UNIVERSITÄT DES SAARLANDES

A First Look At IPv6 Hypergiant Infrastructure

Fahad Hilal, Patrick Sattler, Kevin Vermeulen, Oliver Gasser

CoNEXT'24

MAX PLANCK INSTITUTE FOR INFORMATICS

Technical University of Munich

LAAS CNRS

IPinfo

Introduction-Hypergiants and Traffic Consolidation

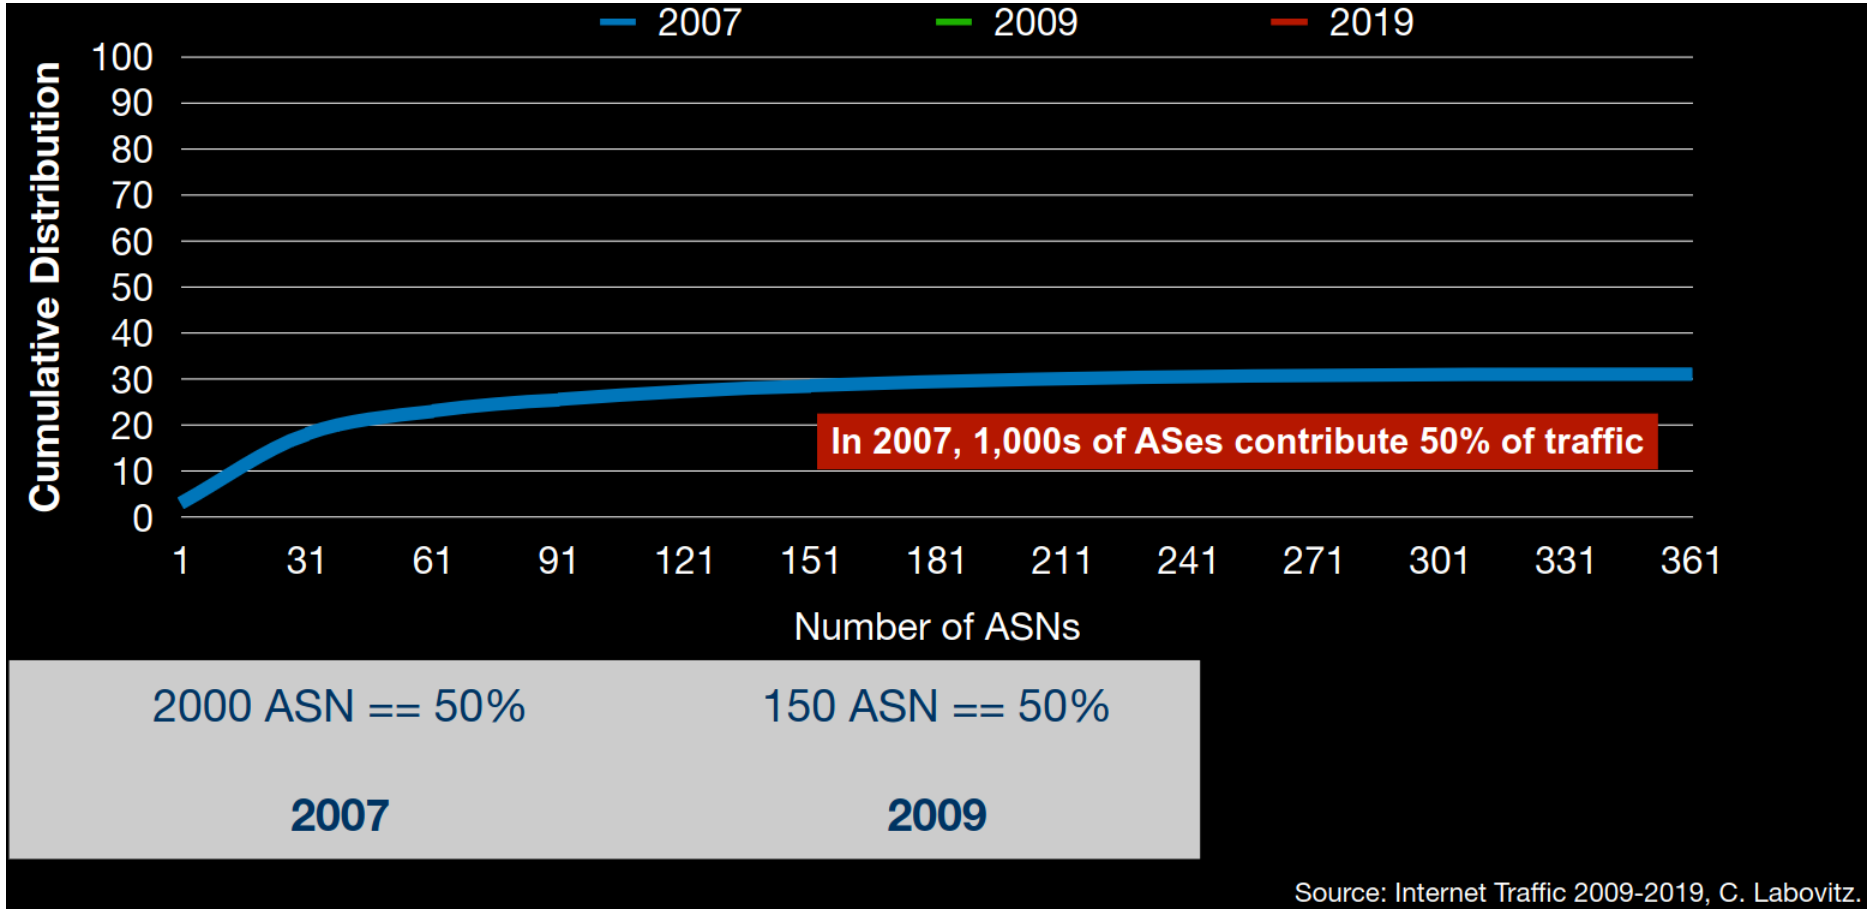

https://pgigis.github.io/hypergiants-offnets/data/pdf/seven_years_in_the_life_of_hypergiants_offnets_slides.pdf

Introduction-Hypergiants and Traffic Consolidation

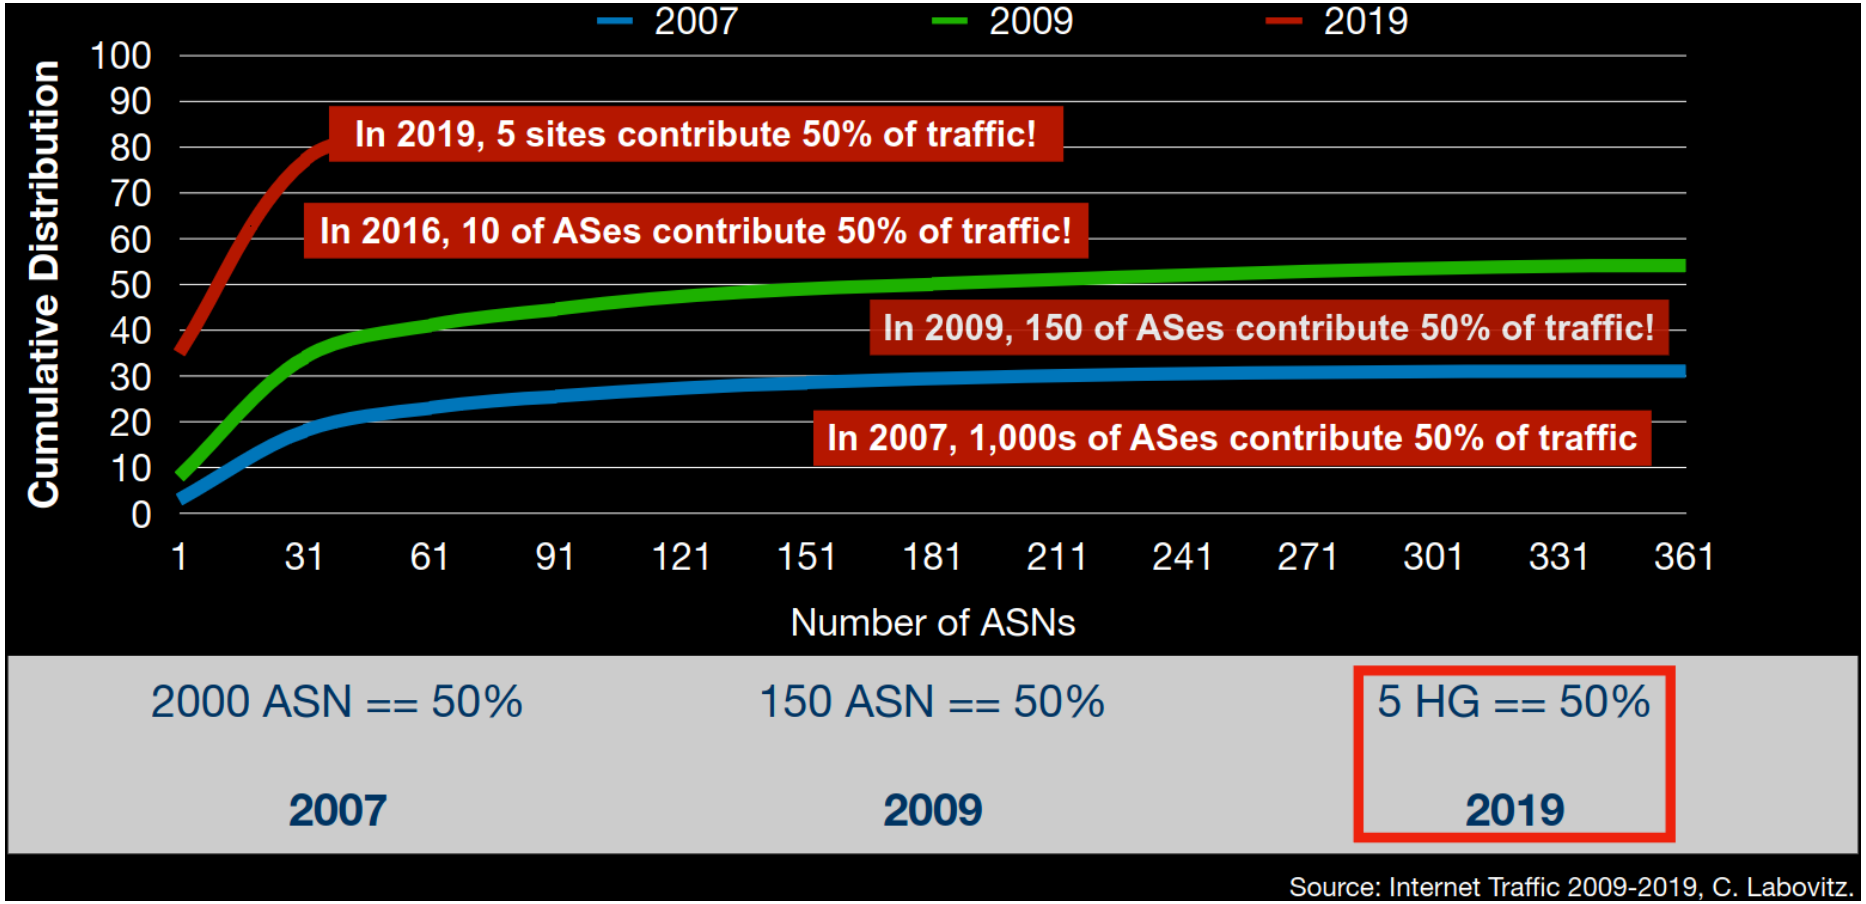

https://pgigis.github.io/hypergiants-offnets/data/pdf/seven_years_in_the_life_of_hypergiants_offnets_slides.pdf

Introduction-Hypergiants and Traffic Consolidation

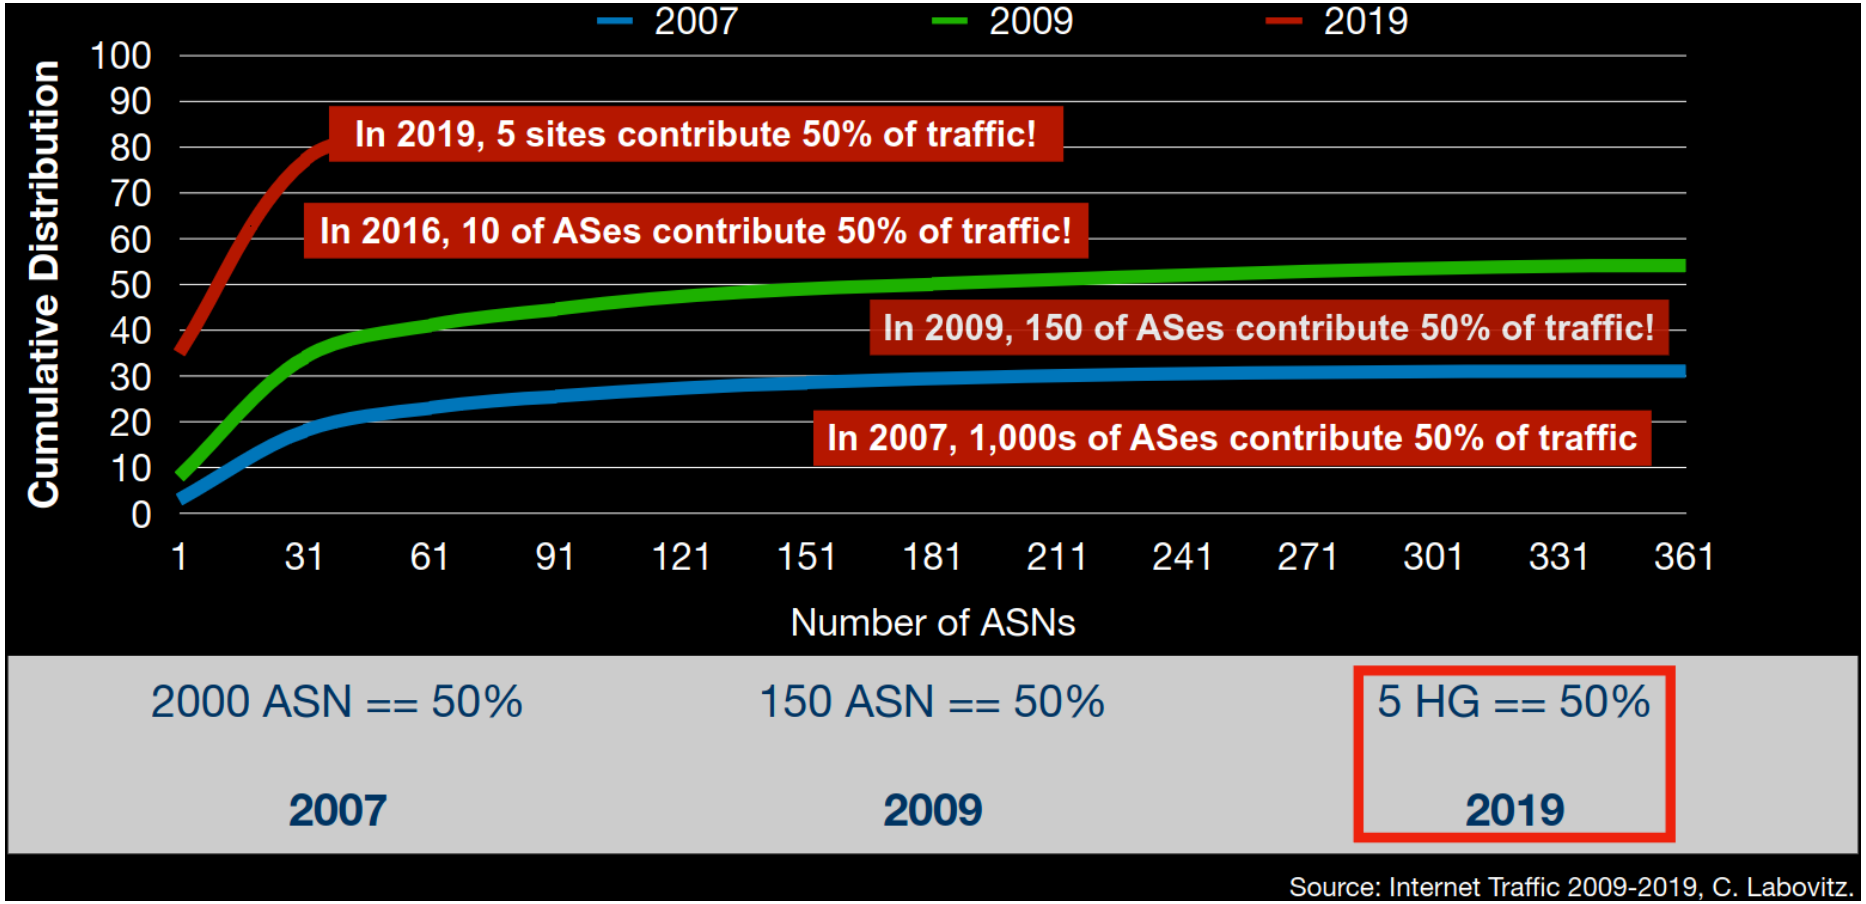

https://pgigis.github.io/hypergiants-offnets/data/pdf/seven_years_in_the_life_of_hypergiants_offnets_slides.pdf

Introduction-Hypergiant Expansion Strategies

UNIVERSITÄT DES SAARLANDES

- Build datacenters
- Roll out fiber to build backbone
- Peer at IXPs and co-location facilities
- Peer directly with eye-balls

Introduction-Hypergiant Expansion Strategies

UNIVERSITÄT DES SAARLANDES

- Build datacenters
- Roll out fiber to build backbone
- Peer at IXPs and co-location facilities
- Peer directly with eye-balls

Also, deploy Off-nets!

Goals

UNIVERSITÄT DES SAARLANDES

- Uncover ***off-nets*** for HGs
  - ***IPv6*** deployment
  - current state in IPv4

Goals

UNIVERSITÄT DES SAARLANDES

- Uncover ***off-nets*** for HGs
  - ***IPv6*** deployment
  - current state in IPv4
- ***Regional*** and ***network-type*** trends
- IPv6 vs IPv4 ***performance***

Motivation-Off-nets

- Impact Internet structure, traffic flows

Motivation-Off-nets

- Impact Internet structure, traffic flows
- Understand HG ***expansion*** strategies
- Serve
  - 70-90% ***Google*** traffic [1]
  - 95% ***Netflix*** traffic [2]

[1] Gill, Phillipa, et al. "M-Lab: User initiated Internet data for the research community." ACM SIGCOMM Computer Communication Review 52.1 (2022): 34-37.
[2] https://about.netflix.com/en/news/red-light-green-light-no-to-network-usage-fees.

Motivation-IPv6

- Increasing IPv6 ***adoption***

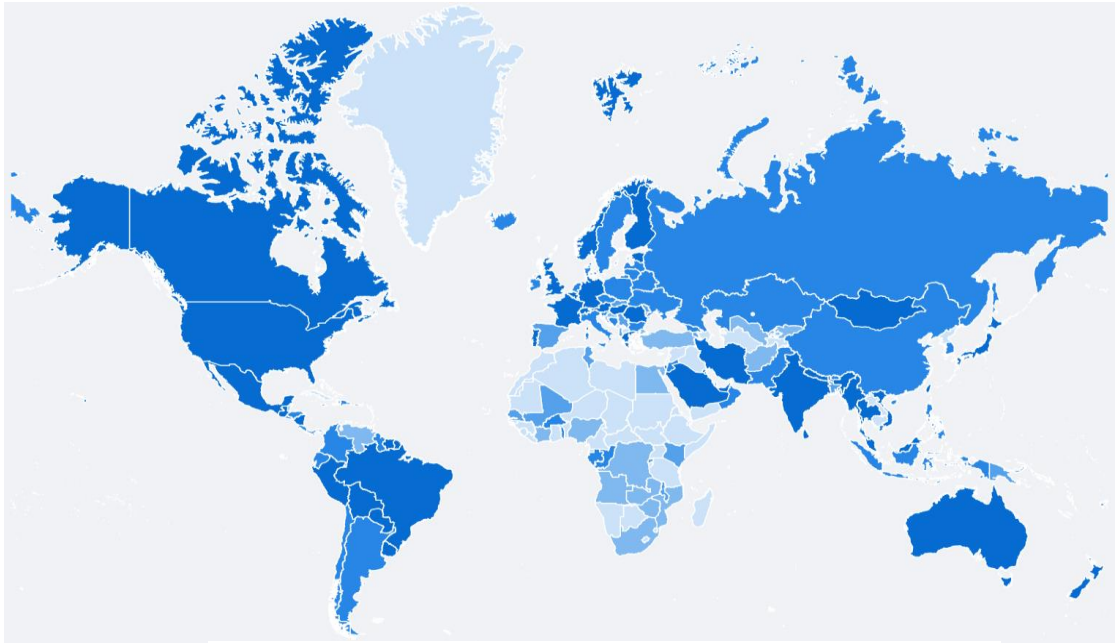

https://www.facebook.com/ipv6/?tab=ipv6_country

Motivation-IPv6

- Increasing IPv6 ***adoption***
- HG IPv6 off-nets ***unexamined***

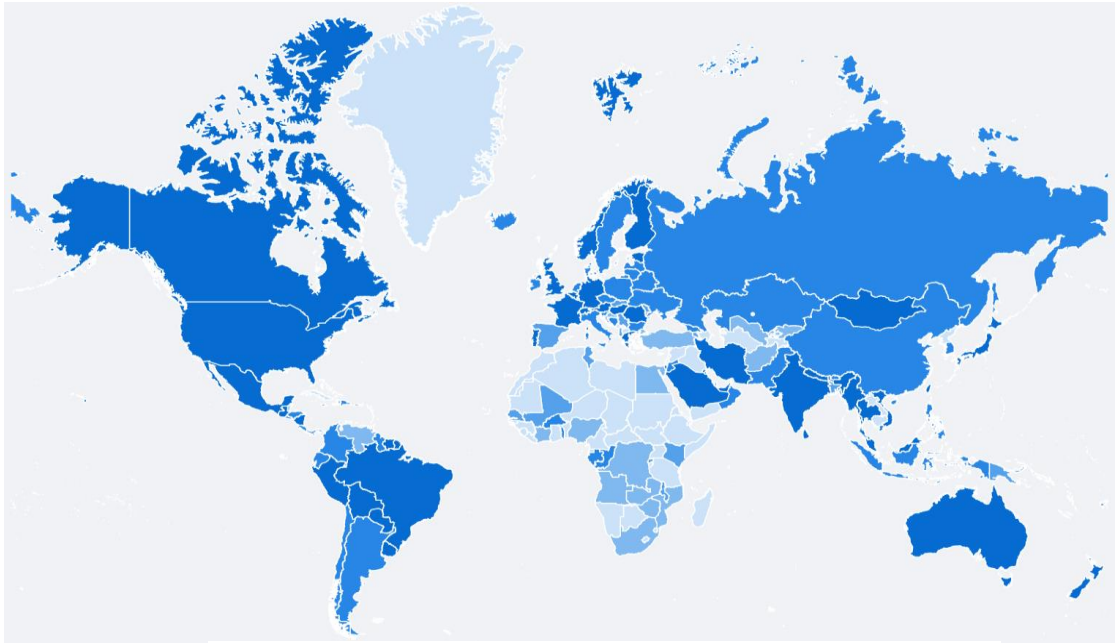

https://www.facebook.com/ipv6/?tab=ipv6_country

Motivation-IPv6

- Increasing IPv6 ***adoption***
- HG IPv6 off-nets ***unexamined***
- Performance
  - benefit or penalty?

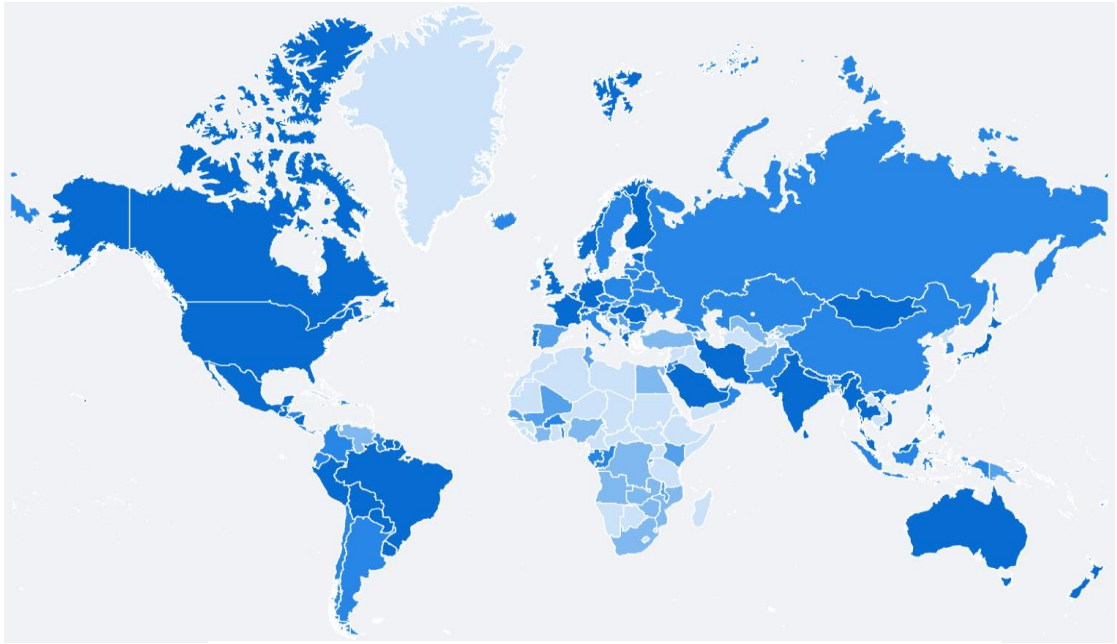

https://www.facebook.com/ipv6/?tab=ipv6_country

Related Work

UNIVERSITÄT DES SAARLANDES

- Prior approaches lack generality
  - Bottger et al. [1], study ***Netflix***
  - Calder et al. [2], study ***Google***

[1] Böttger, Timm, et al. "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the netflix cdn." ACM SIGCOMM Computer Communication Review 48.1 (2018): 28-34.
[2] Calder, Matt, et al. "Mapping the expansion of Google's serving infrastructure." Proceedings of the 2013 conference on Internet measurement conference. 2013.

Related Work

UNIVERSITÄT DES SAARLANDES

- Prior approaches lack generality
  - Bottger et al. [1], study ***Netflix***
  - Calder et al. [2], study ***Google***

General methodology?

[1] Böttger, Timm, et al. "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the netflix cdn." ACM SIGCOMM Computer Communication Review 48.1 (2018): 28-34.
[2] Calder, Matt, et al. "Mapping the expansion of Google's serving infrastructure." Proceedings of the 2013 conference on Internet measurement conference. 2013.

Approach-The TLS Side-Channel

- Most traffic ***encrypted***

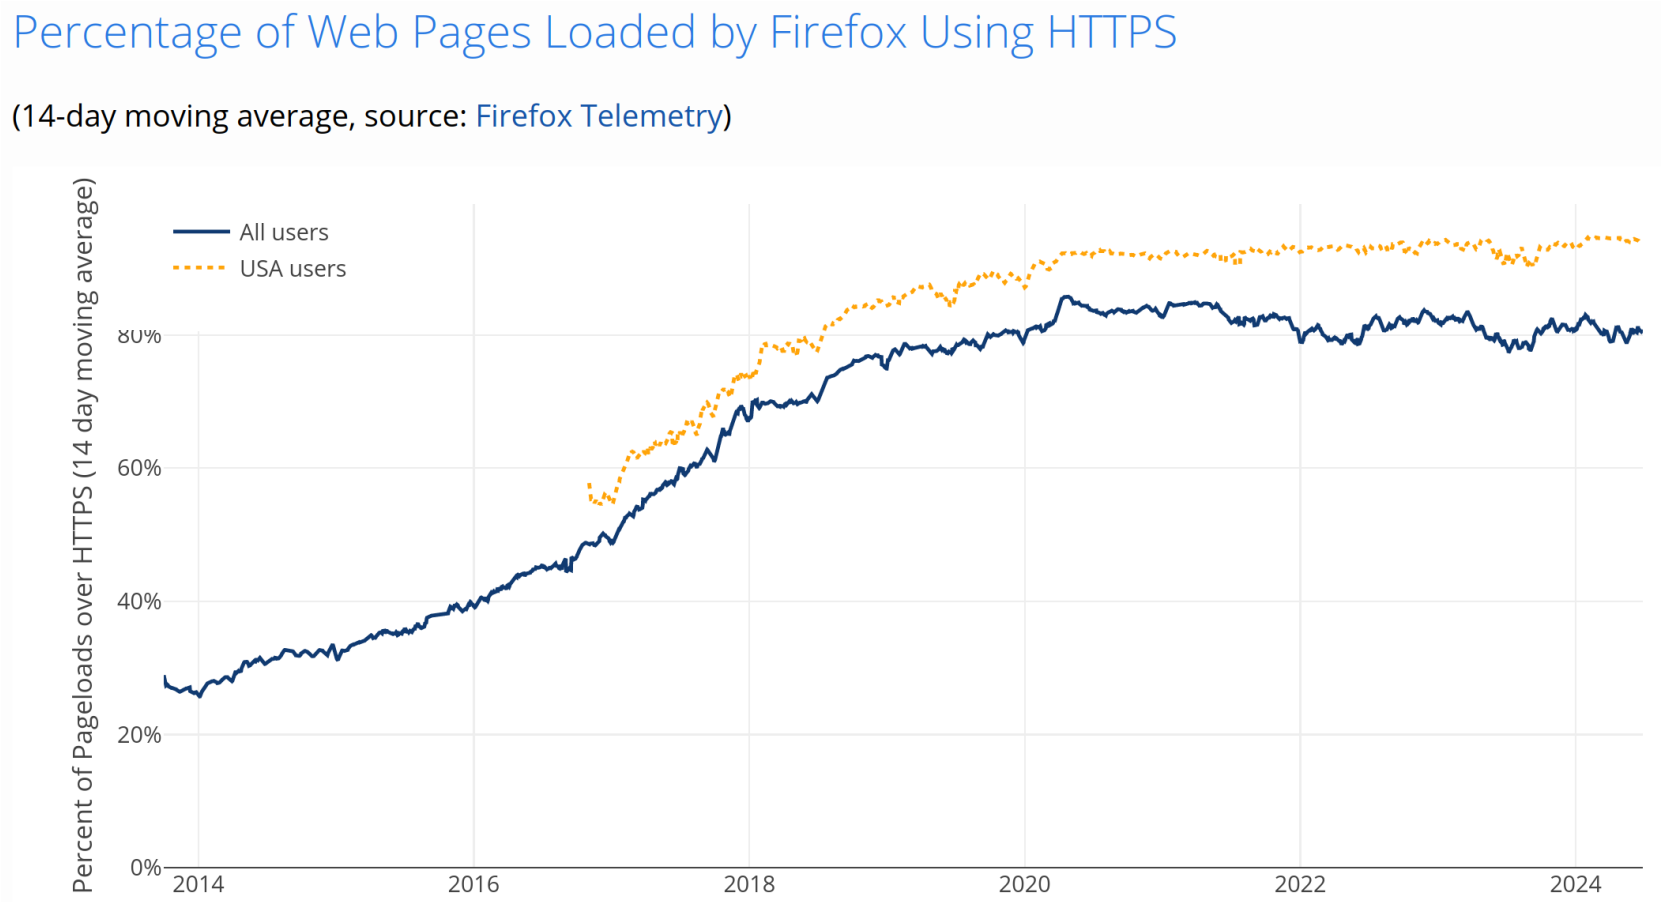

Approach-The TLS Side-Channel

- Most traffic ***encrypted***
  - HGs no *different!*

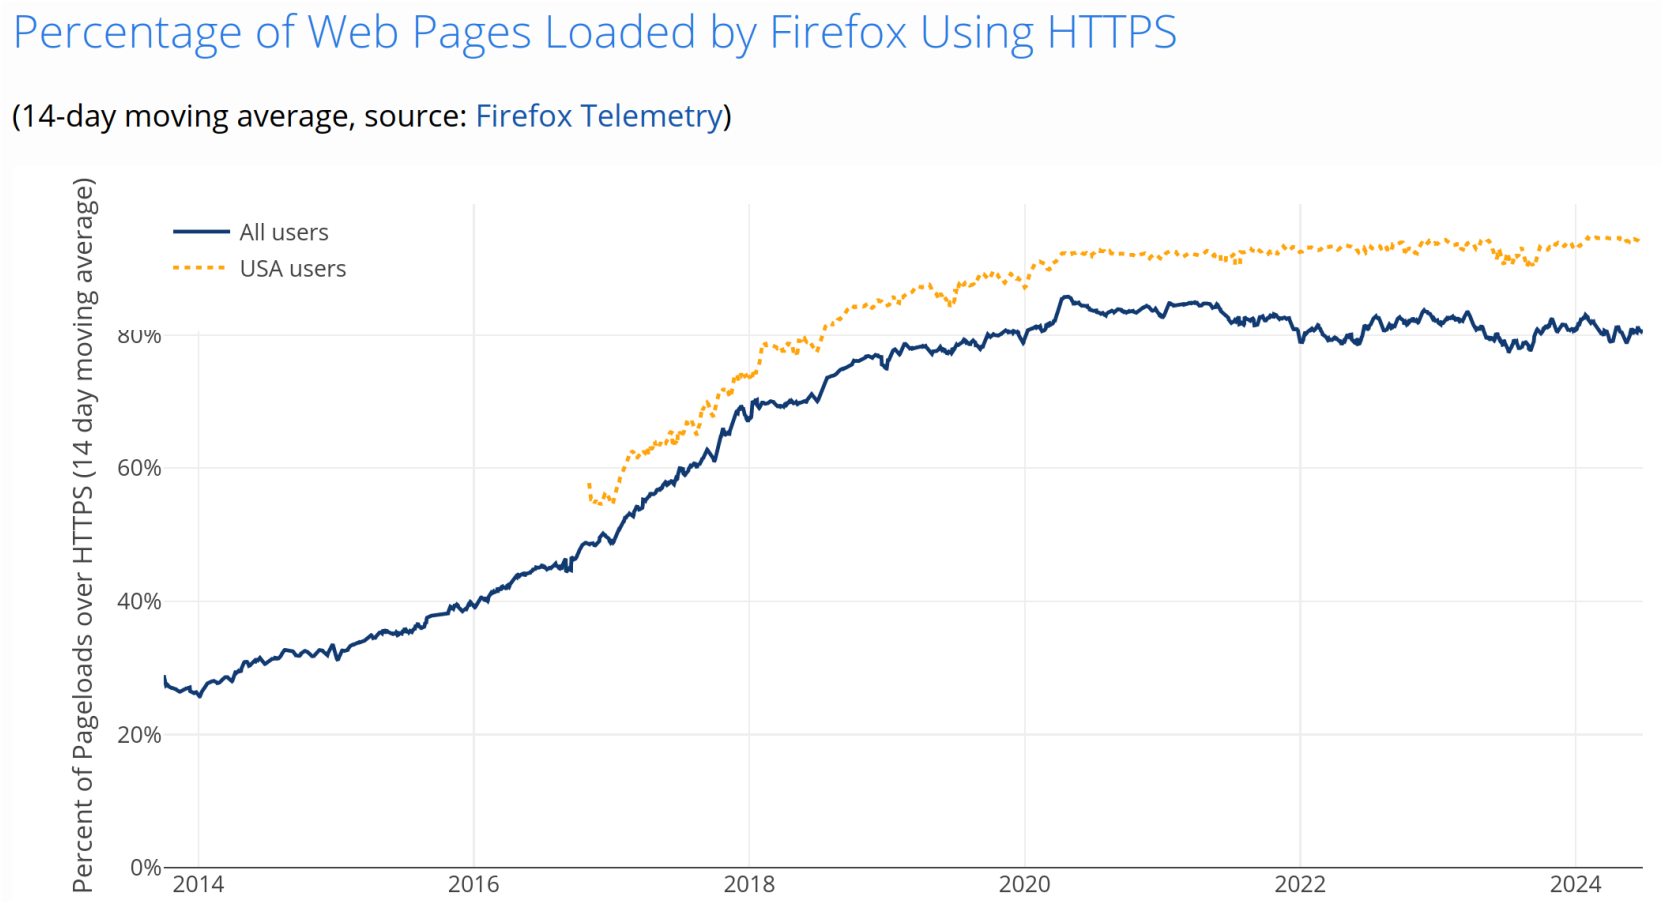

Approach-The TLS Side-Channel

- Most traffic ***encrypted***
  - HGs no *different!*

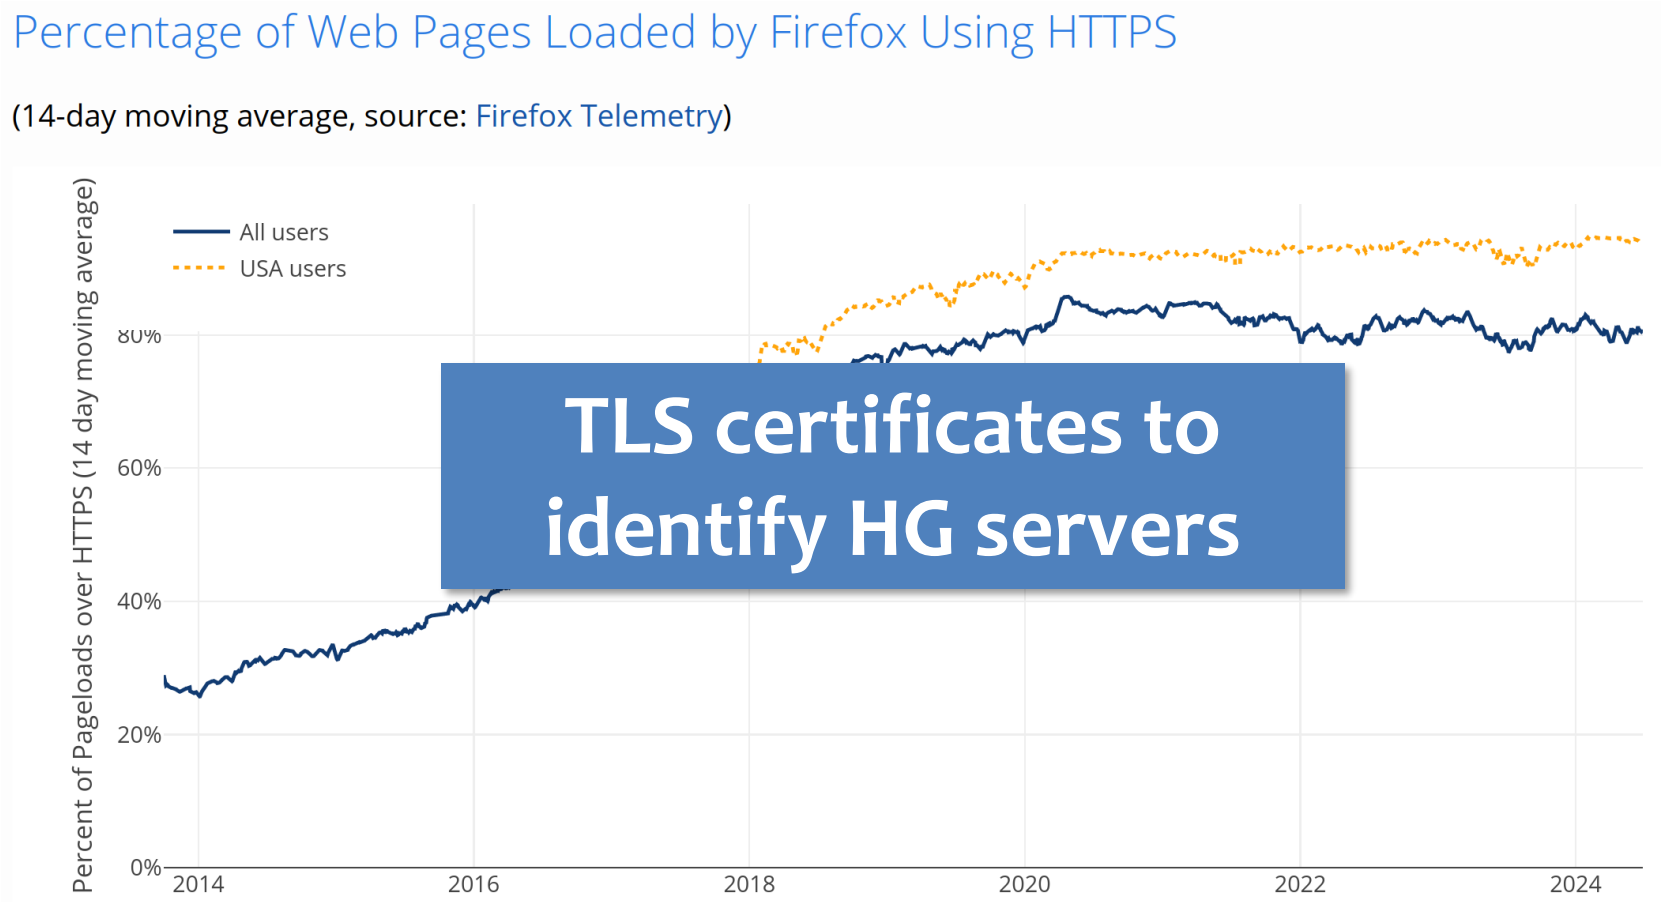

Related Work

UNIVERSITÄT DES SAARLANDES

- Prior approaches lack generality
  - Bottger et al. [1], study ***Netflix***
  - Calder et al. [2], study ***Google***
- General methodology
  - Gigis et al. [3], uncover off-nets of several HGs
    - limited to ***IPv4***

