Anycast service discovery:

Deployment strategies and services deployed

AIMS-19, SAN DIEGO, FEB'26

Remi Hendriks & Gustavo Cesar Luvizotto (remote)

UNIVERSITY
OF TWENTE.

Introduction
What is anycast?

- 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, improve resilience, reduce latency

Introduction
What is anycast?

- Anyc… …f
  Pres…
  - *Ho…*
  - *W…*

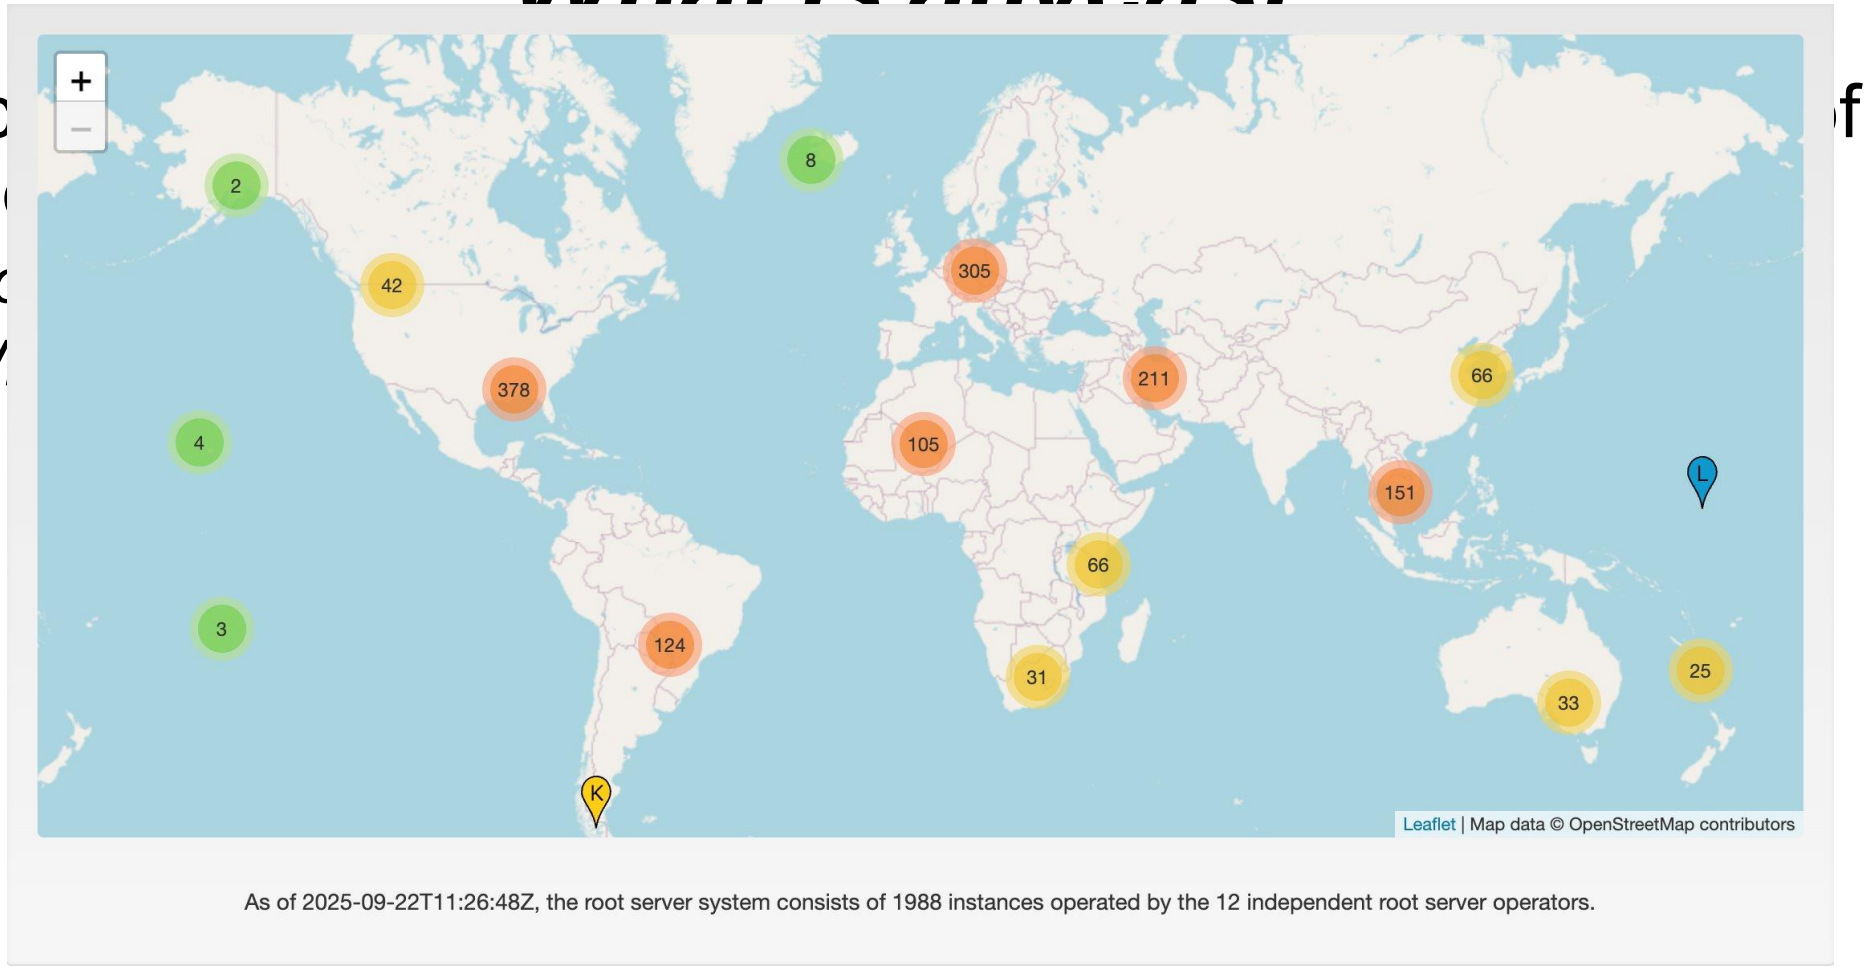

Source: root-servers.org
"All 13 DNS root-letters are deployed using anycast; combined totaling 1,988 sites"

LACeS

- We run a daily anycast census at University of Twente
  - Since Mar'24 (2 years of daily data)
  - Detects /24 and /48 anycast prefixes
  - Includes geolocation using Ark latency measurements
  - Publicly available github.com/ut-dacs/anycast-census

LACeS: An Open, Fast, Responsible, and Efficient Longitudinal Anycast Census System

Remi Hendriks
University of Twente
Enschede, The Netherlands
remi.hendriks@utwente.nl

Matthew Luckie
CAIDA
UC San Diego
La Jolla, CA, USA

Mattijs Jonker
University of Twente
Enschede, The Netherlands

Raffaele Sommese
University of Twente
Enschede, The Netherlands

Roland van Rijswijk-Deij
University of Twente
Enschede, The Netherlands

LACeS

- We run a daily anycast census at University of Twente
  - Since Mar'24 (2 years of daily data)
  - Detects /24 and /48 anycast prefixes
  - Includes geolocation using Ark latency measurements
  - Publicly available github.com/ut-dacs/anycast-census
- Coverage today
  - 14.3k IPv4 /24s (1,078 ASes)
  - 13.0k IPv6 /48s (576 ASes)

Research questions

RQ1 Which operators deploy anycast, and what are the different types of deployment strategies they utilize?

Research questions

RQ1 Which operators deploy anycast, and what are the different types of deployment strategies they utilize?

RQ2 What services do operators replicate using their anycast infrastructure?

Research questions

RQ1 Which operators deploy anycast, and what are the different types of deployment strategies they utilize?

RQ2 What services do operators replicate using their anycast infrastructure?

RQ3 What longitudinal dynamics are visible for anycast?

