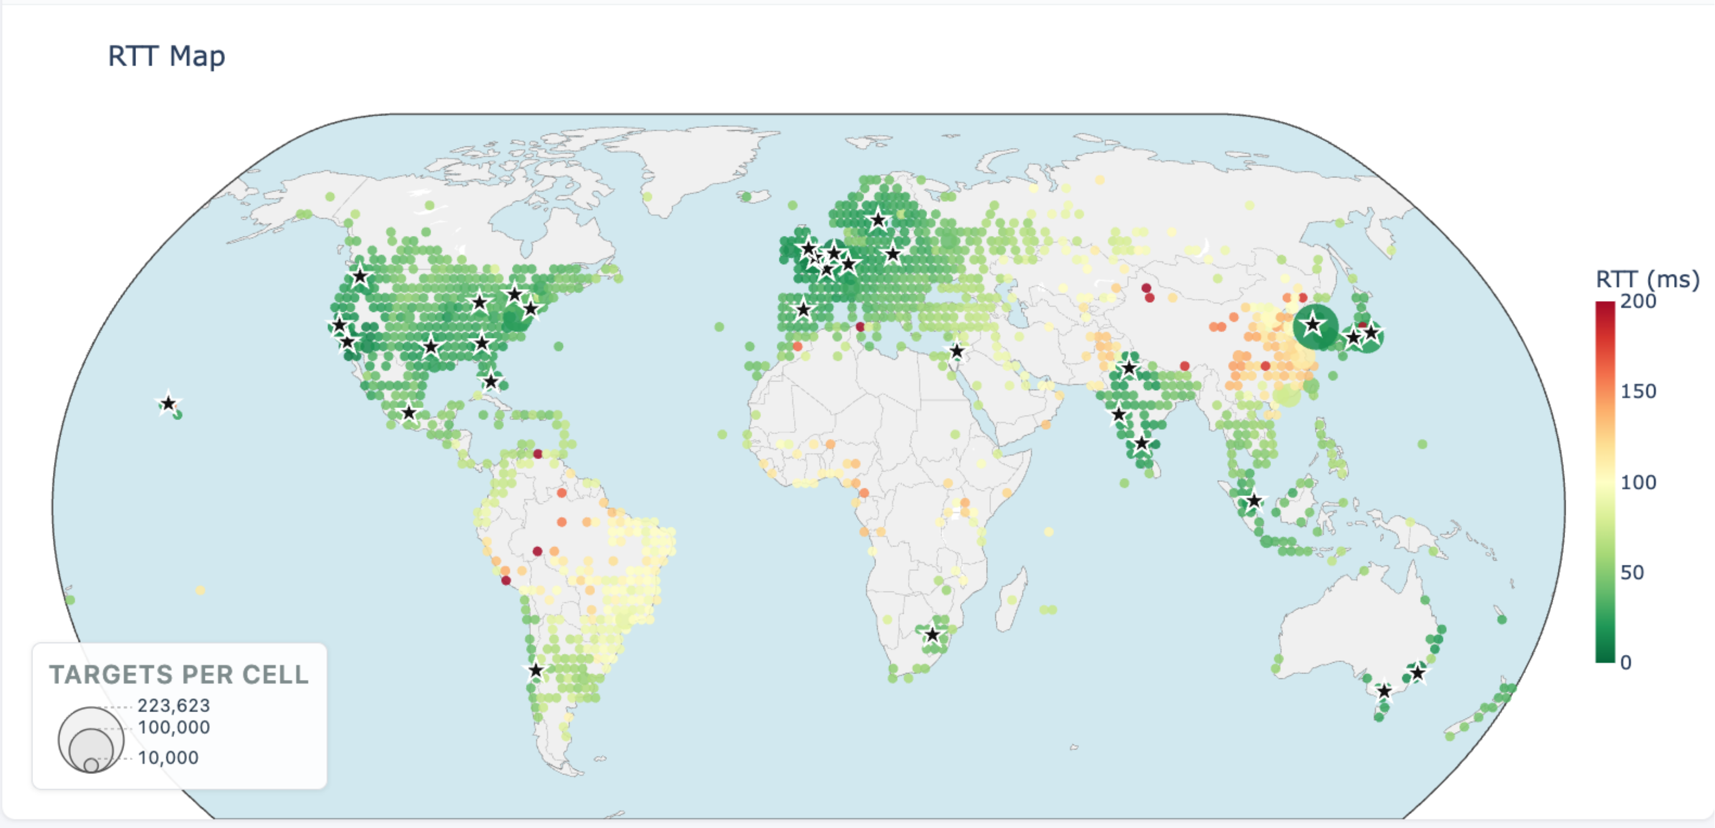

RTTs

- Optimal latencies for networks towards Vultr
- Takeaways:
  - Latencies shows who routes good, who routes bad
  - Optimal latencies show who ‘could’ route good -> TE needed
  - And who ‘cannot’ route good -> PoP placement needed