[1] Böttger, Timm, et al. "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the netflix cdn." ACM SIGCOMM Computer Communication Review 48.1 (2018): 28-34.
[2] Calder, Matt, et al. "Mapping the expansion of Google's serving infrastructure." Proceedings of the 2013 conference on Internet measurement conference. 2013.
[3] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021

Related Work

UNIVERSITÄT DES SAARLANDES

- Prior approaches lack generality
  - Bottger et al. [1], study ***Netflix***
  - Calder et al. [2], study ***Google***
- General methodology
  - Gigis et al. [3], uncover off-nets of several HGs
    - limited to ***IPv4***

No HG IPv6 off-net studies

[1] Böttger, Timm, et al. "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the netflix cdn." ACM SIGCOMM Computer Communication Review 48.1 (2018): 28-34.
[2] Calder, Matt, et al. "Mapping the expansion of Google's serving infrastructure." Proceedings of the 2013 conference on Internet measurement conference. 2013.
[3] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021

Challenges and Tweaks

UNIVERSITÄT DES SAARLANDES

- ***Challenges***
  - **large** IPv6 address space
  - public cert. data ***unreliable***
    - Gigis et al., use Rapid7, Censys (IPv4)

Challenges and Tweaks

- ***Challenges***
  - **large** IPv6 address space
  - public cert. data ***unreliable***
    - Gigis et al., use Rapid7, Censys (IPv4)
- ***Tweaks***
  - scan state of the art ***IPv6 hitlist*** *[1]*
    - ~1B addresses, mostly ISPs

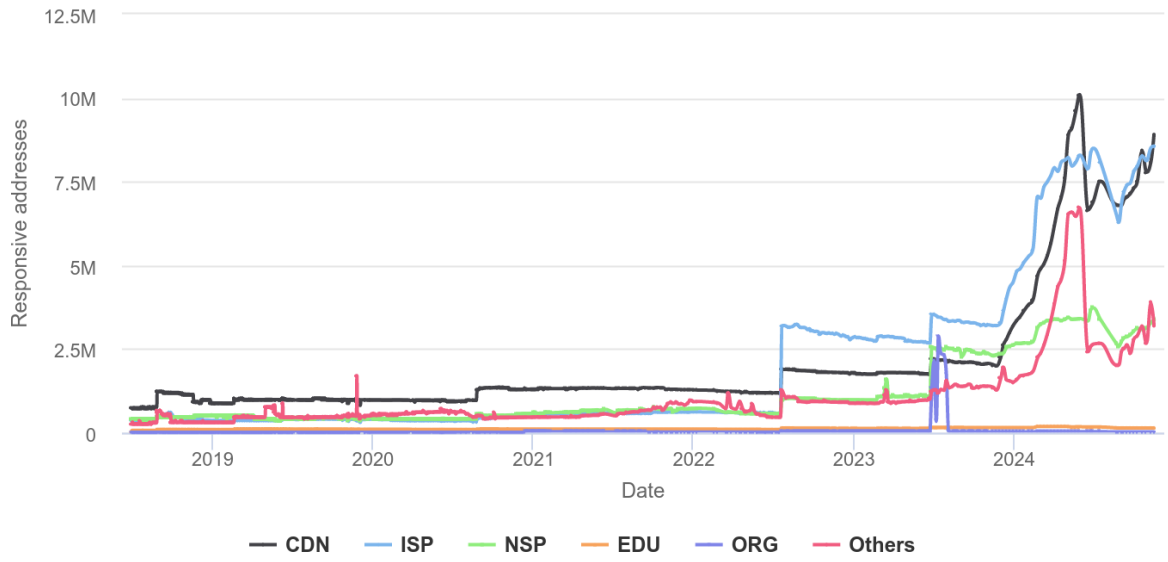

https://ipv6hitlist.github.io/

[1] Zirngibl, Johannes, et al. "Rusty clusters? Dusting an IPv6 research foundation." Proceedings of the 22nd ACM Internet Measurement Conference. 2022.

Challenges and Tweaks

- ***Challenges***
  - **large** IPv6 address space
  - public cert. data ***unreliable***
    - Gigis et al., use Rapid7, Censys (IPv4)
- ***Tweaks***
  - scan state of the art ***IPv6 hitlist*** *[1]*
    - ~1B addresses, mostly ISPs
  - ***SNI*** based TLS measurements
    - top list domains
  - ***merge*** IPv4, IPv6 cert. fingerprints

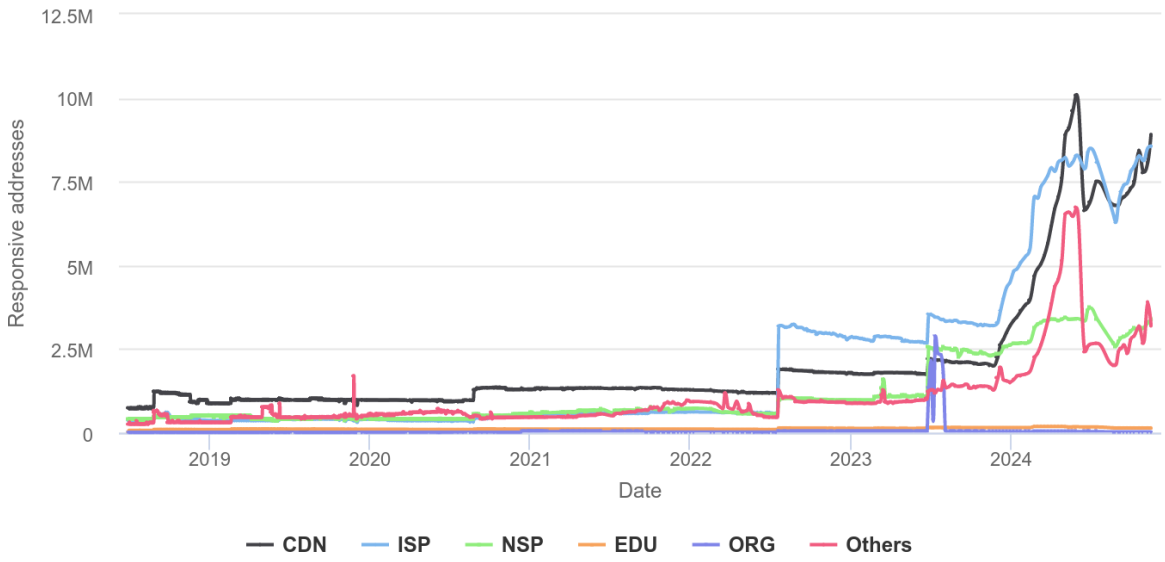

https://ipv6hitlist.github.io/

[1] Zirngibl, Johannes, et al. "Rusty clusters? Dusting an IPv6 research foundation." Proceedings of the 22nd ACM Internet Measurement Conference. 2022.

Challenges and Tweaks

- ***Challenges***
  - **large** IPv6 address space
  - public cert. data ***unreliable***
    - Gigis et al., use Rapid7, Censys (IPv4)
- ***Tweaks***
  - scan state of the art ***IPv6 hitlist*** *[1]*
    - ~1B addresses, mostly ISPs
  - ***SNI*** based TLS measurements
    - top list domains
  - ***merge*** IPv4, IPv6 cert. fingerprints
  - extend coverage
    - dist. ***DNS*** + ***ECS***

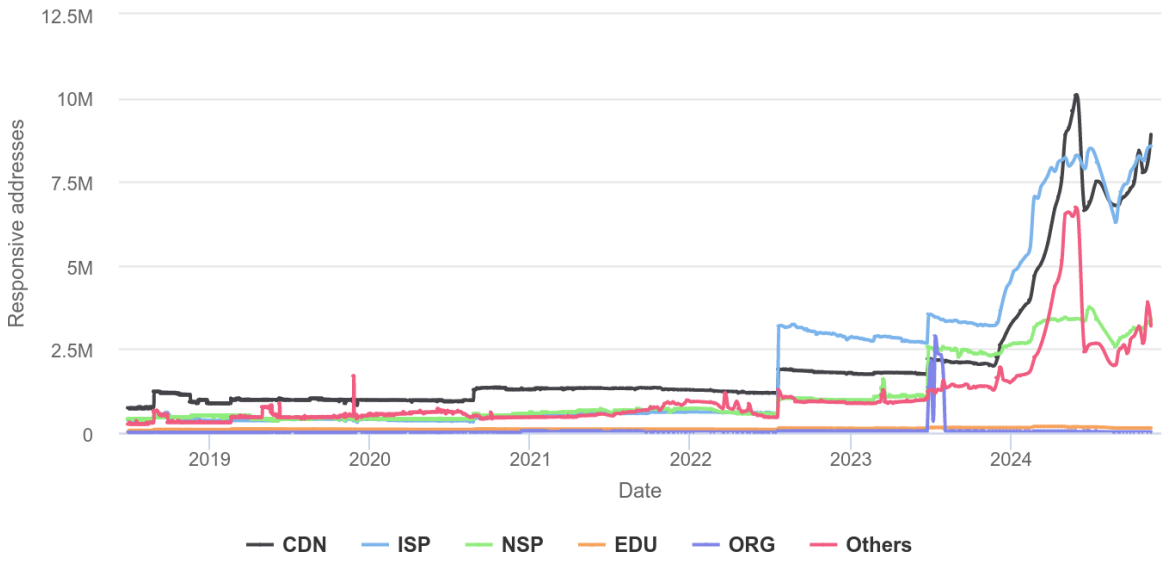

https://ipv6hitlist.github.io/

[1] Zirngibl, Johannes, et al. "Rusty clusters? Dusting an IPv6 research foundation." Proceedings of the 22nd ACM Internet Measurement Conference. 2022.

Methodology

- ***Cert. collection***
  - Step 0: Collect TLS certificates [1, 2]
    - IPv4: full address space + top list domains
    - IPv6: IPv6 hitlist + top list domains

[1] https://zmap.io/
[2] https://github.com/zmap/zgrab2

Methodology

- ***Cert. validation***
  - Step 1: remove self-signed, expired and certificates with non-verified chain

Collected TLS certificates

Methodology

- ***Cert. validation***
  - Step 1: remove self-signed, expired and certificates with non-verified chain

Collected TLS certificates

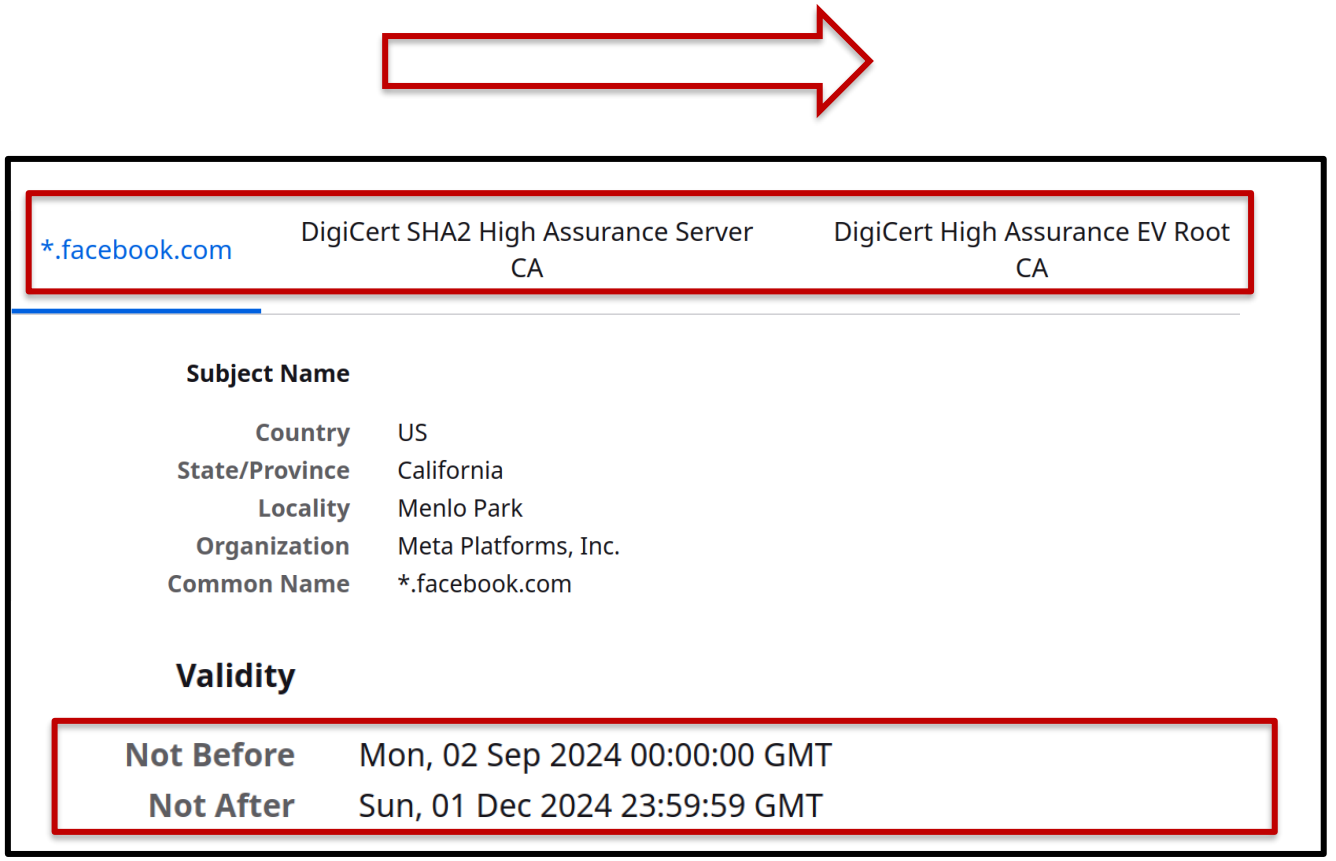

Valid TLS certificates

Methodology

- ***Hypergiant (HG) TLS Fingerprints (FPs)***
  - Step 2: Build per-HG TLS FPs

Valid TLS certificates

Methodology

- ***Hypergiant (HG) TLS Fingerprints (FPs)***
  - Step 2: Build per-HG TLS FPs

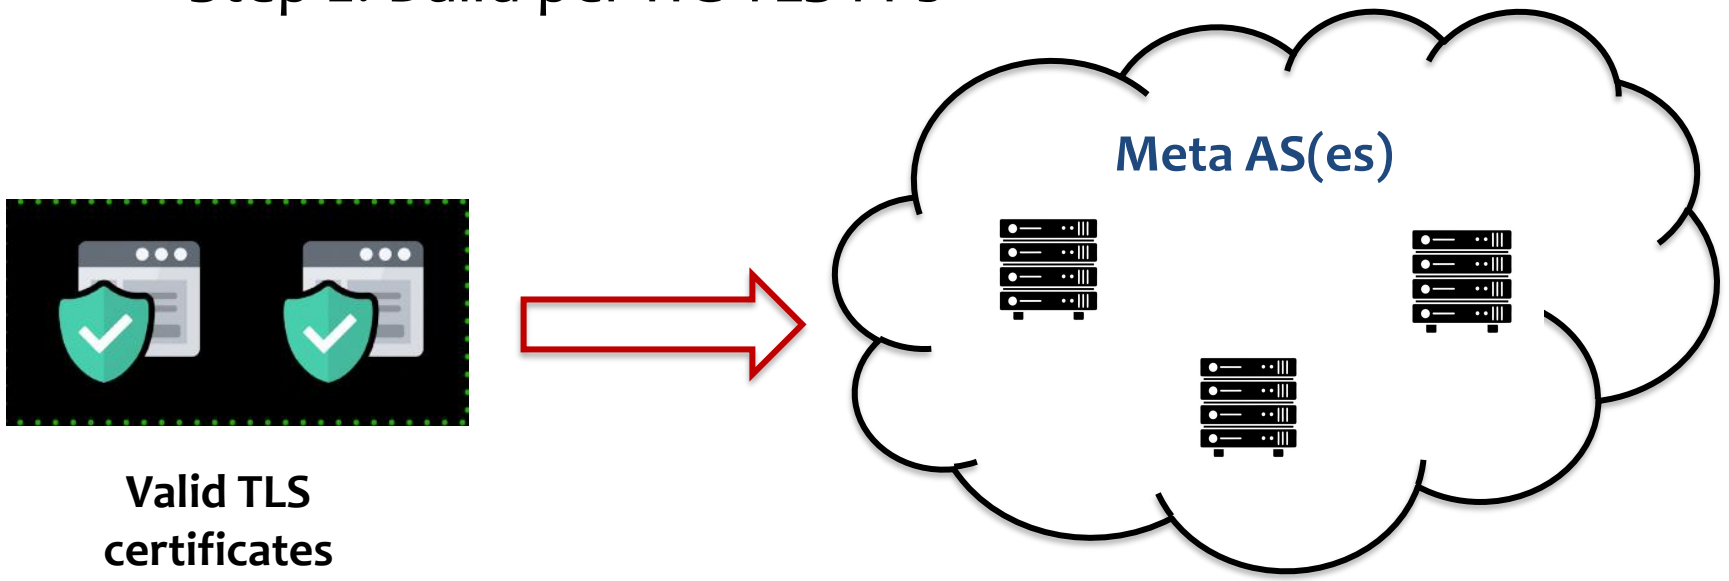

Methodology

- ***Hypergiant (HG) TLS Fingerprints (FPs)***
  - Step 2: Build per-HG TLS FPs

Valid TLS certificates

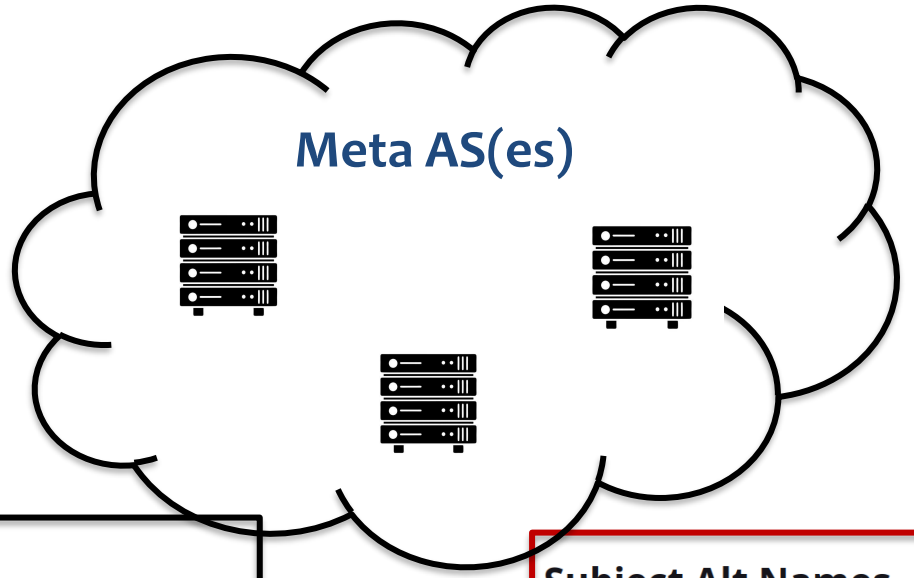

Subject Name	
Country	US
State/Province	California
Locality	Menlo Park
Organization	Meta Platforms, Inc.
Common Name	*.facebook.com

Subject Alt Names	
DNS Name	*.facebook.com
DNS Name	*.facebook.net
DNS Name	*.fbcdn.net
DNS Name	*.fbsbx.com

Methodology

- ***Hypergiant (HG) TLS Fingerprints (FPs)***
  - Step 2: Build per-HG TLS FPs
    - merge IPv6, IPv4 FPs

Valid TLS certificates

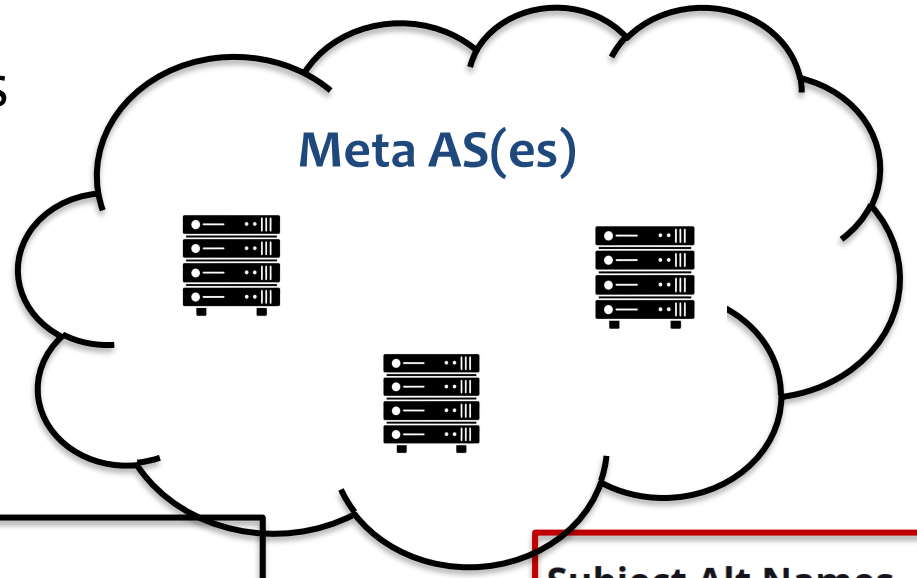

On-net TLS Fingerprints

Subject Name	
Country	US
State/Province	California
Locality	Menlo Park
Organization	Meta Platforms, Inc.
Common Name	*.facebook.com

Subject Alt Names	
DNS Name	*.facebook.com
DNS Name	*.facebook.net
DNS Name	*.fbcdn.net
DNS Name	*.fbsbx.com

Methodology

- ***Candidate Off-nets***
  - Step 3: Apply TLS FPs

On-net TLS Fingerprints

Subject Name	
Country	US
State/Province	California
Locality	Menlo Park
Organization	Meta Platforms, Inc.
Common Name	*.facebook.com

Subject Alt Names	
DNS Name	*.facebook.com
DNS Name	*.facebook.net
DNS Name	*.fbcdn.net
DNS Name	*.fbsbx.com

Methodology

- ***Candidate Off-nets***
  - Step 3: Apply TLS FPs

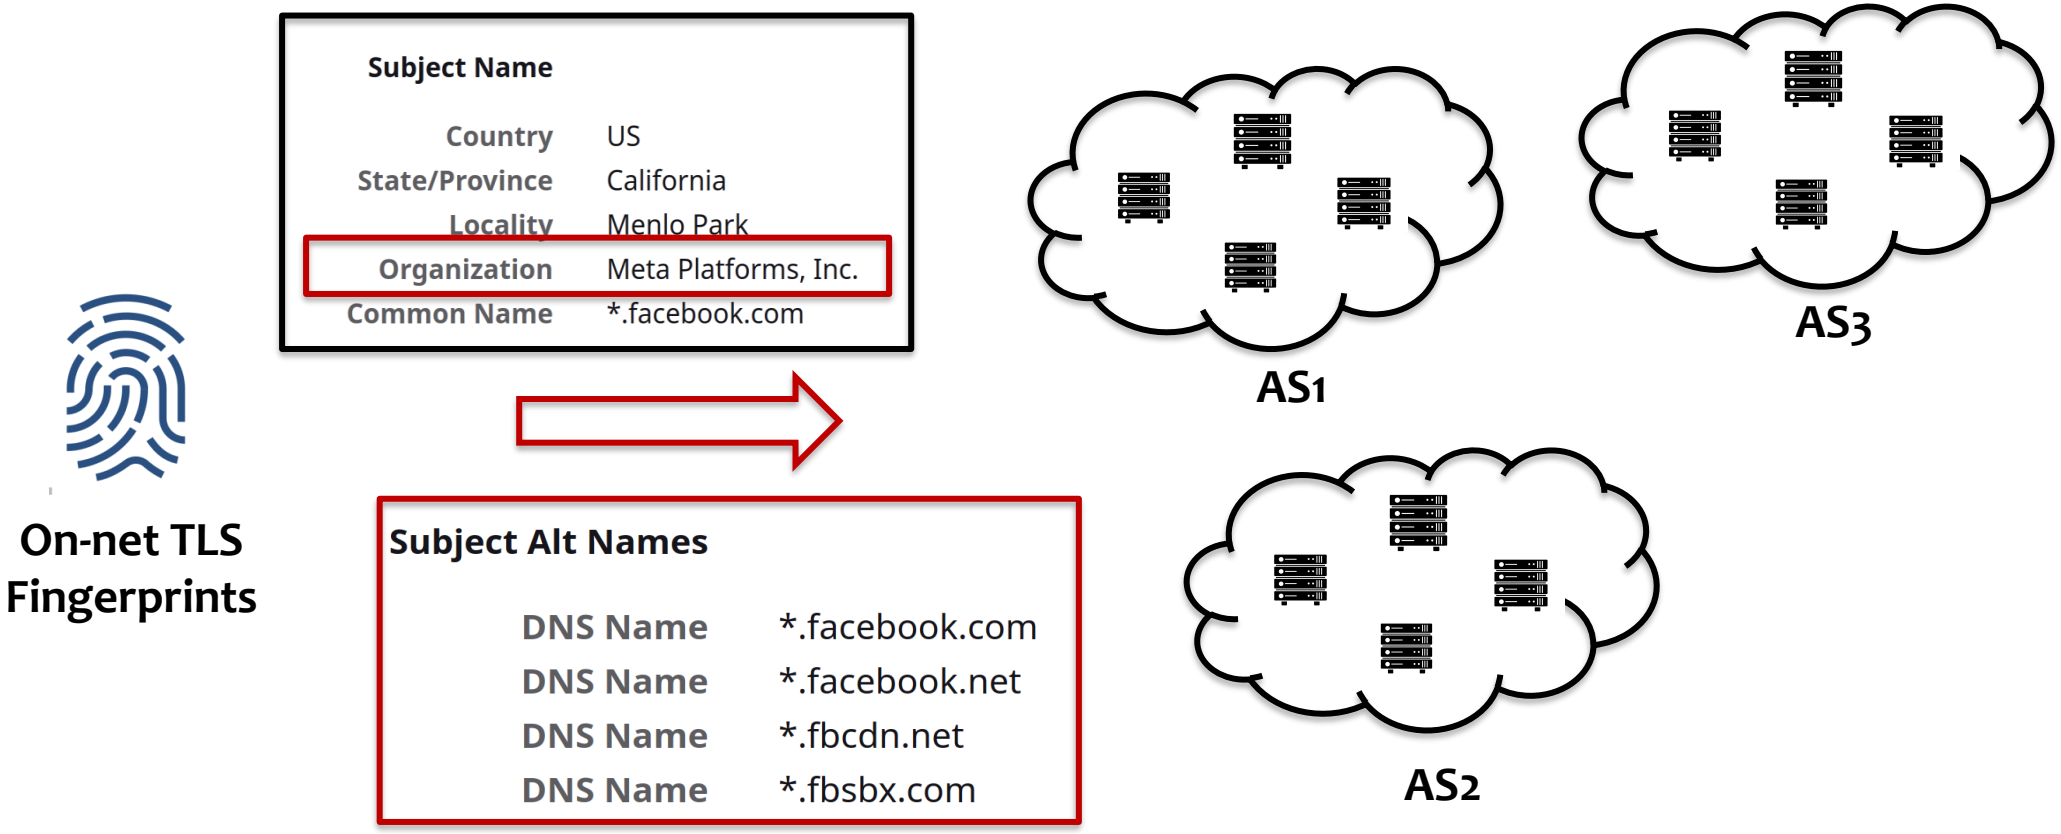

Methodology

- ***Candidate Off-nets***
  - Step 3: Apply TLS FPs
    - match oustside HG ASes - > HG candidate off-net