Background
Anycast

- Operators are known to deploy regional anycast
  - Confined within a continent, or sometimes even a country
  - Used by ccTLDs (like .nz in NZ, or .be in BE & NL)
  - Used by CDNs (EU users -> EU region, NA users -> NA region)

Background
Anycast

- Operators are known to deploy regional anycast
  - Confined within a continent, or sometimes even a country
  - Used by ccTLDs (like .nz in NZ, or .be in BE & NL)
  - Used by CDNs (EU users -> EU region, NA users -> NA region)
- Hypergiants (like AWS, Google) deploy partial anycast
  - Same routed prefix has both unicast and anycast
  - Example: NTT announces a /16 at all their PoPs
    - Single address used for their anycasted DNS resolver
    - All other addresses routed internally to single location

Background & related work
Scanning tools

- ZMap (USENIX Security ’13)
  - Well-known and widely used
  - Says TCP/443 is open on these hosts

Background & related work
Scanning tools

- ZMap (USENIX Security ’13)
  - Well-known and widely used
  - Says TCP/443 is open on these hosts
- LZR (USENIX Security’21)
  - State-of-the-art
  - Fingerprint with fewer packets (and computation)
  - Can fingerprint 36 protocols
  - Says TCP/443 is running HTTPS on these hosts

Background & related work
Scanning tools

- ZMap (USENIX Security ’13)
  - Well-known and widely used
  - Says TCP/443 is open on these hosts
- LZR (USENIX Security’21)
  - State-of-the-art
  - Fingerprint with fewer packets (and computation)
  - Can fingerprint 36 application-layer protocols
  - Says TCP/443 is running HTTPS on these hosts
- We use Highly Responsive Prefixes (HRP) (CoNext’23)
  - /24s with more than 231 addresses responding to TCP SYN
  - Widely observed at CDNs

Measurement methodology

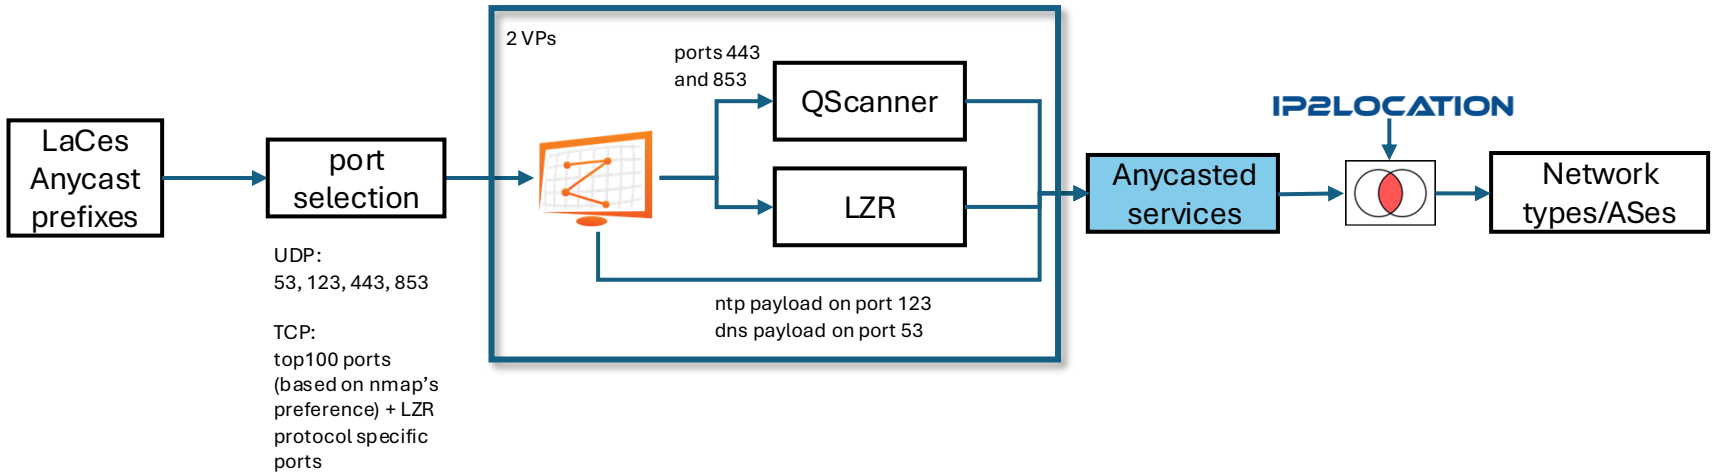

Scan	IANA-Assigned Ports		Ephemeral Ports	
Order	Protocol	Δ Coverage	Protocol	Δ Coverage
1	*wait*	51.3%	*wait*	66.3%
2	TLS	29.0%	HTTP	17.1%
3	HTTP	13.6%	TLS	15.9%
4	DNS	3.4%	Oracle DB	0.23%
5	PPTP	1.8%	PPTP	0.14%

Table 2: **Optimal Handshake Order**—For IANA-assigned ports, waiting and then sending a TLS Client Hello discovers 80.3% of unexpected services. Five handshakes can identify over 99% of identifiable unexpected services.

we are also performing a 7-day daily measurement

Results

RQ1 *Which operators deploy anycast, and what deployment strategies do they use?*

Anycast operators
Hypergiants

- Hypergiants dominate anycast
  - More than 2/3rds of /24-prefixes

Org (ASN)	/24s	Active IPs	Active IP ratio	HRP ratio
Cloudflare (13335)	2,866	657,587	0.896	0.989
Google (396982)	4,319	262,845	0.238	0.974
Fastly (54113)	830	210,443	0.990	0.990
AWS (16509)	1,589	130,649	0.321	0.245
Cloudflare (209242)	268	59,822	0.872	0.978
Fly.io (40509)	223	21,711	0.380	1.000
GoDaddy (398787)	64	16,383	1.000	1.000
AceVille (139341)	115	12,943	0.440	0.835
Nazwa.pl (15967)	195	11,741	0.235	0.959
Afilias (207266)	82	11,142	0.531	0.061

TABLE I

LARGEST ASES BY NUMBER OF ACTIVE IPS, ALONGSIDE THE NUMBER OF /24S AND ACTIVE IP & HRP RATIOS.

Anycast operators Hypergiants

- Hypergiants dominate anycast
  - More than 2/3rds of /24-prefixes
  - 86% of active IPs

Org (ASN)	/24s	Active IPs	Active IP ratio	HRP ratio
Cloudflare (13335)	2,866	657,587	0.896	0.989
Google (396982)	4,319	262,845	0.238	0.974
Fastly (54113)	830	210,443	0.990	0.990
AWS (16509)	1,589	130,649	0.321	0.245
Cloudflare (209242)	268	59,822	0.872	0.978
Fly.io (40509)	223	21,711	0.380	1.000
GoDaddy (398787)	64	16,383	1.000	1.000
AceVille (139341)	115	12,943	0.440	0.835
Nazwa.pl (15967)	195	11,741	0.235	0.959
Afilias (207266)	82	11,142	0.531	0.061

TABLE I

LARGEST ASES BY NUMBER OF ACTIVE IPS, ALONGSIDE THE NUMBER OF /24S AND ACTIVE IP & HRP RATIOS.

Anycast operators
Hypergiants

- Hypergiants dominate anycast
  - More than 2/3rds of /24-prefixes
  - 86% of active IPs
  - High active IP ratios
  - High HRP ratios

Org (ASN)	/24s	Active IPs	Active IP ratio	HRP ratio
Cloudflare (13335)	2,866	657,587	0.896	0.989
Google (396982)	4,319	262,845	0.238	0.974
Fastly (54113)	830	210,443	0.990	0.990
AWS (16509)	1,589	130,649	0.321	0.245
Cloudflare (209242)	268	59,822	0.872	0.978
Fly.io (40509)	223	21,711	0.380	1.000
GoDaddy (398787)	64	16,383	1.000	1.000
AceVille (139341)	115	12,943	0.440	0.835
Nazwa.pl (15967)	195	11,741	0.235	0.959
Afilias (207266)	82	11,142	0.531	0.061

TABLE I
LARGEST ASES BY NUMBER OF ACTIVE IPS, ALONGSIDE THE NUMBER OF /24S AND ACTIVE IP & HRP RATIOS.

