

Measuring Anycast Performance

Catchment, RTT, and Optimal Site

maprg, IETF, MADRID, JUL25

Remi Hendriks

**UNIVERSITY
OF TWENTE.**

What is anycast?

- Anycast: replicating a service at multiple locations using a single shared IP address
 - Querying 1.1.1.1 from New Zealand -> reach server in New Zealand
 - Querying 1.1.1.1 here -> reach Cloudflare server in Madrid



What is anycast?

- Anycast: replicating a service at multiple locations using a single shared IP address
 - Querying 1.1.1.1 from New Zealand -> reach server in New Zealand
 - Querying 1.1.1.1 here -> reach Cloudflare server in Madrid
- Used for critical Internet infrastructure (e.g., DNS)
- Used by CDNs for a large variety of services
- Used to provide DDoS mitigation services



Cloudflare's anycast network

What is anycast?

- Anycast: replicating a service at multiple locations using a single shared IP address
 - Querying 1.1.1.1 from New Zealand -> reach server in New Zealand
 - Querying 1.1.1.1 here -> reach Cloudflare server in Madrid
- Used for critical Internet infrastructure (e.g., DNS)
- Used by CDNs for a large variety of services
- Used to provide DDoS mitigation services
- Why?
 - Proven technique
 - Reduces latency, load-balances traffic
 - Most importantly, improves resilience



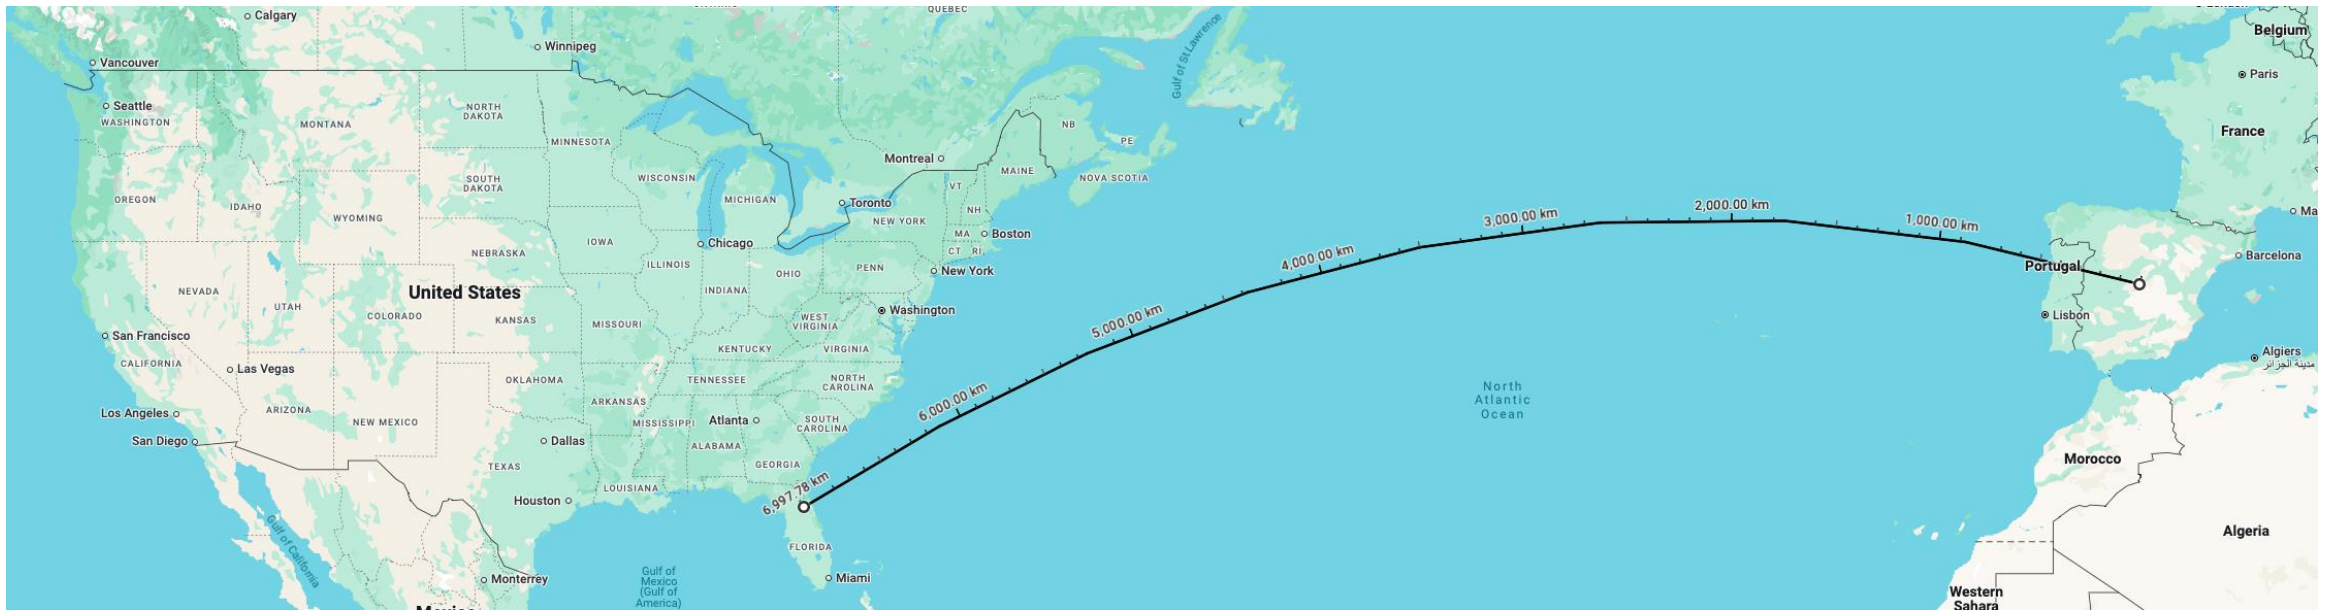
Cloudflare's anycast network

Anycast problems

- Anycast relies on BGP to route clients to the nearest site
 - BGP is not performance aware
- Sub-optimal anycast routing (client reaching a distant site when a nearby one is available)
 - Known issue for anycast operators
 - Causes high RTTs
 - May violate privacy regulations
- Clients may experience anycast site flipping
 - Results in failure of stateful services