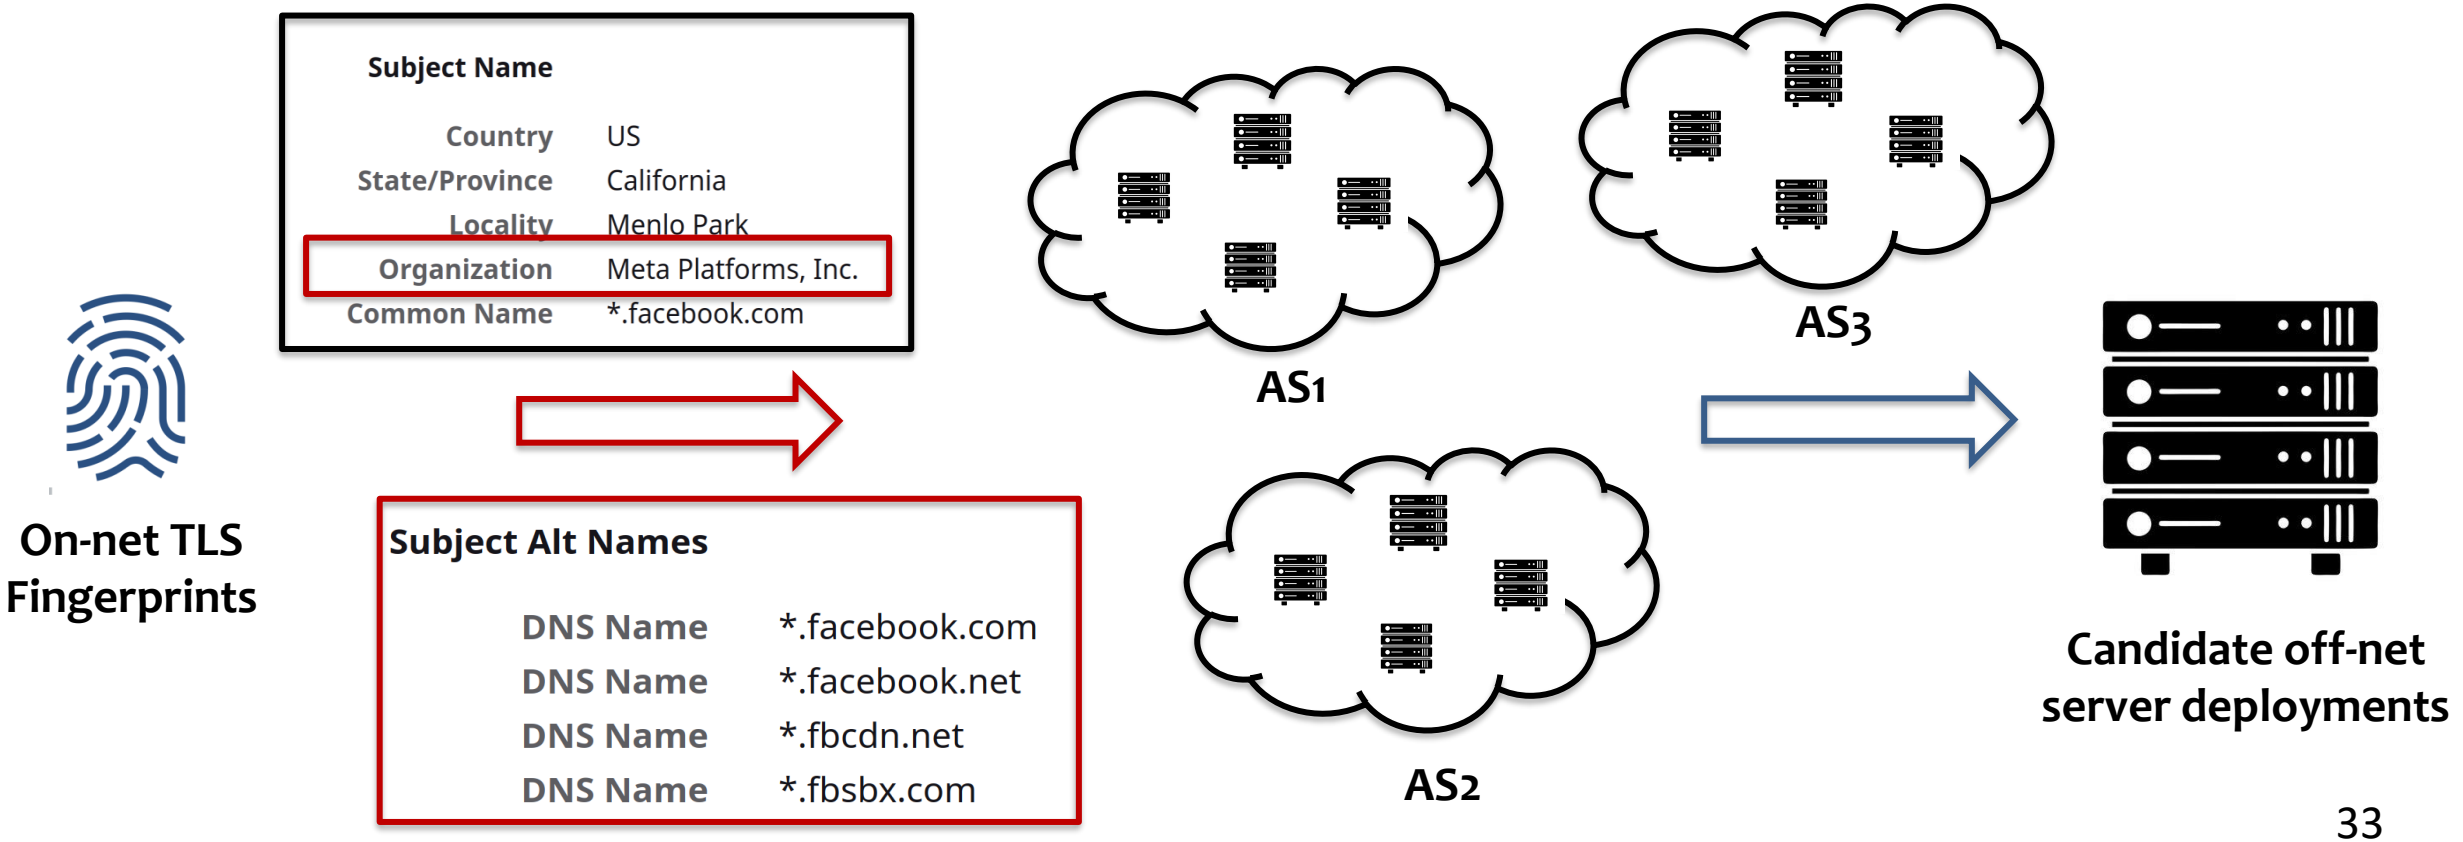

Methodology

- ***Candidate Off-nets***
  - Step 3: Apply TLS FPs
    - match oustside HG ASes - > HG candidate off-net

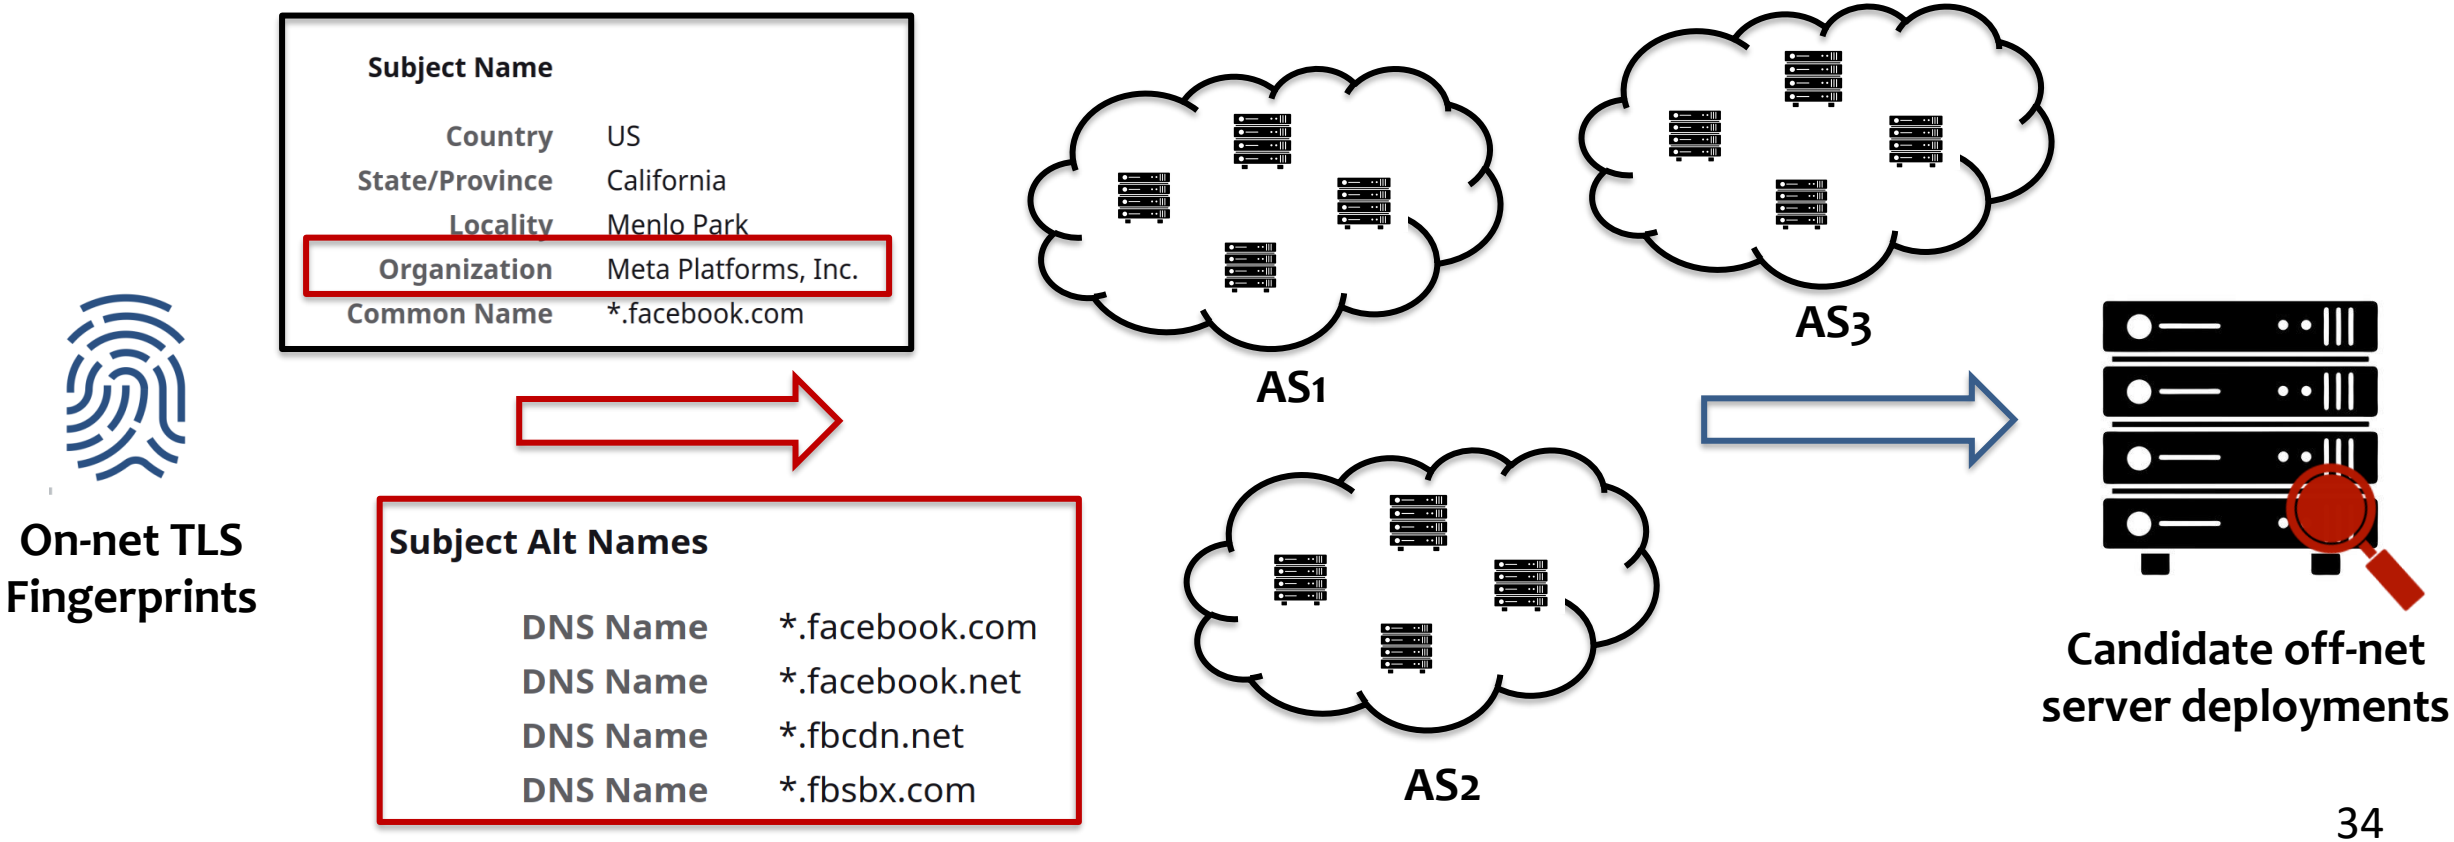

Confirming Off-nets

- ***HG cert. outside HG network***
  - no guarantee of off-net
    - eg. X images from Akamai, verizon other from own infra. [1]
    - NETFLIX uses aws for front-end [2]
  - only ***confirms HG service*** outside HG network

[1] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021
[2] Adhikari, Vijay Kumar, et al. "Unreeling netflix: Understanding and improving multi-cdn movie delivery." 2012 Proceedings IEEE Infocom. IEEE, 2012.

Confirming Off-nets

- ***HG cert. outside HG network***
  - no guarantee of off-net
    - eg. X images from Akamai, verizon other from own infra. [1]
    - NETFLIX uses aws for front-end [2]
  - only ***confirms HG service*** outside HG network
- ***Response headers***
  - used by large providers, CDNs for debugging
    - eg: Server: AkamaiGHost, ***x-fb-debug***

location: https://www.facebook.com/
strict-transport-security: max-age=15552000; preload
content-type: text/html; charset="utf-8"
x-fb-debug: cIYNJ6/AajqLhpnPqcwaUlwui/FPHnhJFg
content-length: 0
date: Sun, 24 Nov 2024 12:01:28 GMT
x-fb-connection-quality: EXCELLENT; q=0.9, rtt=11, rtx=0, c=10, mss=1380, tbw=3535, tp=-1, tpl=-1, uplat=28, ullat=0
alt-svc: h3=":443"; ma=86400

[1] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021
[2] Adhikari, Vijay Kumar, et al. "Unreeling netflix: Understanding and improving multi-cdn movie delivery." 2012 Proceedings IEEE Infocom. IEEE, 2012.

Methodology

- ***Header Fingerprints***
  - Step 4: Learn HG HTTP(S) FPs using headers
    - IPv4 full address space, IPv6 hitlist

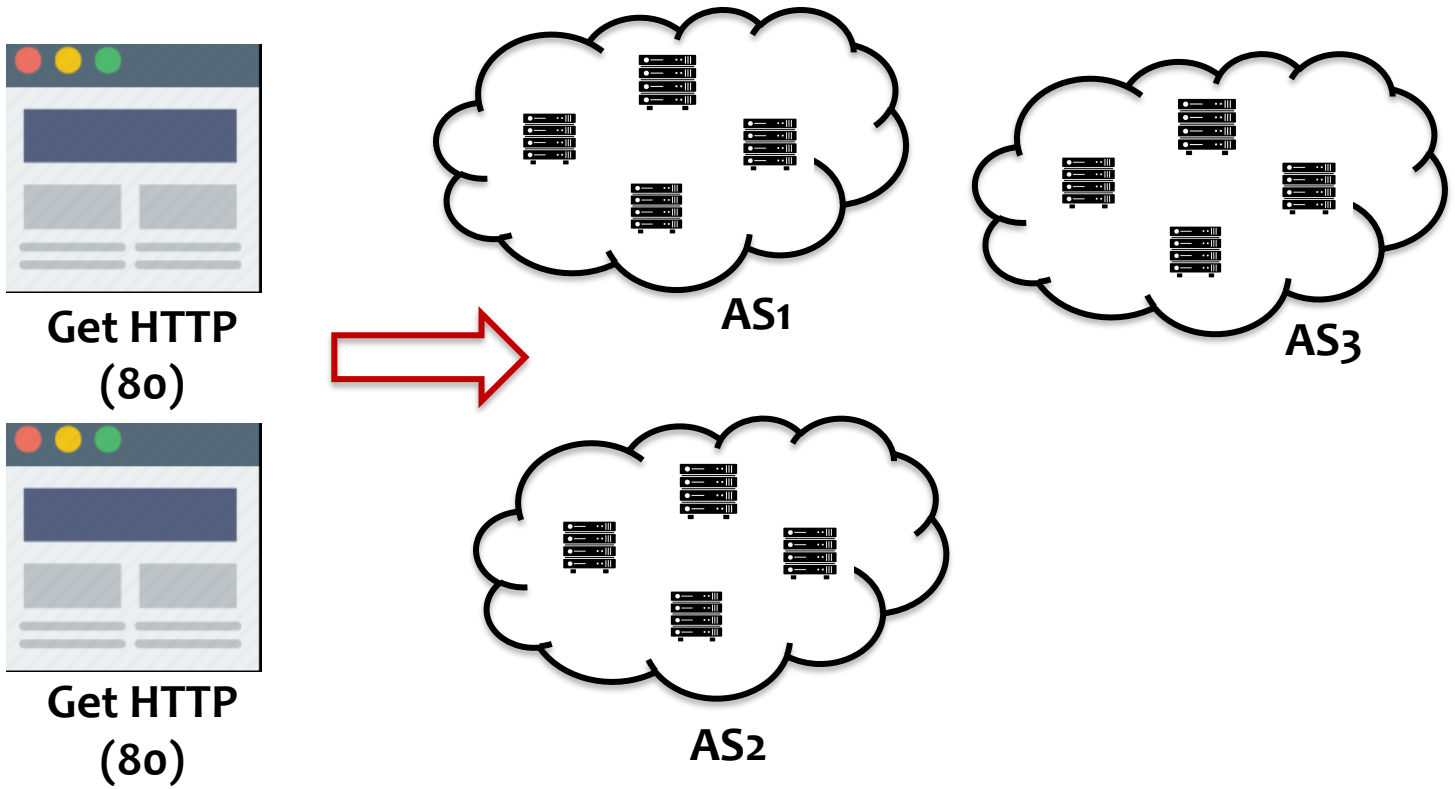

Methodology

- ***Header Fingerprints***
  - Step 4: Learn HG HTTP(S) FPs using headers
    - IPv4 full address space, IPv6 hitlist

Get HTTP (80)

Get HTTP (80)

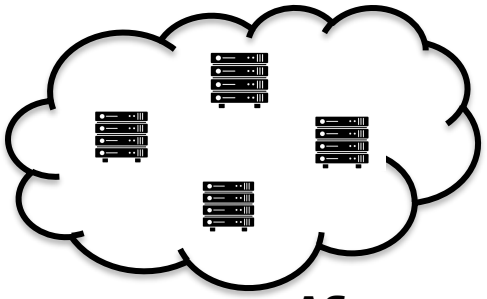

AS1

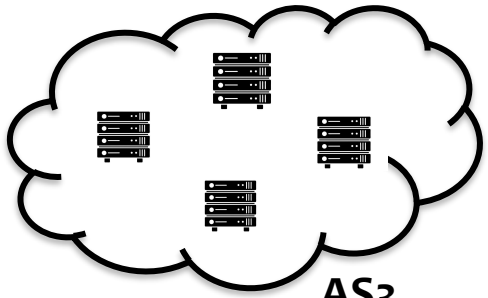

AS3

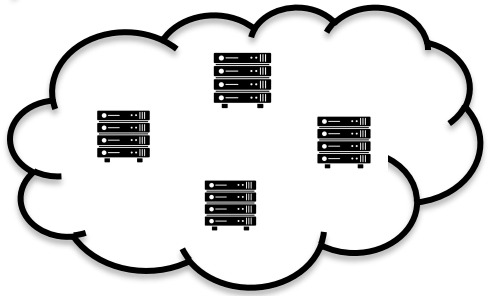

AS2

location: https://www.facebook.com/
strict-transport-security: max-age=15552000; preload
content-type: text/html; charset="utf-8"
x-fb-debug: cIYNJ6/AajqLhpnPqcwaUlwui/FPHnhJFg
content-length: 0
date: Sun, 24 Nov 2024 12:01:28 GMT
x-fb-connection-quality: EXCELLENT; q=0.9, rtt=11, rtx=0, c=10, mss=1380, tbw=3535, tp=-1, tpl=-1, uplat=28, ullat=0
alt-svc: h3=":443"; ma=86400

Methodology

- ***Header Fingerprints***
  - Step 4: Learn HG HTTP(S) FPs using headers
    - IPv4 full address space, IPv6 hitlist

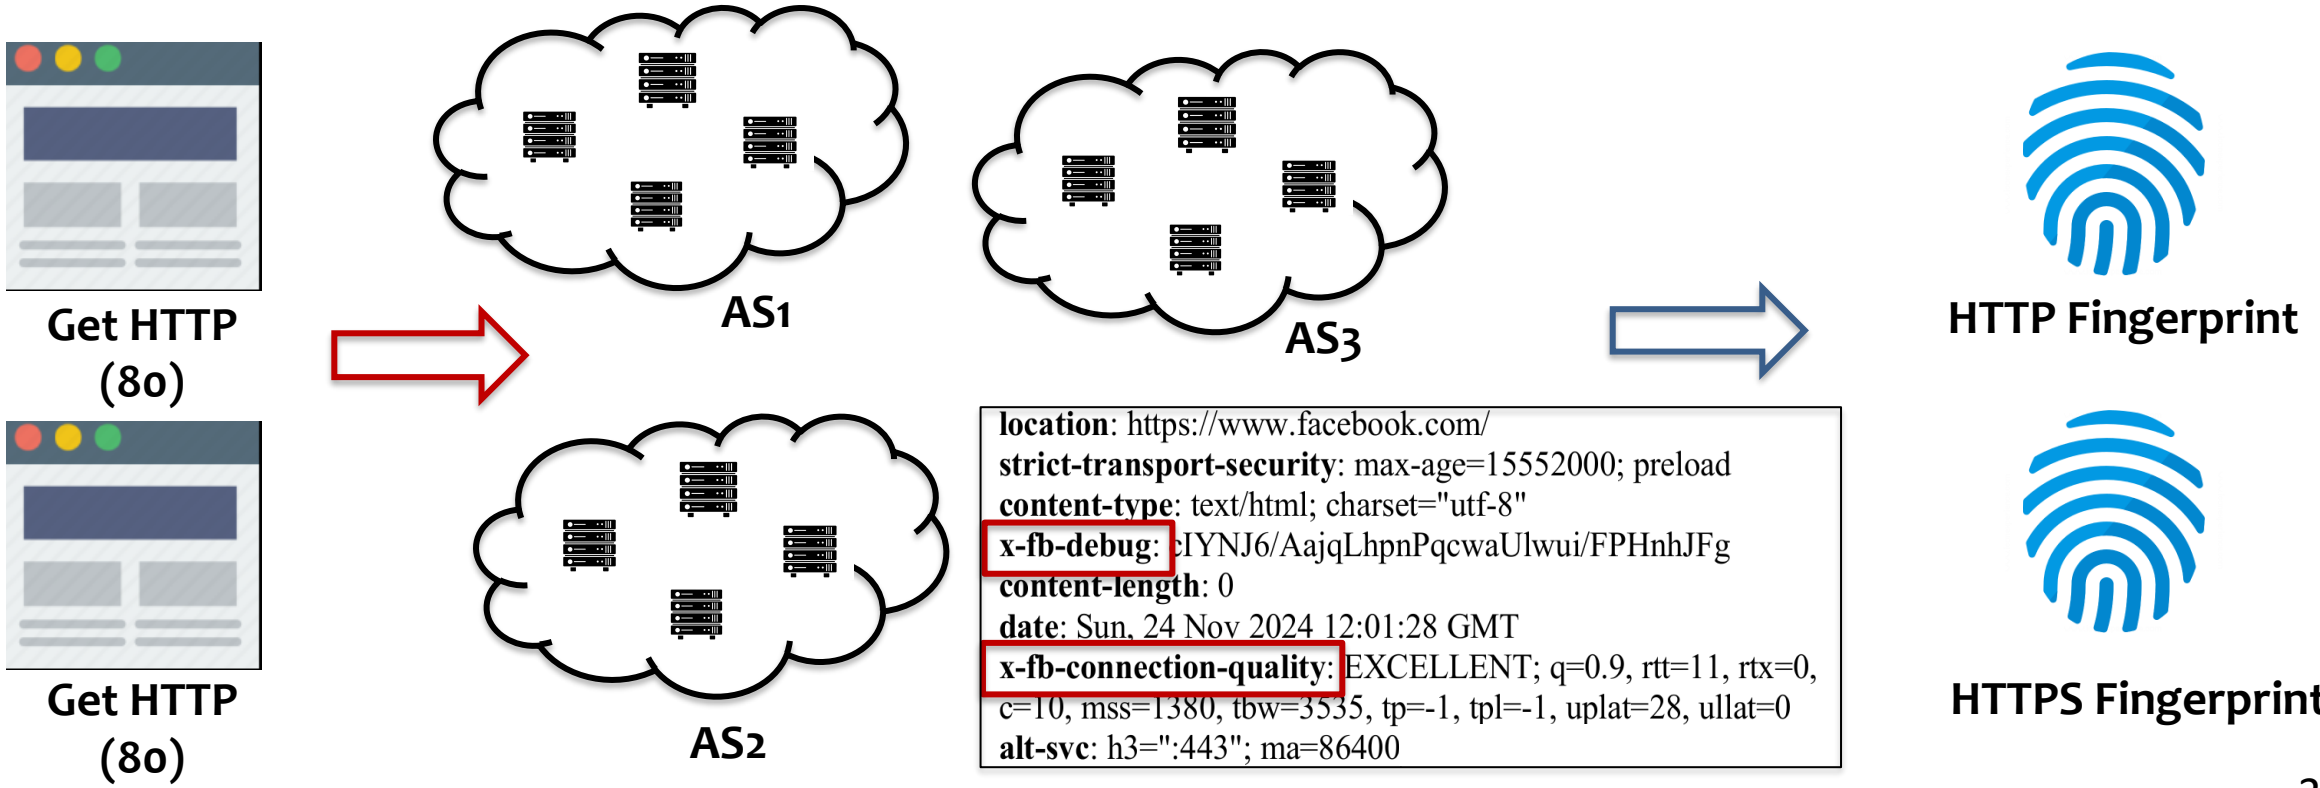

Methodology

- ***Off-net confirmation***
  - Step 4: Confirm Candidates Using HTTP(S)
    - match HG hdr. FPs -> classify as off-nets

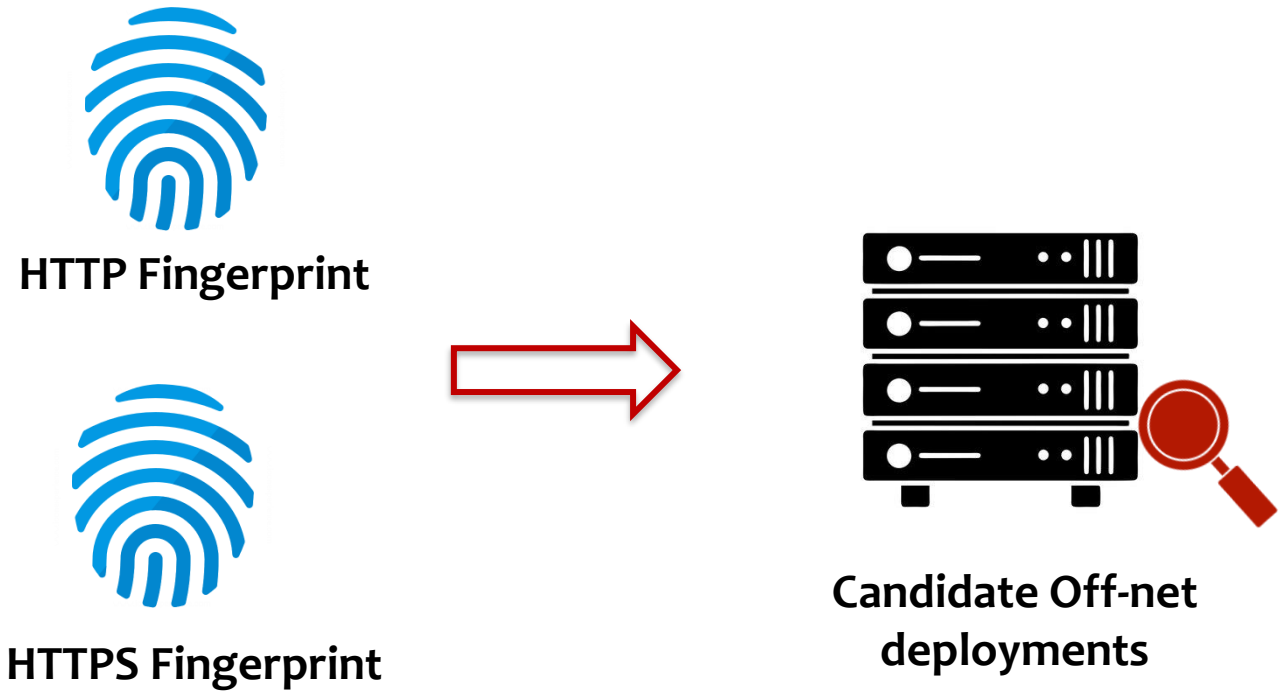

Methodology

- ***Off-net confirmation***
  - Step 4: Confirm Candidates Using HTTP(S)
    - match HG hdr. FPs -> classify as off-nets

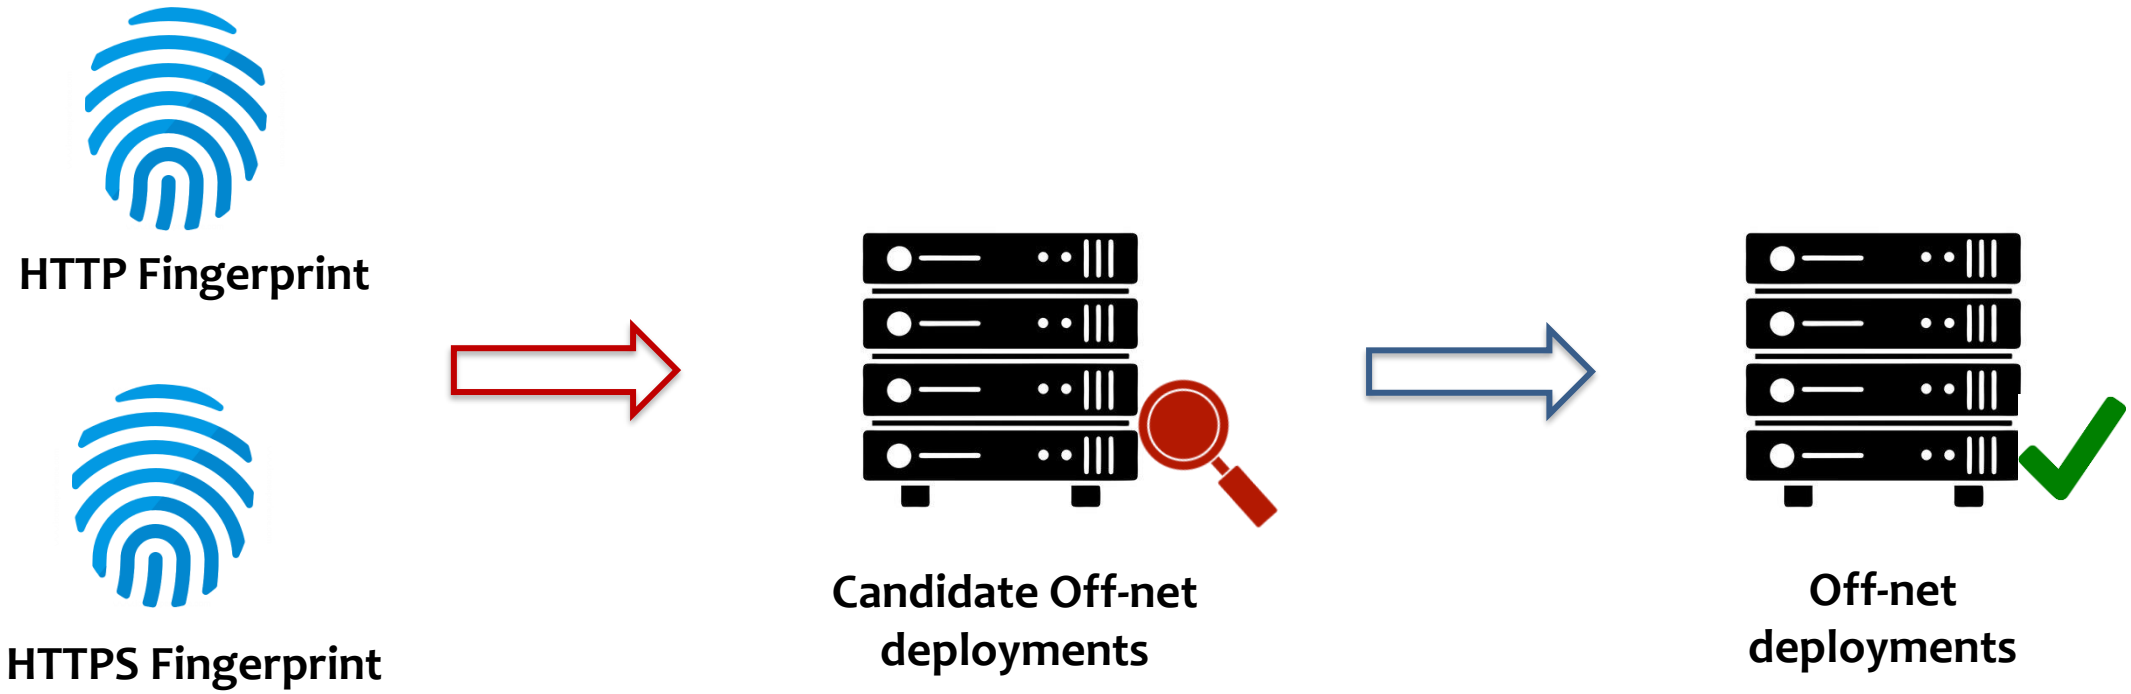

Limitations

- Full IPv6 address space infeasible
  - coverage lower bound

Limitations

- Full IPv6 address space infeasible
  - coverage lower bound
- Anycast deployments

Contributions

UNIVERSITÄT DES SAARLANDES

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k networks

Contributions

UNIVERSITÄT DES SAARLANDES

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k networks
- Investigate regional deployment
  - ***aggressive*** deployment in South America

Contributions

UNIVERSITÄT DES SAARLANDES

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k networks
- Investigate regional deployment
  - ***aggressive*** deployment in South America
- Find off-nets to be at ***risk***
- Find IPv6 ***latency*** to be ***at par*** to IPv4

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k networks
- Investigate regional deployment
  - ***aggressive*** deployment in South America
- Find off-nets to be at ***risk***
- Find IPv6 ***latency*** to be ***at par*** to IPv4

Results-Regional Deployment

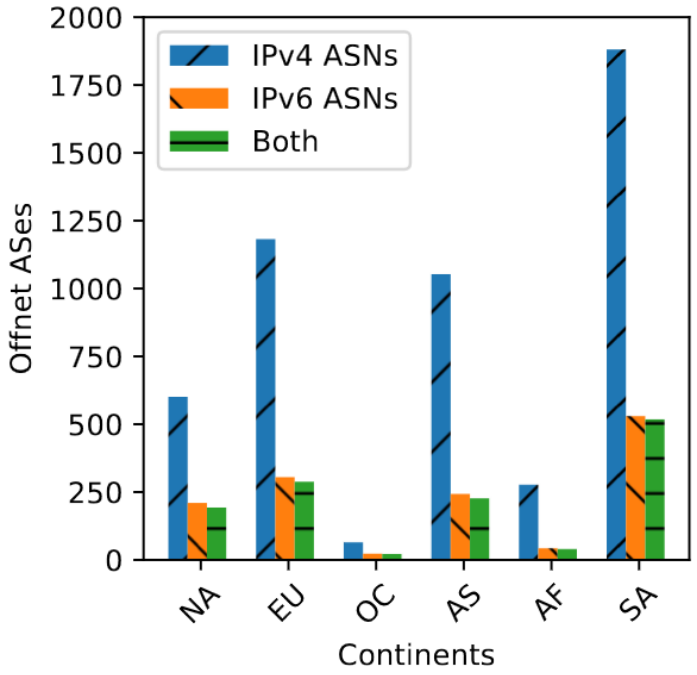

(a) Google.

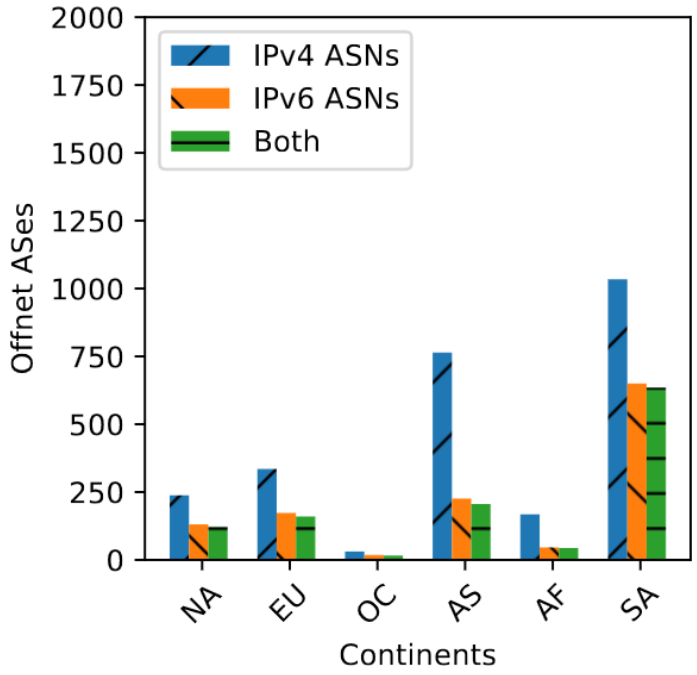

(b) Meta.