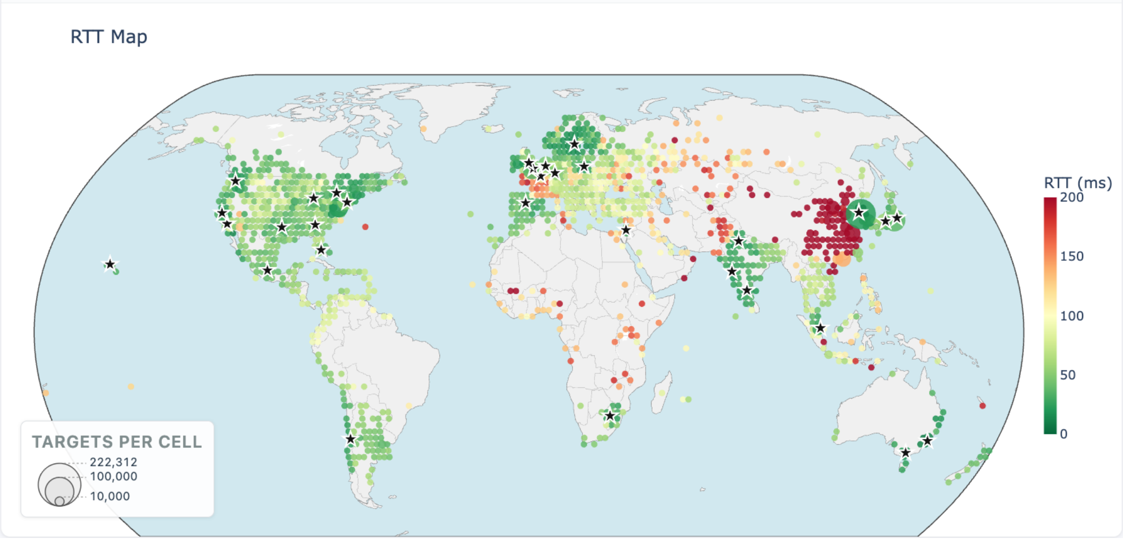

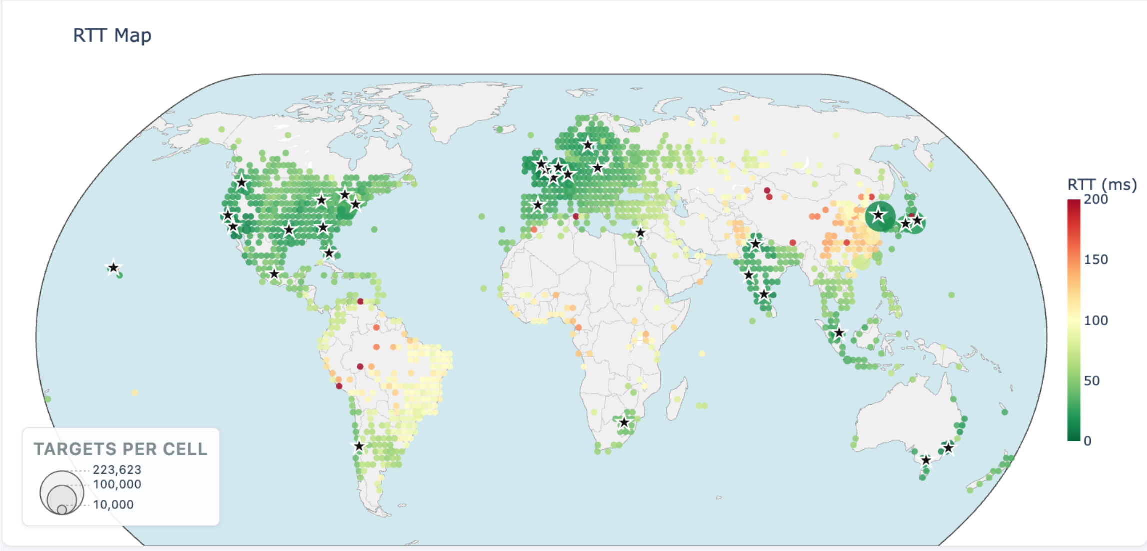

Quantifying Improvement

- Zooming in on China
  - Horrible median RTT (215)
  - Optimal can lower it to 110
  - Still bad however..
- If you anycast using Vultr
  - Expect poor connectivity with China

Geographic Lookup

Draw a bounding box on the map to explore all targets in that region — their anycast routing, latencies, and where they *should* route under optimal unicast.

△ **Draw polygon** `6-vertex polygon` × Clear ⛶ Fullscreen

432,331 targets *(showing first 500 by anycast RTT)*

ANYCAST RTT	
MEDIAN	215.4ms
P90	260.9ms
P25	185.8ms
MIN	0.9ms
MAX	4903.6ms
SAMPLES	425,313

BEST RTT (UNICAST)	
MEDIAN	109.8ms
P90	148.4ms
P25	82.1ms
MIN	0.2ms
MAX	52488.7ms
SAMPLES	431,685

Quantifying Improvement

- Zooming in on China
  - Horrible median RTT (215)
  - Optimal can lower it to 110
  - Still bad however..

▼ ASNs 200

ASN	OPERATOR	TARGETS	TOP POP	ANYCAST RTT	BEST RTT
AS4134	CHINANET-BAC...	111,647	de-fra-manycast	228.4ms	134.0ms
AS4837	CHINA169-BAC...	90,880	de-fra-manycast	219.0ms	90.3ms
AS37963	ALIBABA-CN-N...	24,915	gb-lhr-manycast	218.1ms	97.8ms
AS9808	CHINAMOBILE-...	21,340	gb-lhr-manycast	237.9ms	129.6ms
AS45090	TENCENT-NET-...	17,739	de-fra-manycast	217.9ms	123.3ms

Geographic Lookup

Draw a bounding box on the map to explore all targets in that region — their anycast routing, latencies, and where they *should* route under optimal unicast.

△ Draw polygon `6-vertex polygon` × Clear ⛶ Fullscreen

432,331 targets *(showing first 500 by anycast RTT)*

ANYCAST RTT	
MEDIAN	215.4ms
P90	260.9ms
P25	185.8ms
MIN	0.9ms
MAX	4903.6ms
SAMPLES	425,313

BEST RTT (UNICAST)	
MEDIAN	109.8ms
P90	148.4ms
P25	82.1ms
MIN	0.2ms
MAX	52488.7ms
SAMPLES	431,685

RTTs
mx-mex

- Actual (left) and optimal (right) RTTs for our second-largest PoP by catchment
- RTTs are optimal for Mexican networks
- But south US networks have good RTTs
  - RTT gain probably not worth the traffic engineering

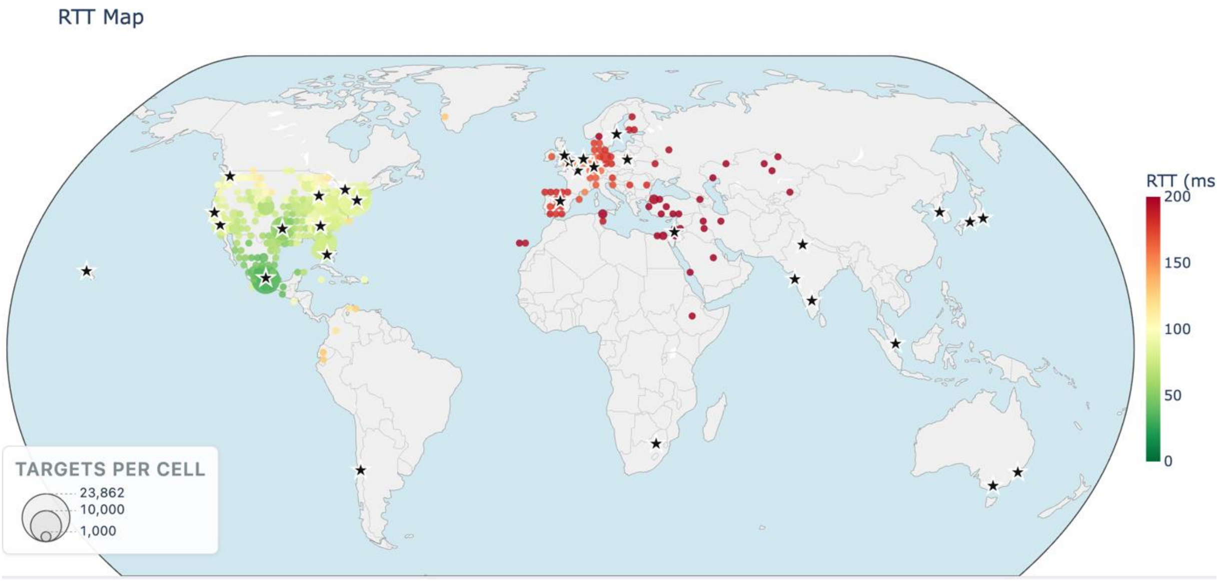

RTT against Catchment

- Measured anycast RTTs (y-axes)
- Distance from catching PoP (x-axis)
  - Using IPInfo