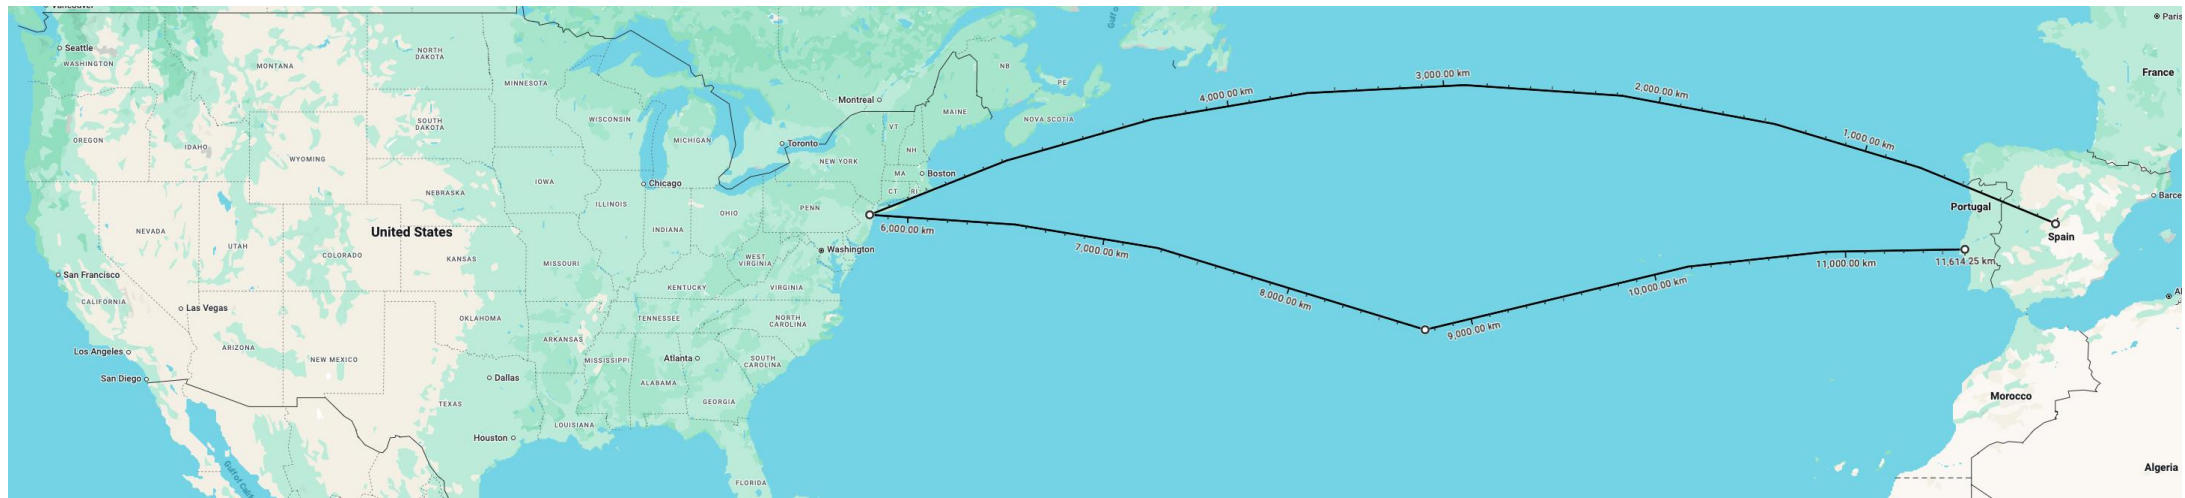
Anycast problems

- Example 1) BGP routes you to 1.1.1.1 in Miami instead of Madrid



Anycast problems

- Example 2) Long path to 1.1.1.1 through New York to reach Lisbon



Motivation

- Sub-optimal anycast routing hurts performance
 - Affected areas change over time due to dynamic nature of Internet
- Anycast requires active Traffic Engineering (TE)
- To make these TE decisions, performance metrics are needed

Motivation

- Sub-optimal anycast routing hurts performance
 - Affected areas change over time due to dynamic nature of Internet
- Anycast requires active Traffic Engineering (TE)
- To make these TE decisions, performance metrics are needed

- Option 1) Passive traffic analysis
 - Cannot prevent issues pro-actively
 - RTT unknown from e.g., passive DNS traffic

Motivation

- Sub-optimal anycast routing hurts performance
 - Affected areas change over time due to dynamic nature of Internet
- Anycast requires active Traffic Engineering (TE)
- To make these TE decisions, performance metrics are needed
- Option 2) Probing platform (e.g., RIPE Atlas)
 - Limited to coverage of probing platform
 - Rate-limiting and credit costs



RIPE Atlas probe locations

Motivation

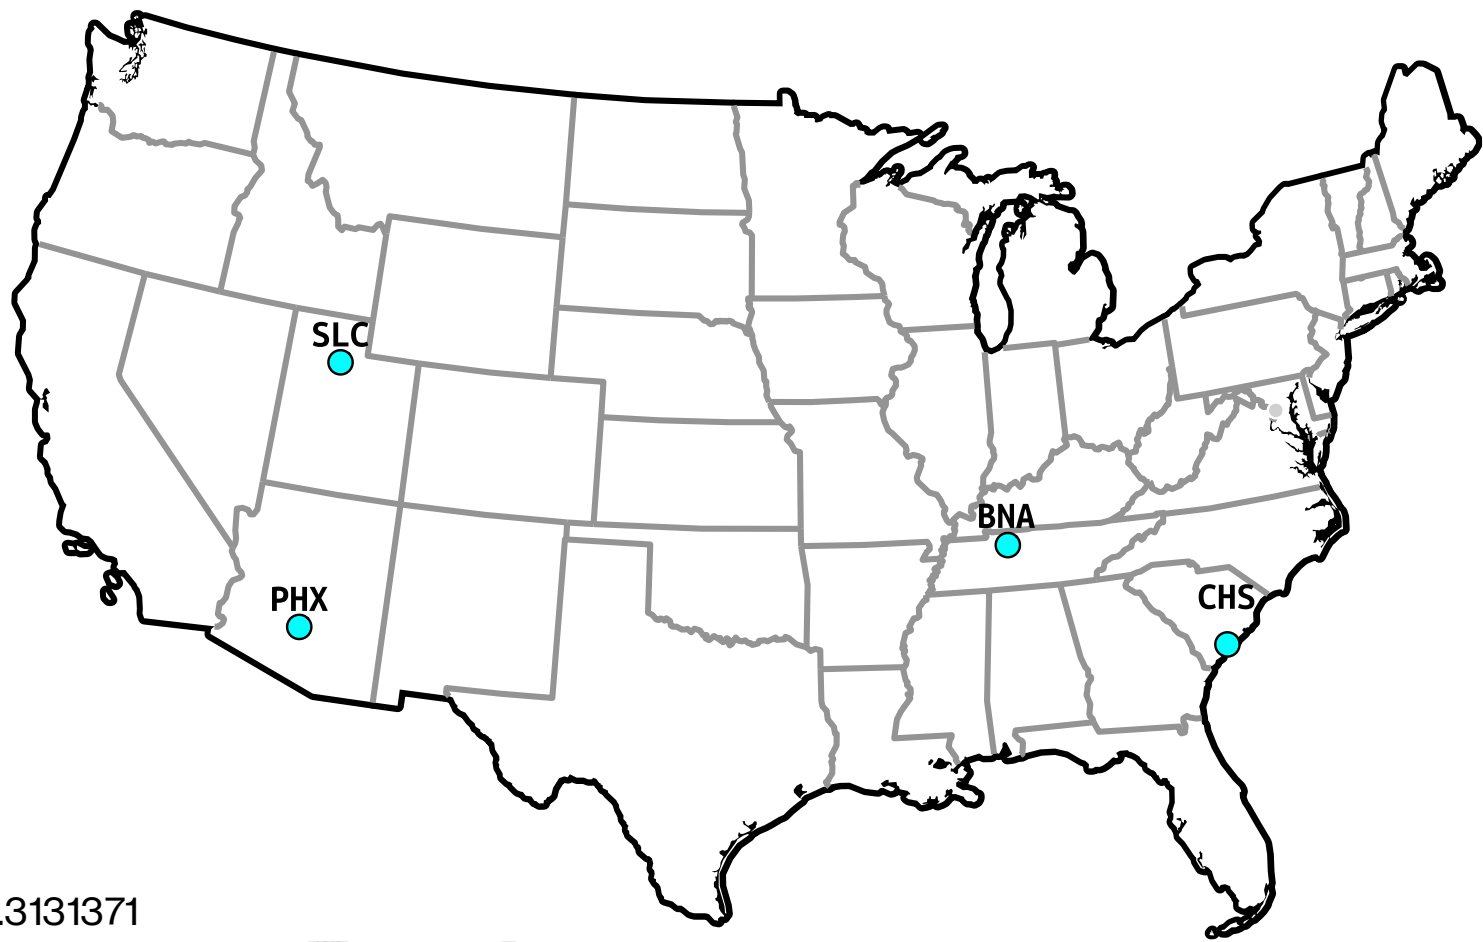
- Sub-optimal anycast routing hurts performance
 - Affected areas change over time due to dynamic nature of Internet
- Anycast requires active Traffic Engineering (TE)
- To make these TE decisions, performance metrics are needed