Anycast operators Hypergiants

- Hypergiants dominate anycast
  - More than 2/3rds of /24-prefixes
  - 86% of active IPs
  - High active IP ratios
  - High HRP ratios

Org (ASN)	/24s	Active IPs	Active IP ratio	HRP ratio
Cloudflare (13335)	2,866	657,587	0.896	0.989
Google (396982)	4,319	262,845	0.238	0.974
Fastly (54113)	830	210,443	0.990	0.990
AWS (16509)	1,589	130,649	0.321	0.245
Cloudflare (209242)	268	59,822	0.872	0.978
Fly.io (40509)	223	21,711	0.380	1.000
GoDaddy (398787)	64	16,383	1.000	1.000
AceVille (139341)	115	12,943	0.440	0.835
Nazwa.pl (15967)	195	11,741	0.235	0.959
Afilias (207266)	82	11,142	0.531	0.061

TABLE I

LARGEST ASES BY NUMBER OF ACTIVE IPS, ALONGSIDE THE NUMBER OF /24S AND ACTIVE IP & HRP RATIOS.

Anycast operators
The long tail

- Most ASes announce few anycast /24-prefixes

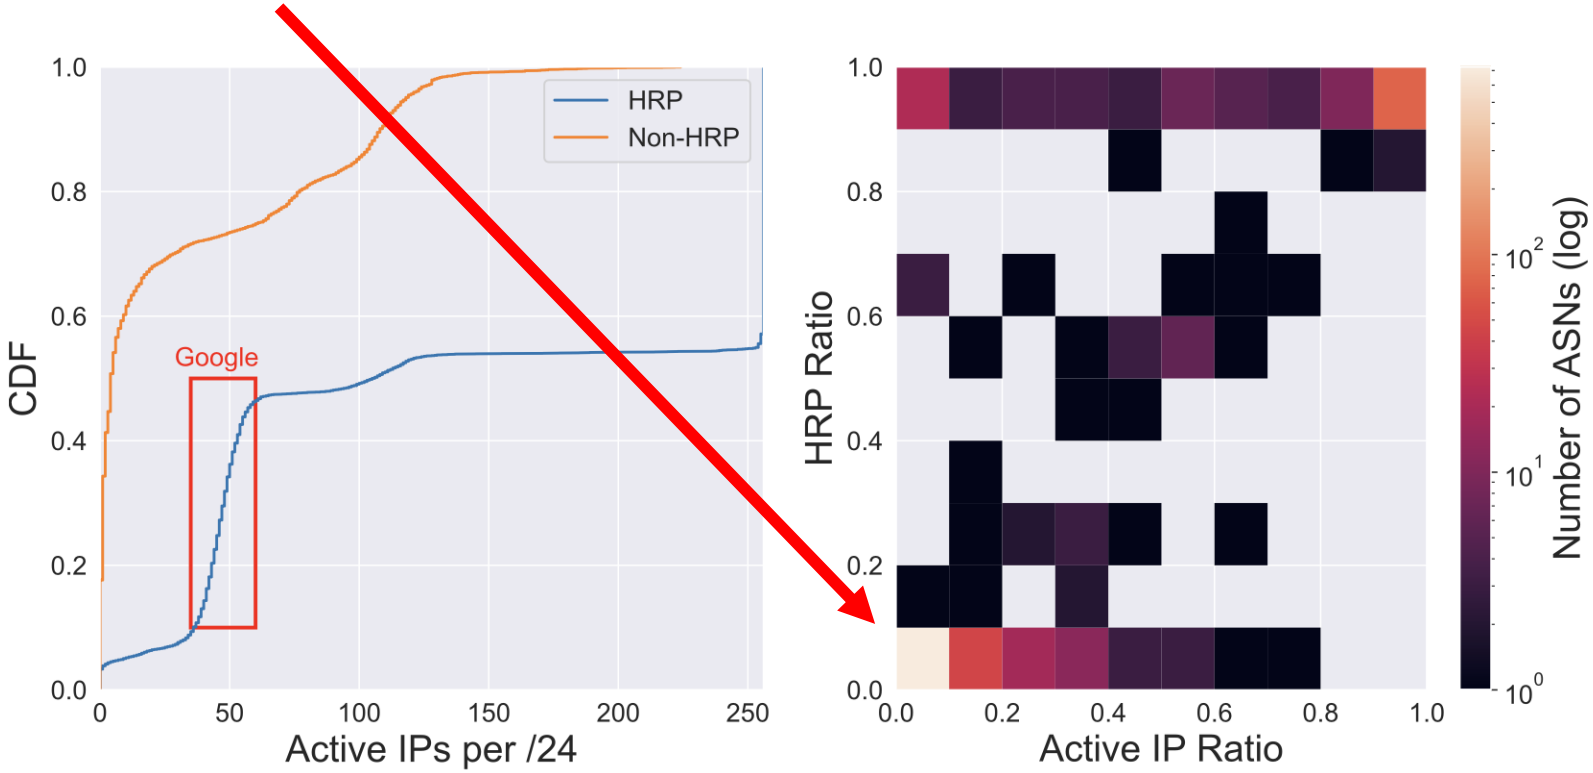

Fig. 2. Distribution of number of active IPs per /24 (CDF, left) and HRP ratio against Active IP ratio for ASes (heatmap, right).

Anycast operators
The long tail

- Most ASes announce few anycast /24-prefixes
  - Often a few active IPs
  - No HRP prefixes

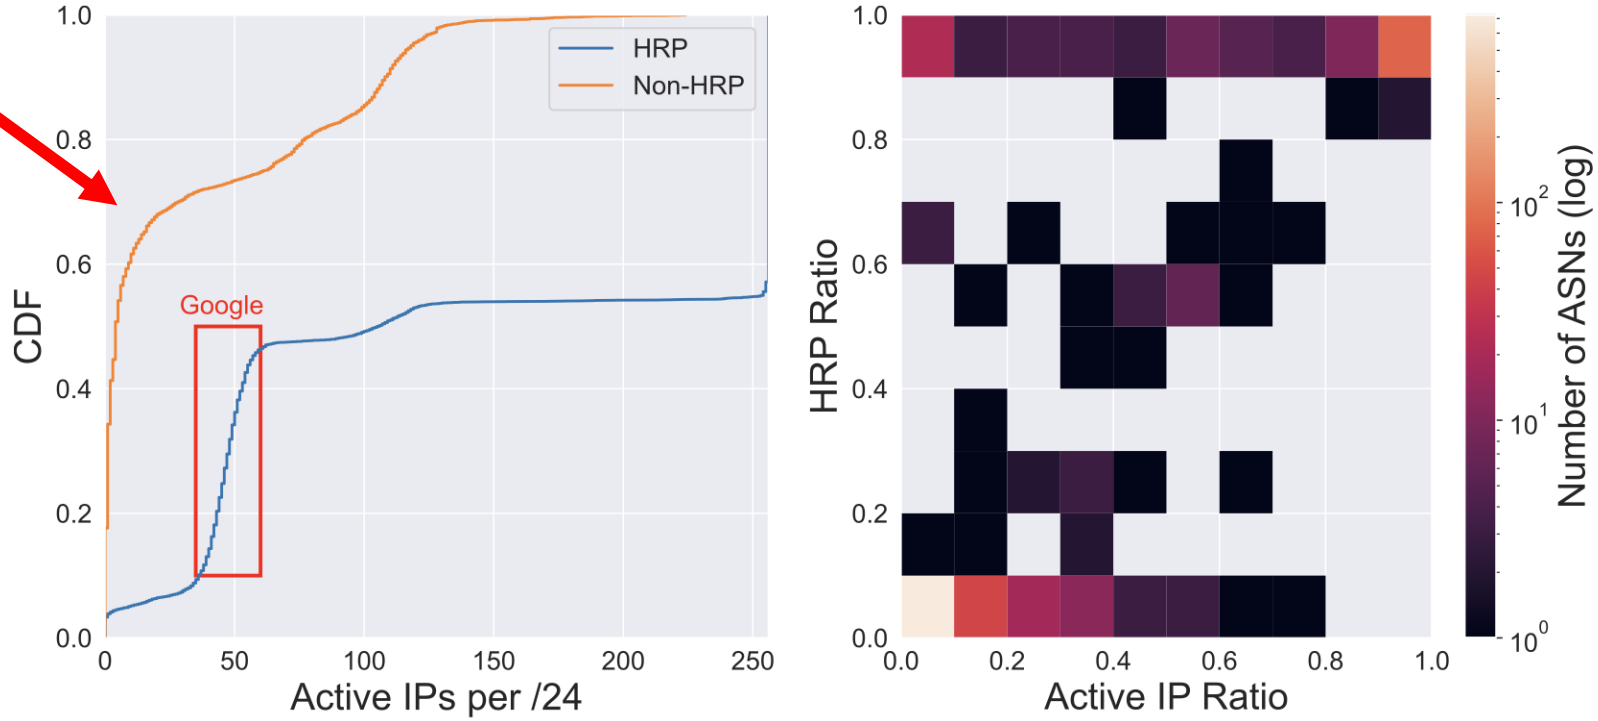

Fig. 2. Distribution of number of active IPs per /24 (CDF, left) and HRP ratio against Active IP ratio for ASes (heatmap, right).