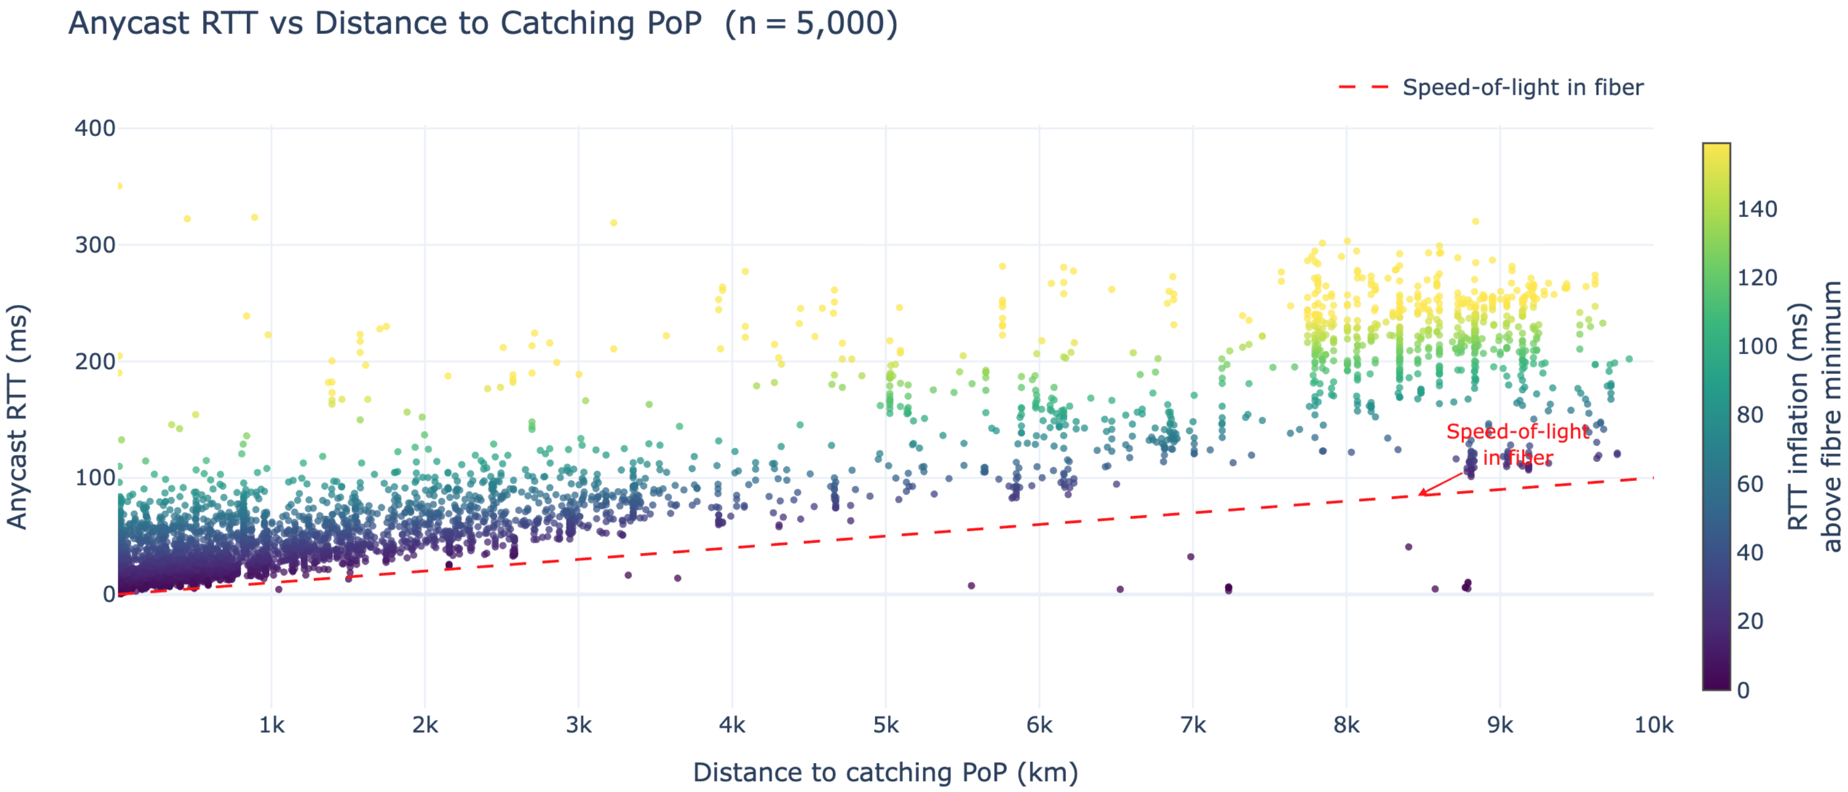

RTT against Catchment

- Measured anycast RTTs (y-axes)
- Distance from catching PoP (x-axis)
  - Using IPInfo

Takeaways

- Catchments can be misleading
  - Catchment inferred distance =/= path length
  - *E.g.*, boomerang routes
- RTT is a better indicator of performance
  - Accounts for congested links
  - Accounts for long paths (to nearby targets)