- Option 3) Probing responsive hosts with anycast source IP

Previous work

Verfploeter

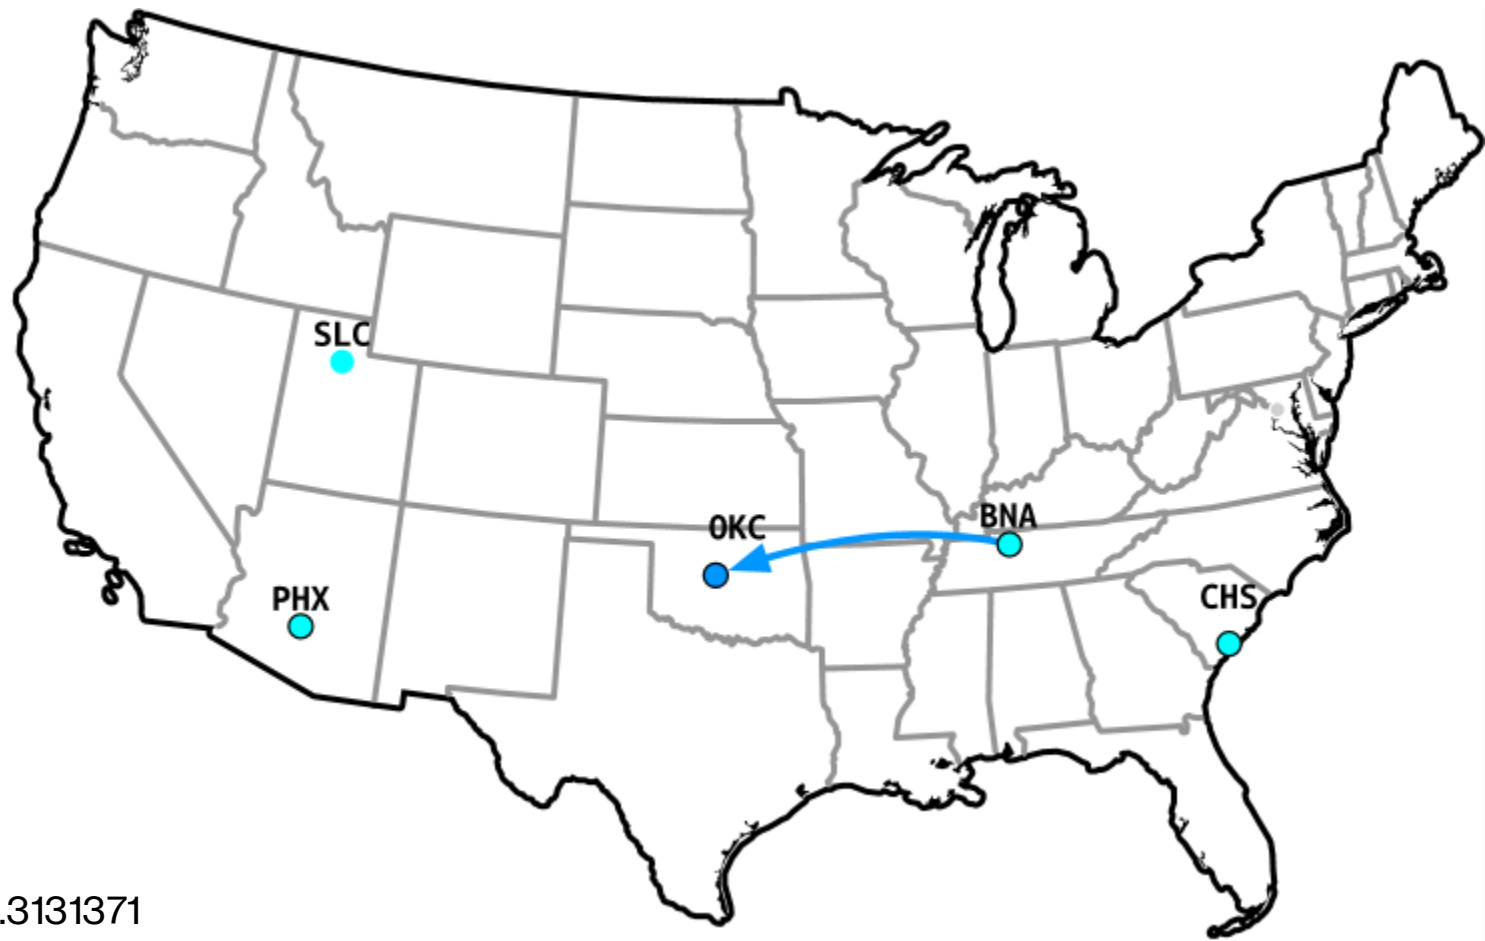
- Verfploeter maps anycast catchments [1]
 - I.e., which anycast site 'catches' which traffic
- How?
 - Ping host with anycast source IP
 - See where the ping reply comes back