Anycast operators
The long tail

- Most ASes announce few anycast /24-prefixes
  - Often a few active IPs
  - No HRP prefixes

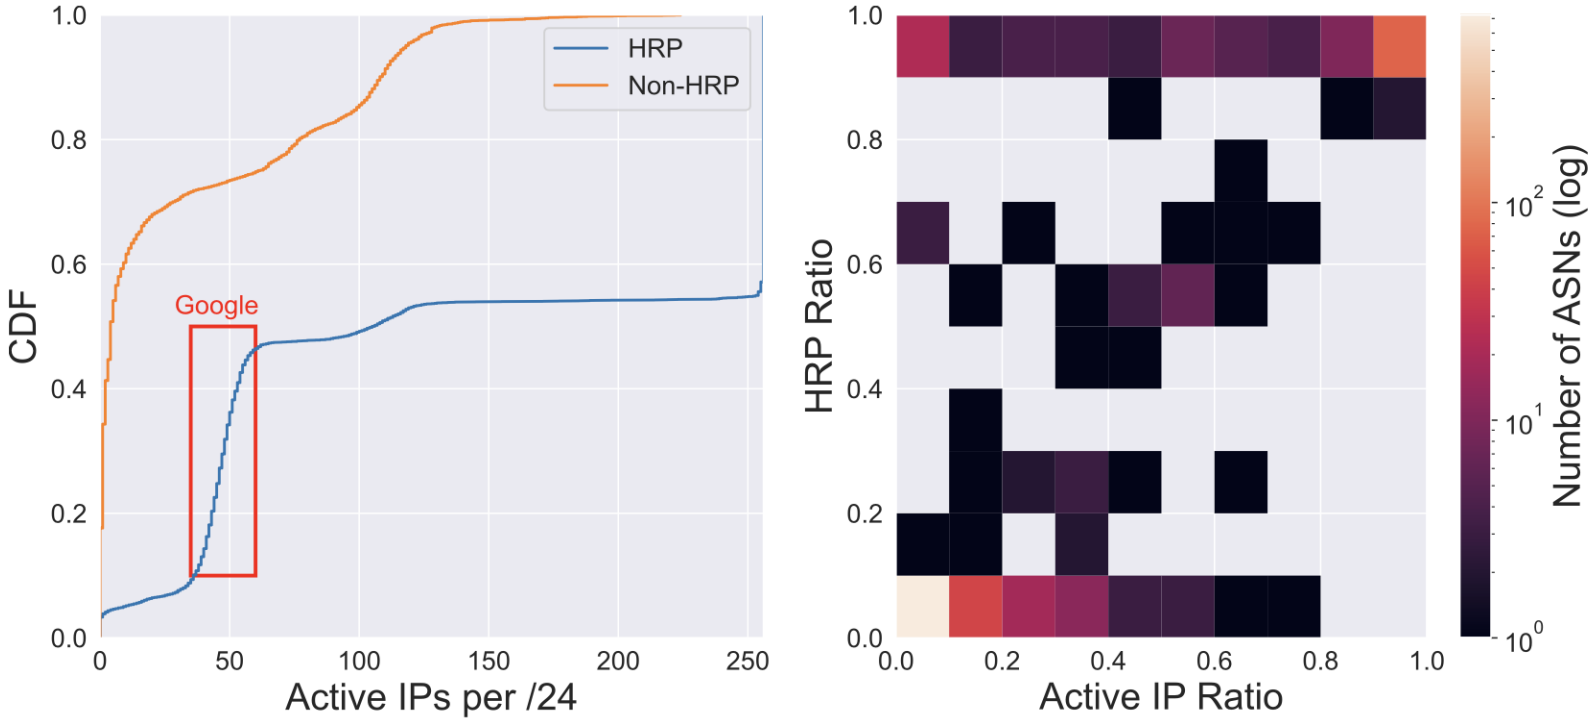

Fig. 2. Distribution of number of active IPs per /24 (CDF, left) and HRP ratio against Active IP ratio for ASes (heatmap, right).

Anycast upstreams

- Most anycast prefixes have a single upstream

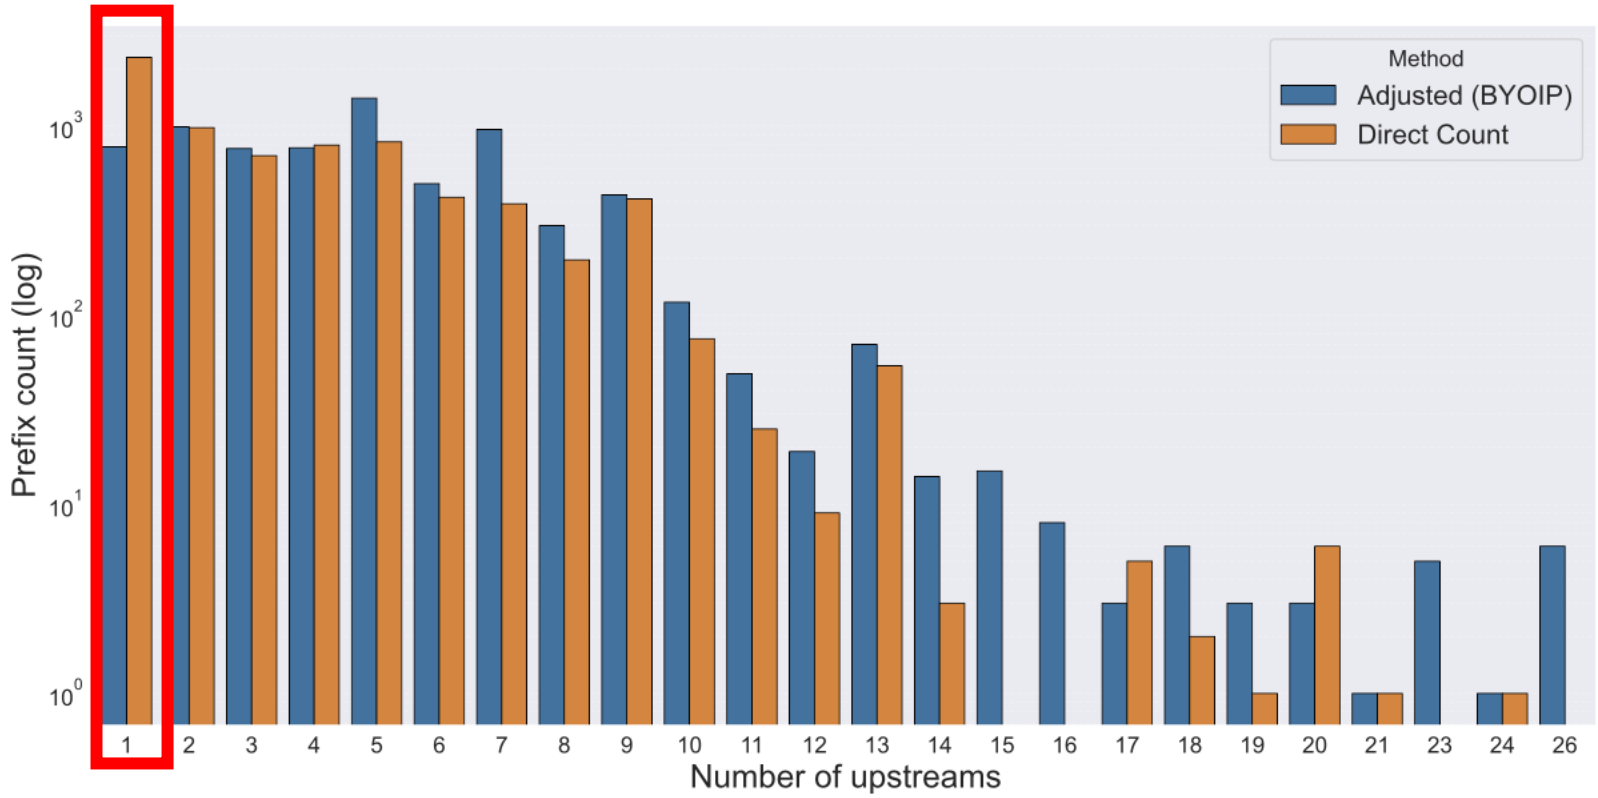

Anycast upstreams

- Most anycast prefixes have a single upstream
  - Often BYOIP provider relationships
  - E.g., Vultr, Cloudflare