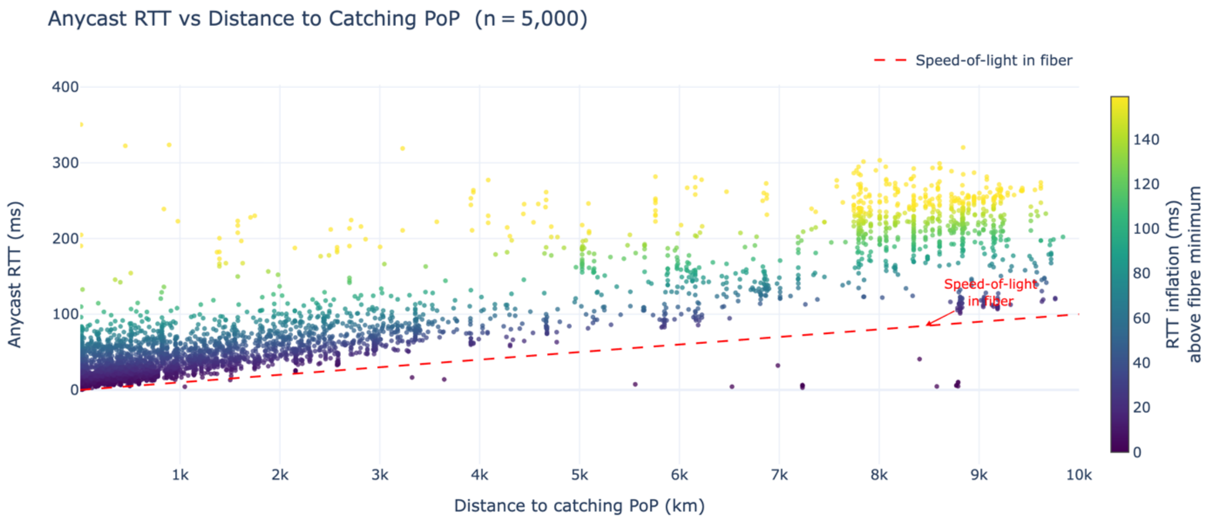

TTL as performance indicator

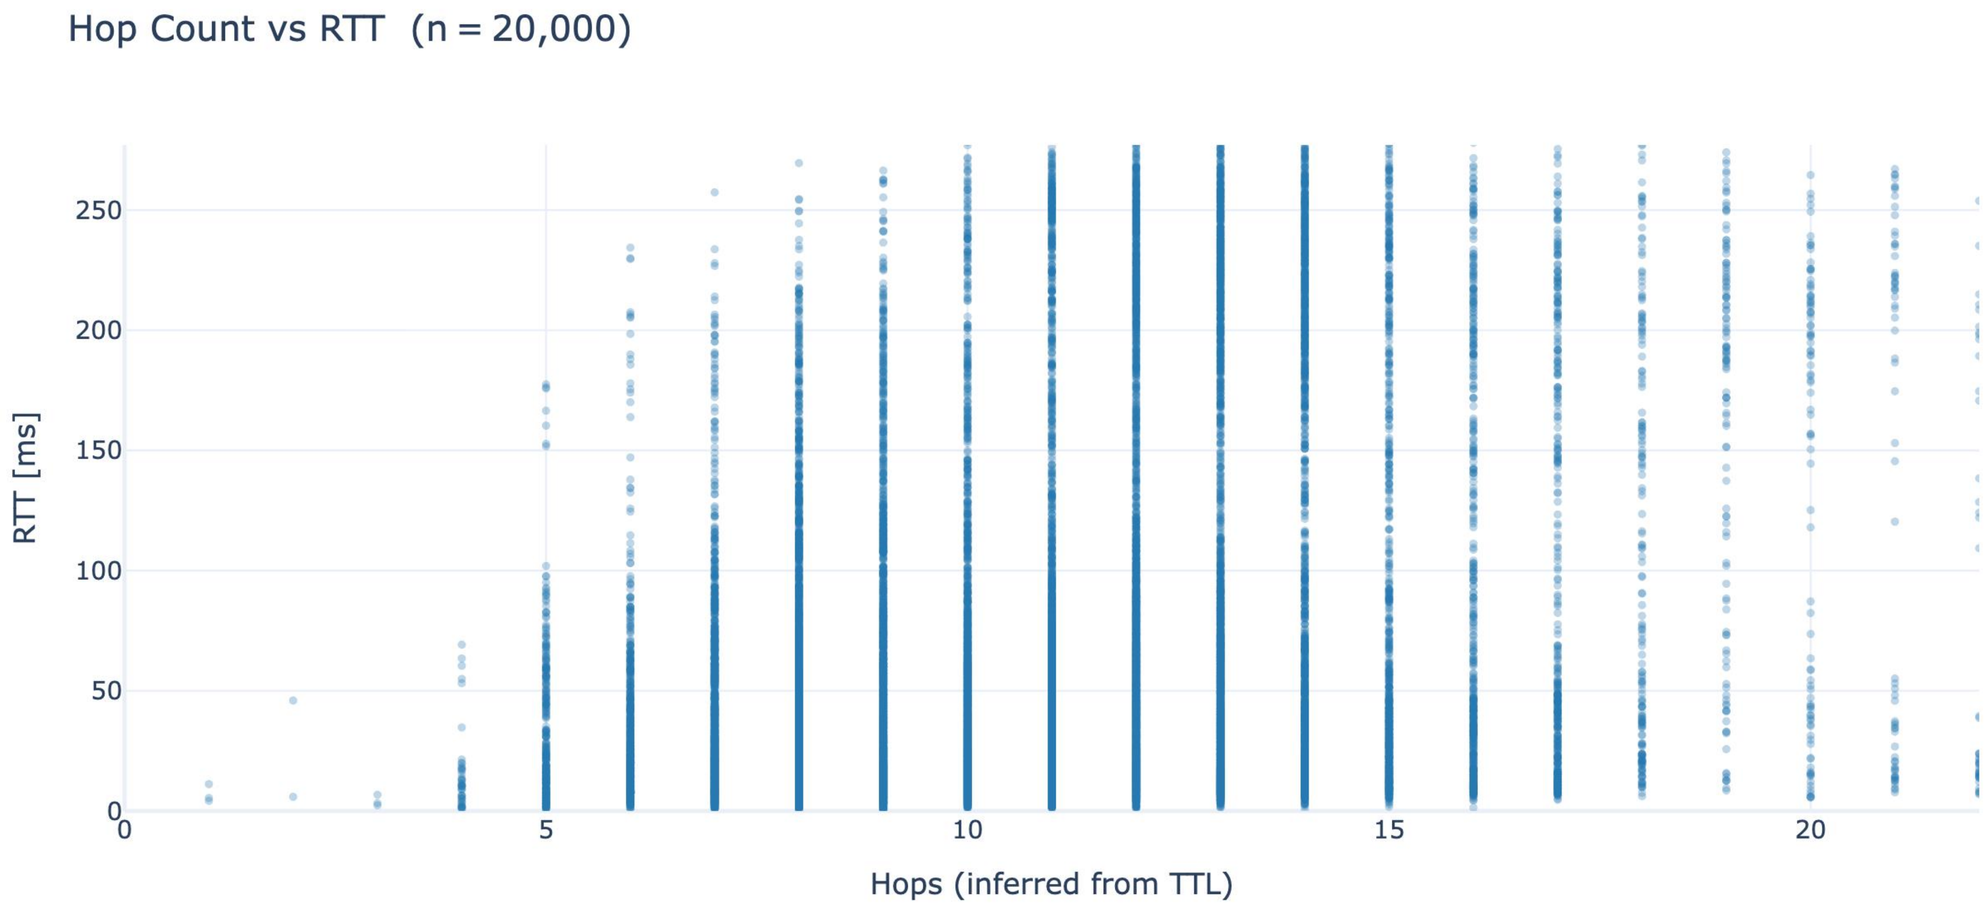

TTL as performance indicator

- TTL is a poor indicator of performance
  - E.g., tunneling (single hop -> large distance & RTT inflation)
- If you optimize your DNS deployment for TTL
  - This will not improve performance for clients
- Tooling gives better metrics
  - RTTs and catchments