Previous work

Verfploeter

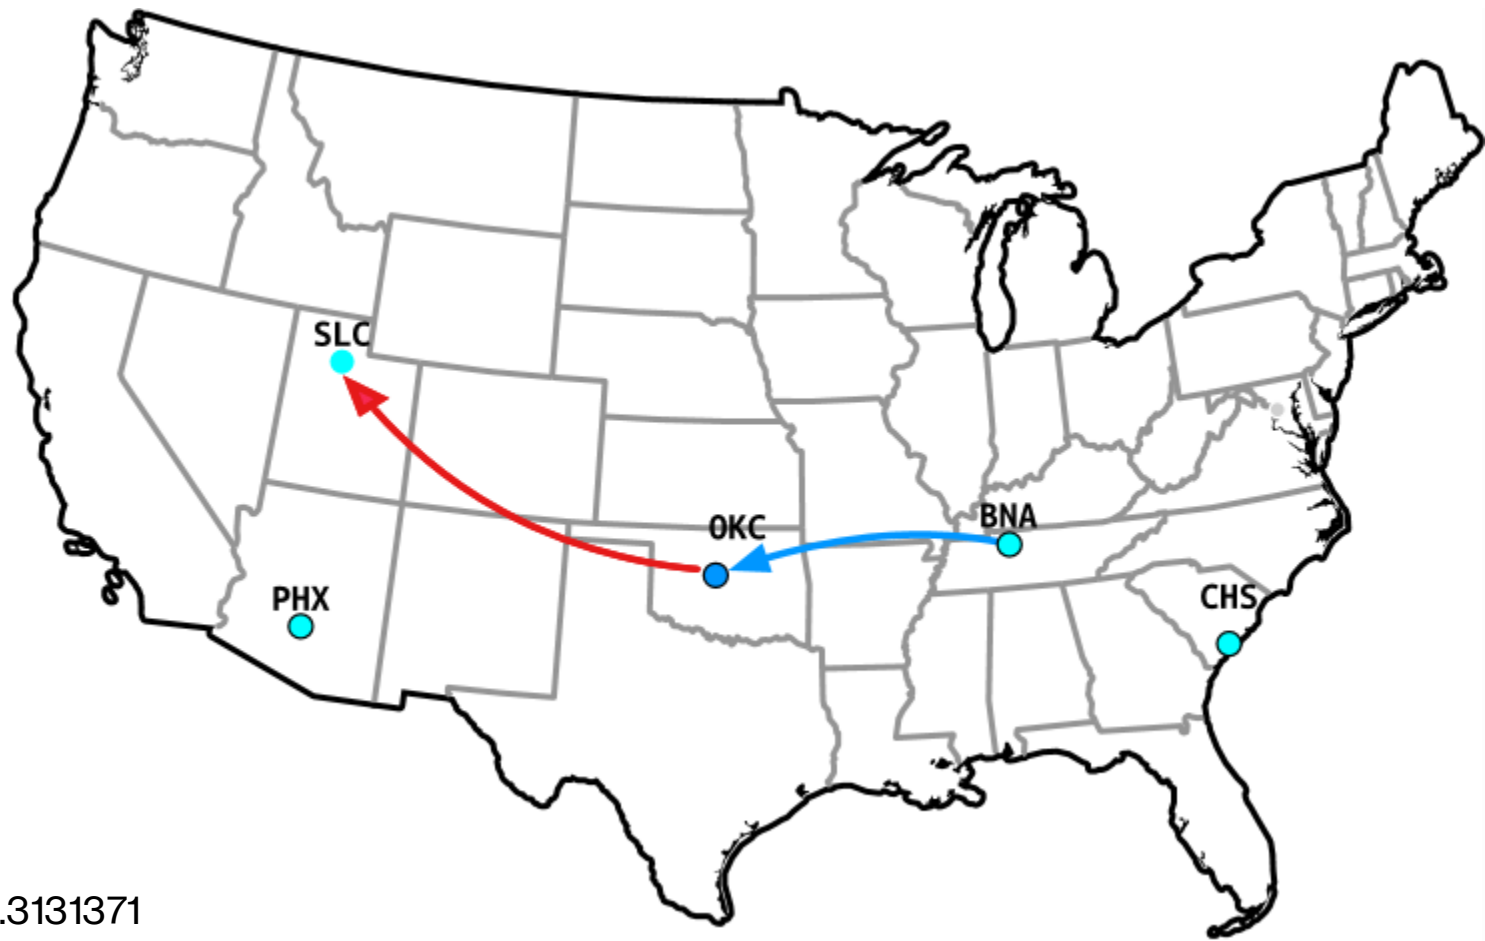
- Verfploeter maps anycast catchments [1]
 - I.e., which anycast site 'catches' which traffic
- How?
 - Ping host with anycast source IP
 - See where the ping reply comes back



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



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
 - Can measure pro-actively
- Disadvantages
 - Cannot measure RTT
 - Not always clear what is the ‘optimal’ site
 - Catchment can be deceiving (RTT better metric)



Our tooling

- Builds on Verfploeter
 - anycast RTT
 - unicast probing
- Designed as a 'Swiss knife'
 - Configuration files (for complex measurements)
 - Large variety of supported measurements
- Low resource usage
- 8 MB static binary or docker image