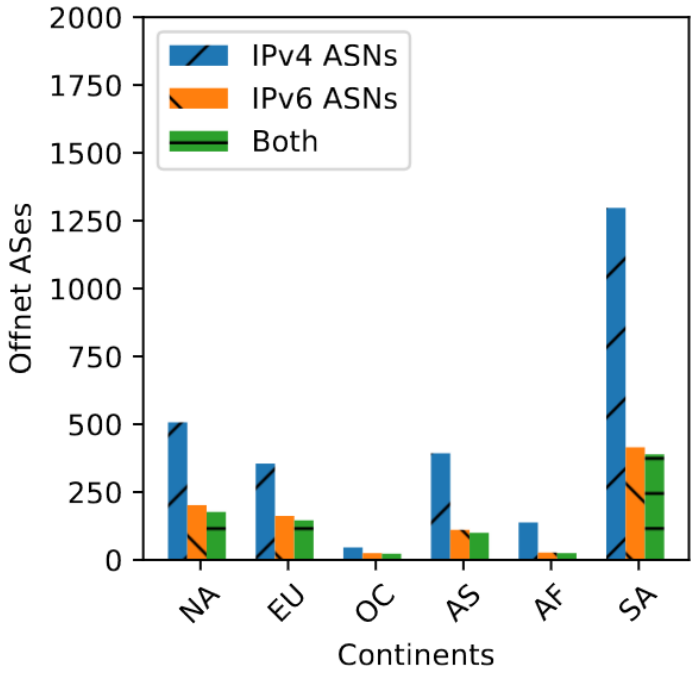

(c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Regional Deployment

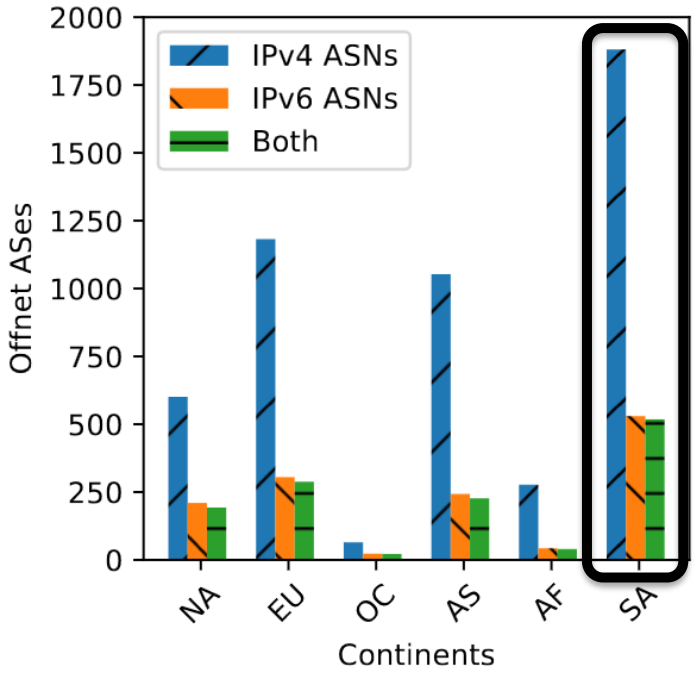

(a) Google.

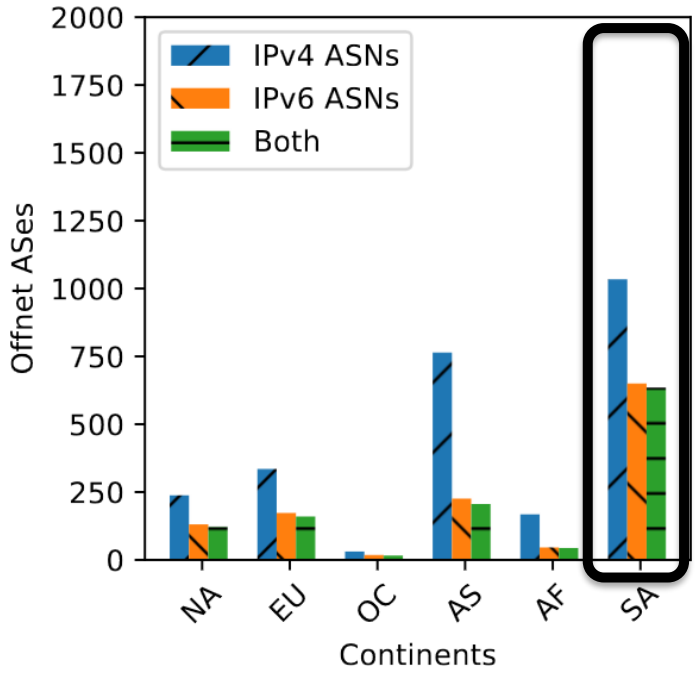

(b) Meta.

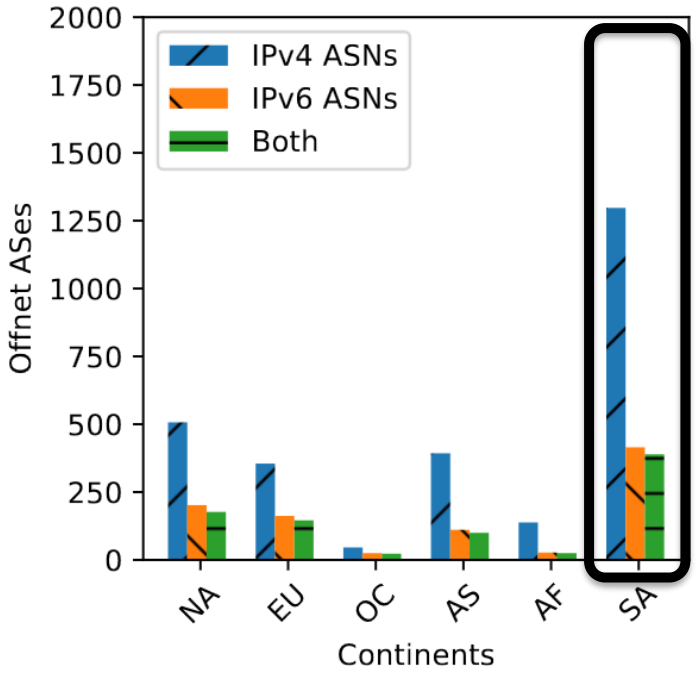

(c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Regional Deployment

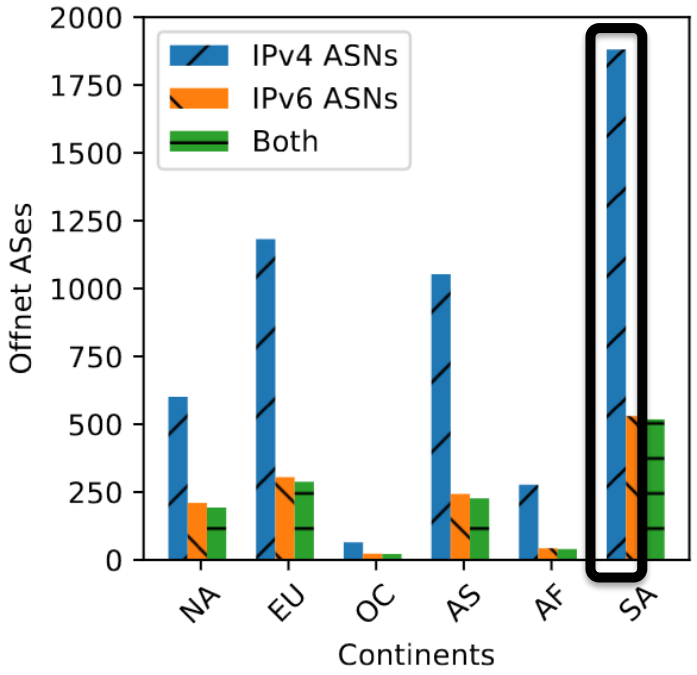

(a) Google.

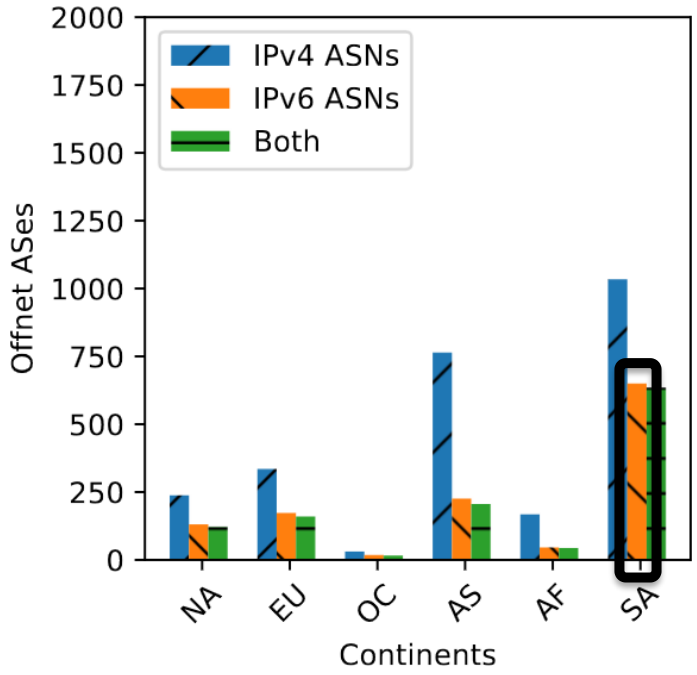

(b) Meta.

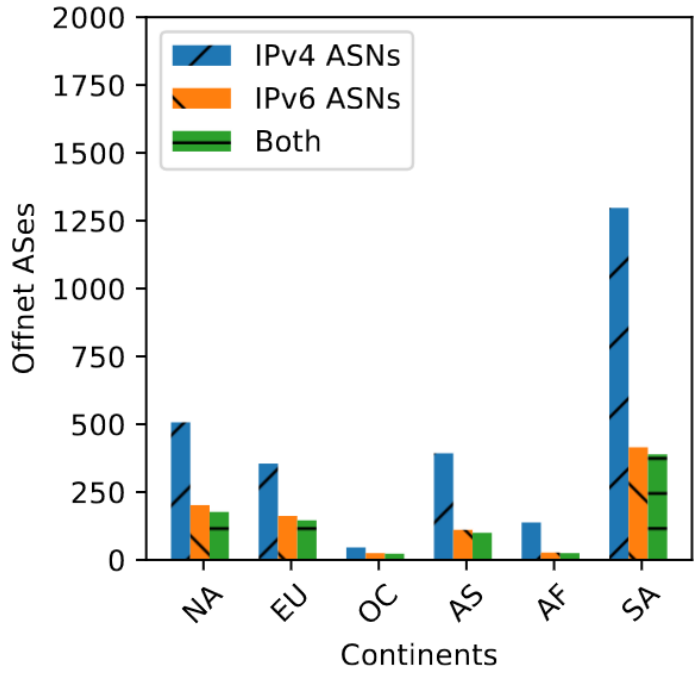

(c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Regional Deployment

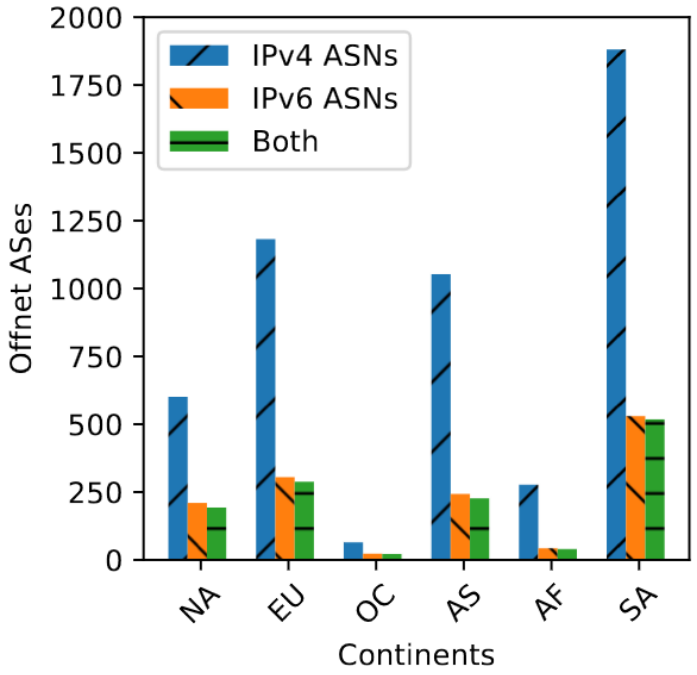

(a) Google.

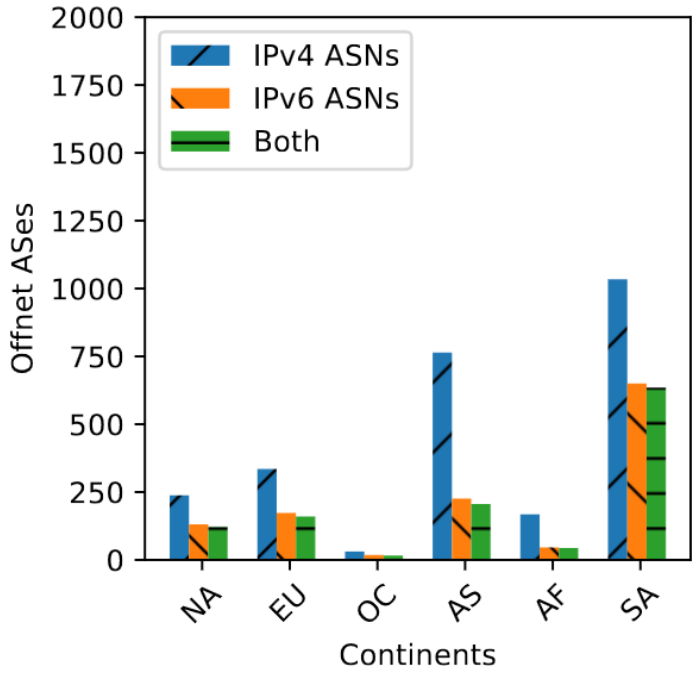

(b) Meta.

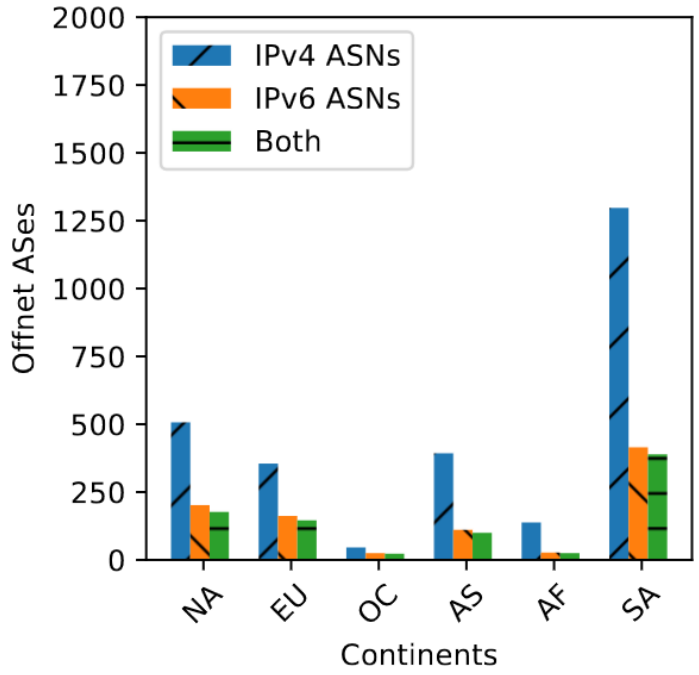

(c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google: 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Regional Deployment

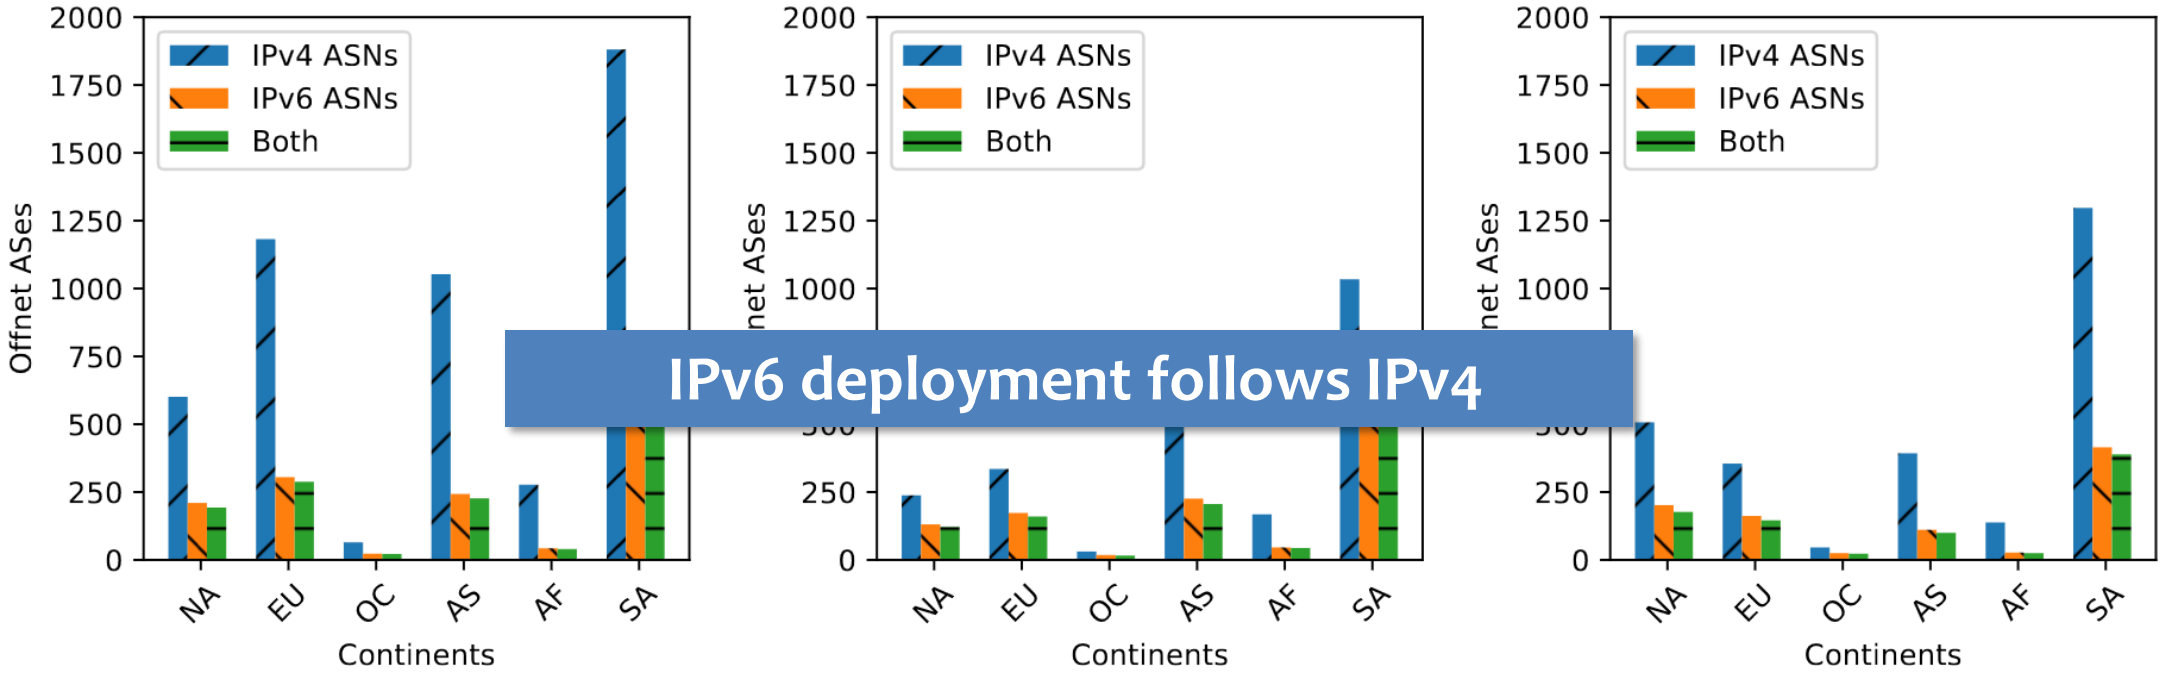

(a) Google. (b) Meta. (c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- ~93% of IPv6 off-nets where IPv4 off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Network Types Hosting Off-nets

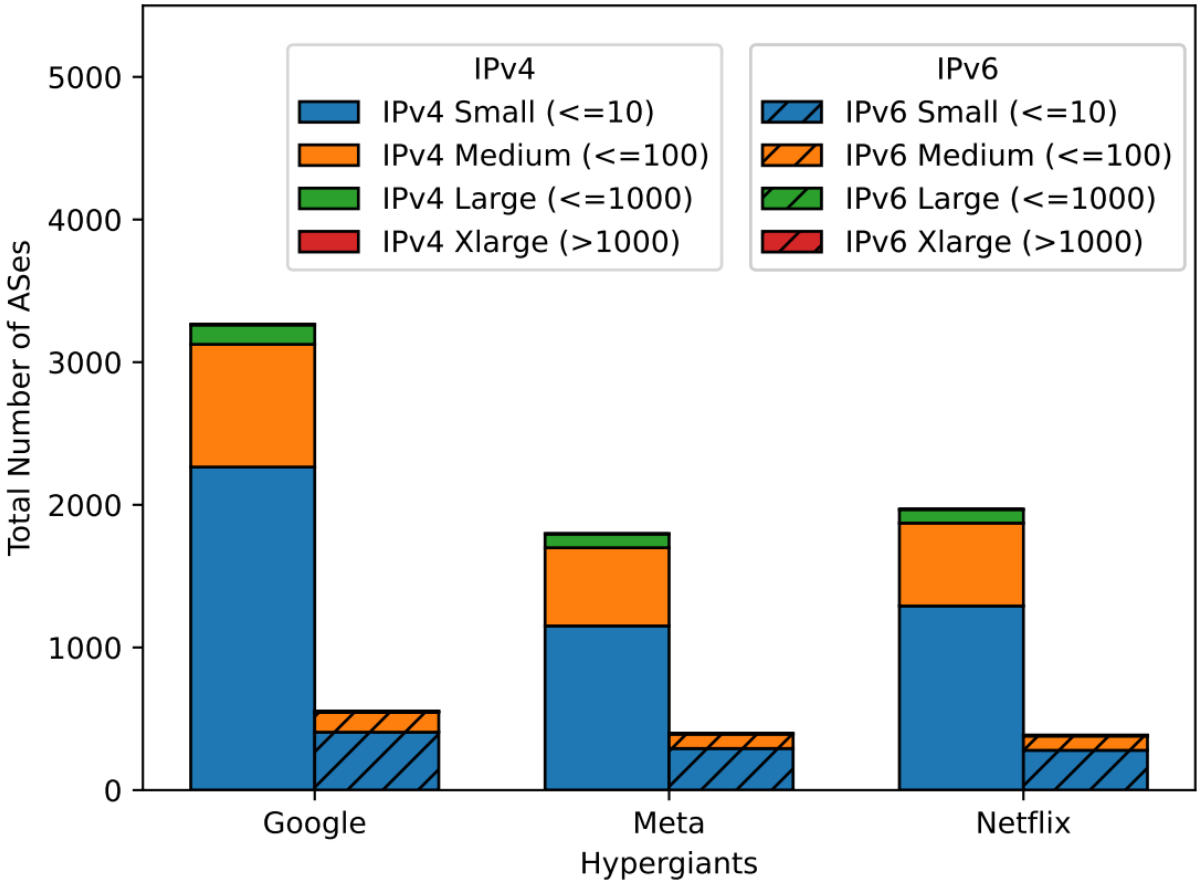

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Network Types Hosting Off-nets

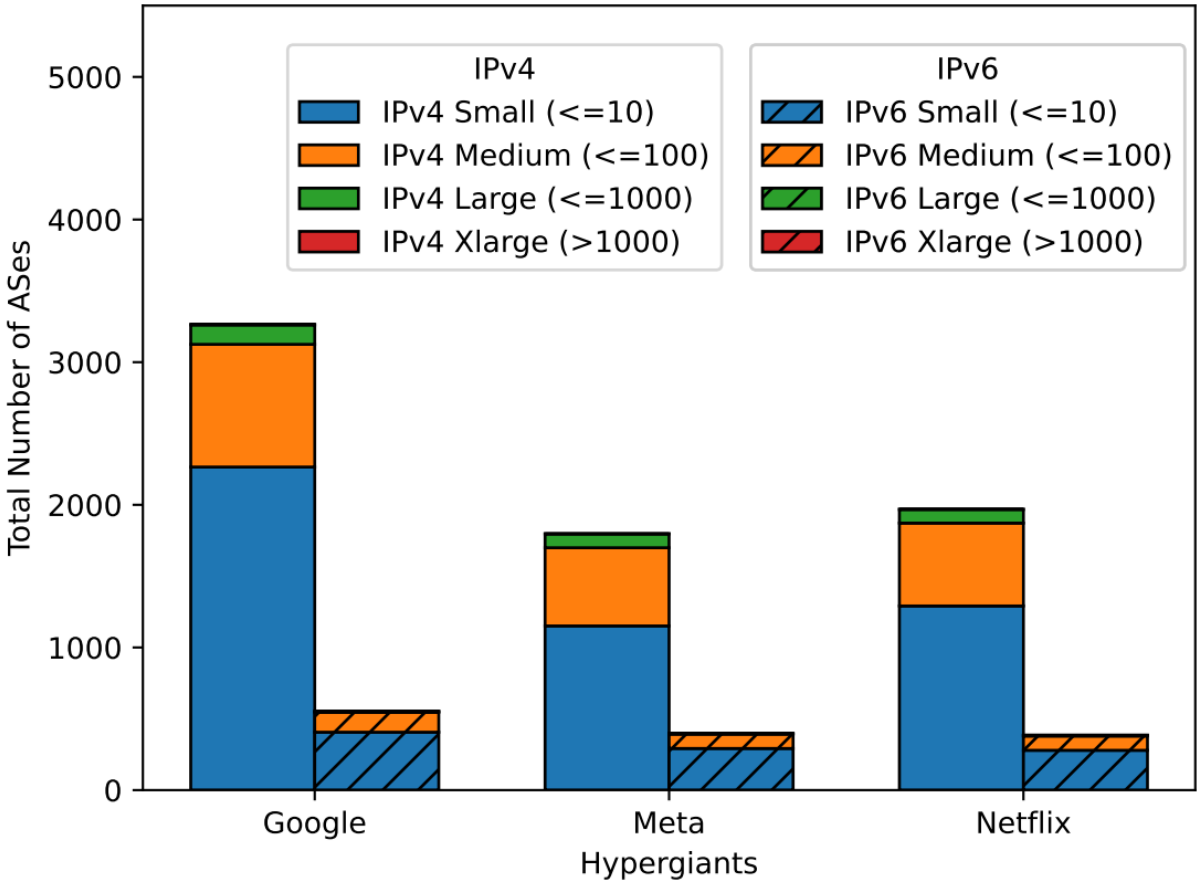

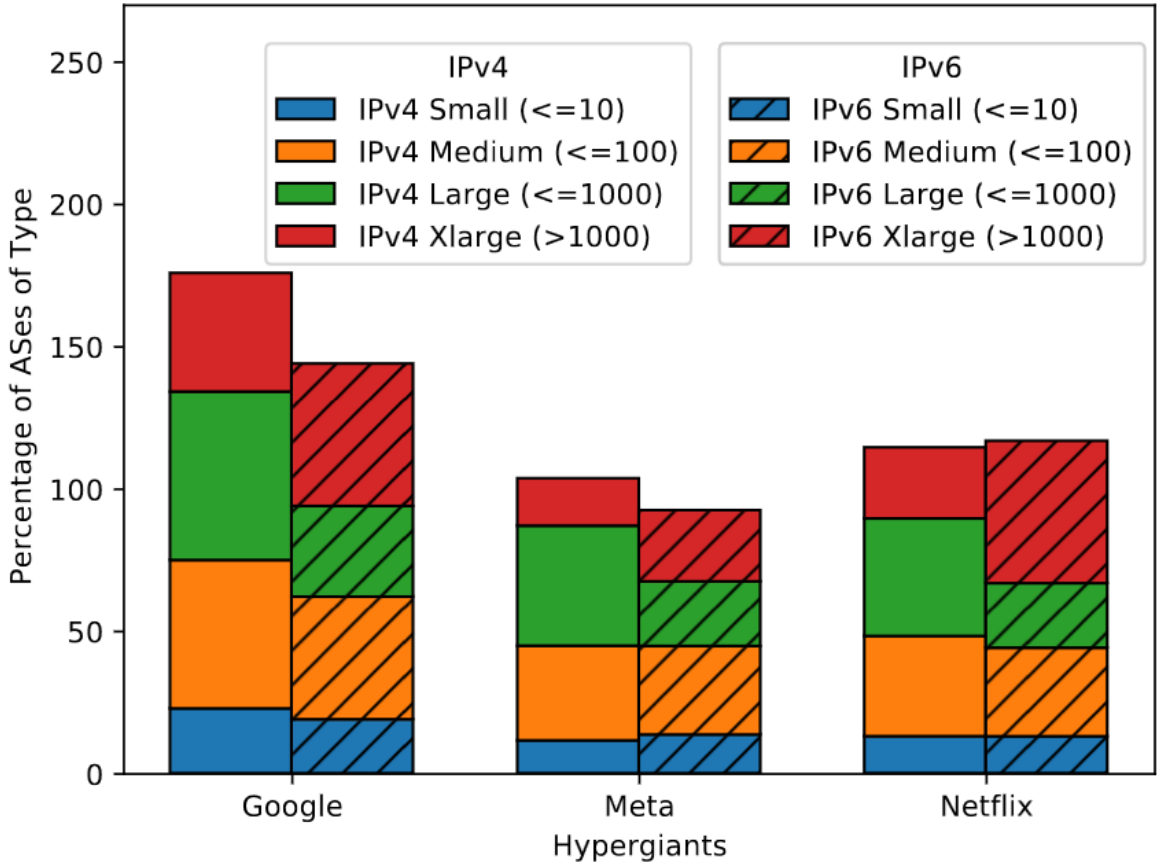

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Network Types Hosting Off-nets

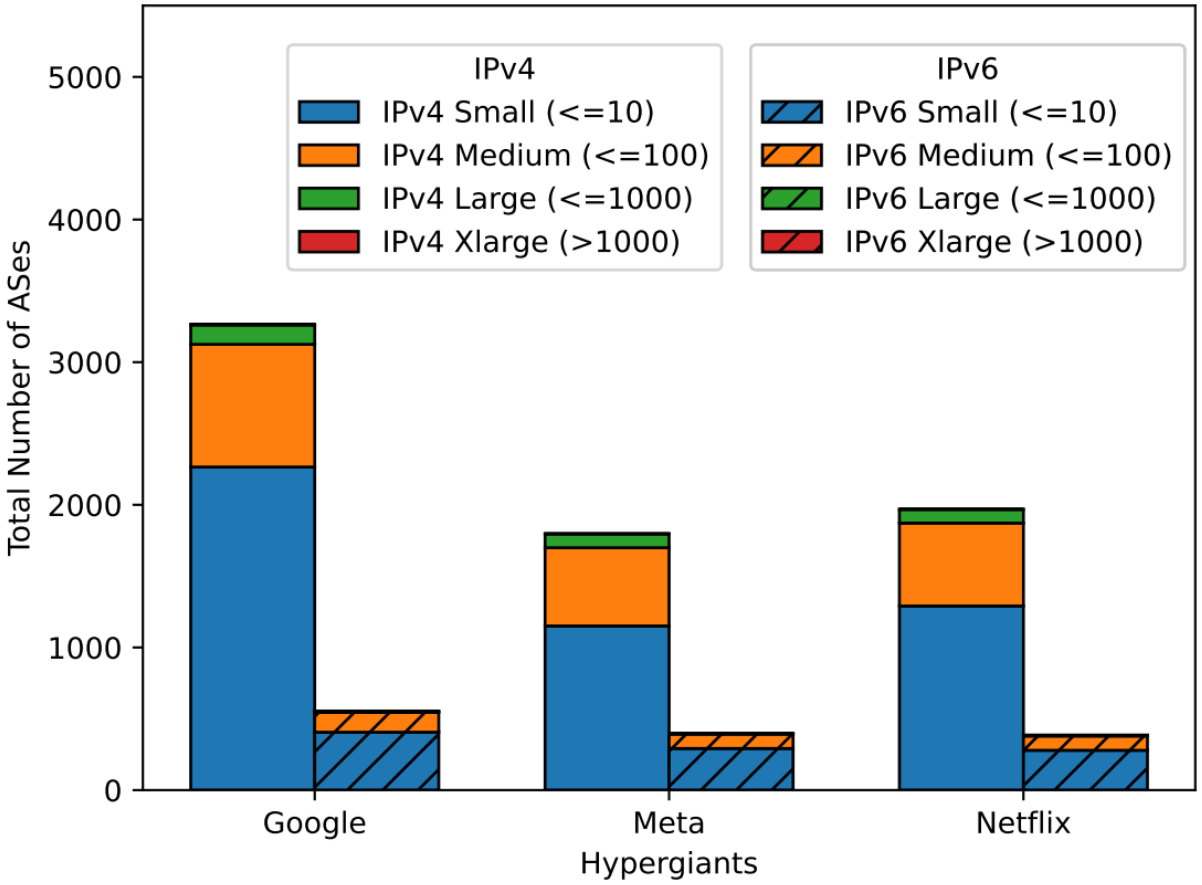

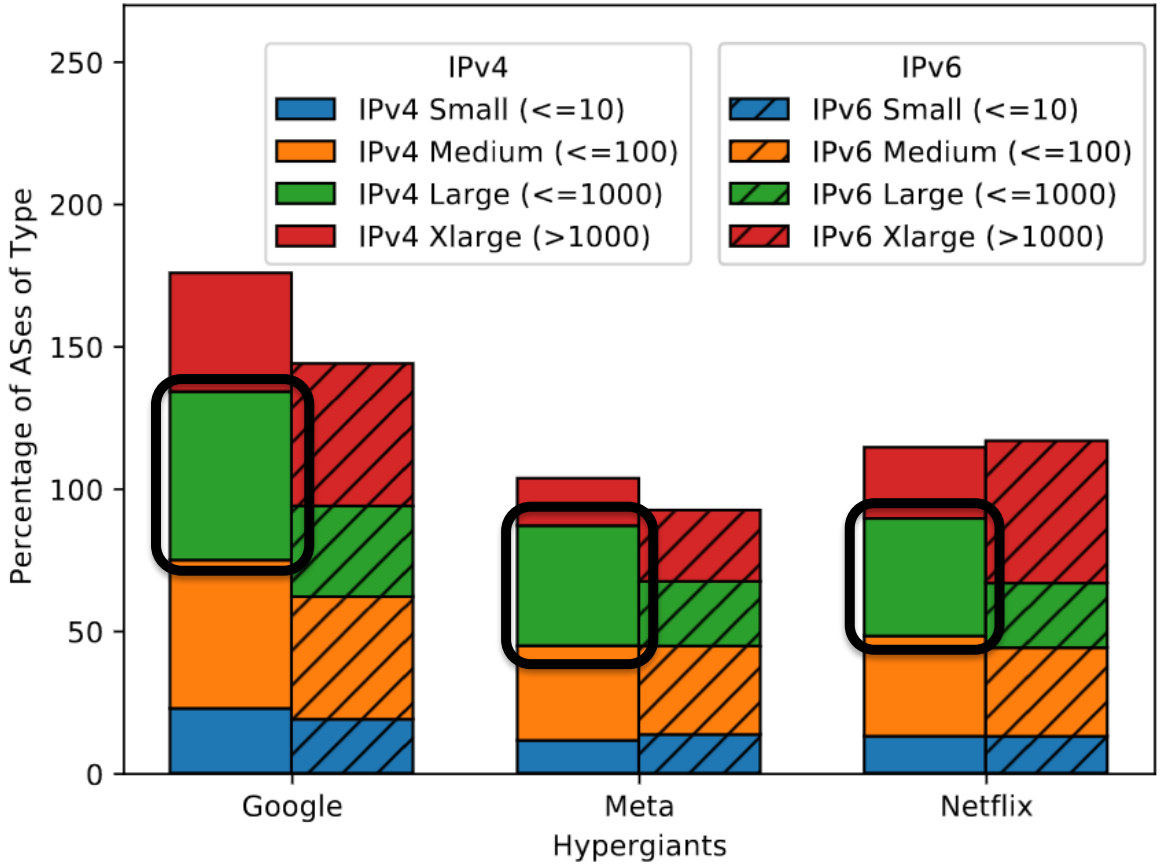

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Network Types Hosting Off-nets