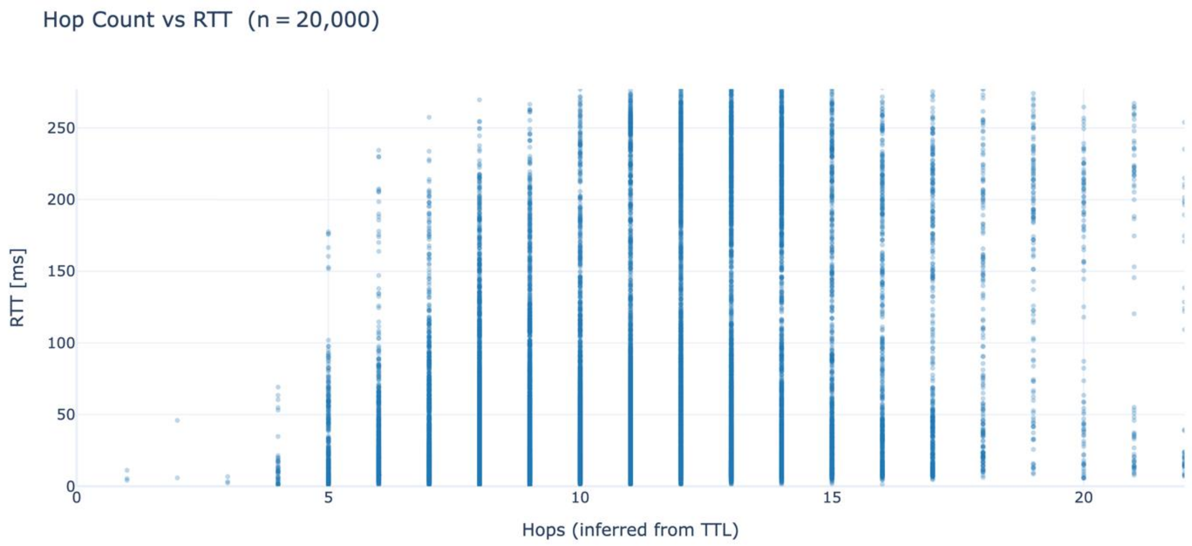

Other improvements

- Anycast traceroute
  - Time exceeded replies can map additional networks
  - Helps with troubleshooting suboptimal paths
- Measure networks experiencing anycast site flipping

Conclusion

- We built MAnycastR (publicly available at github.com/rhendriks/MAnycastR)
- Originally designed for research purposes
  - Anycast census
- Showcases highlight its operational value
  - Catchment and latency mappings
  - Can help with TE decisions
  - And with PoP placement strategizing
- Tooling might be helpful in attack mitigation
  - Measure traffic redirection results after applying them in production

Q&A / Discussion

- What do you (DNS operators) care about?
  - Does RTT matter for DNS anycast?
  - Do you optimize for short paths?
  - Is load-distribution more important?
- Does this tooling help with realizing these goals?
  - What requirements do you have for using it?
  - What features would you like to see?
- We are open to feedback from the community
- Interested in active collaborations
  - Especially for measuring anycast routing in developing regions
- Contact: remi.hendriks@utwente.nl

Appendix slides

Catchment Analysis

- Grouping catchments by ASN / announced BGP prefix
  - Most ASes/prefixes route consistently to the same PoP
  - Long tail -> ASes with private backbone / global presence