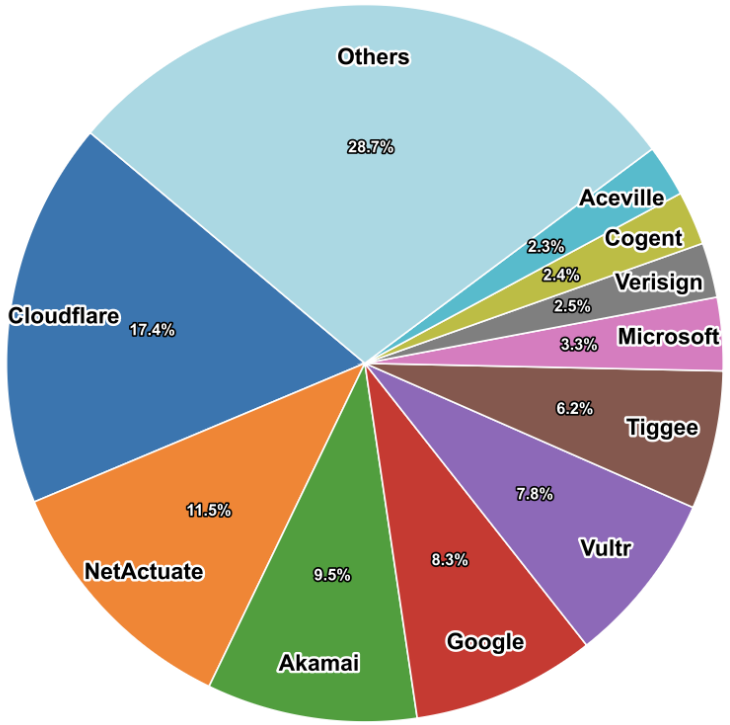

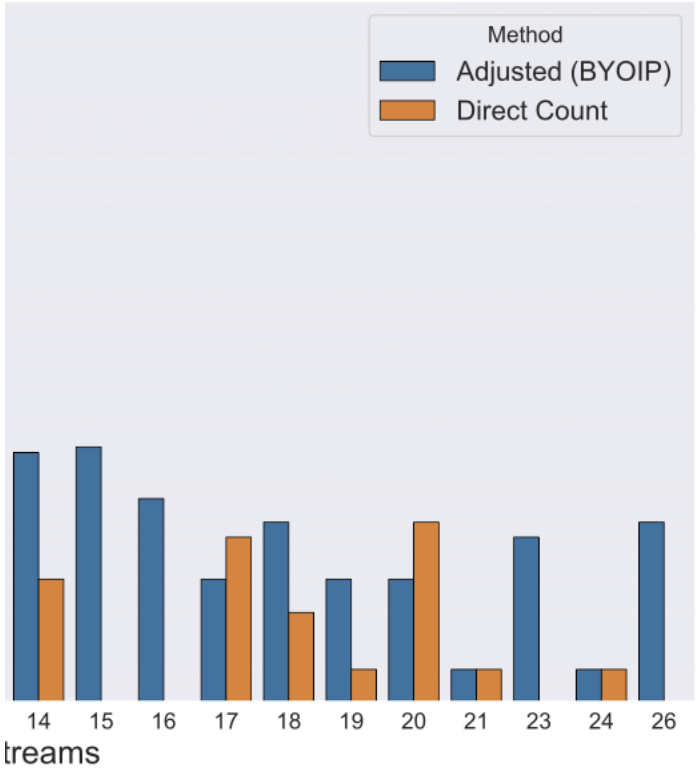

Fig. 3. Popular single-upstream ASes by announced anycast prefix share.

Anycast upstreams

- Most anycast prefixes have a single upstream
  - Often BYOIP provider relationships
  - E.g., Vultr, Cloudflare
- Adjusting for BYOIP
  - Often 1 to 7 upstream