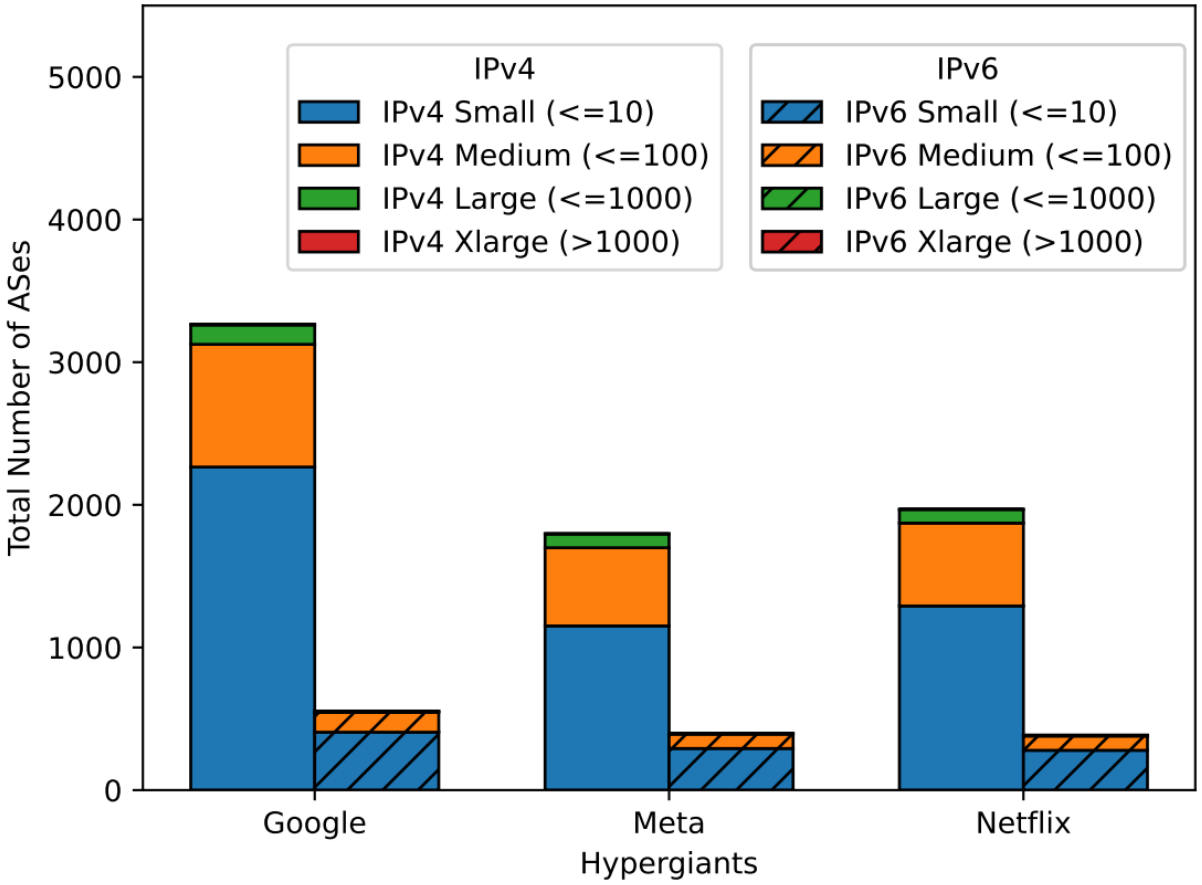

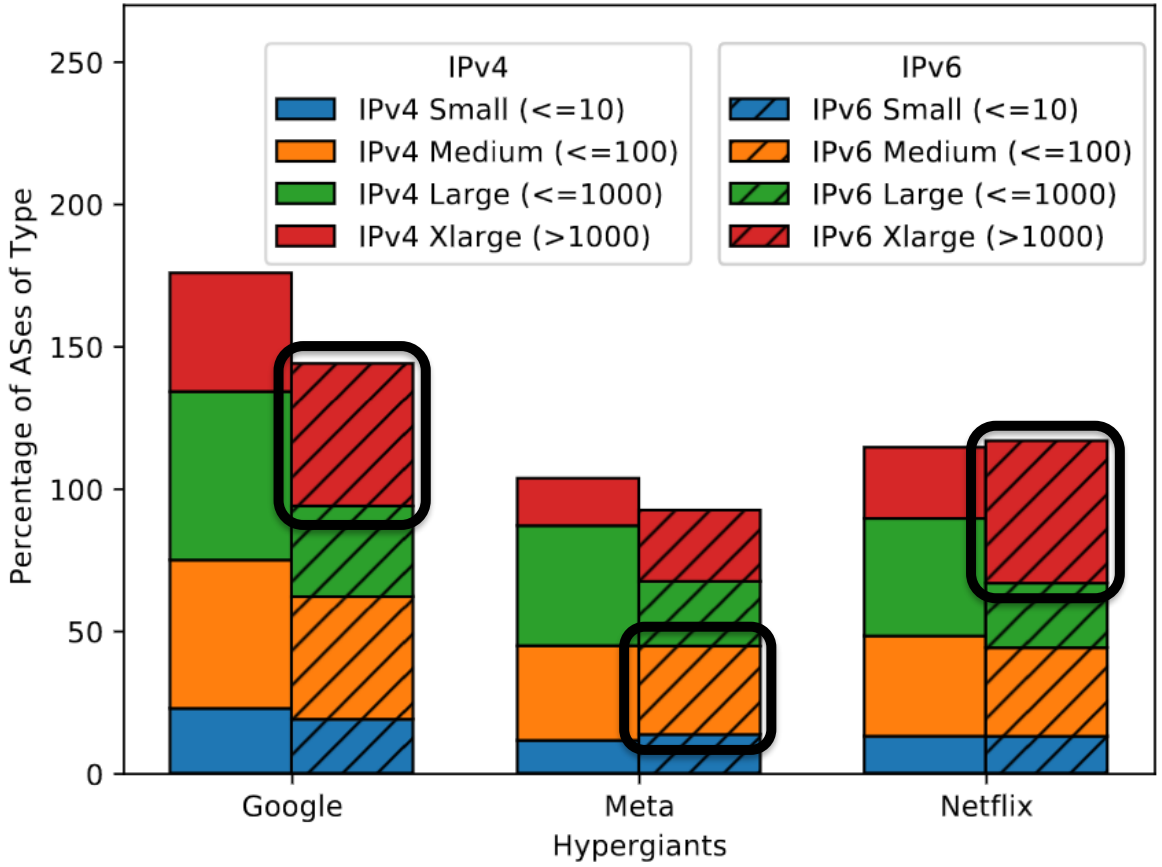

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Network Types Hosting Off-nets

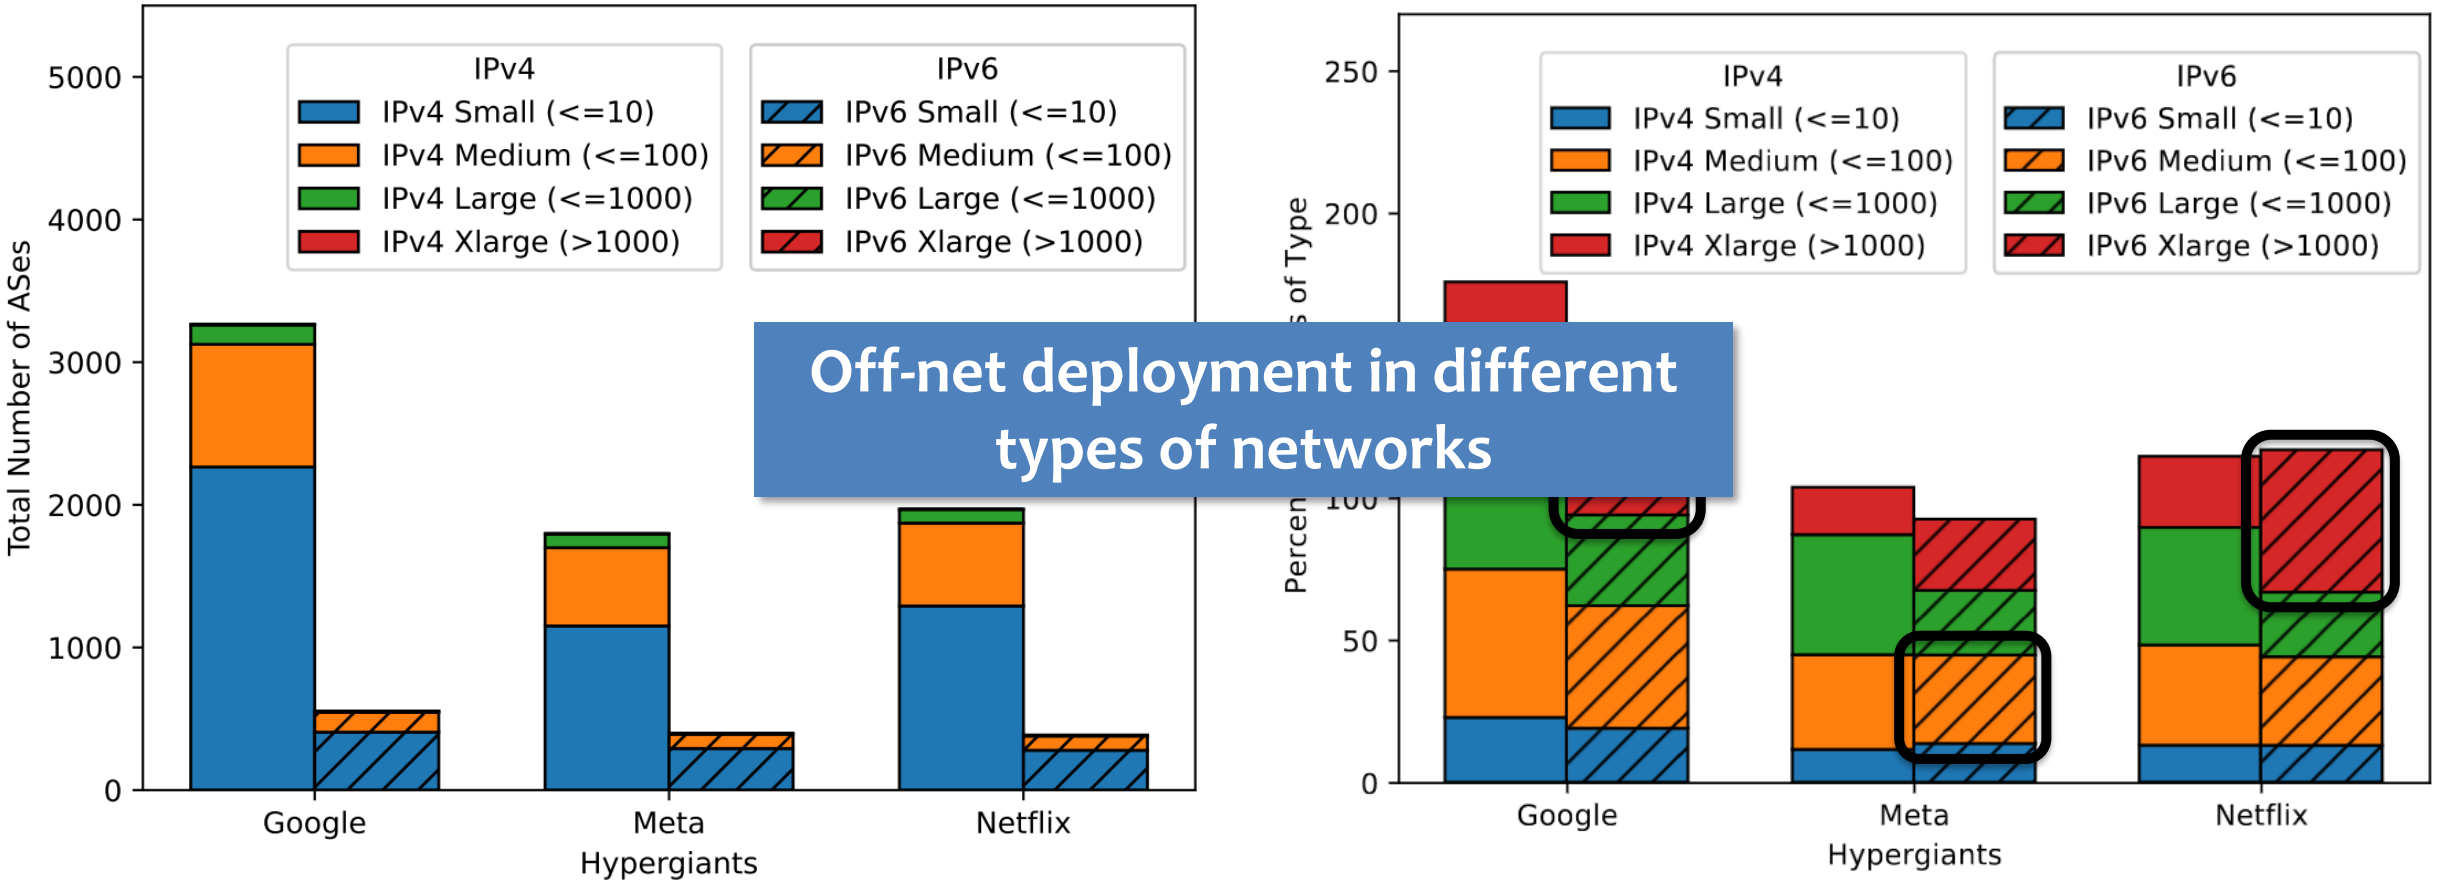

Off-net deployment in different types of networks

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

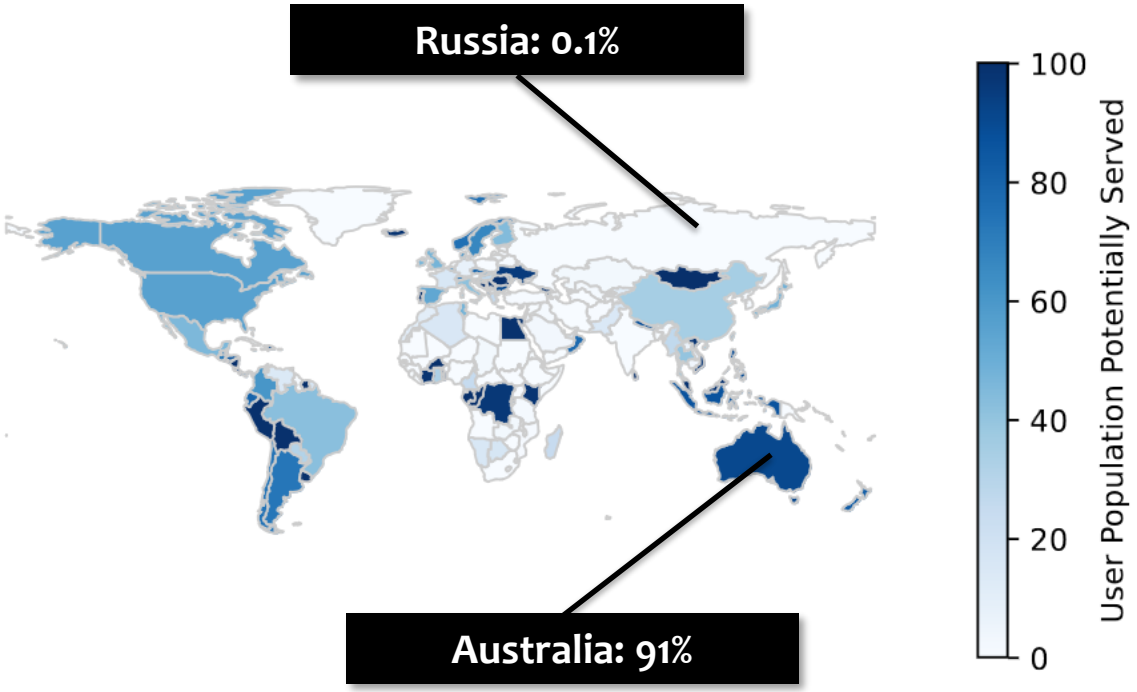

(a) IPv6.

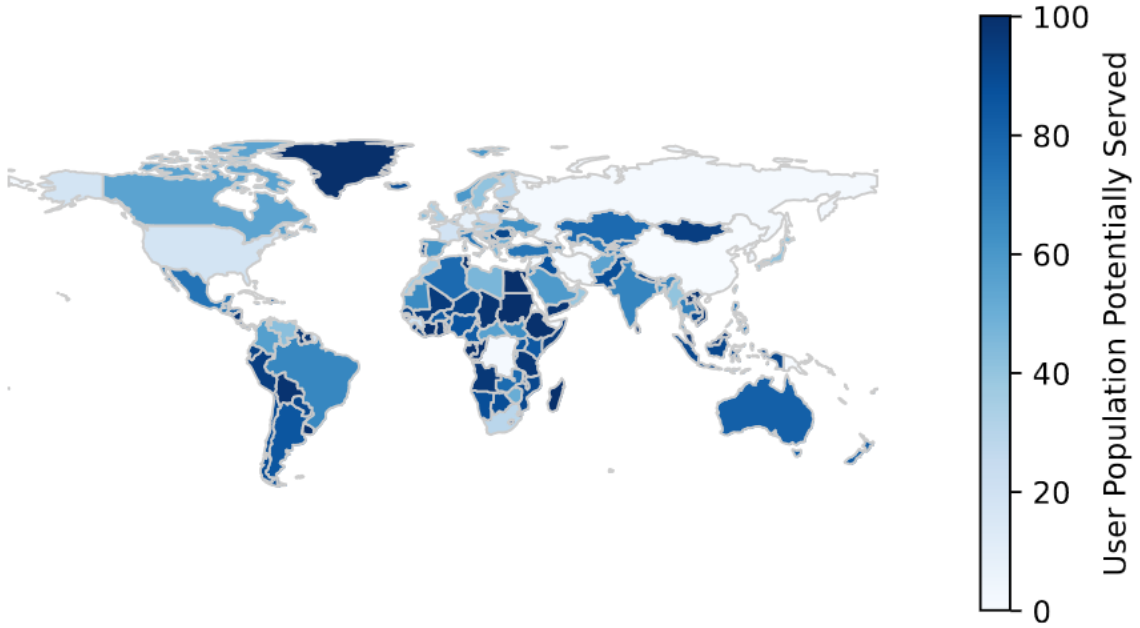

(b) IPv4.

Results-Access to Off-nets (Meta)

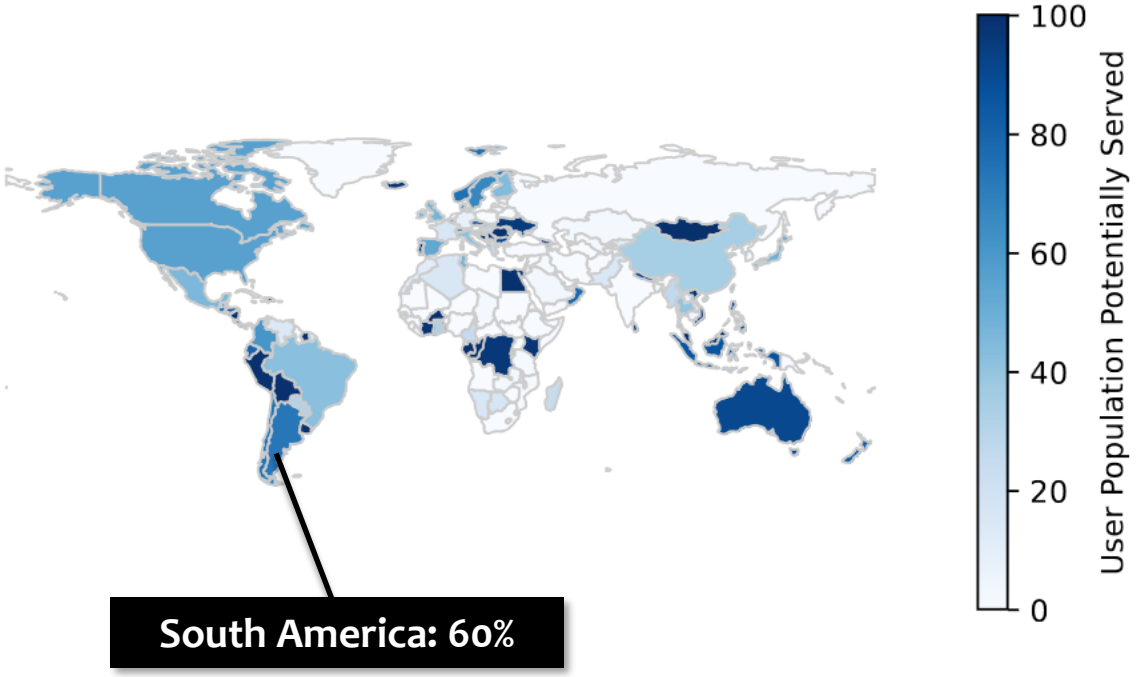

(a) IPv6.

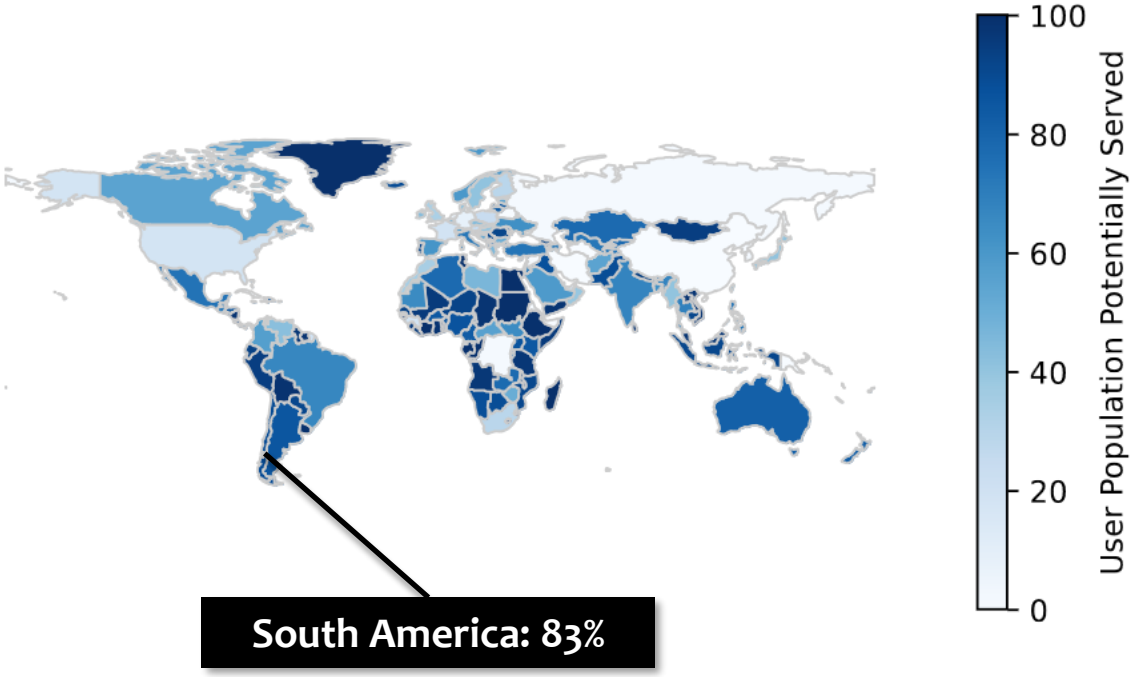

(b) IPv4.

Results-Access to Off-nets (Meta)

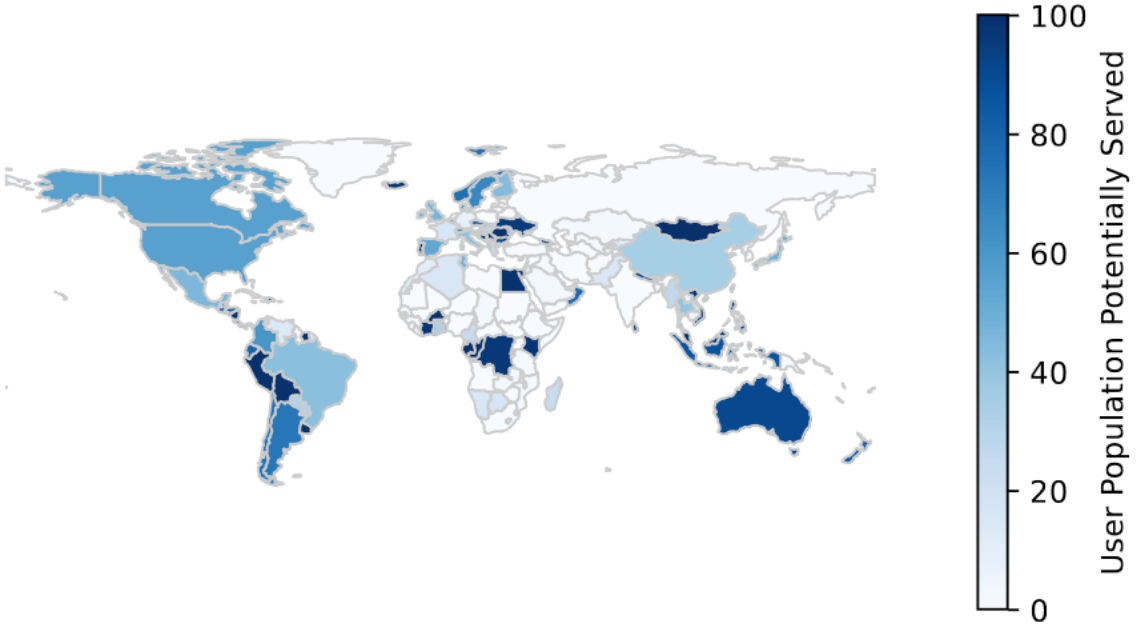

(a) IPv6.

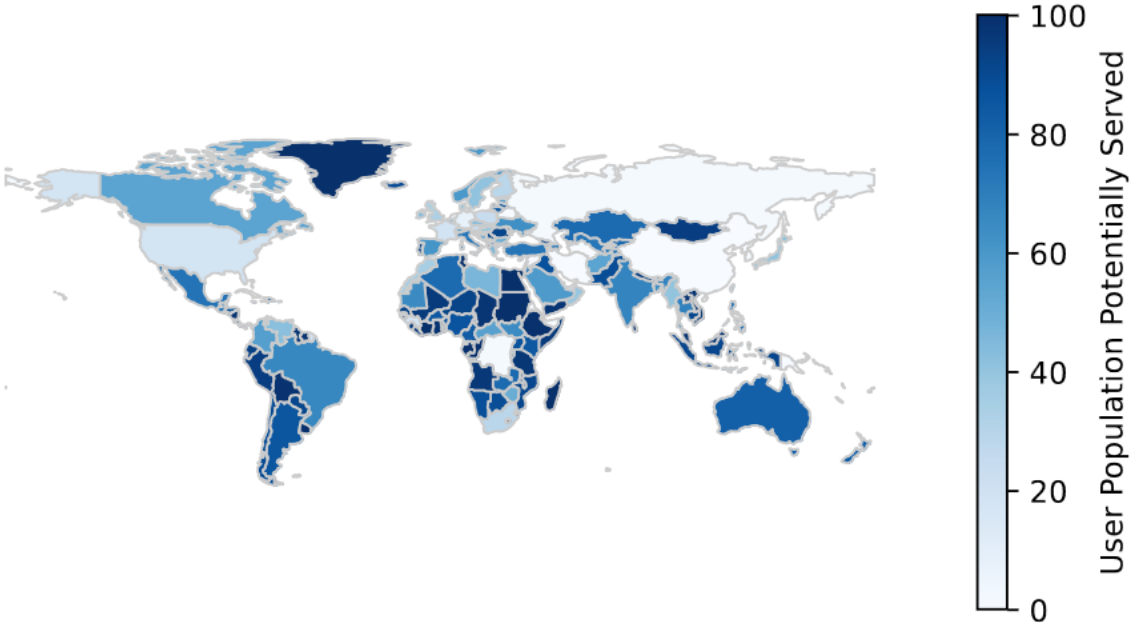

(b) IPv4.

Results-Access to Off-nets (Meta)

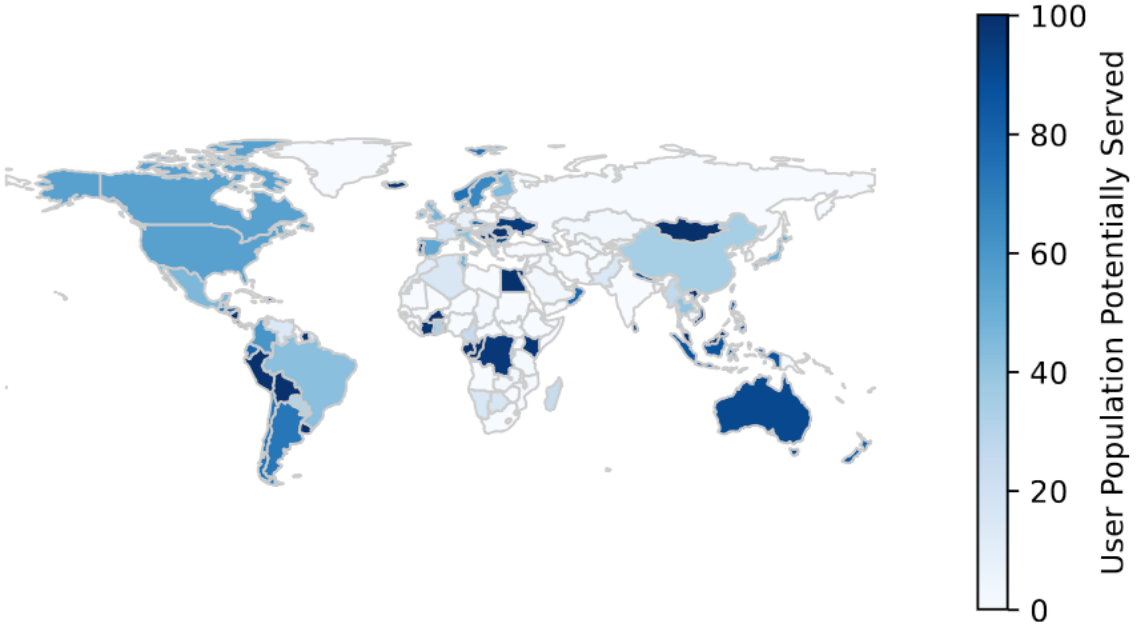

(a) IPv6.

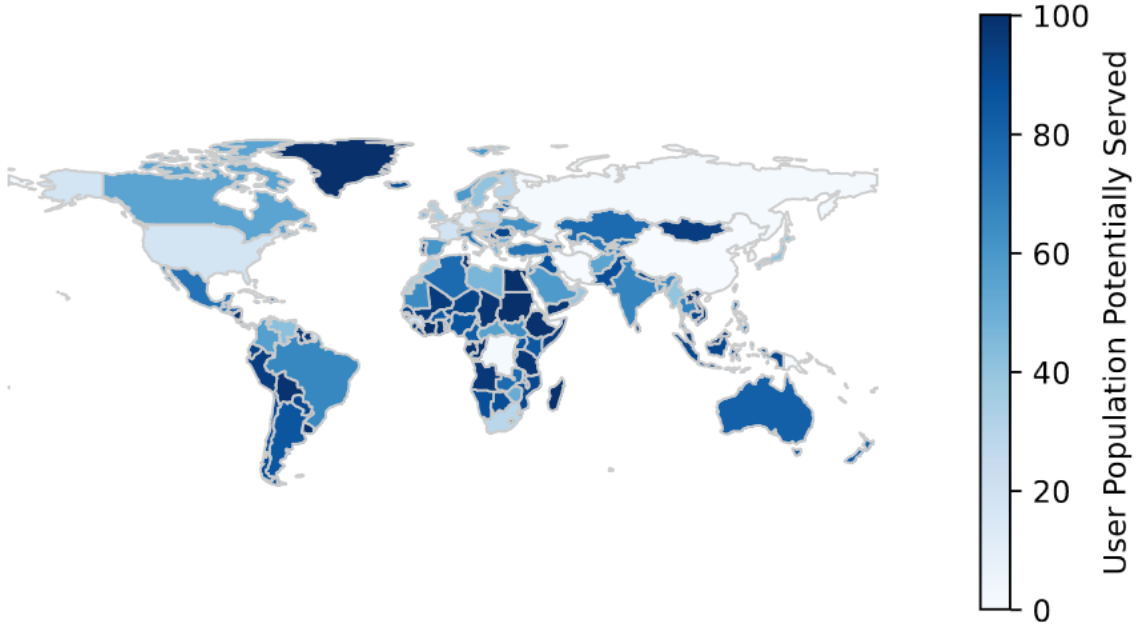

(b) IPv4.

- Africa
  - ~4% traffic to facebook over IPv6 [1]

[1] https://www.facebook.com/ipv6/?tab=ipv6_country

Results-Access to Off-nets (Meta)

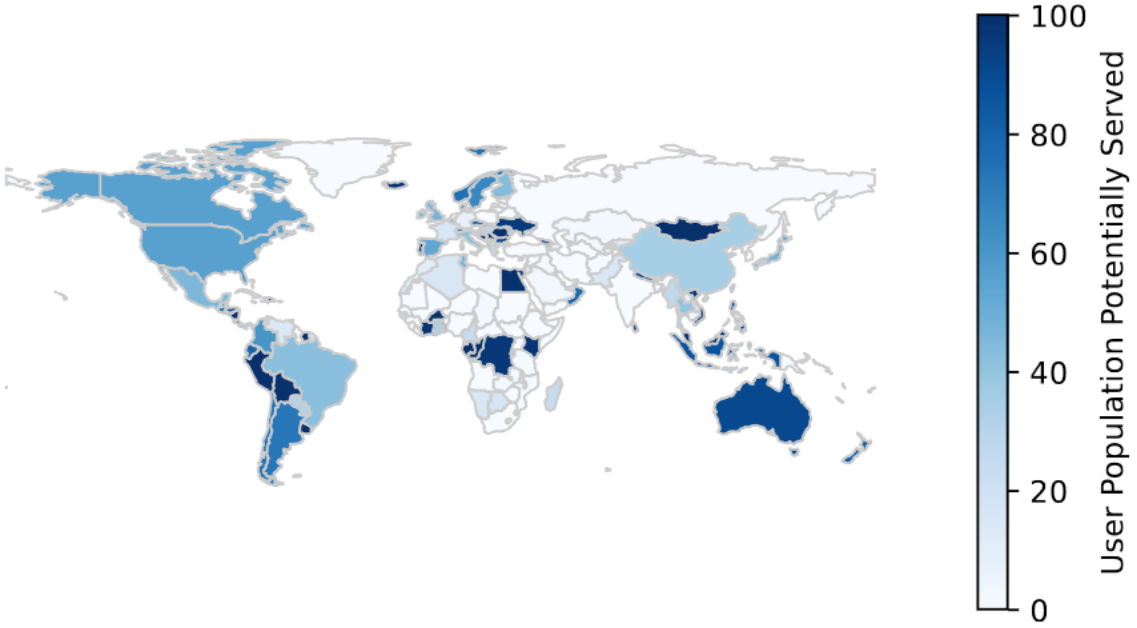

(a) IPv6.

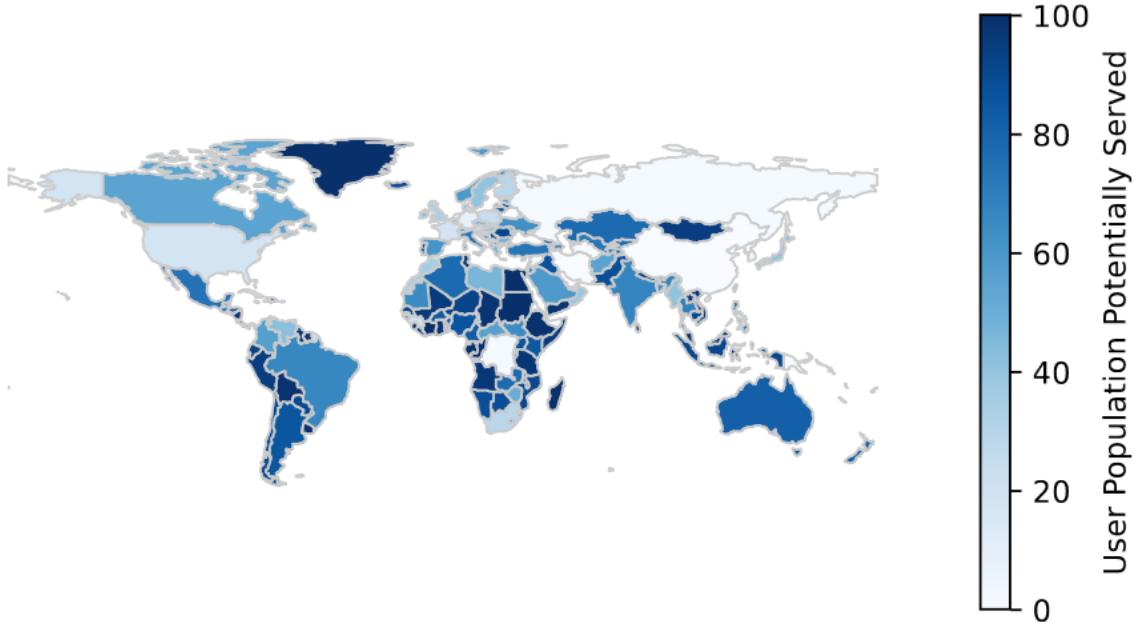

(b) IPv4.