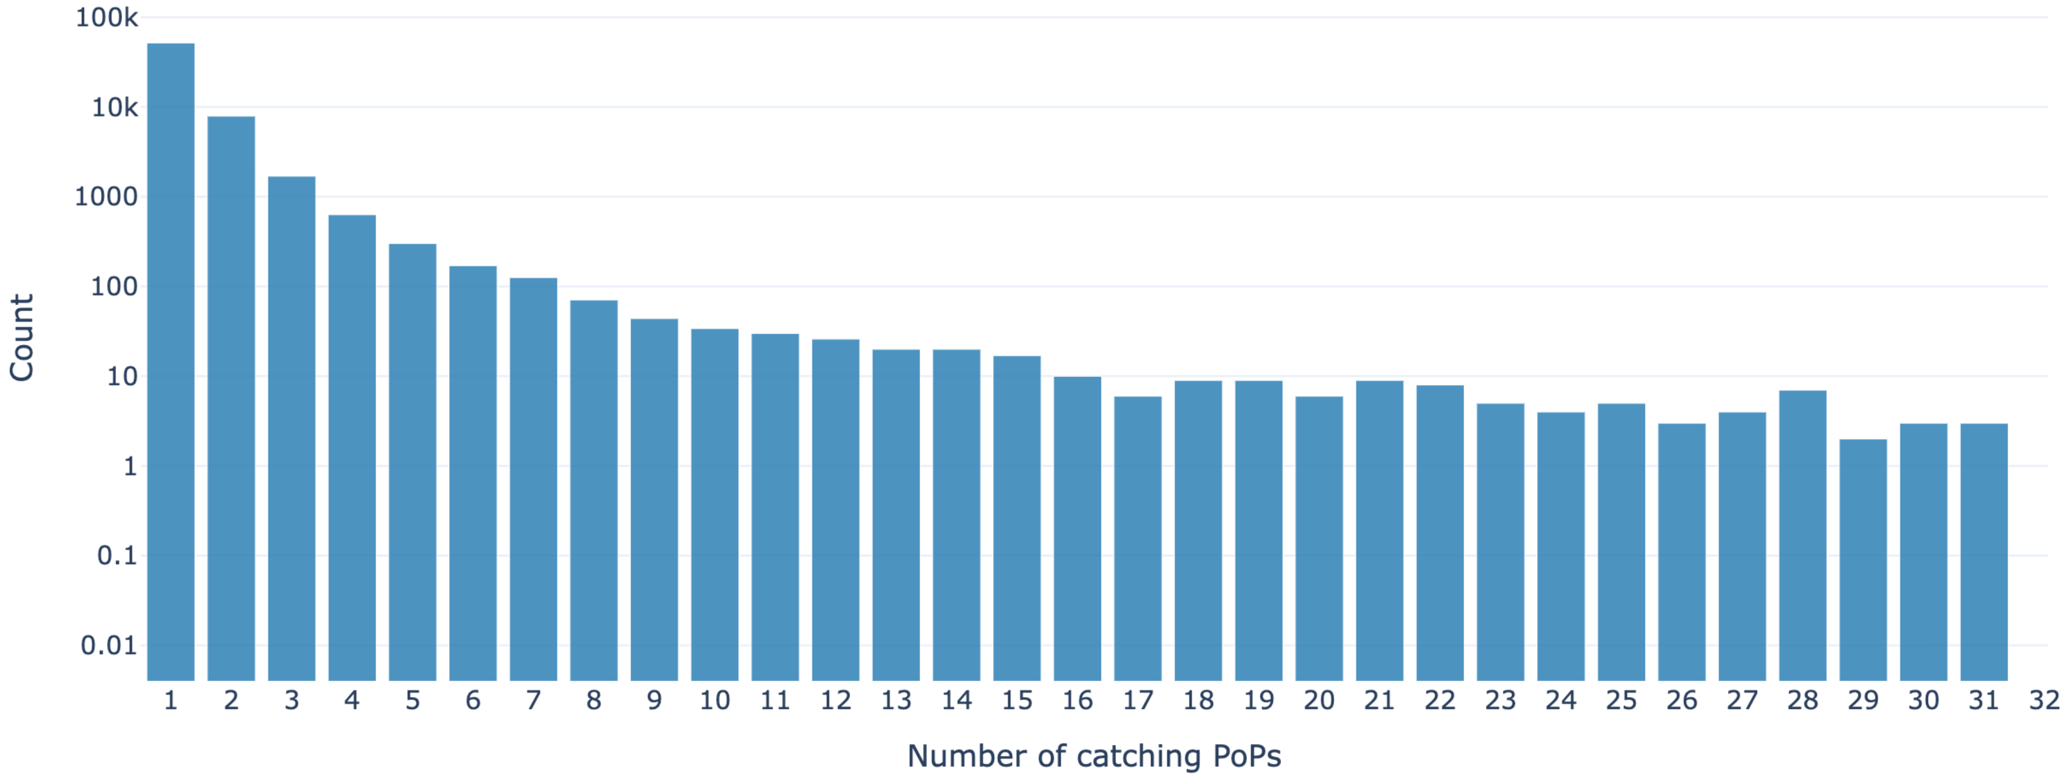

Catchment Analysis

- Grouping catchments by ASN / announced BGP prefix
  - Most ASes/prefixes route consistently to the same PoP
  - Long tail -> ASes with private backbone / global presence
- Takeaway
  - 82% of ASes have homogenous routing properties (same catchment for all networks)
  - 18% require multiple TE decisions
  - Long tail (CDNs)

Recommendation

DNS operators must expect

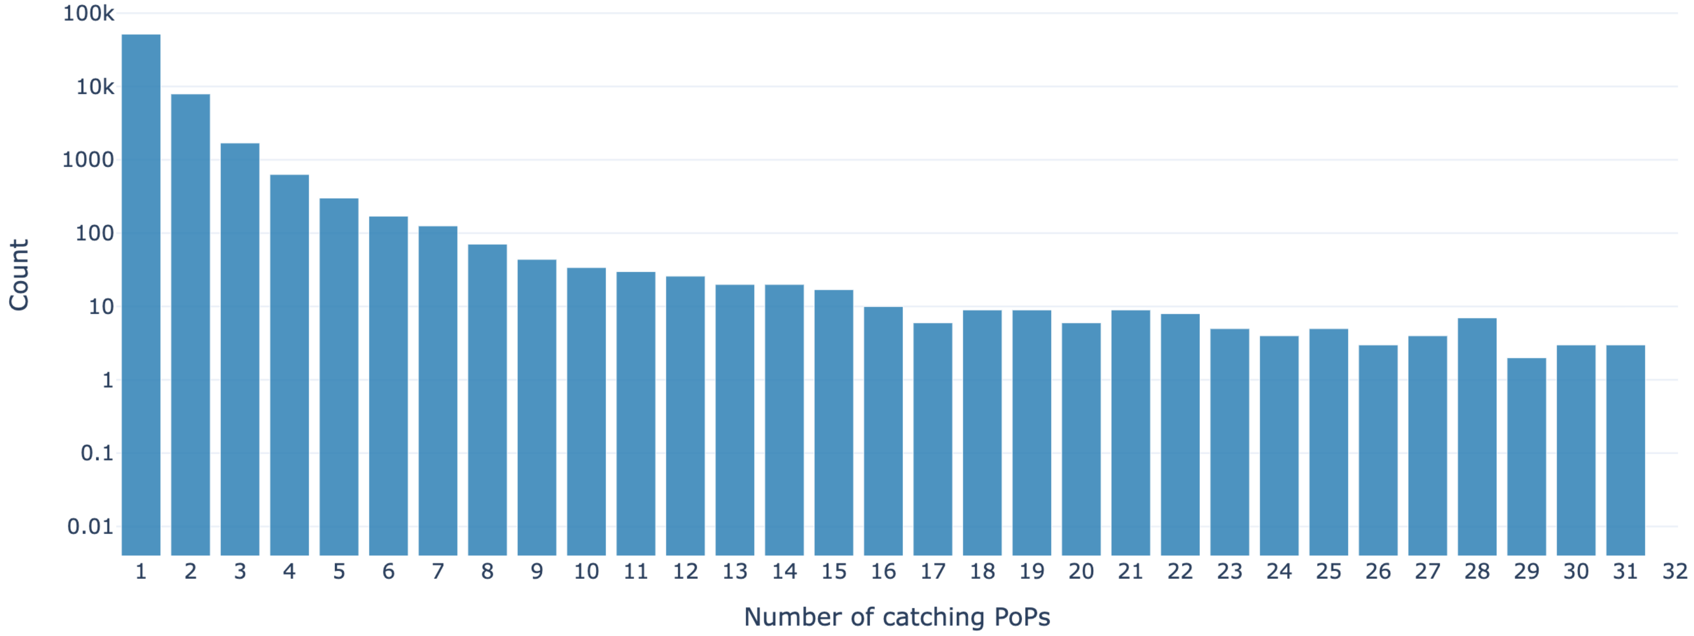

Catchment Analysis

- Grouping catchments by ASN / announced BGP prefix
  - Most ASes/prefixes route consistently to the same PoP
  - Long tail -> ASes with private backbone / global presence
- Takeaway
  - 82% of ASes consistently route to a single PoP
  - 18% to multiple
  - Long tail (CDNs)