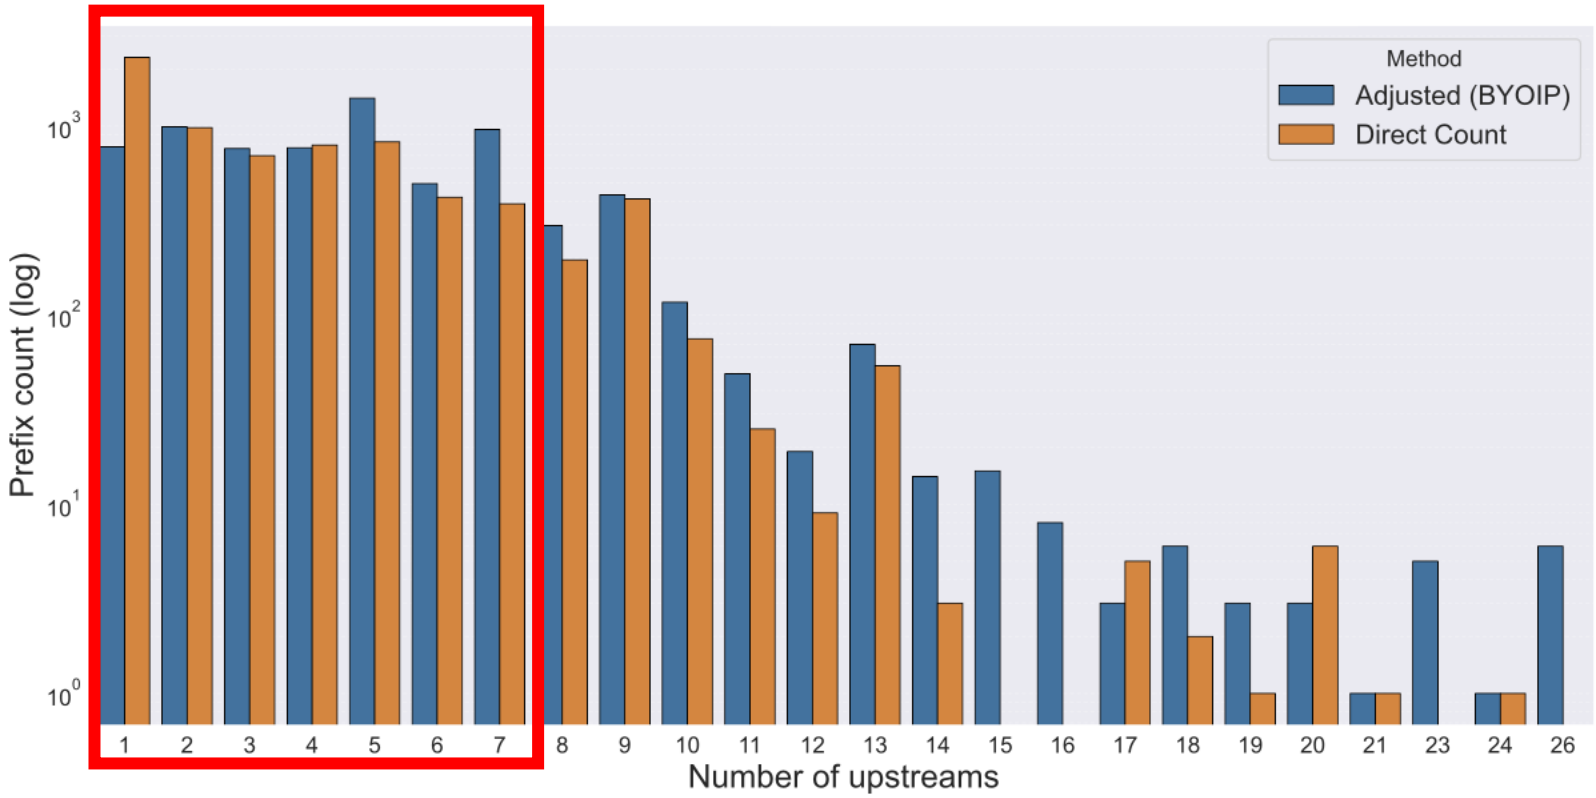

PoP geographic placement

- Operators mostly deploy in NA, EU
- Rarely in Africa, Central America

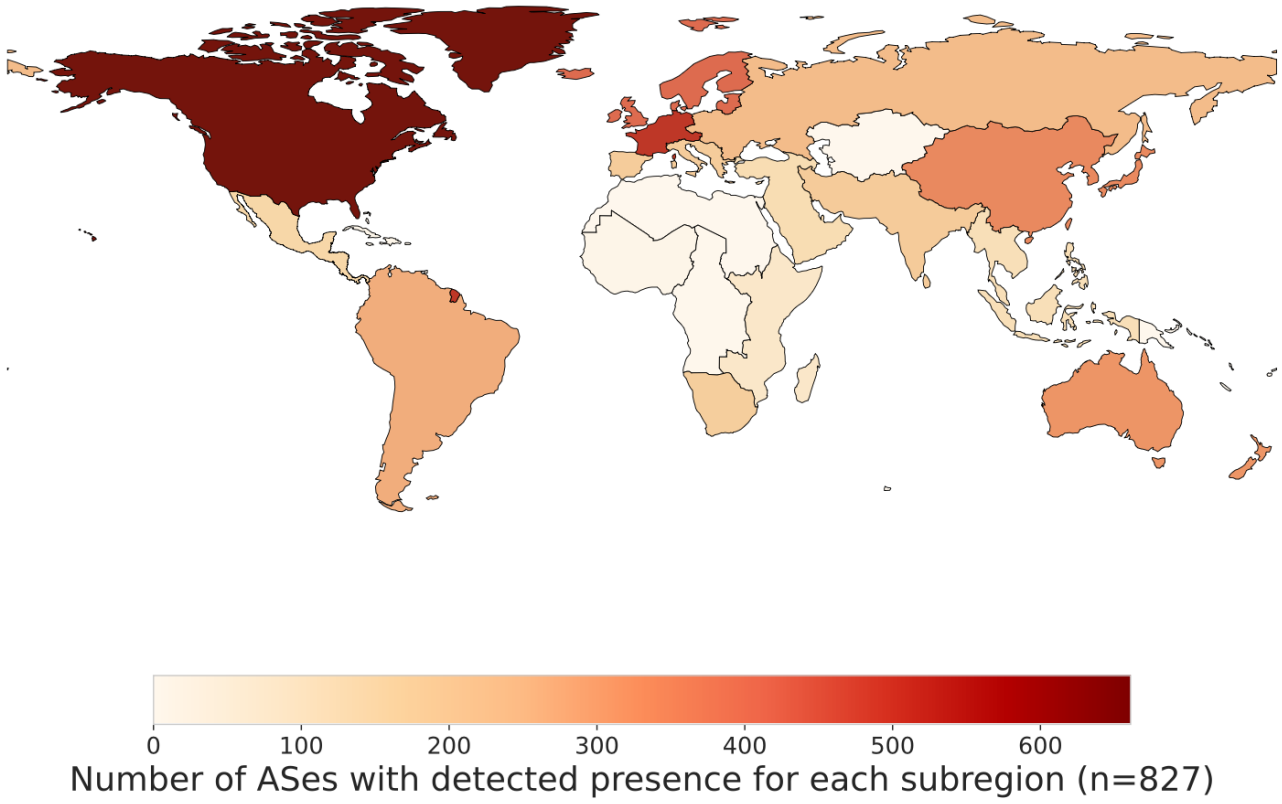

PoP geographic placement

- Operators mostly deploy in NA, EU
- Rarely in Africa, Central America
- 319 ASes deploy within a continent (regional)
  - NA, EU, Oceania, Asia most frequent

PoP geographic placement

- Operators mostly deploy in NA, EU
- Rarely in Africa, Central America
- 319 ASes deploy within a continent (regional)
  - NA, EU, Oceania, Asia most frequent
- 192 ASes deploy within a country (regional)
  - Mostly USA
  - Australia also seen often

Results

RQ2 *What services do operators replicate using anycast?*

Measurement overview

- HRP are more predominant on web, email, but also seen on databases (5432, 6379) and kubernetes (6443)
- non-HRP are more predominant on DNS and NTP, and a few other appearances
- But which services do we actually see?

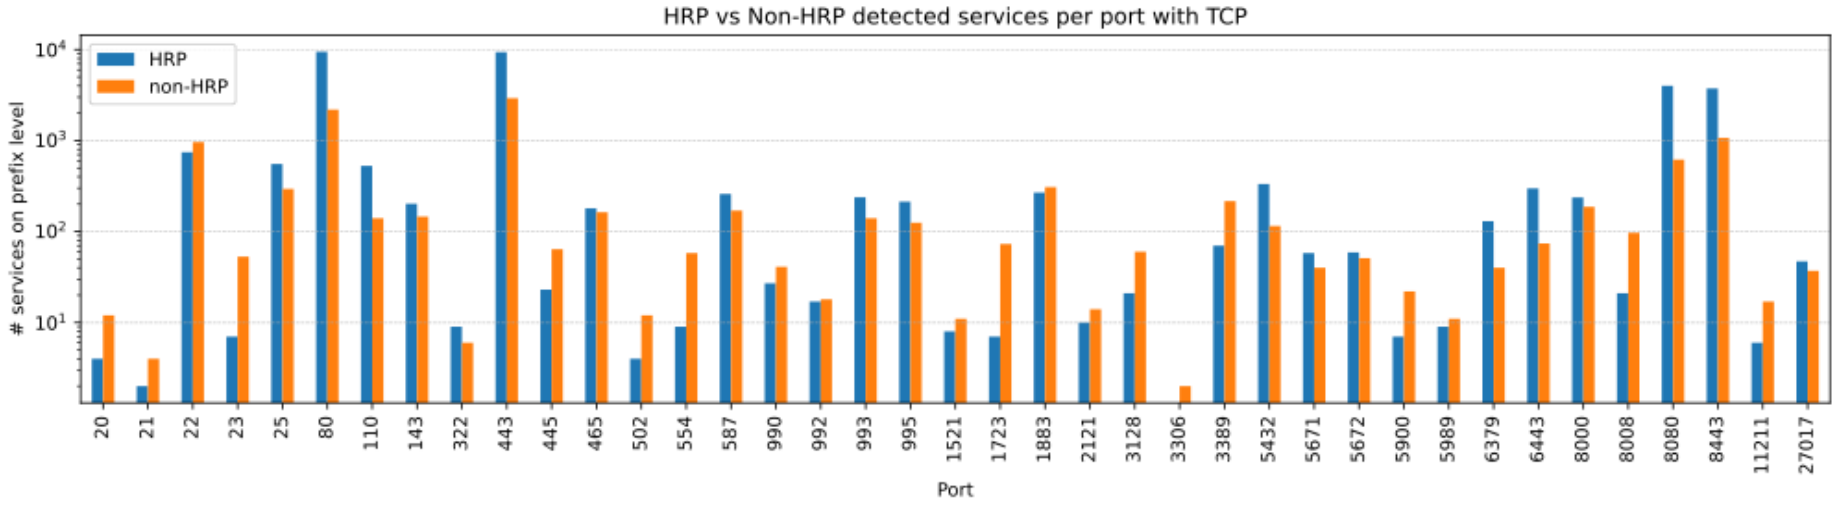

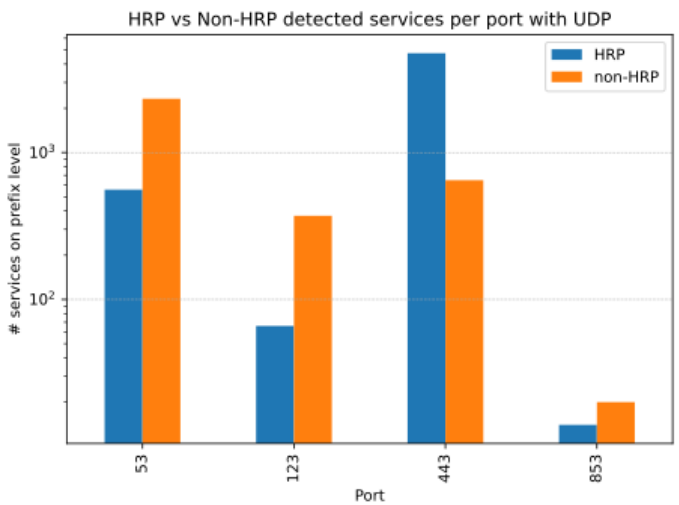

What do hypergiants replicate

- HTTP and QUIC (HTTP/3) are mostly on hypergiants
- Amazon's prefixes are unneven: some are higly responsive
- DNS is proeminent on non-HRP