Results-Access to Off-nets (Meta)

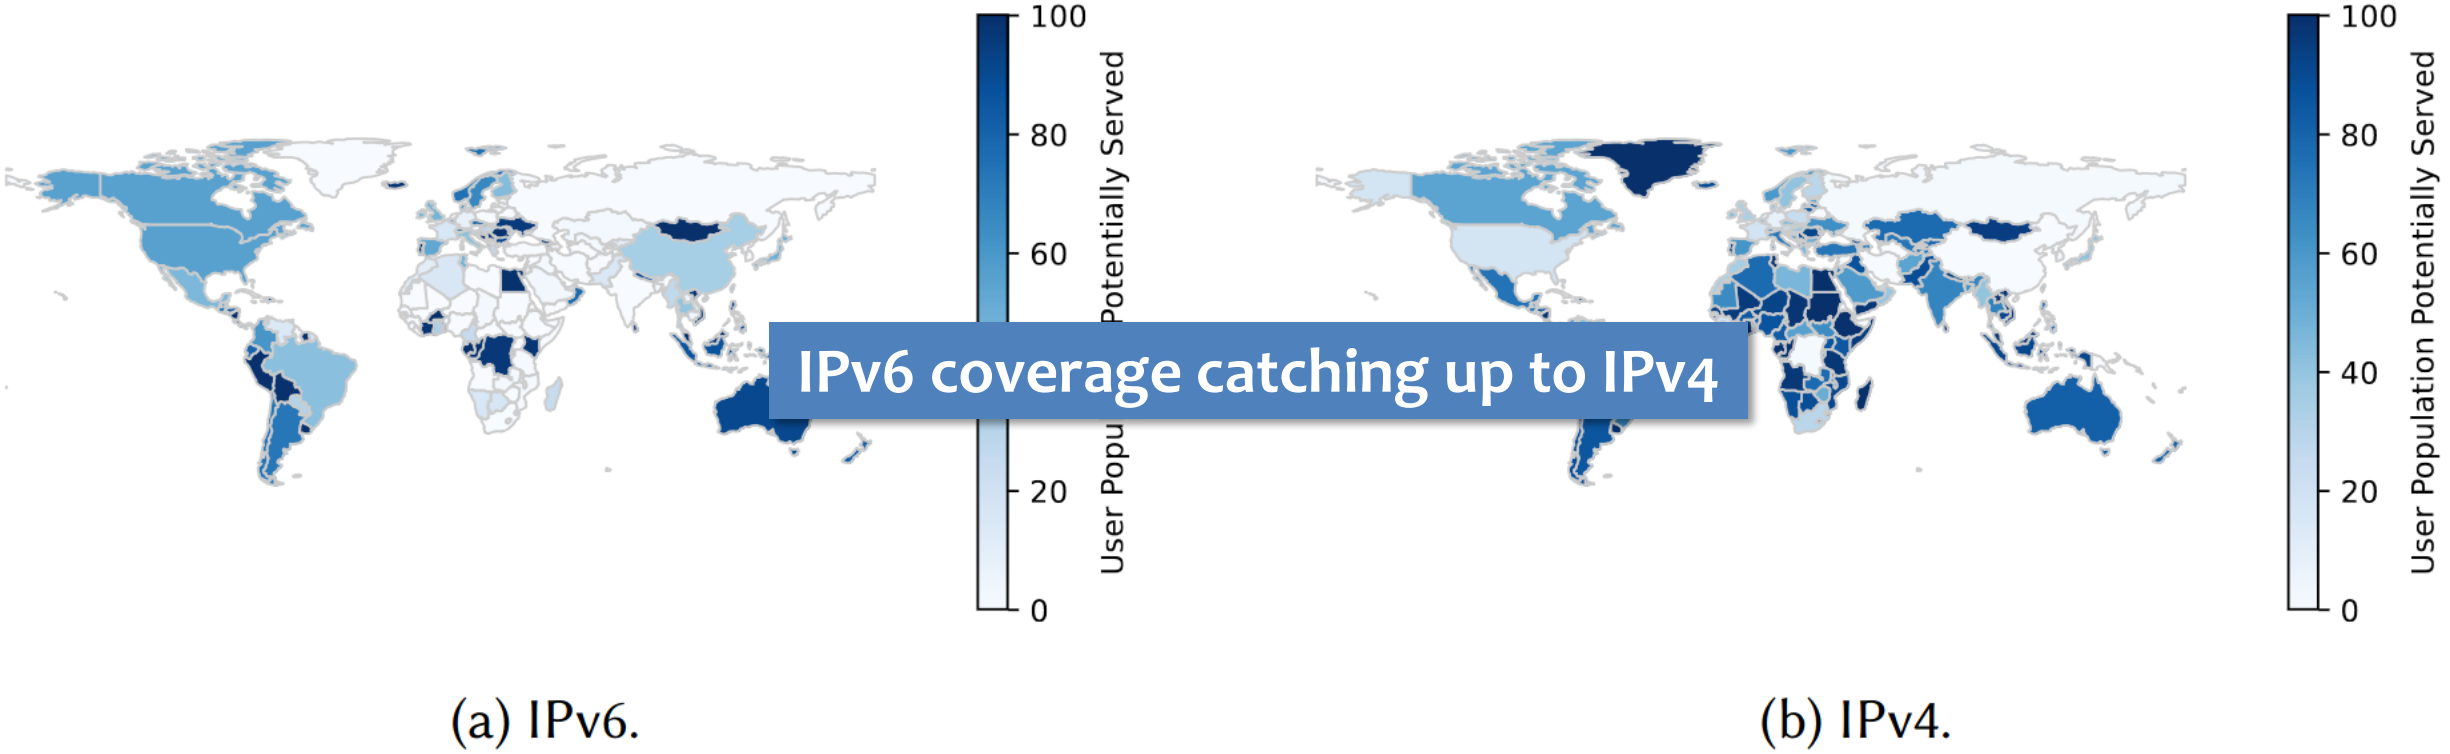

(a) IPv6.

(b) IPv4.

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k networks
- Investigate regional depl.
  - ***aggressive*** depl. in South America
- Find off-nets to be at ***risk***
- Find IPv6 ***latency*** to be ***at par*** to IPv4

Results-Regional Deployment

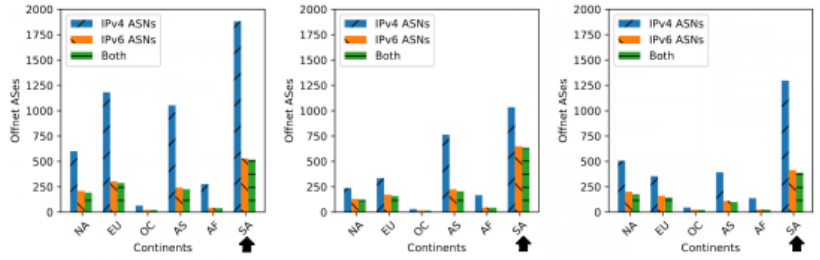

(a) Google. (b) Meta. (c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google : 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Network Types Hosting Off-nets

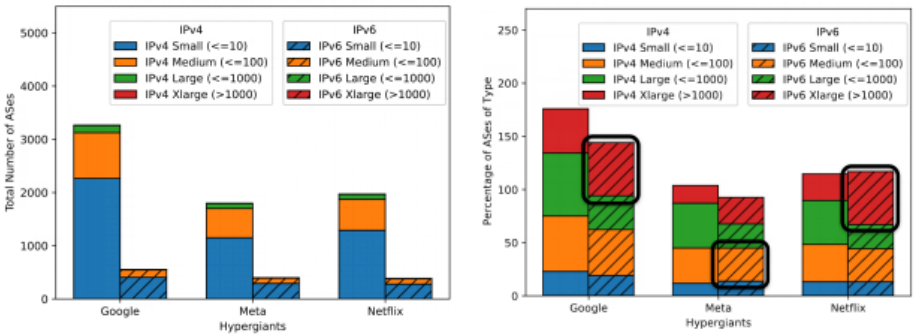

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

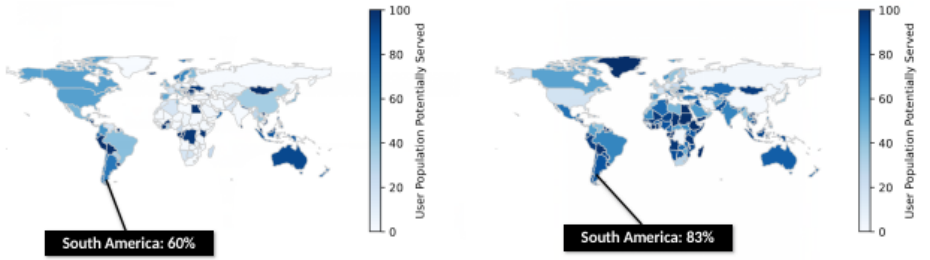

(a) IPv6. (b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k
- Invest
  - ***agg***
- Find o
- Find IPv6 ***latency*** to be ***at par*** to IPv4

First study into HGs' IPv6 off-nets

Results-Regional Deployment

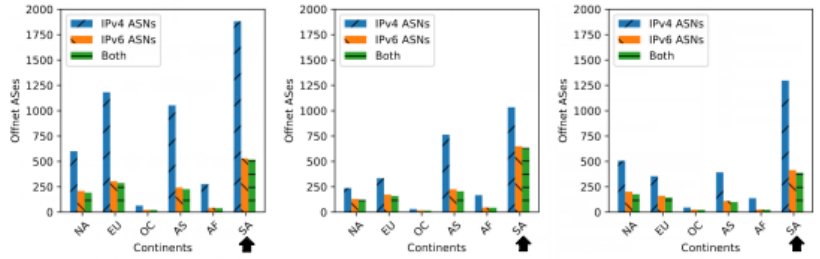

(a) Google. (b) Meta. (c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google : 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Network Types Hosting Off-nets

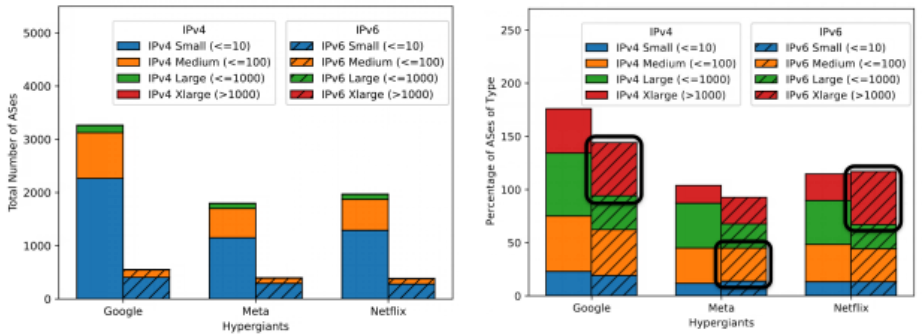

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

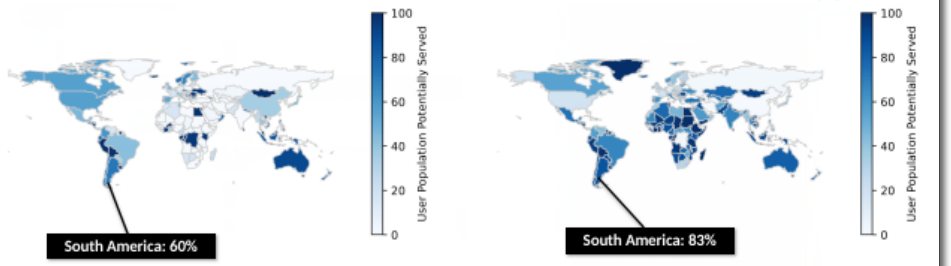

(a) IPv6. (b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k
- Invest
  - ***agg***
- Find o
- Find IPv6 ***latency*** to be ***at par*** to IPv4

First study into HGs' IPv6 off-nets

Results-Regional Deployment

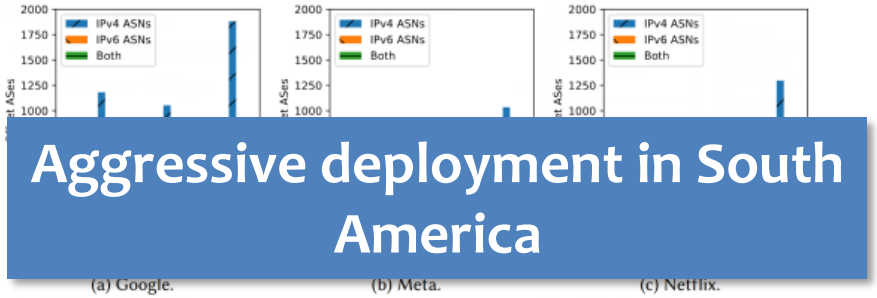

(a) Google. (b) Meta. (c) Netflix.

Aggressive deployment in South America

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google : 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Network Types Hosting Off-nets

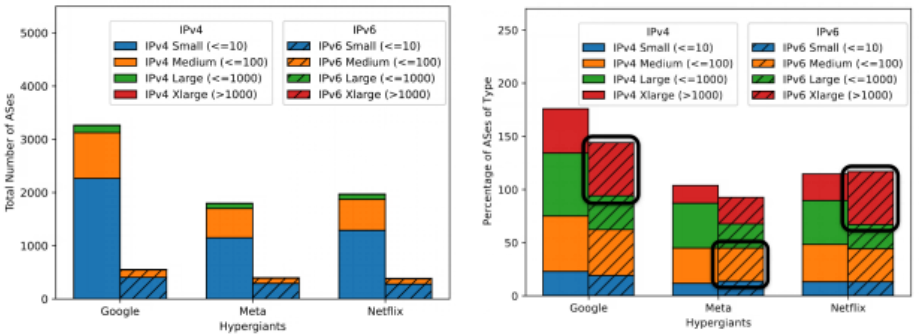

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

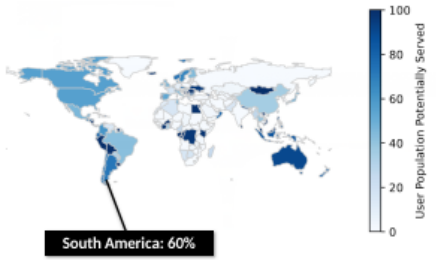

(a) IPv6.

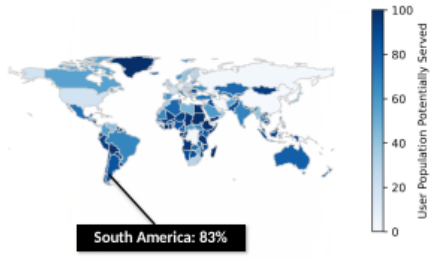

(b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k
- Invest
  - ***agg***
- Find o
- Find IPv6 ***latency*** to be ***at par*** to IPv4

First study into HGs' IPv6 off-nets

Results-Regional Deployment

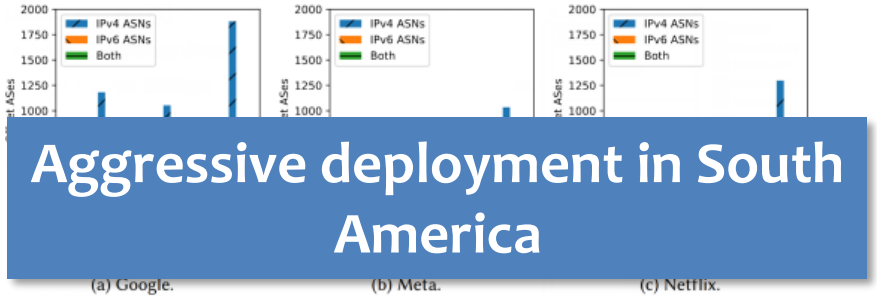

(a) Google. (b) Meta. (c) Netflix.

Aggressive deployment in South America

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google : 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Network Types Hosting Off-nets

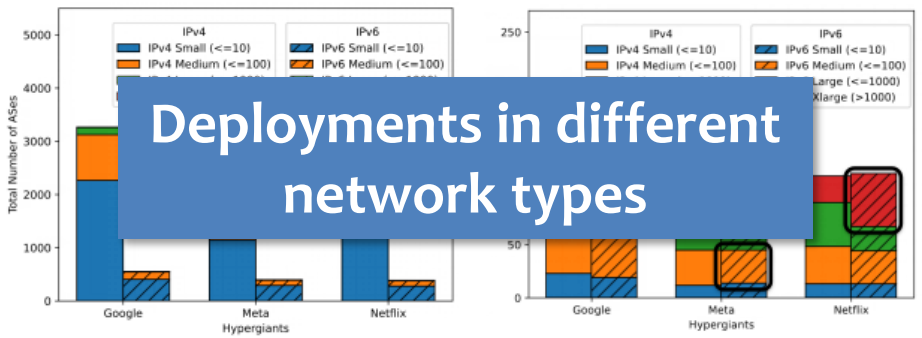

Deployments in different network types

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

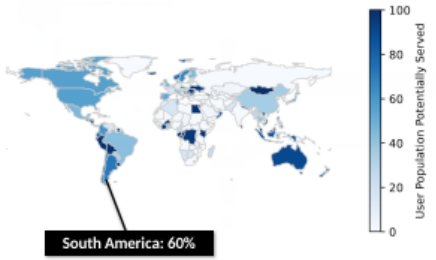

(a) IPv6.

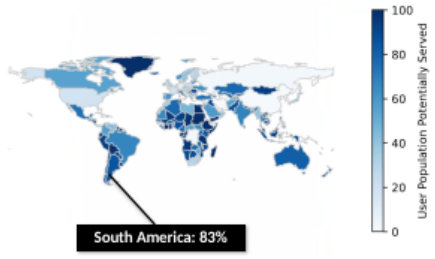

(b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k
- Invest
  - ***agg***
- Find o
- Find IPv6 ***latency*** to be ***at par*** to IPv4

First study into HGs' IPv6 off-nets

Results-Regional Deployment

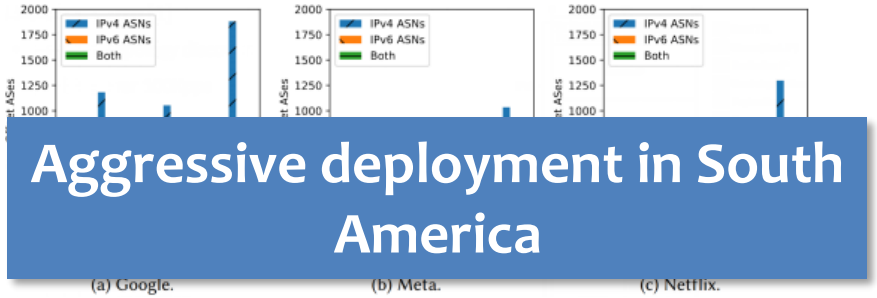

(a) Google. (b) Meta. (c) Netflix.

Aggressive deployment in South America

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google : 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Network Types Hosting Off-nets

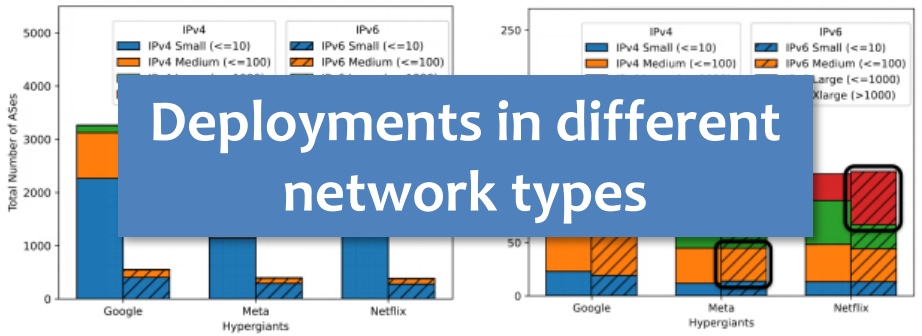

Deployments in different network types

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

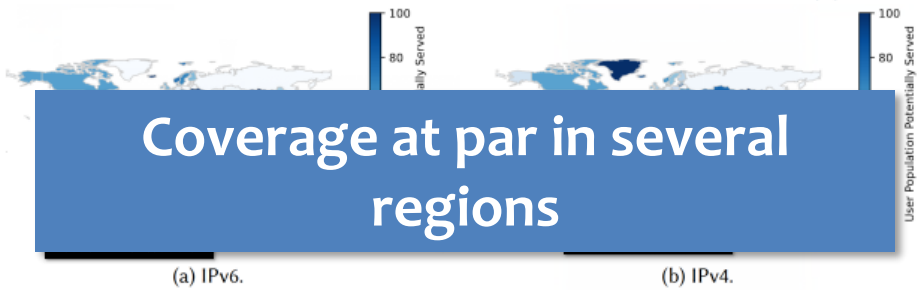

(a) IPv6. (b) IPv4.

Coverage at par in several regions

[1] https://stats.labs.apnic.net/aspop/

fhilal@mpi-inf.mpg.de

UNIVERSITÄT DES SAARLANDES

Contributions

- ***First*** study into HGs IPv6 expansion
  - 2k networks across 14 HGs
- Reveal ***current*** state of IPv4 depl.
  - 6k
- Invest
  - ***agg***
- Find o
- Find IPv6 ***latency*** to be ***at par*** to IPv4

First study into HGs' IPv6 off-nets

Results-Regional Deployment

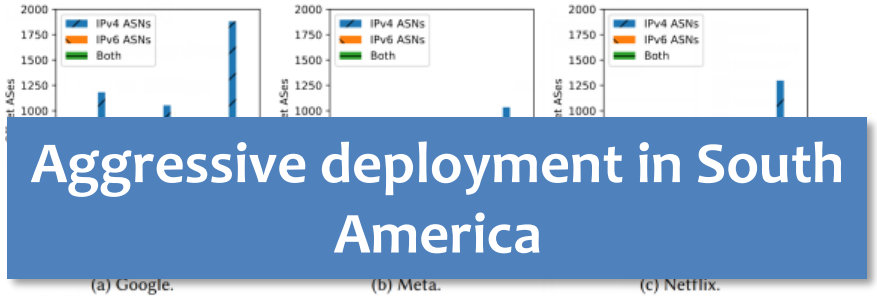

(a) Google. (b) Meta. (c) Netflix.

Aggressive deployment in South America

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google : 0.02% of all on-nets, 21% of all off-nets

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Questions?

Results-Network Types Hosting Off-nets

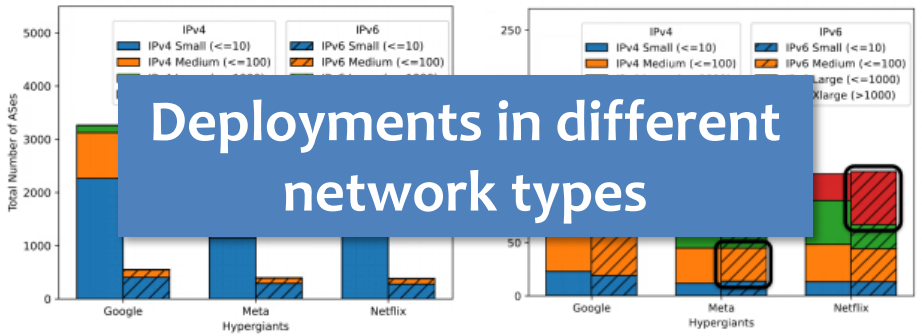

Deployments in different network types

- Based on ***customer cone*** size [1]

[1] https://publicdata.caida.org/datasets/as-relationships/

Results-Access to Off-nets (Meta)

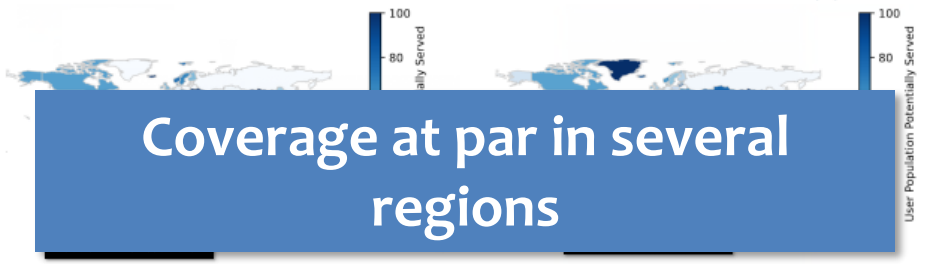

(a) IPv6. (b) IPv4.

Coverage at par in several regions

[1] https://stats.labs.apnic.net/aspop/

Additional Slides

Introduction-Server Deployments

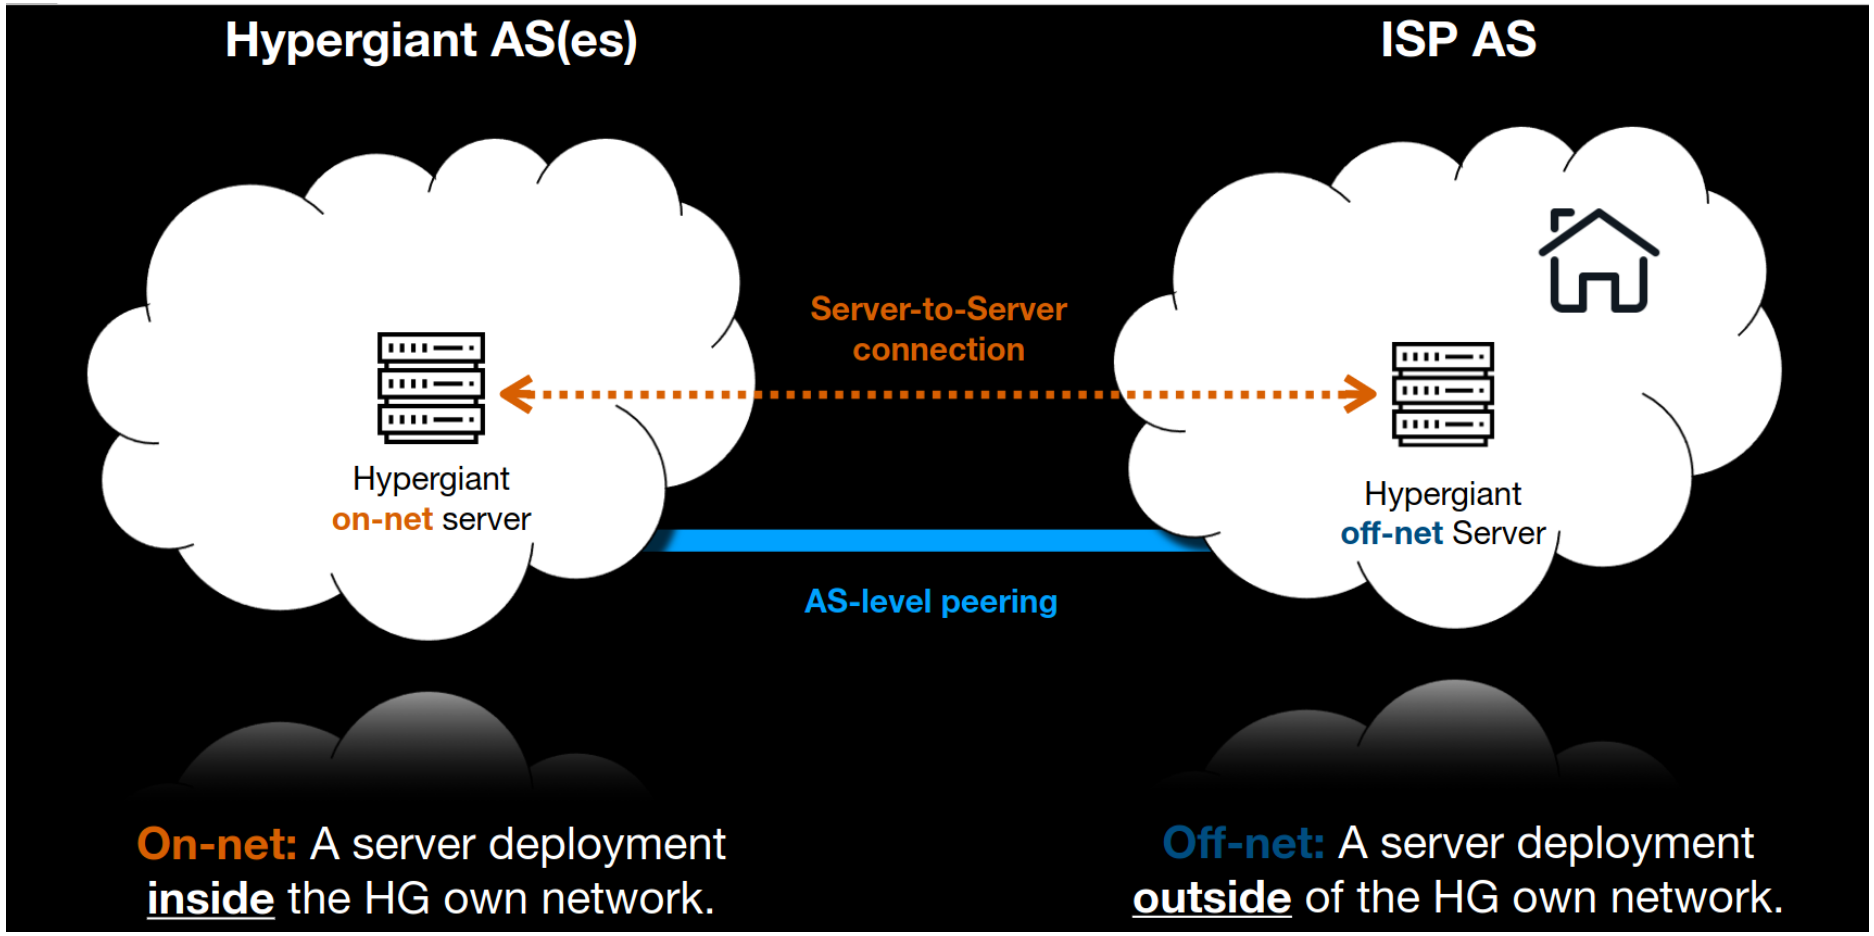

https://pgigis.github.io/hypergiants-offnets/data/pdf/seven_years_in_the_life_of_hypergiants_offnets_slides.pdf

Technique Outline

UNIVERSITÄT DES SAARLANDES

- ***Terminology***
  - TLS certificate reveals if IP hosts a HG service
  - HTTP(S) header reveals who operates server
  - IP address reveals if on-net or off-net server
- ***Detection***
  - For server to be HG off-net
    - TLS certificate + HTTP(S) headers map HG
    - IP not part of HG network

Results-Requirements for Off-net Hosting

UNIVERSITÄT DES SAARLANDES

NETWORK AND DATA CENTER CRITERIA

Criteria	Description
Network	Your ISP should have a **public** autonomous system number (ASN) that you are able to use for peering
Netflix Viewer Traffic	For offload efficiency and system scale, we evaluate the level of current Netflix traffic to determine whether embedded OCA(s) can provide sufficient offload
Network Capacity	Each site must have the capacity to handle **1.2 Gbps of inbound traffic daily** per appliance for **fill and updates**. The majority of fill occurs during low traffic times to maximize offload.
Interconnection	You must connect to Netflix via SFI (peering) at one or more Netflix **points of presence (PoPs)** if it is possible to do so.
AWS Connectivity	OCAs and clients must be able to communicate with the control plane services running in Amazon Web Services (AWS). If connectivity to AWS is lost, OCAs will stop serving traffic. **Note**: You can confirm the status of required inbound/outbound OCA connectivity **in the Partner Portal**.
Installation	You must be able to **physically install** OCAs at your assigned site within **10 business days of receipt**, or as soon as possible thereafter
Consistent room temperatures	No higher than **78°F (26°C)**
Maximum room temperatures	No higher than **104°F (40°C)** for very short periods of time only
Physical Connection	You must be able to provision **1-2 X 100G** or **2-6 x 10 Gbps optical ethernet ports in a LACP LAG** per OCA. The exact quantity depends on the OCA hardware type.
Maximum Power Supply Draw	350-700W per OCA, varies by hardware type and utilization

https://openconnect.zendesk.com/hc/en-us/articles/360034538352-Requirements-for-deploying-embedded-appliances

Results-TLS vs TLS + HTTP(S)

HG	Off-net ASes Dropped (IPv6)	Off-net ASes Dropped (IPv4)
Google	2.5%	1.3%
Meta	0.6%	1.3%
Netflix	5%	6.4%

- Drop in Off-net ASes after applying HTTP(S) validation ***small***

Validation

Method	Off-net ASes
Ours (2023)	2611
Brute-Force (2022)	1560
Intersection	**92%**

- brute-force airport codes in FB urls [1]
  - scontent.fXXXY-Z.fna.fbcdn.net.
  - XXX: airport code, Y: 1st ISP, Z:1st cluster.
  - eg: scontent.f***FRA***1-1.fna.fbcdn.net - ***Frankfurt***
- ***89%-95%*** HG IPv4 footprint uncovered by Gigis et al

[1] https://anuragbhatia.com/2022/07/networking/isp-column/facebook-cache-fna-updates-july-2022/

Results-Networks with Off-nets

UNIVERSITÄT DES SAARLANDES

HG	# Off-net ASes (IPv6)	# Off-net ASes (IPv4)	% Off-net ASes (Both)
Google	1.3k	5k	**96.2%**
Meta	1.2k	2.6k	**96.3%**
Netflix	928	2.7k	**92.7%**

- **15%-30%** increase in IPv4 off-net footprint since 2021 [1]
- ***no*** IPv6 prefixes announced by **~30%** IPv4 off-net ASes

[1] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021.

Results-Networks with Off-nets

HG	# Off-net ASes (IPv6)	# Off-net ASes (IPv4)	% Off-net ASes (Both)
Google	1.3k	5k	96.2%
Meta	1.2k	2.6k	96.3%
Netflix	928	2.7k	92.7%

IPv6 only Off-net networks rare

- **15%-30%** increase in IPv4 off-net footprint since 2021 [1]
- ***no*** IPv6 prefixes announced by **~30%** IPv4 off-net ASes

[1] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021.