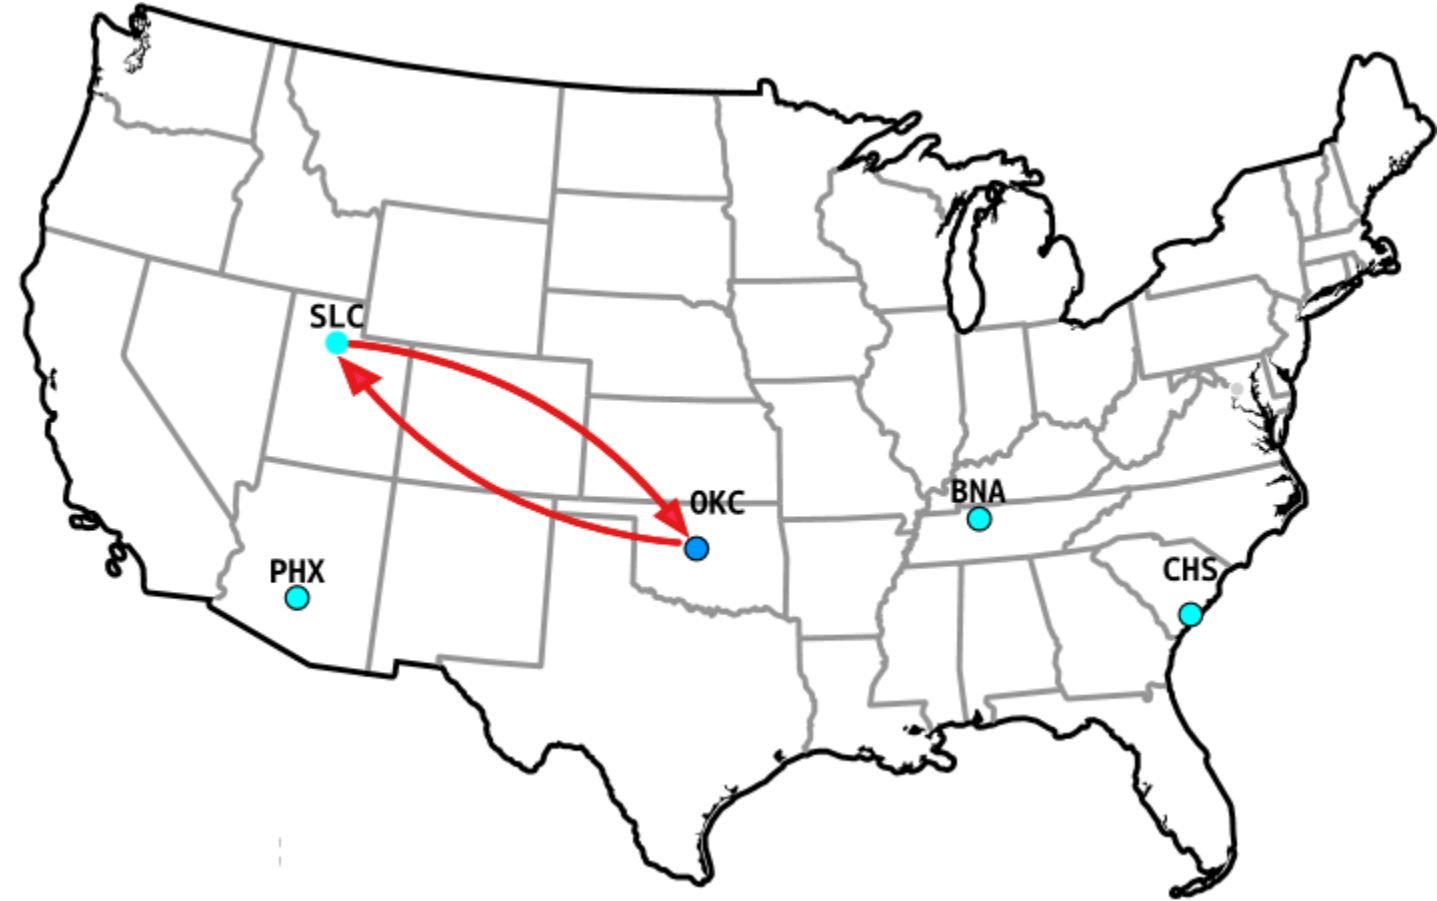
Improvements Anycast RTT

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



Improvements Anycast RTT

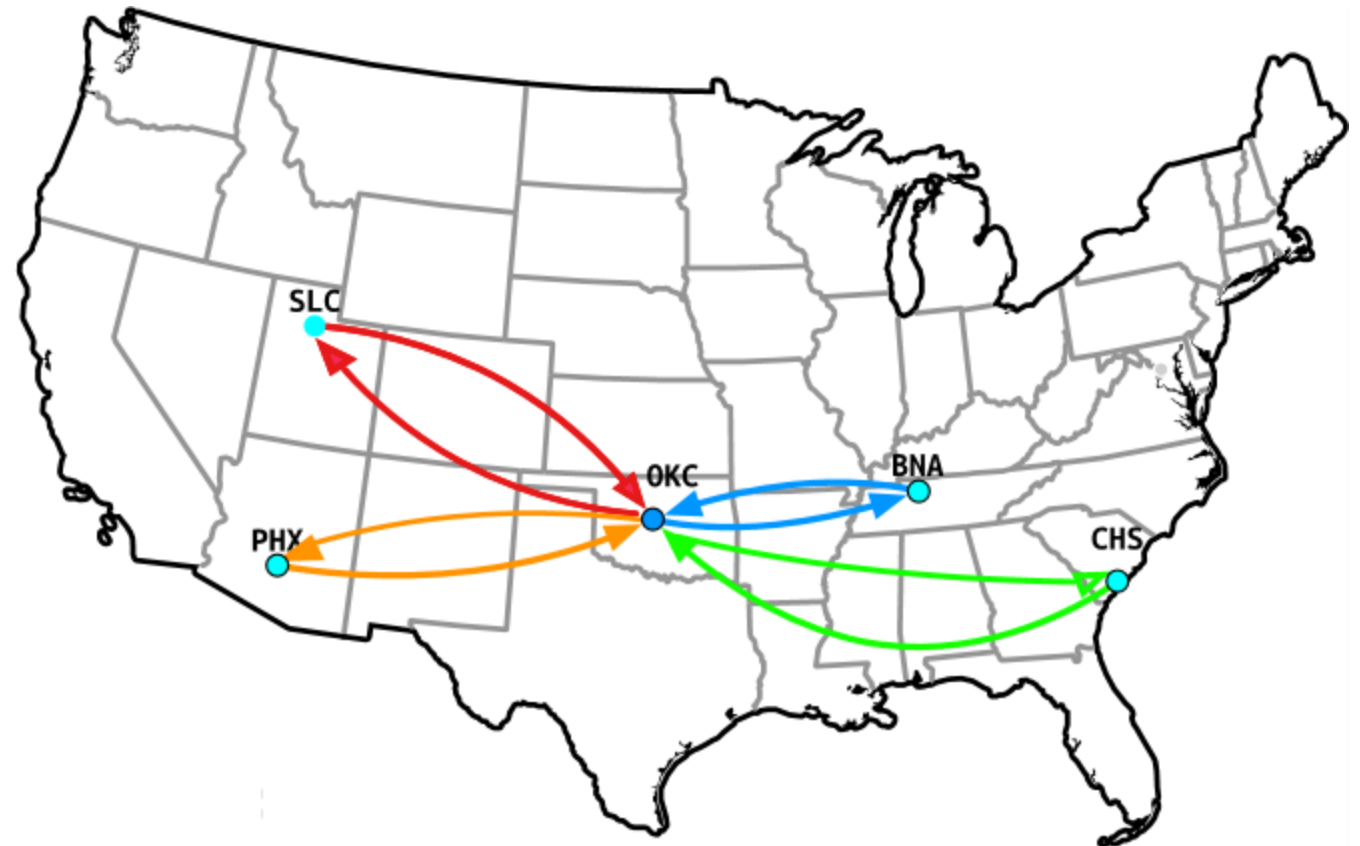
- Probe 1: assess catchment (discovery probe)
- **Probe 2: probe from catcher (measurement probe)**
 - $RTT = \text{receive time} - \text{transmit time}$