Recommendation

DNS operators must expect some networks (especially CDNs) to reach multiple PoPs

Catchments

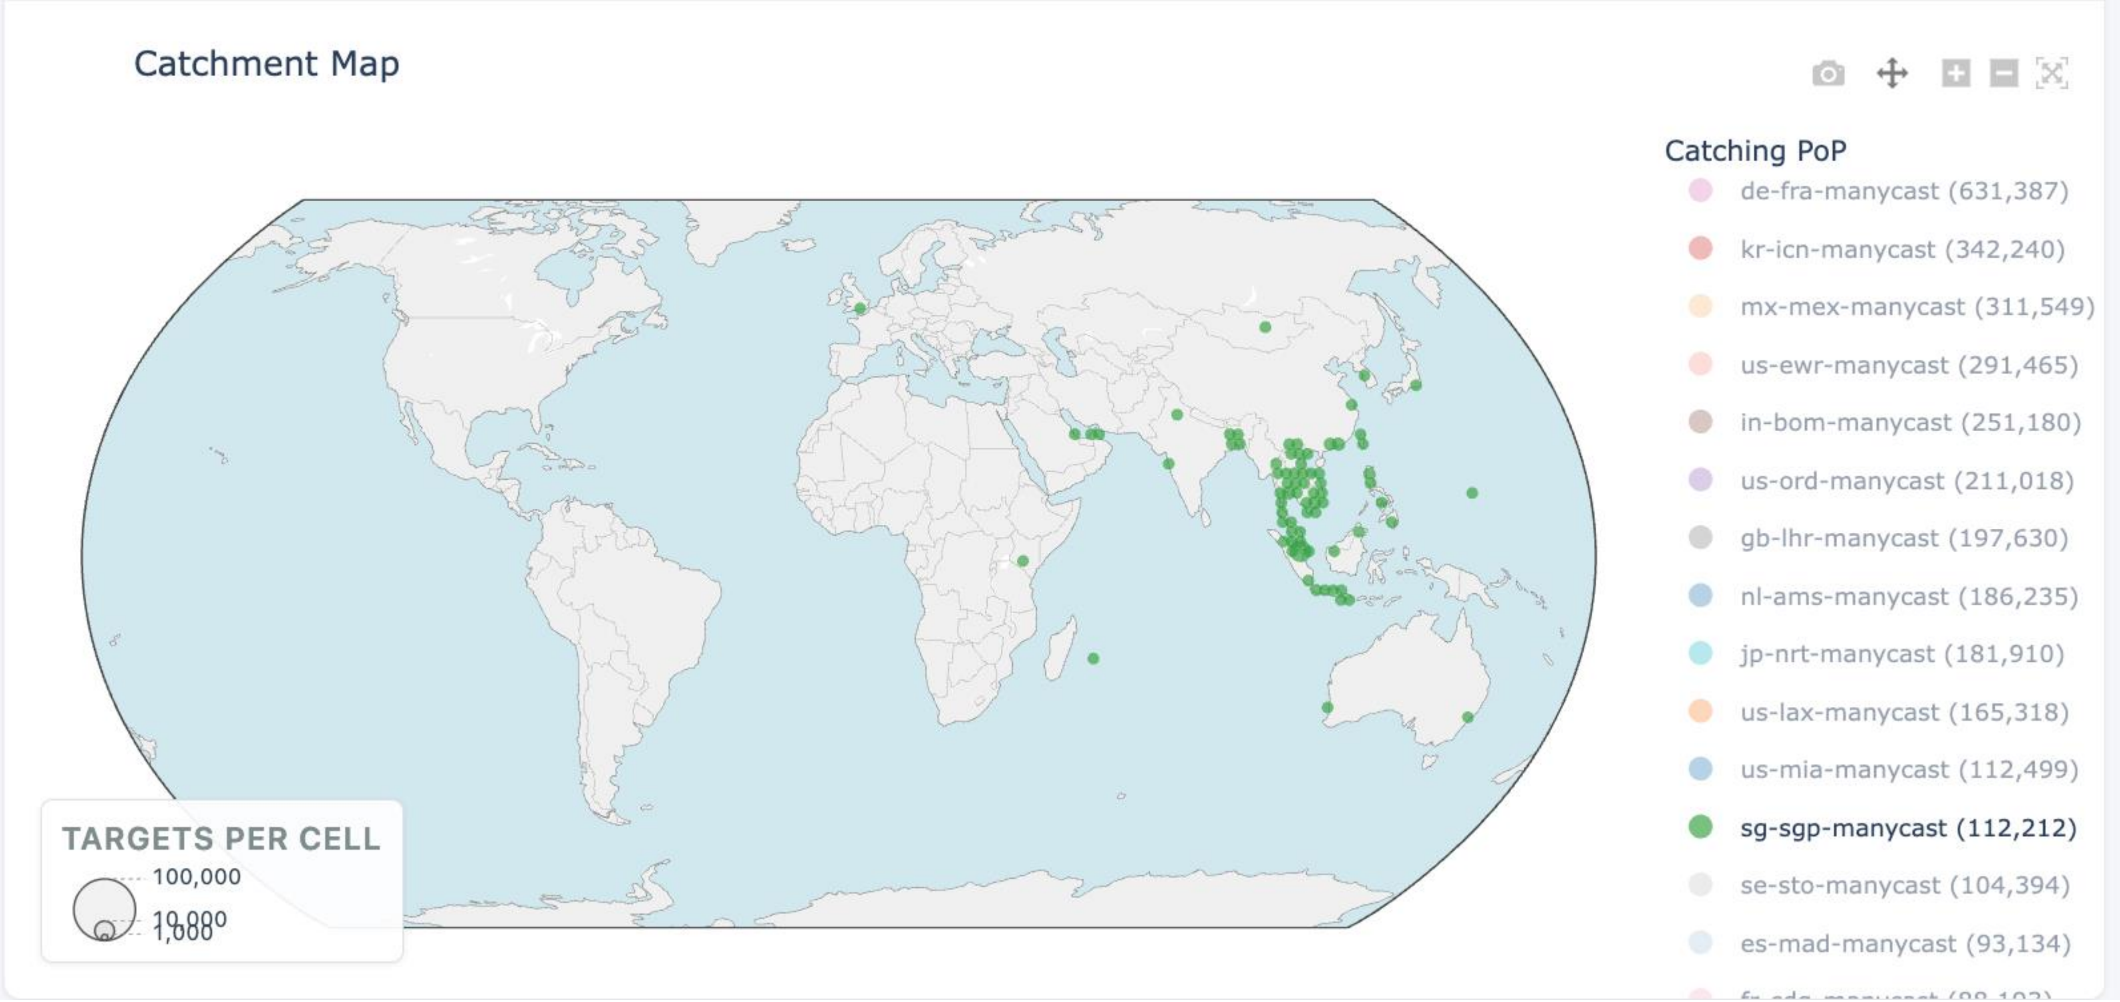
Catchment Map
Catching PoP
de-fra-manycast (631,387)
kr-icn-manycast (342,240)
mx-mex-manycast (311,549)
us-ewr-manycast (291,465)
in-bom-manycast (251,180)
us-ord-manycast (211,018)
gb-lhr-manycast (197,630)
nl-ams-manycast (186,235)
jp-nrt-manycast (181,910)
us-lax-manycast (165,318)
us-mia-manycast (112,499)
sg-sgp-manycast (112,212)
se-sto-manycast (104,394)
es-mad-manycast (93,134)
TARGETS PER CELL
100,000
10,000
1,000