Results-Google Data Centers

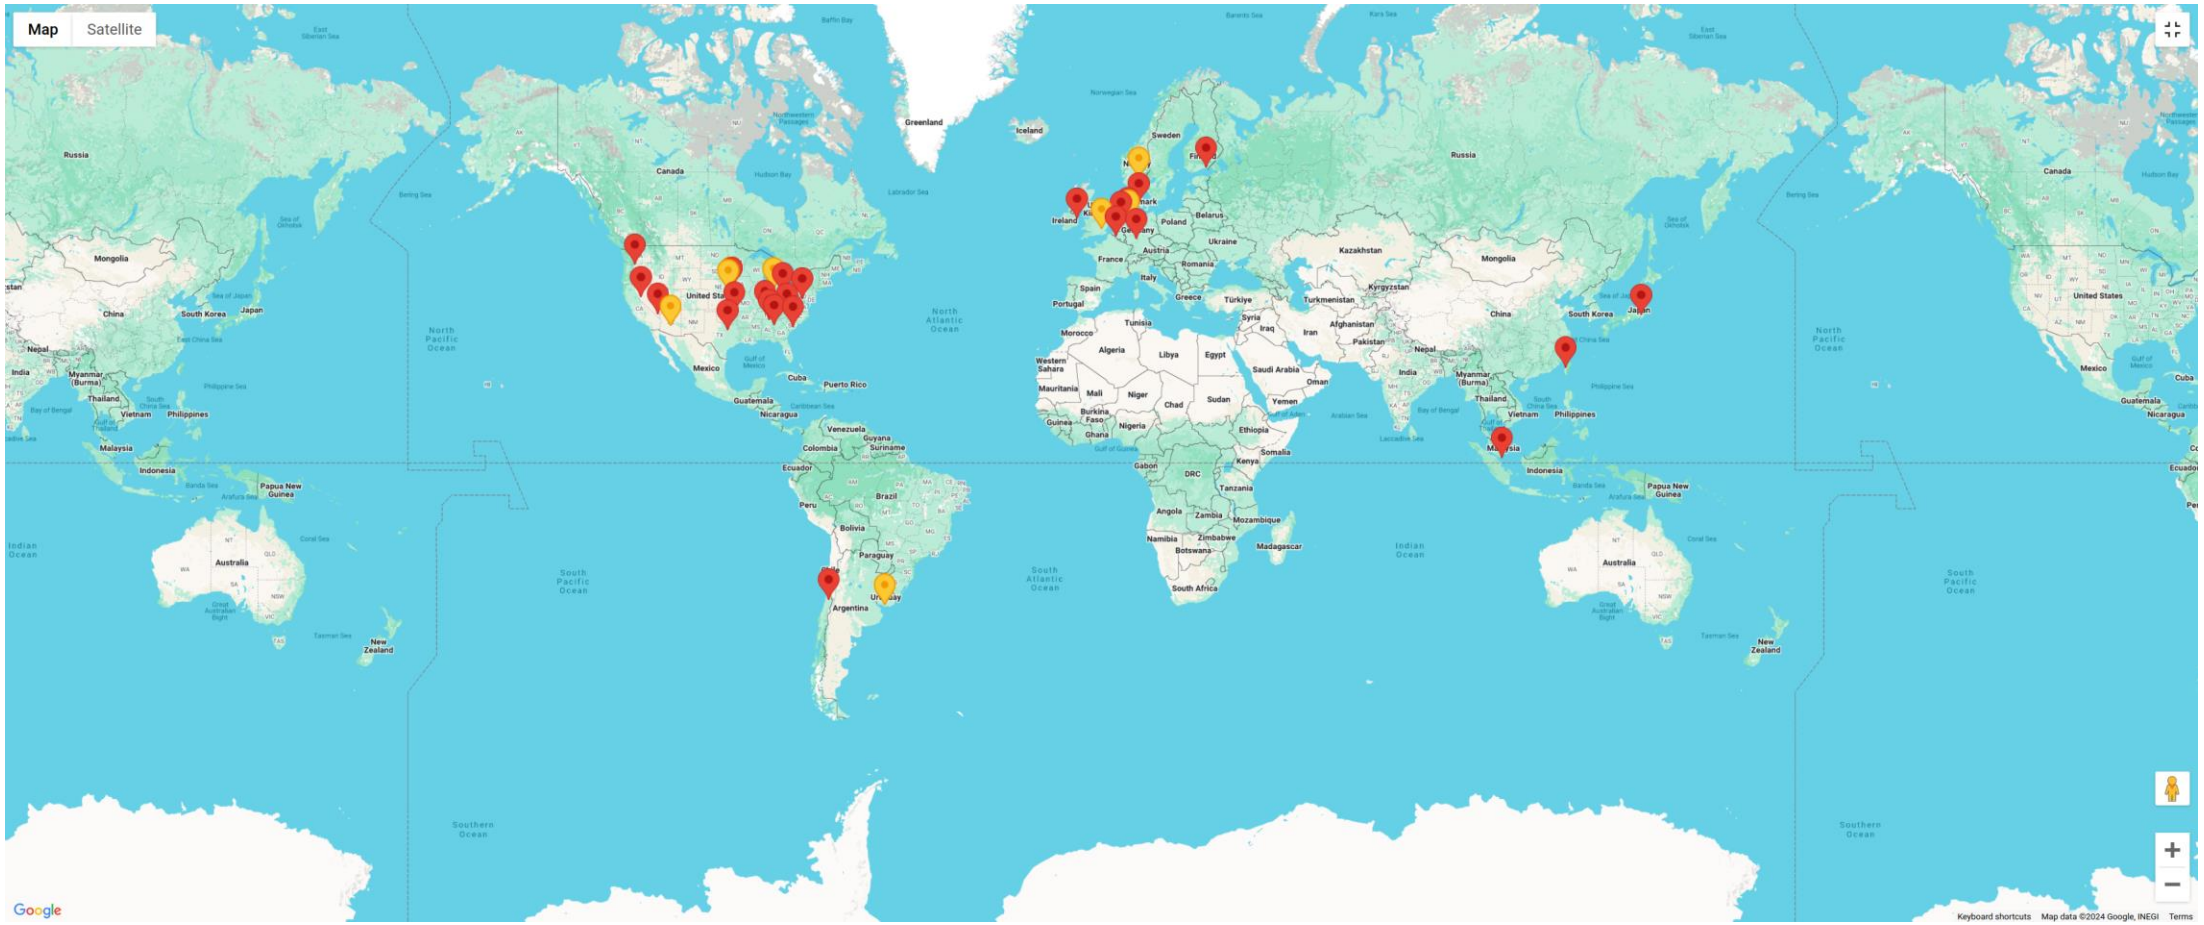

https://www.google.com/about/datacenters/locations/

Results-Networks with Off-nets

HG	# Off-net ASes (IPv6)	# Off-net ASes (IPv4)	Both
Google	1,342	4,976	1,291
Meta	1,231	2,565	1,185
Netflix	928	2,731	860
Akamai	241	881	223
Apple	**117**	**219**	**104**
Alibaba	37	175	26
Amazon	11	171	7
Microsoft	**2**	**174**	**0**
Fastly	2	6	0

- Apple, Microsoft: ***no*** off-nets in 2021, Amazon: 175% ***growth*** [1]
- Akamai: ~20% ***decrease***

[1] Gigis, Petros, et al. "Seven years in the life of Hypergiants' off-nets." Proceedings of the 2021 ACM SIGCOMM 2021 Conference. 2021.

Results-Off-net Coverage Per Region

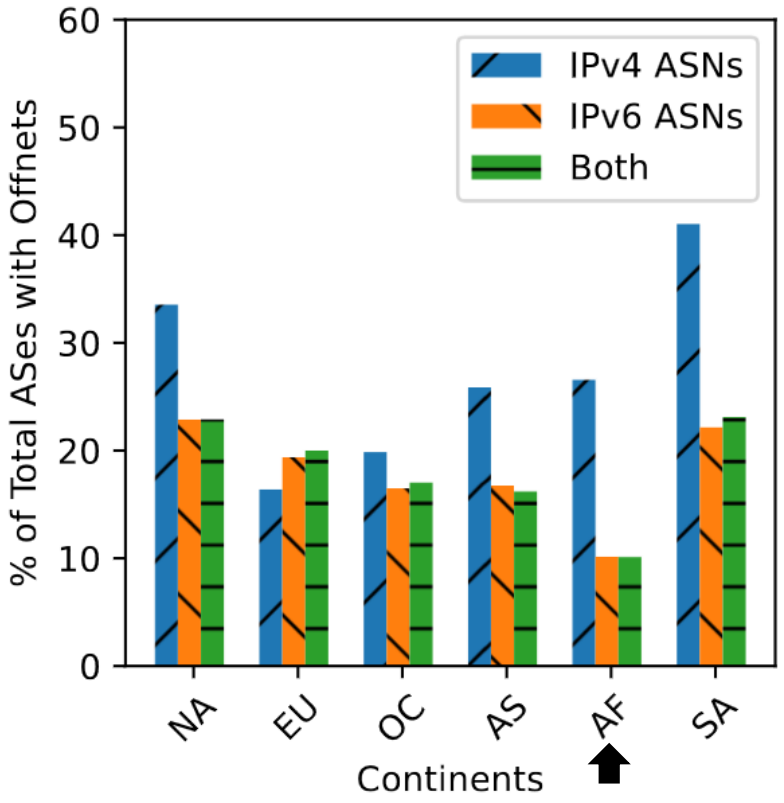

(a) Google.

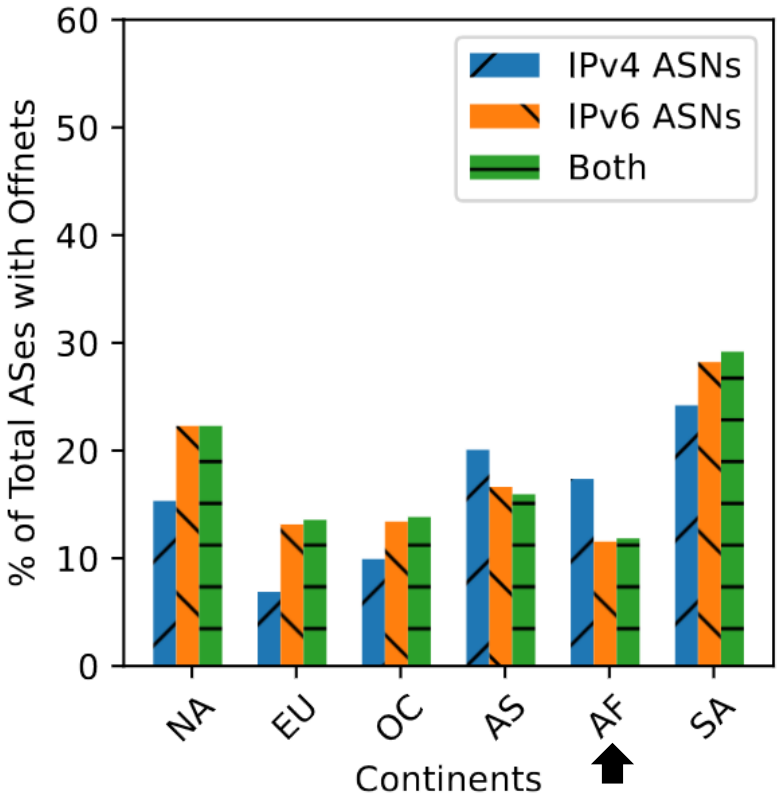

(b) Meta.

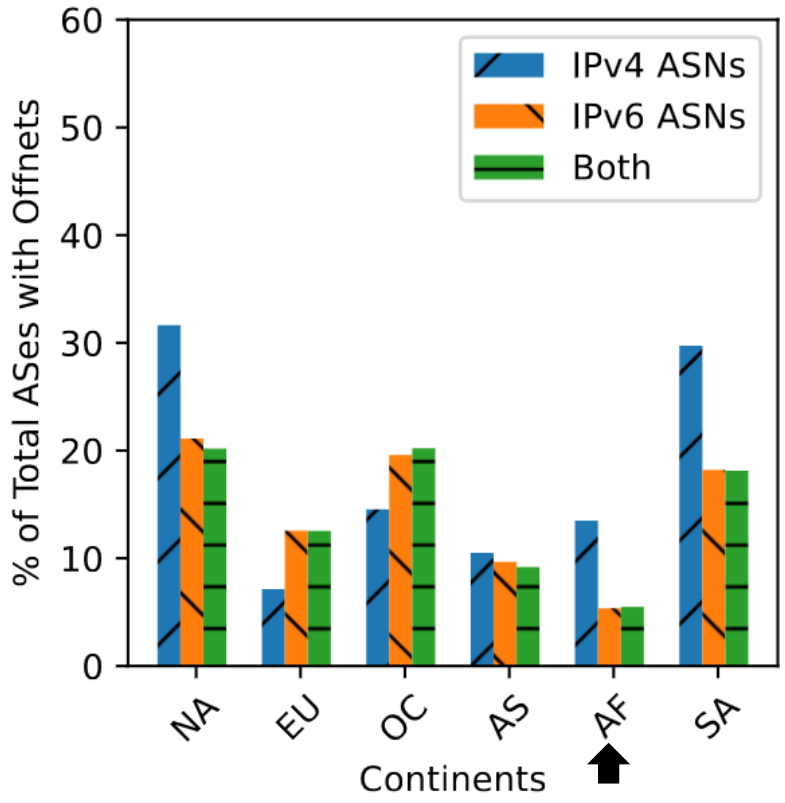

(c) Netflix.

- Out of ASes serving non-zero user base per continent [1]

[1] https://labs.apnic.net/measurements/

Results-Off-net Coverage Per Region

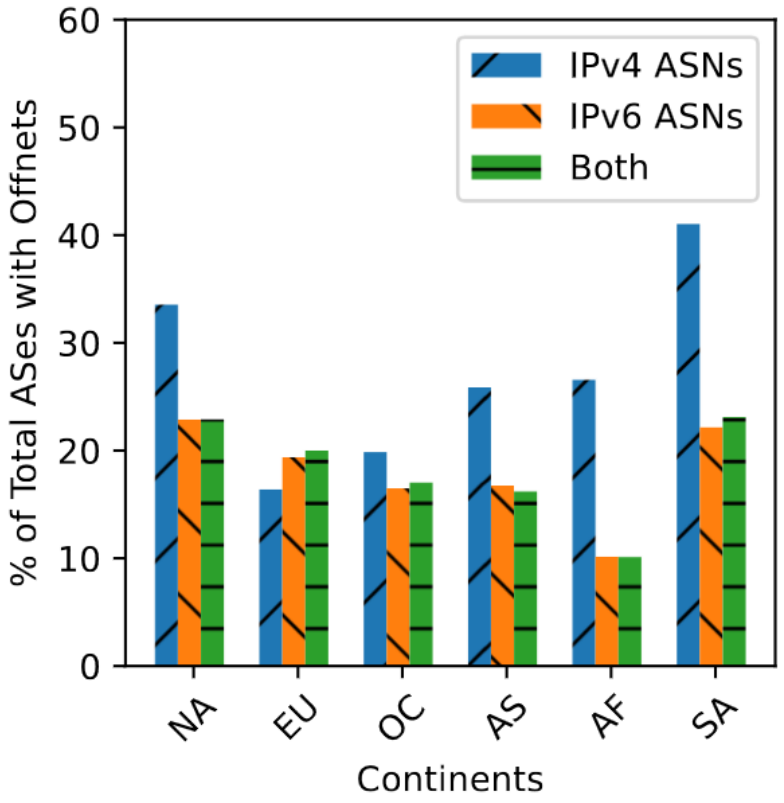

(a) Google.

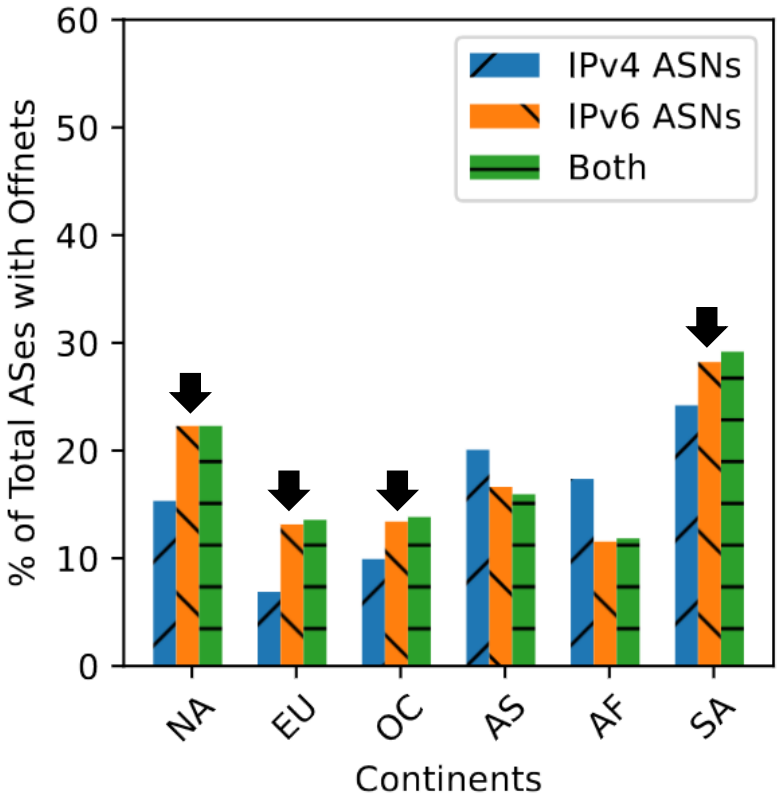

(b) Meta.

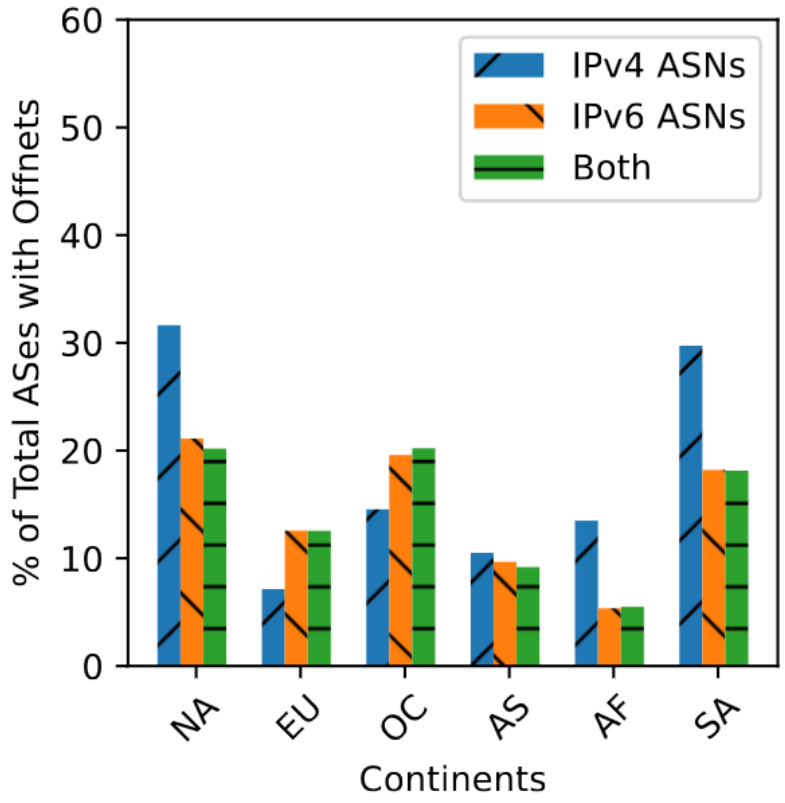

(c) Netflix.

- Out of ASes serving non-zero user base per continent [1]

[1] https://labs.apnic.net/measurements/

Results-Off-net Coverage Per Region

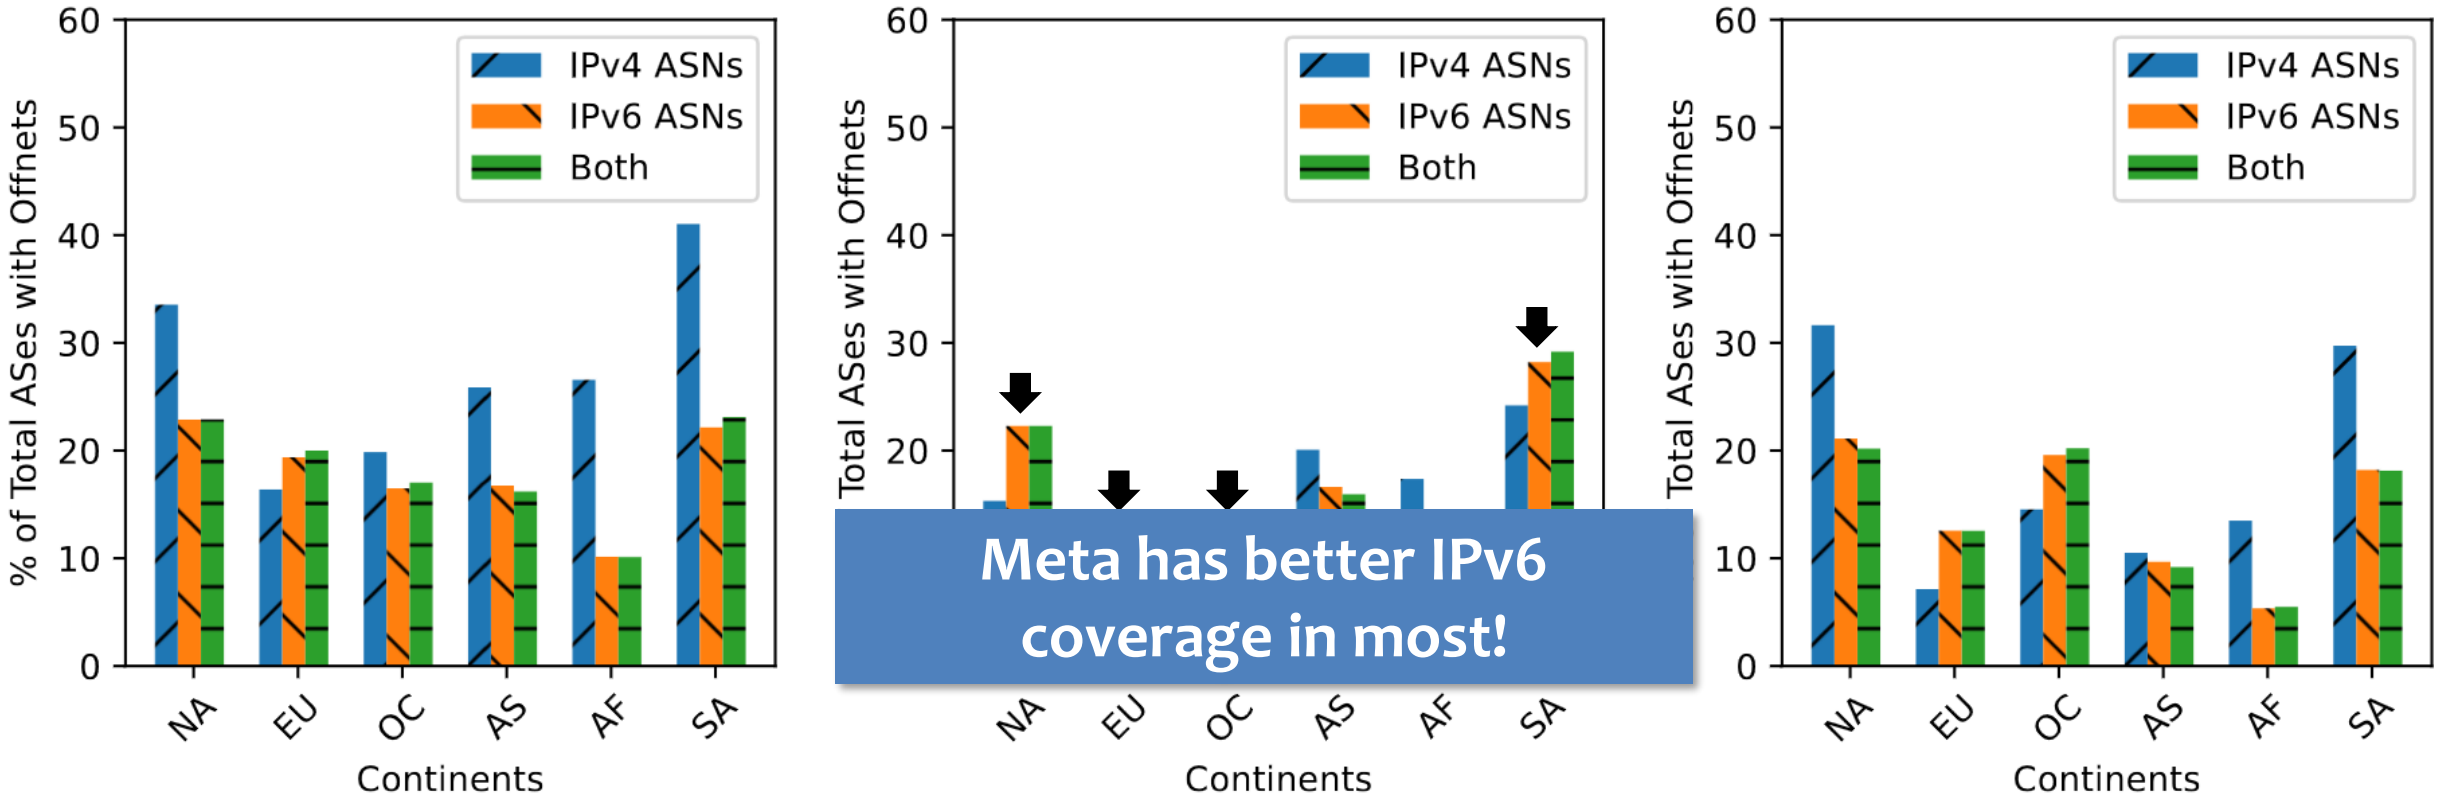

(a) Google. (b) Meta. (c) Netflix.

- Out of ASes serving non-zero user base per continent [1]

[1] https://labs.apnic.net/measurements/

Results-Regional Deployment (Alibaba)

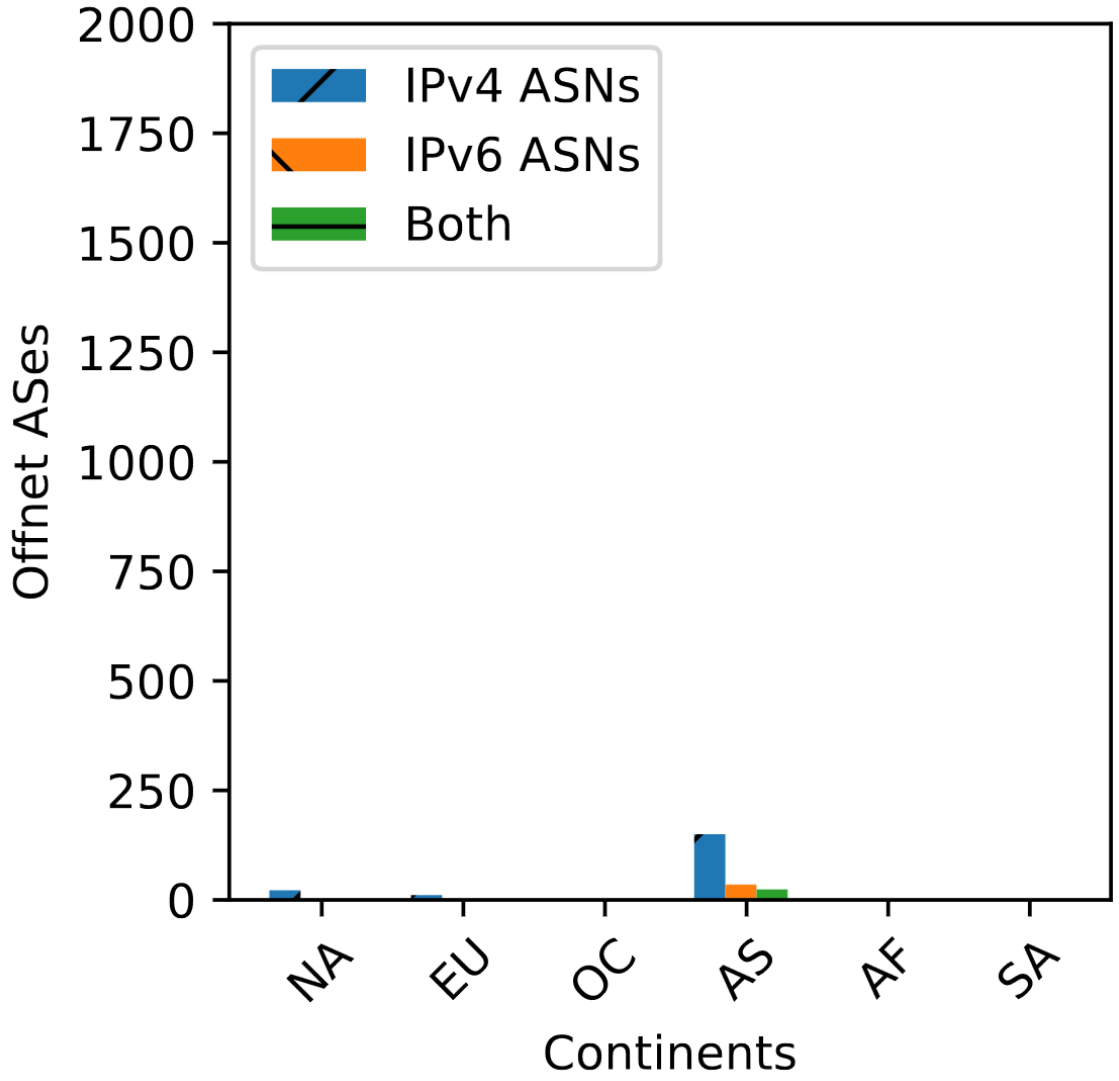

Results-Access to Off-nets (Google)

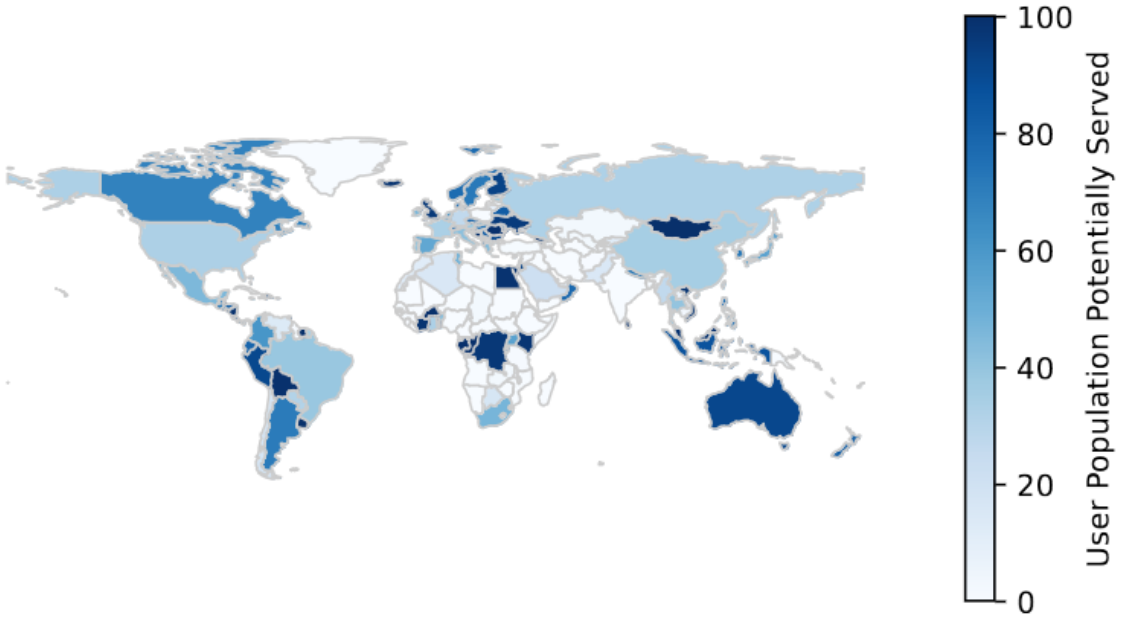

(a) IPv6.

(b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Results-Access to Off-nets (Netflix)

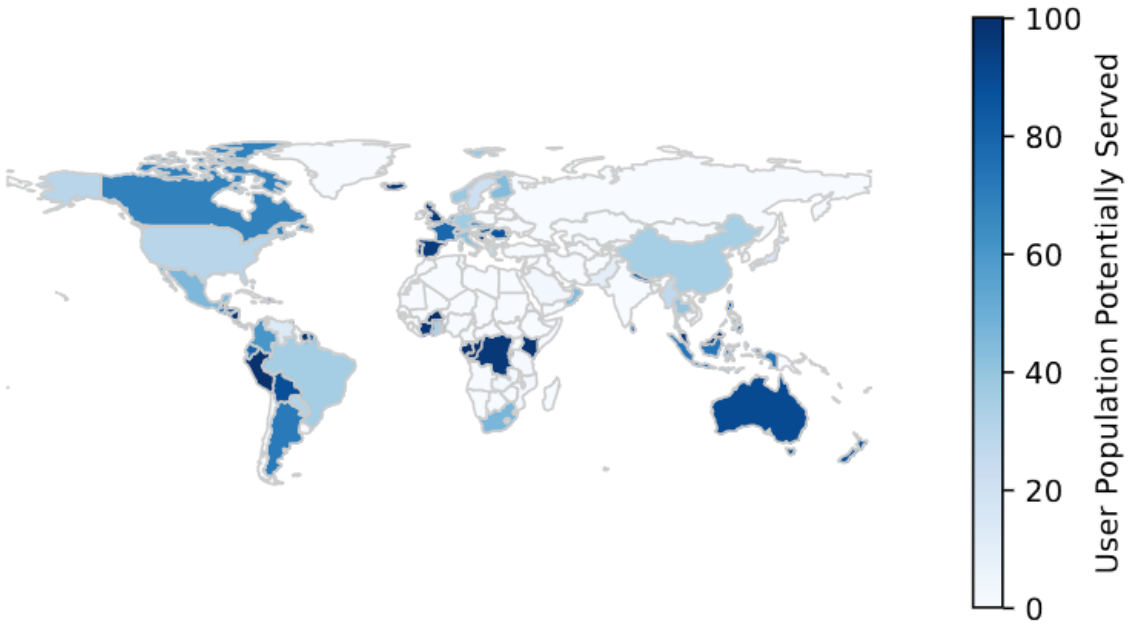

(a) IPv6.