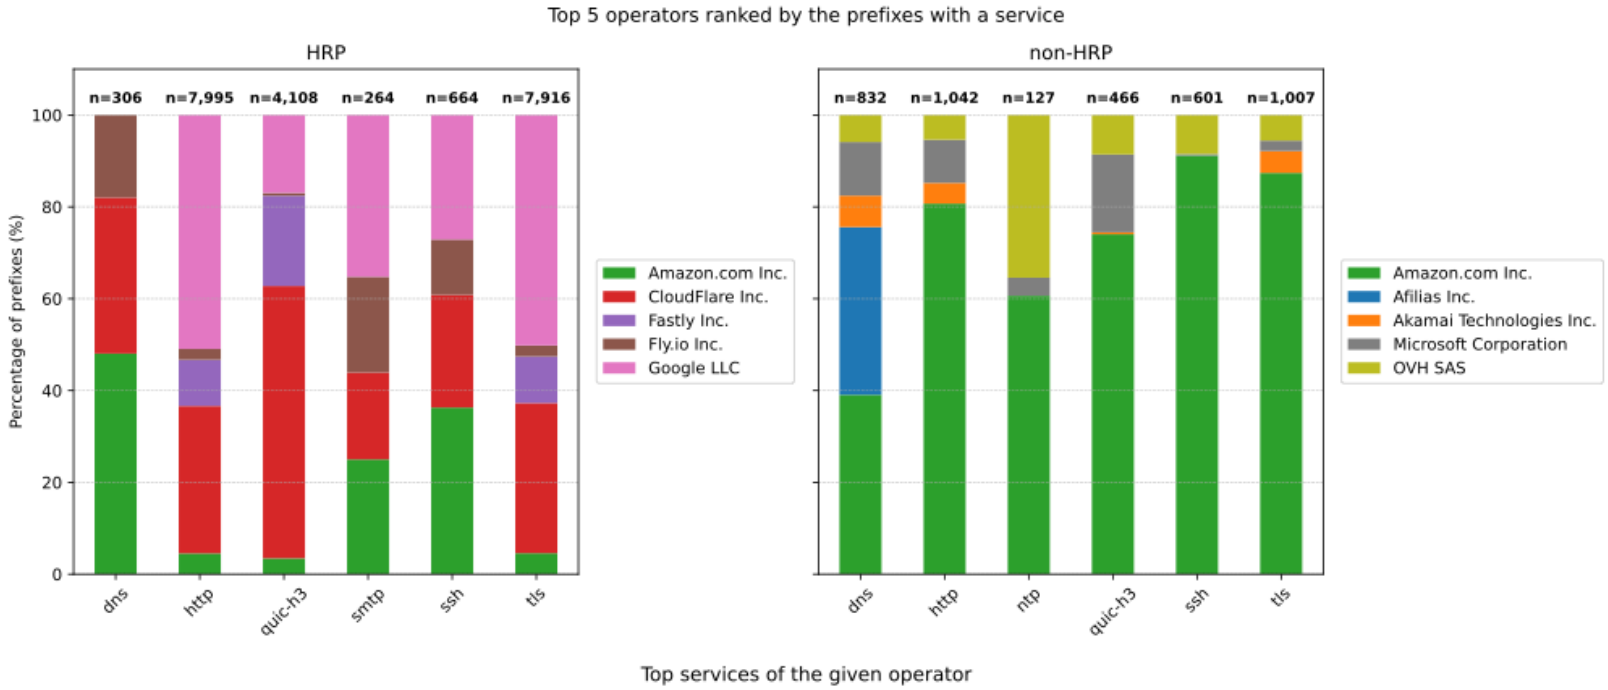

Single purpose prefixes

- Avoid interference from other services hosted on the same prefix or even IP
  - HRP 229 vs. 575 non-HRP
  - DNS has the most dedicated prefixes
  - Email services also visible on dedicated prefixes