Improvements

Optimal routing

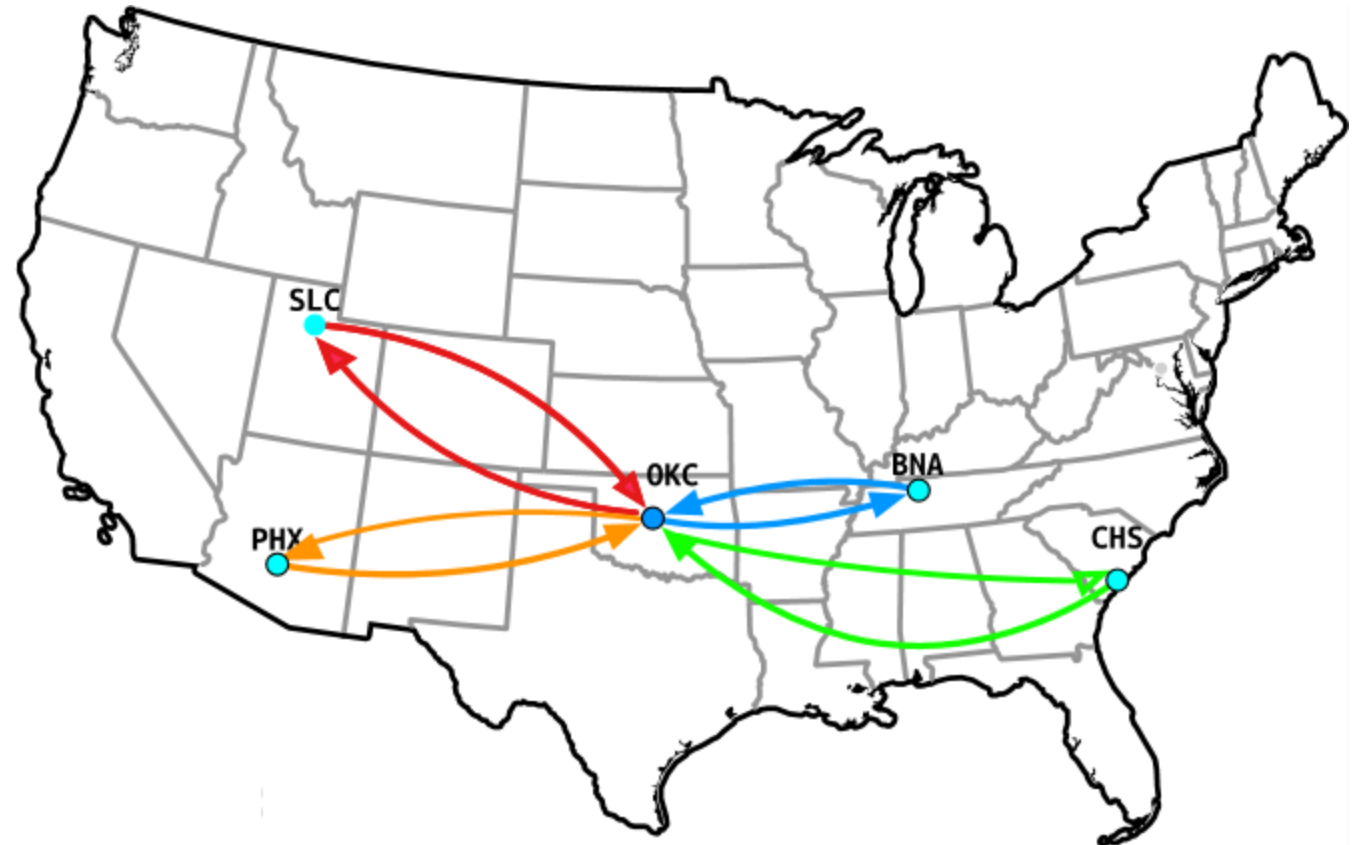
- Probe target from all VPs using unicast IPs
 - Measures RTT with all possible sites
 - Site with lowest RTT -> 'optimal' site



Improvements

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 improvement



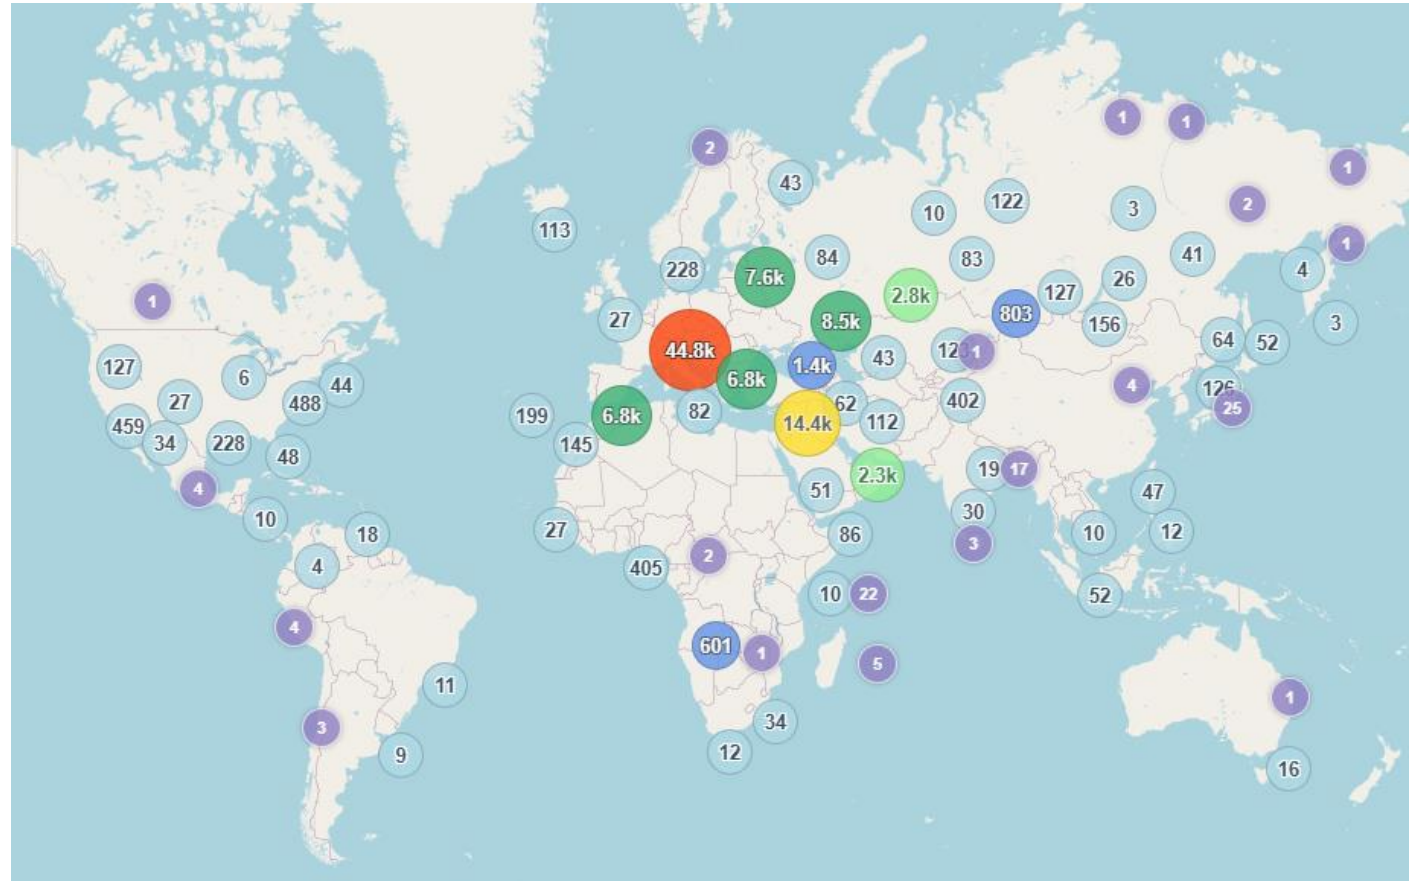
Other improvements

- Extended protocol support
 - ICMP, TCP, UDP
 - IPv4, IPv6
- Divide-and-conquer style anycast measurements
- Measure multiple anycast prefixes simultaneously
- Measure load-balancing induced routing instability
- And more ...