Catchments

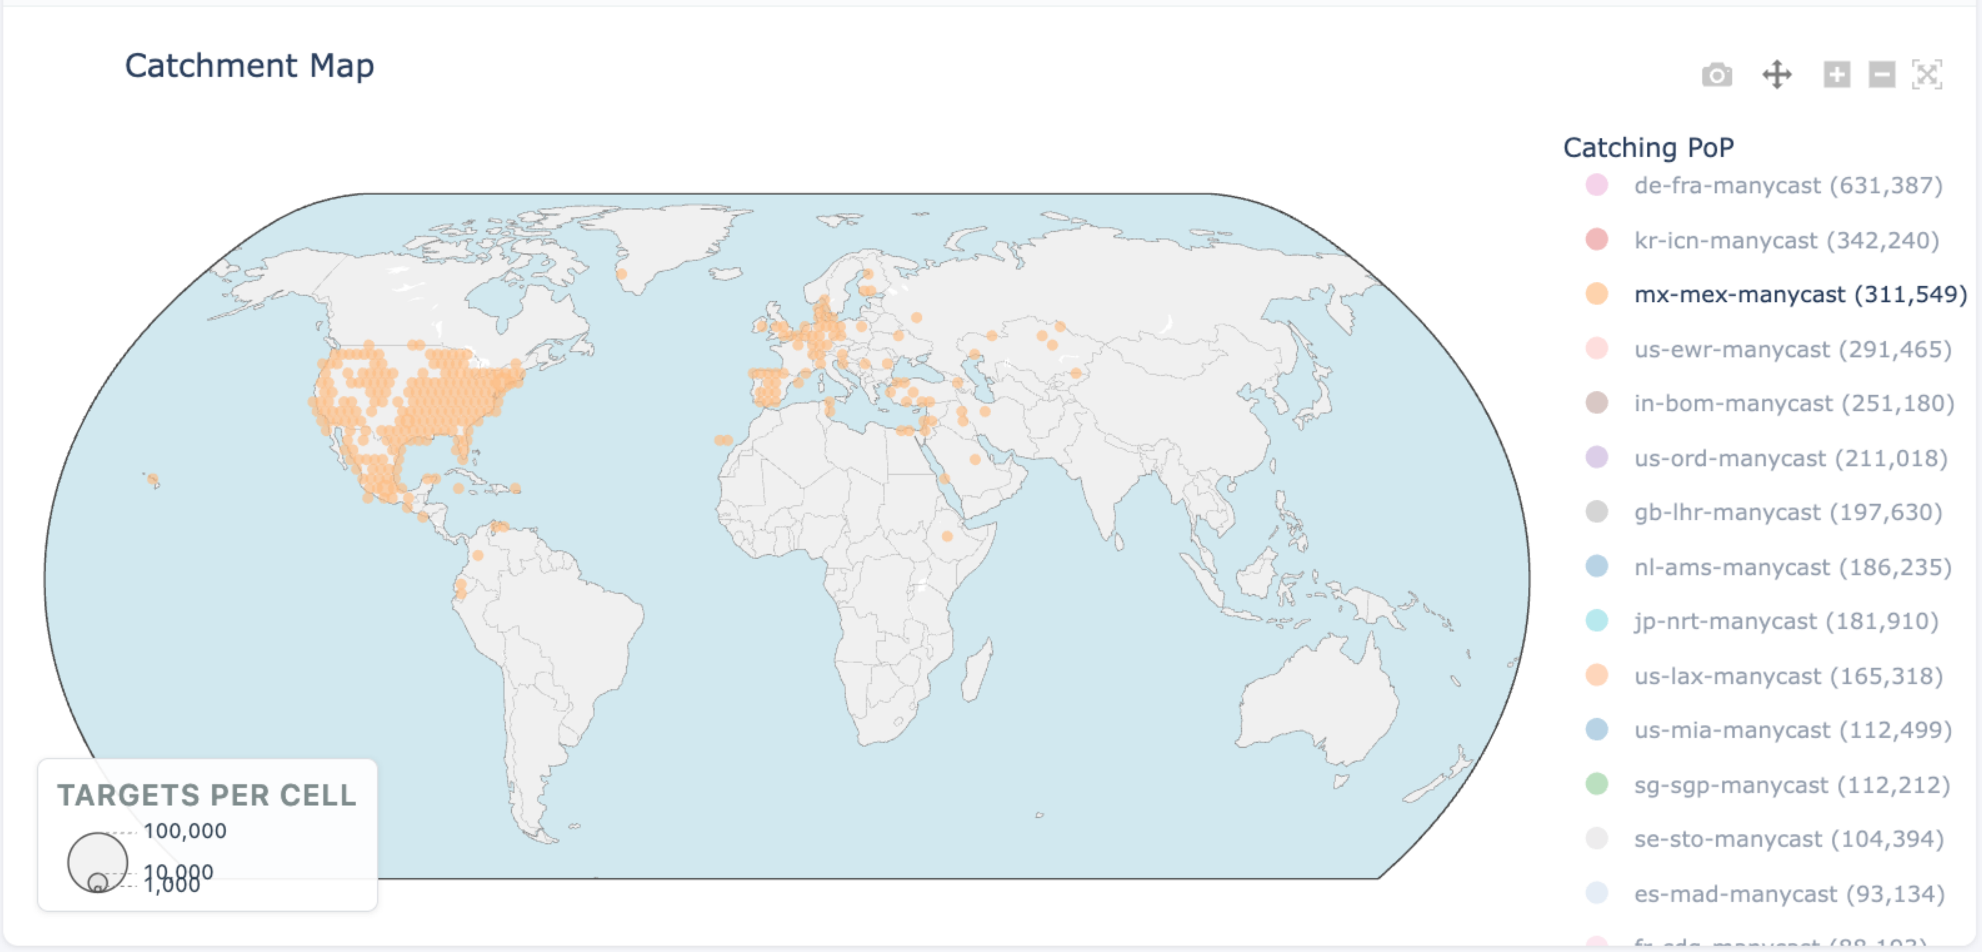
Catchment Map
Catching PoP
de-fra-manycast (631,387)
kr-icn-manycast (342,240)
mx-mex-manycast (311,549)
us-ewr-manycast (291,465)
in-bom-manycast (251,180)
us-ord-manycast (211,018)
gb-lhr-manycast (197,630)
nl-ams-manycast (186,235)
jp-nrt-manycast (181,910)
us-lax-manycast (165,318)
us-mia-manycast (112,499)
sg-sgp-manycast (112,212)
se-sto-manycast (104,394)
es-mad-manycast (93,134)
TARGETS PER CELL
100,000
10,000
1,000