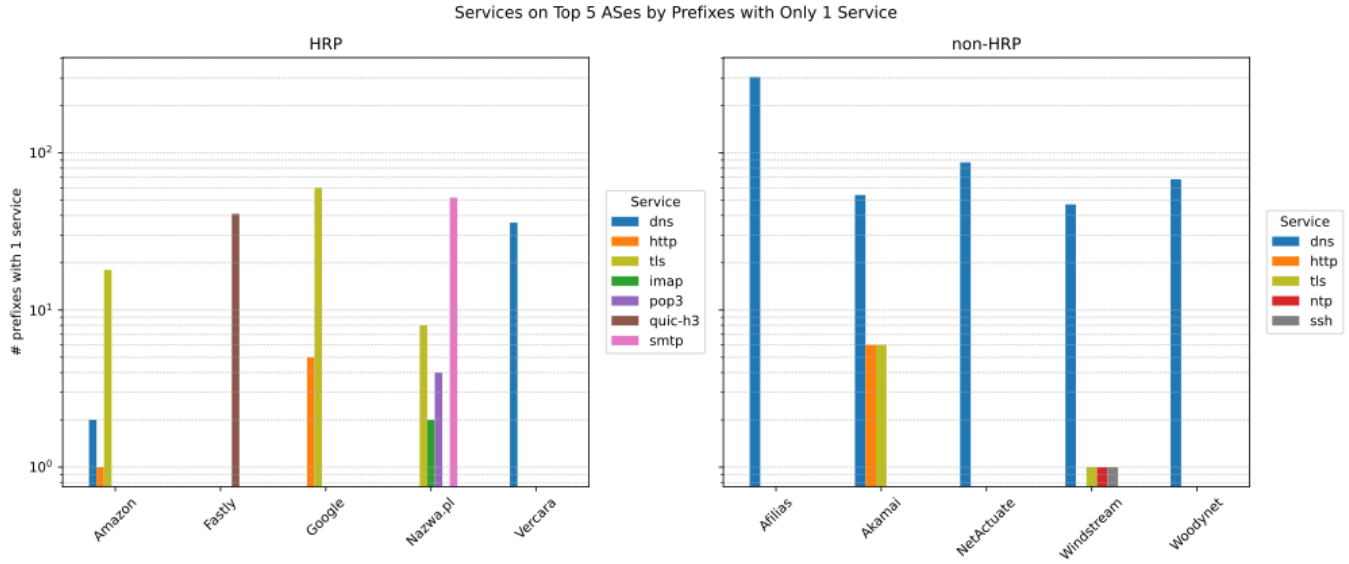

Prefixes with dedicated purpose

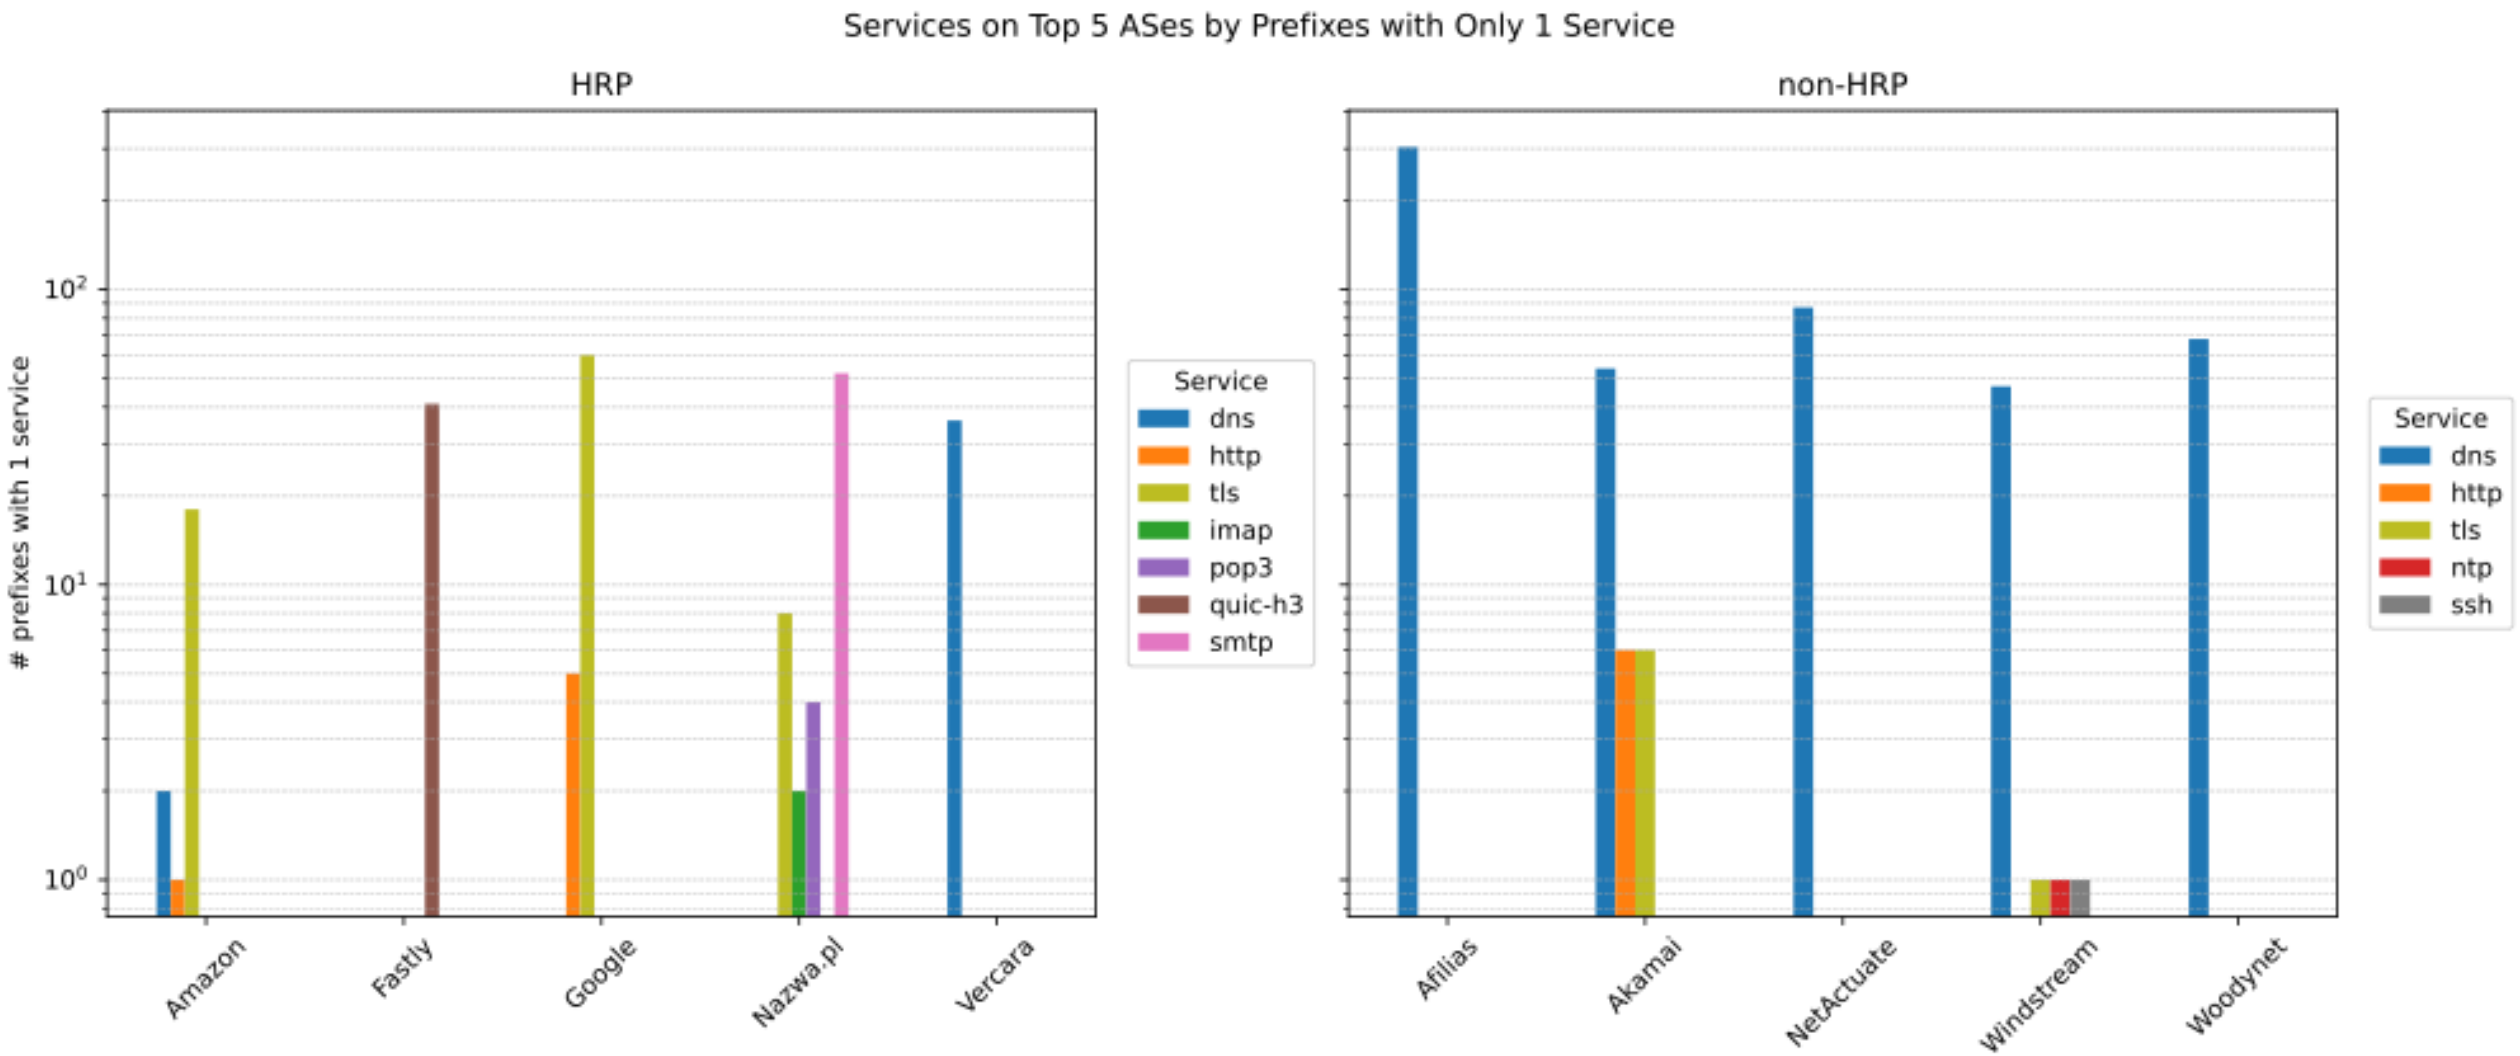

Conclusion

- Anycast is widely used
  - Mostly used by CDNs for HTTP (prefix, IP granularity)
  - Mostly used for DNS (AS granularity)
- Deployment strategies
  - Regional anycast widespread
  - Many ASes anycast using BYOIP providers (Vultr, Cloudflare, ..)

Next steps

- What is behind hosts with TLS? Follow-up ZGrab scan on such hosts
- Which DNS records point to anycast? (update view from 2021)
  - A/AAAA (web)
  - NS (nameservers)
  - MX (mail)
- RQ3 longitudinal analysis
  - Monthly granularity (2+ years) using Censys
  - Daily granularity using our scanning infrastructure (short-lived anycast)