Dashboard

Catchment mapping (nl-ams)

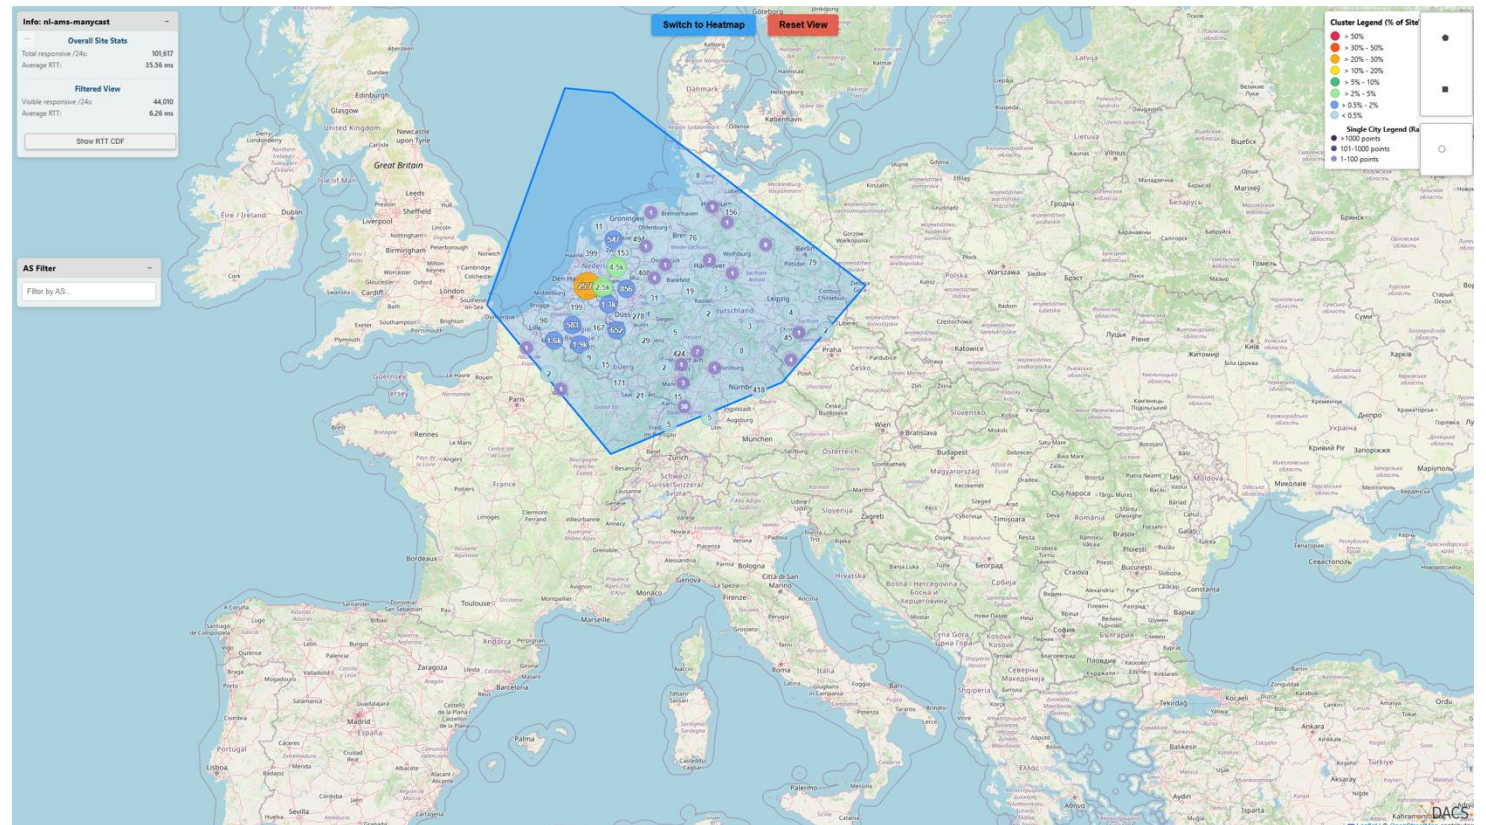
- Ongoing work by Nikos Fiotakis
 - Visualize catchments



Dashboard

Catchment mapping (nl-ams)

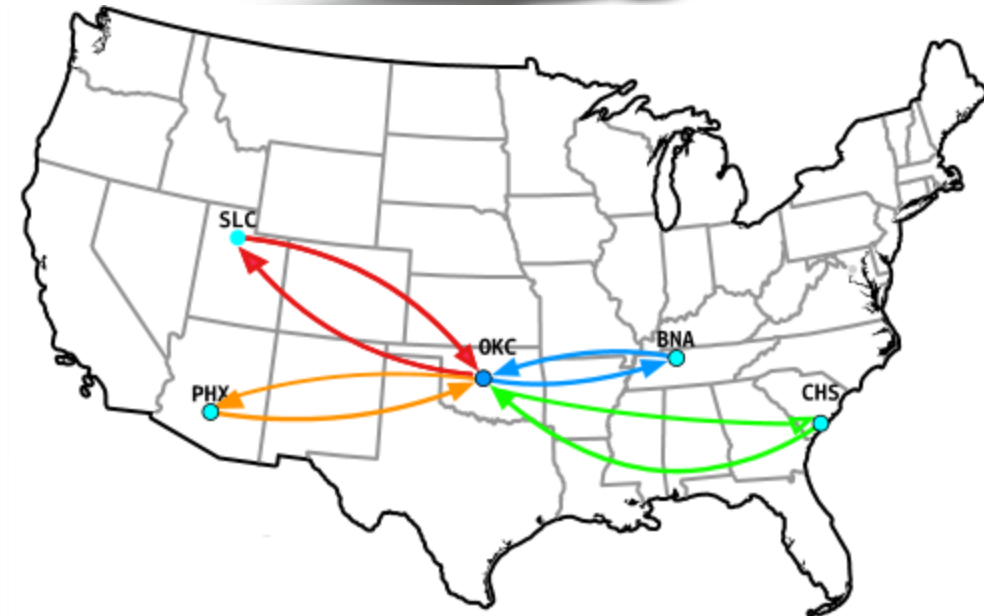
- Ongoing work by Nikos Fiotakis
- Visualize catchments
- Selection tool
- Visualize latencies
- Filter by AS
- ...



Summary

- Tooling allows for...
 - Measuring anycast RTTs
 - Unicast RTTs to infer optimal site
 - And much more...
- Currently used in production by a ccTLD
- Ongoing talks with root operators

- We are open to collaborate with anycast operators
 - Tooling available on request :-)
 - (we would like data in return)
 - Contact (remi.hendriks@utwente.nl)

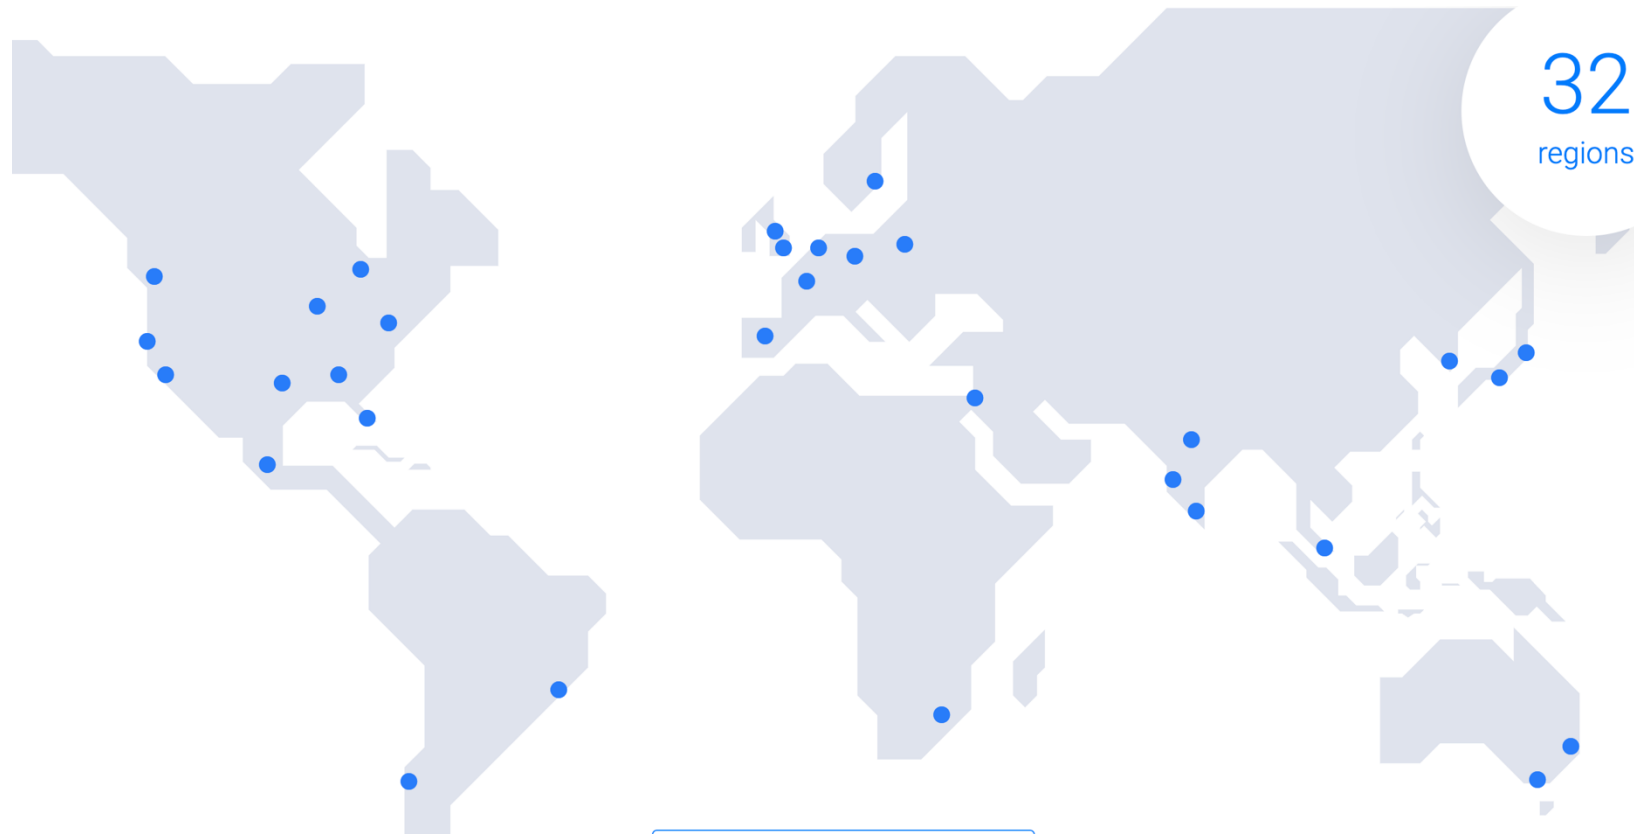


Appendix slides

Measurement setup

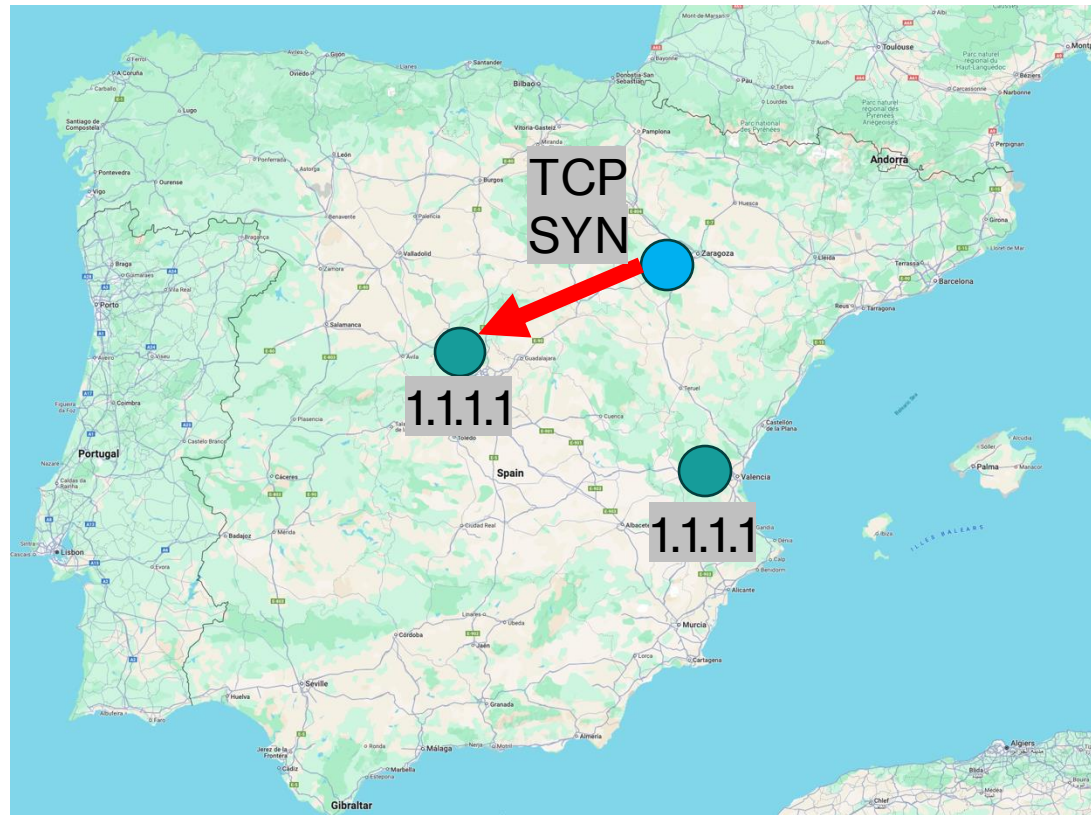
Deployed using Vultr (32 PoPs)

5.9 million /24-prefix targets (ISI hitlist)



Anycast problems

- Example 3) Load-balancers send subsequent packets to different 1.1.1.1 locations



Anycast problems

- Example 3) Load-balancers send subsequent packets to different 1.1.1.1 locations



Anycast problems

- Example 3) Load-balancers send subsequent packets to different 1.1.1.1 locations



Anycast problems

- Example 3) Load-balancers send subsequent packets to different 1.1.1.1 locations



Measuring site flipping

- Tooling can measure site flipping
- How?
 - Probe each target multiple times
 - Vary flow header fields in probes to the same target
 - Triggers load-balancers

Our tooling System design

8 MB image size
Low CPU/RAM usage

All computation offloaded
to orchestrator