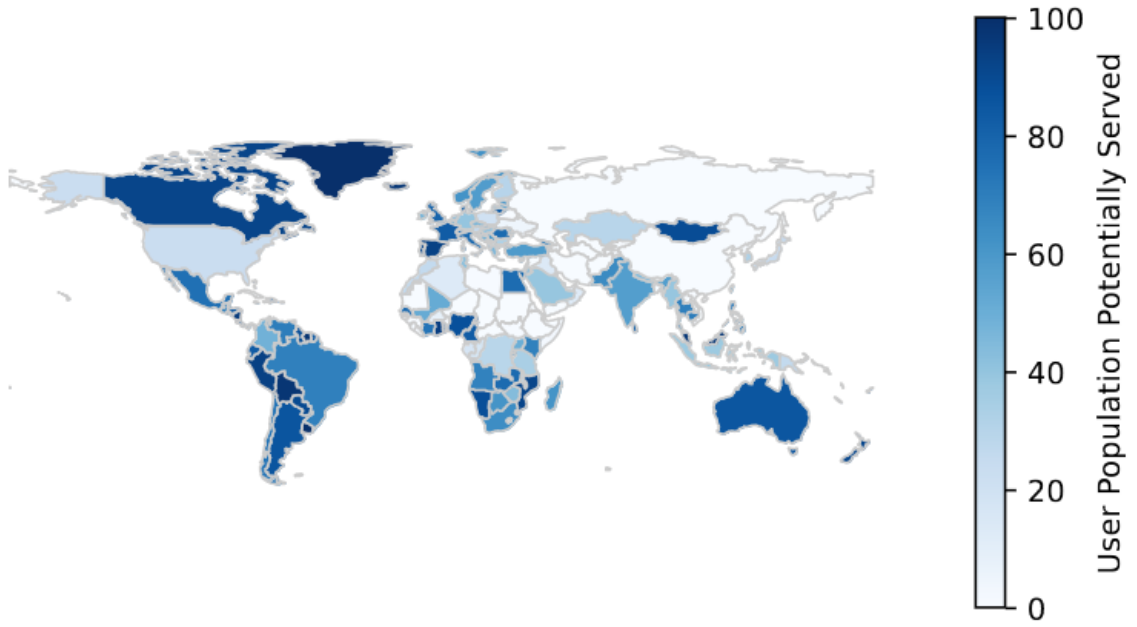

(b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Results-Access to Off-nets (Akamai)

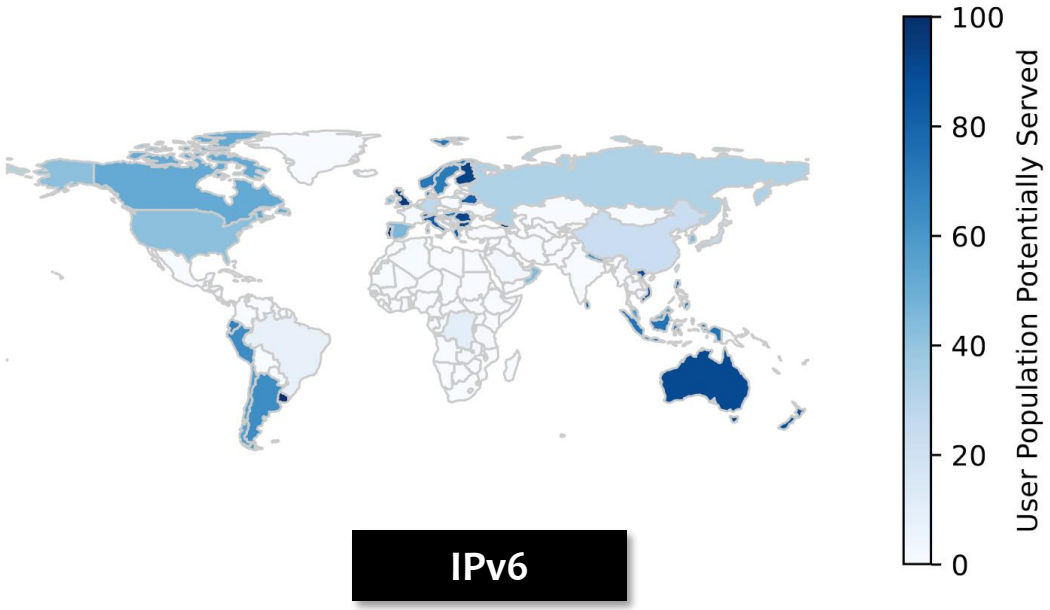

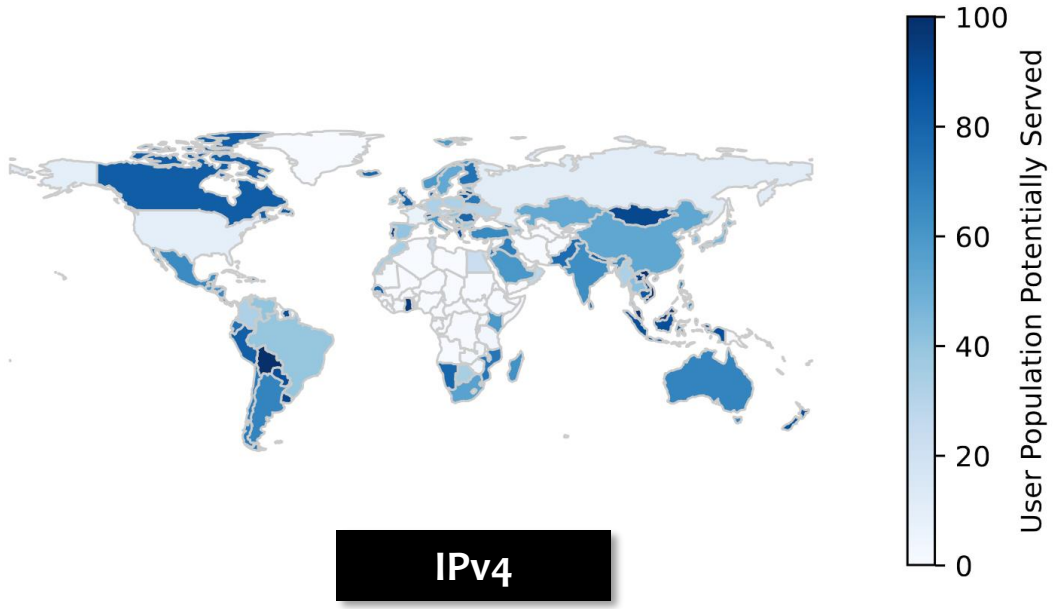

- ~20% drop in IPv4 off-net footprint since 2021

[1] https://stats.labs.apnic.net/aspop/

Results-Access to Off-nets (Alibaba)

- Most deployment in China
  - 40% ASes have off-nets
- Third-party servers outside

[1] https://stats.labs.apnic.net/aspop/

Results-Access to Off-nets (IPv6)

UNIVERSITÄT DES SAARLANDES

Google

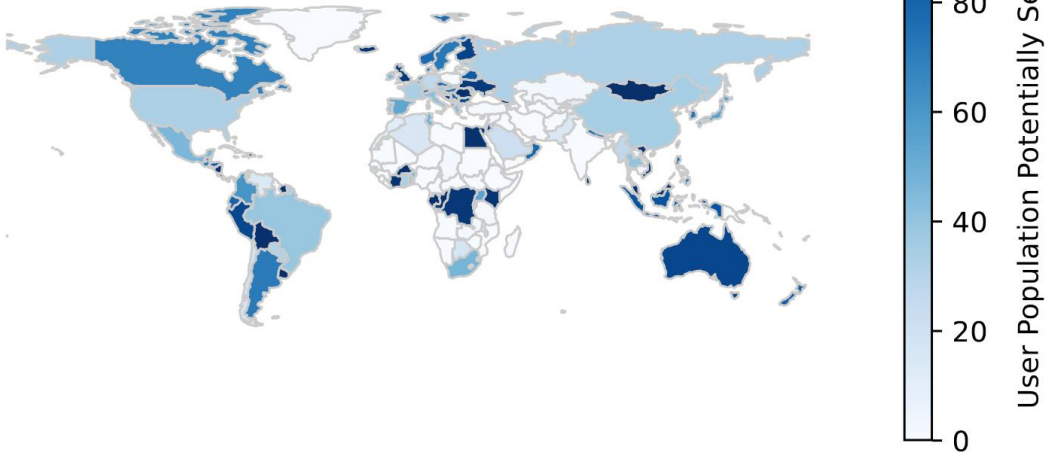

Meta

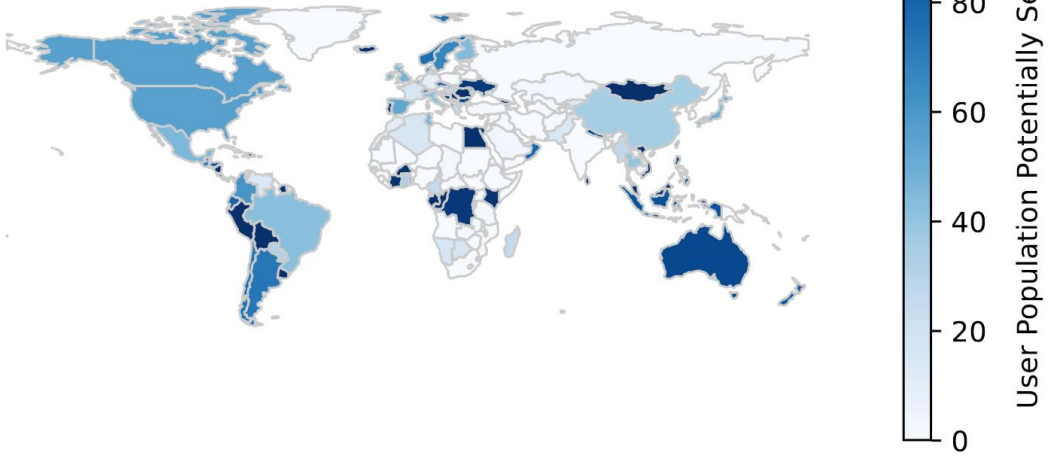

Netflix

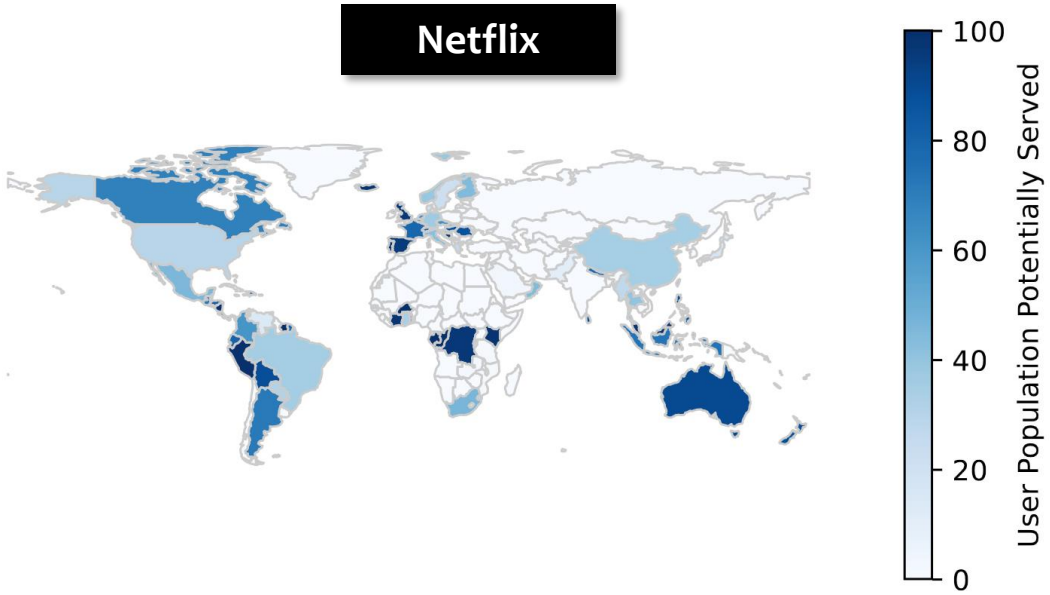

[1] https://stats.labs.apnic.net/aspop/

Results-Hosting Multiple Hypergiants

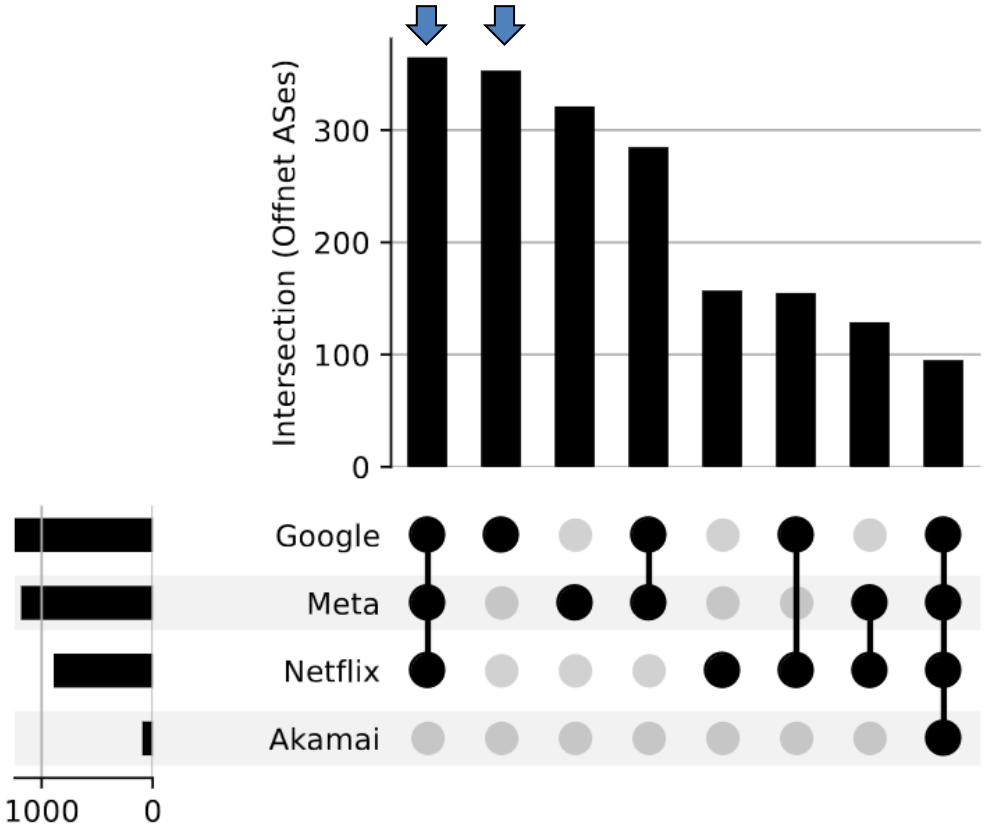

(a) IPv6.

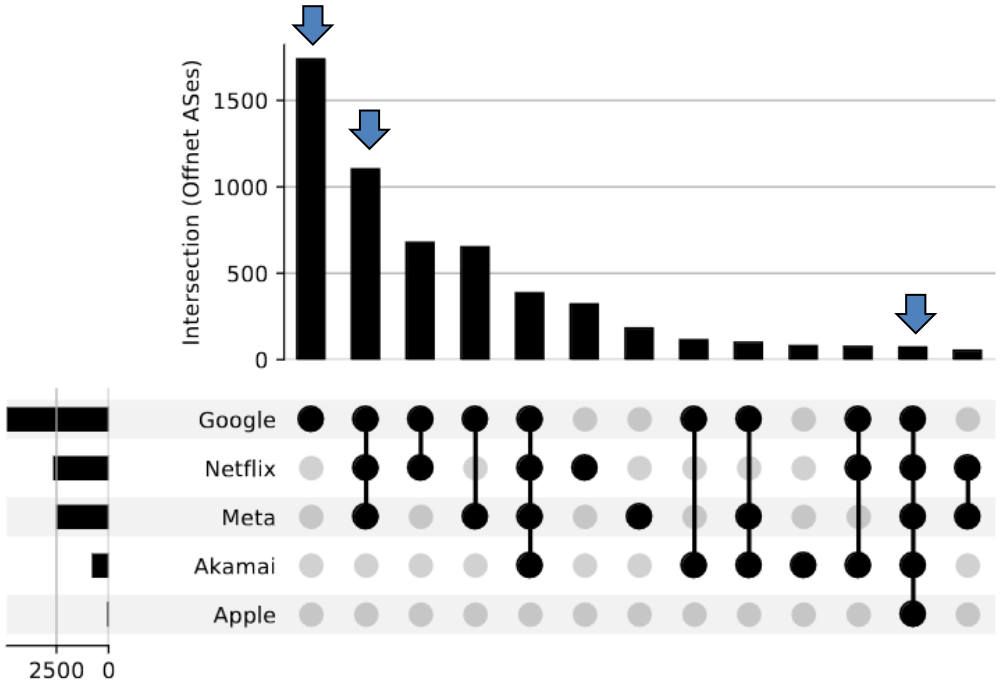

(b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Results-Hosting Multiple Hypergiants

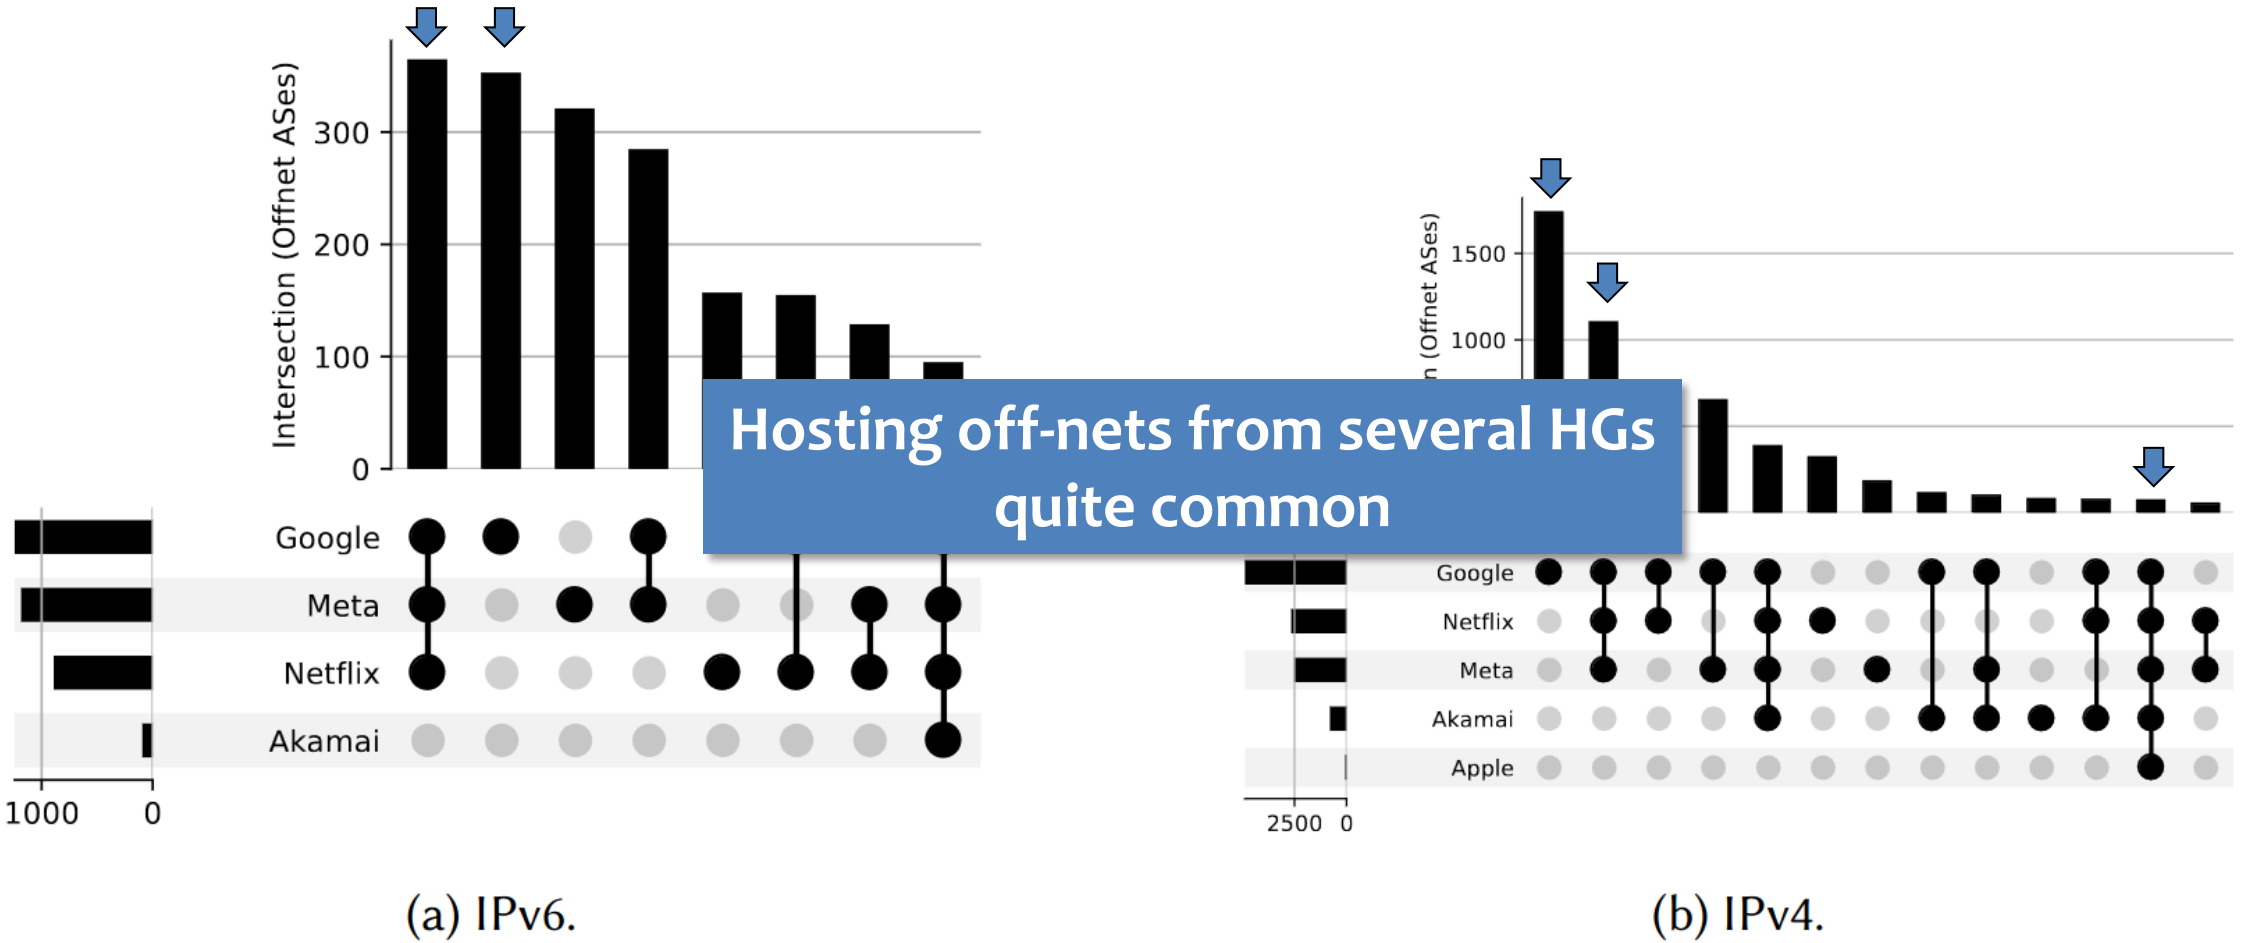

(a) IPv6.

(b) IPv4.

[1] https://stats.labs.apnic.net/aspop/

Future Work

UNIVERSITÄT DES SAARLANDES

- ***Enhance*** target list for IPv6 off-net discovery
- Track growth
- Off-net ***services***

Results-Off-nets at Risk

HG	% of Off-net ASes (IPv6)	% of Off-net ASes (IPv4)
Google	32.2%	50.7%
Meta	29.8%	51.6%
Netflix	**35.6%**	**58%**

- ROA protected prefixes:
  - Google: 99.5% (on-net IPv6), 35.6% (off-net IPv6)

Results-Off-nets at Risk

HG	% of Off-net ASes (IPv6)	% of Off-net ASes (IPv4)
Google	32.2%	50.7%
Meta	29.8%	51.6%
Netflix	**35.6%**	**58%**

- ROA protected prefixes:
  - Google: 99.5% (on-net IPv6), 35.6% (off-net IPv6)
- ROA covered prefixes (ROA BGP)
- ROA covered off-net prefixes (ROA Off.)
- For off-net AS,
  - atleast: ROA Off. > ROA BGP
    - special treatment

Results-Off-nets at Risk

HG	% of Off-net ASes (IPv6)	% of Off-net ASes (IPv4)
Google	32.2%	50.7%
Meta	29.8%	51.6%
Netflix	**35.6%**	**58%**

- ROA protected prefixes:
  - Google: 99.5% (on-net IPv6), 35.
- ROA covered prefixes (ROA BGP)
- ROA covered off-net prefixes (ROA Off.)
- For off-net AS,
  - atleast: ROA Off. > ROA BGP
    - special treatment

Majority of networks do not treat off-net prefixes differently

Results-Per Continent Deployment

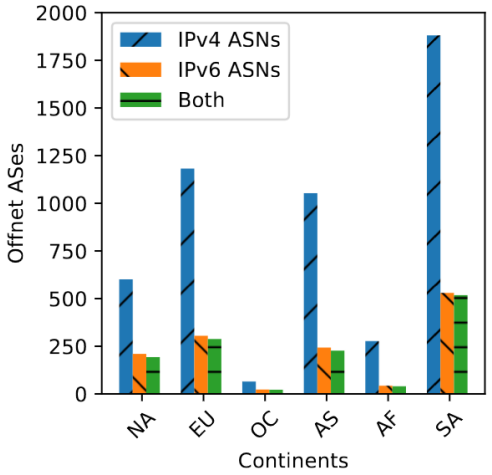

(a) Google.

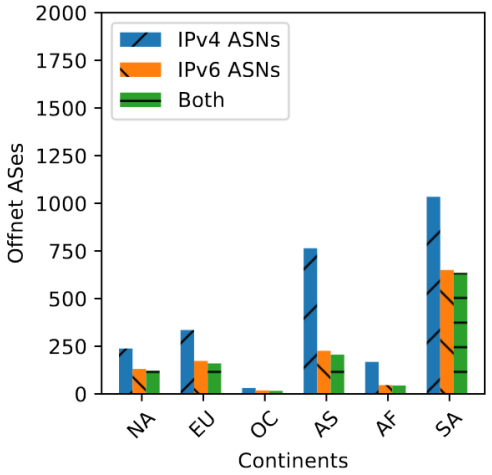

(b) Meta.

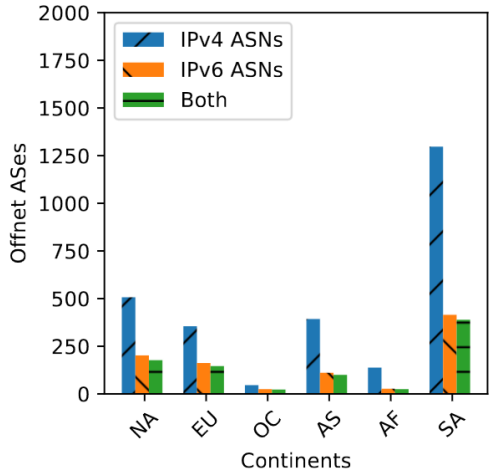

(c) Netflix.

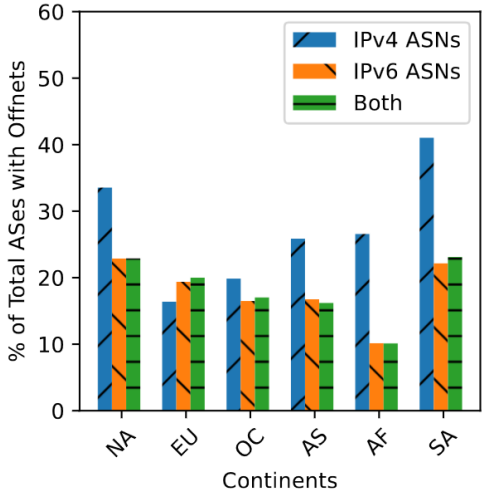

(a) Google.

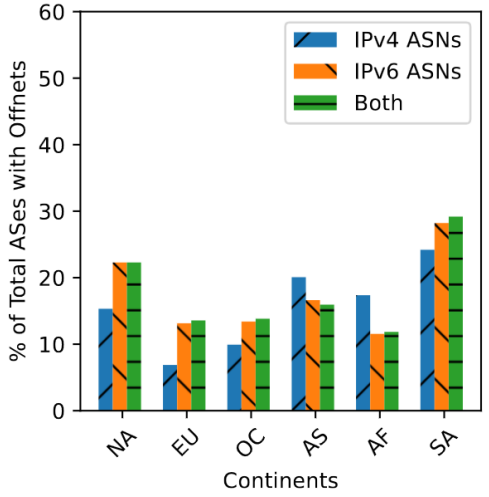

(b) Meta.

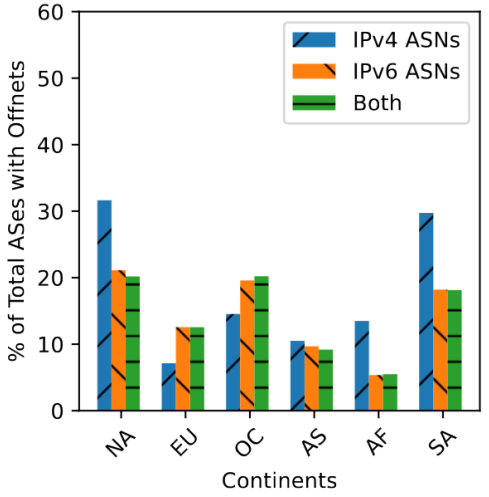

(c) Netflix.

Results-Potential Censorship Implications

HG	China	Iran	US	Scandinavia
Google	7%, 3%	8%, 0%	**41%, 18%**	22%, 21%
Meta	1%, 3%	0%, 0%	9%, 11%	6%, 10%

- China, Iran known to censor Facebook, Google [1,2,3]
- Find ASes with non-zero user base in region
  - check if off-nets deployed

[1] Farnan, Oliver, Alexander Darer, and Joss Wright. "Poisoning the well: Exploring the great firewall's poisoned dns responses." Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society. 2016.
[2] Hoang, Nguyen Phong, et al. "How Great is the Great Firewall? Measuring China's {DNS} Censorship." 30th USENIX Security Symposium (USENIX Security 21). 2021.
[3] Aryan, Simurgh, Homa Aryan, and J. Alex Halderman. "Internet censorship in Iran: A first look." 3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13). 2013.

Results-Internet User Population Coverage

- APNIC AS POP
  - APNIC online ad-based (Google ads) measurements to estimate user population per AS
  - ITU data to normalize findings

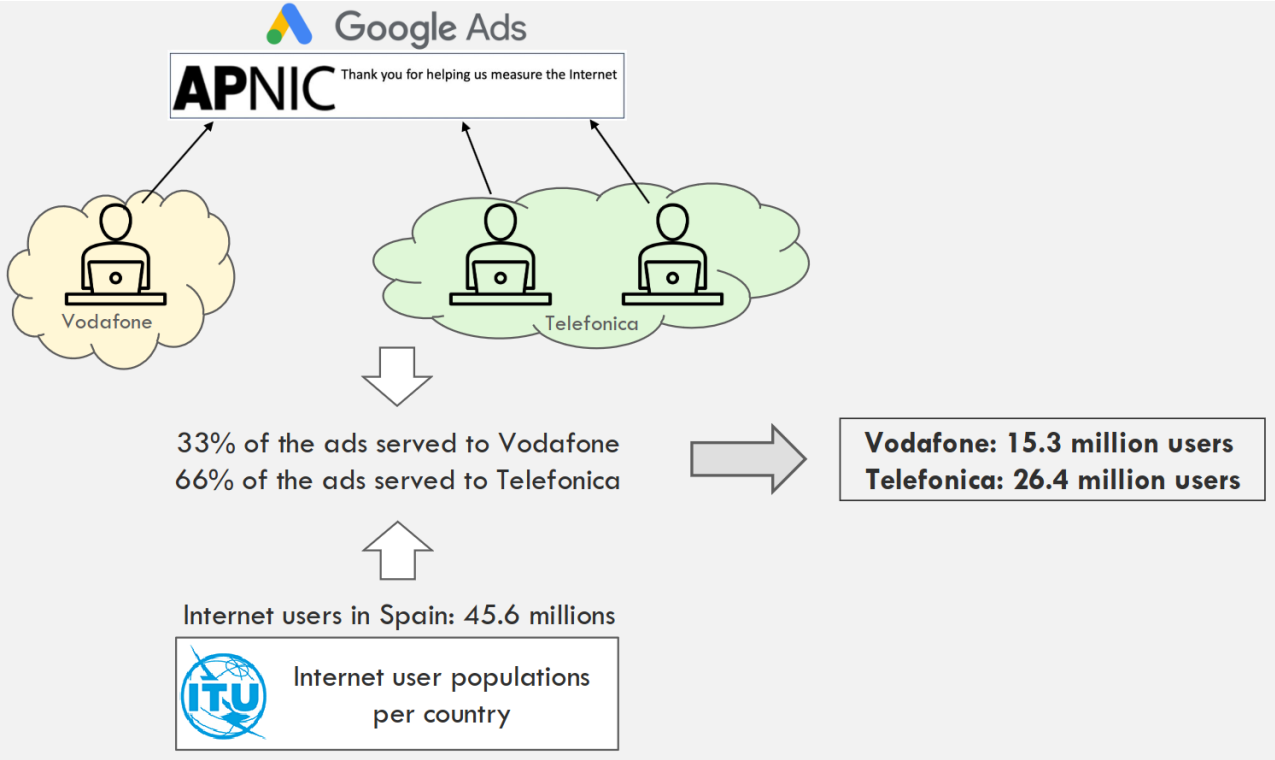

https://ripe89.ripe.net/presentations/77-UnboxingAPNIC-RIPE89.pdf

Results-Regional Deployment

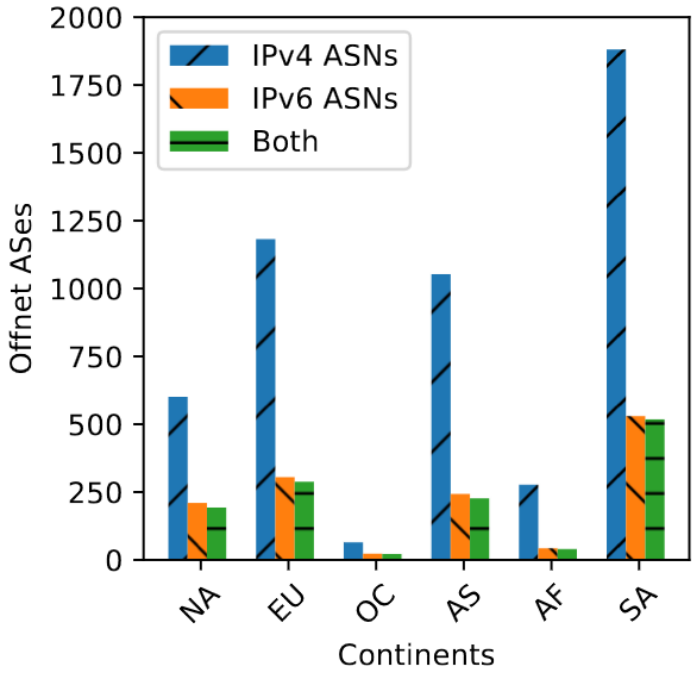

(a) Google.

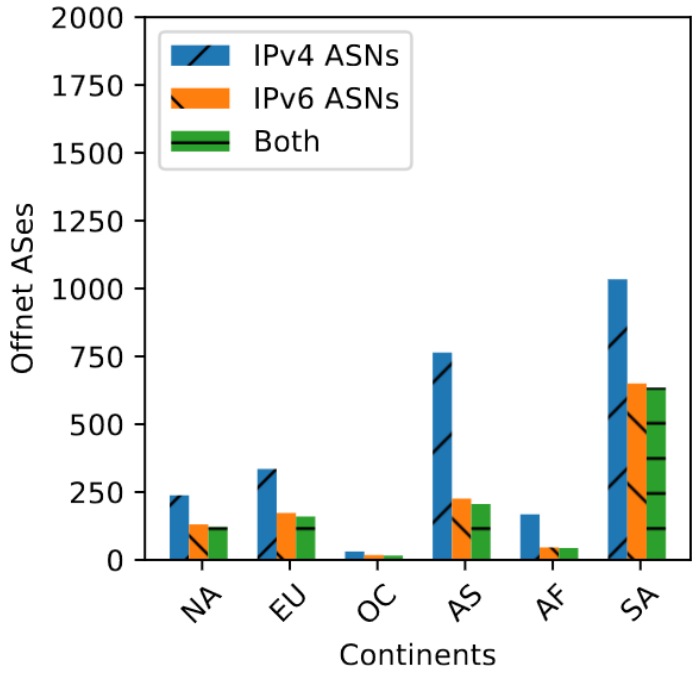

(b) Meta.

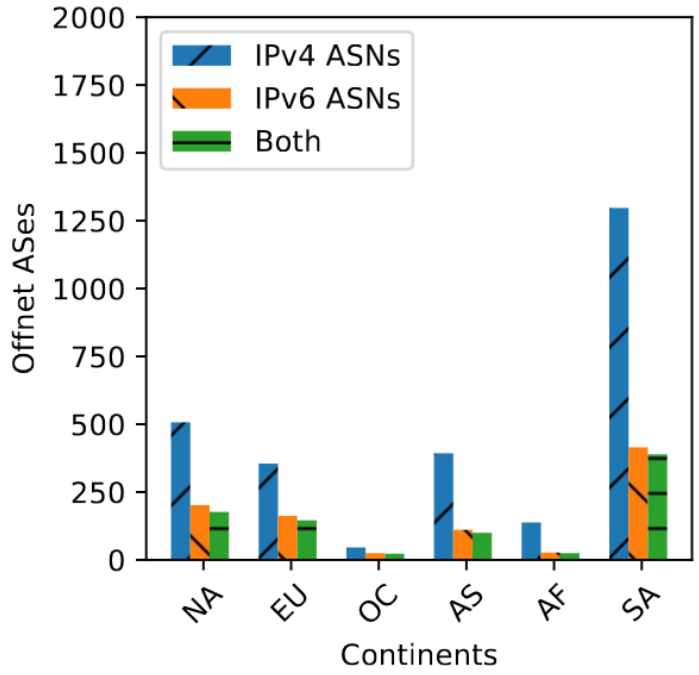

(c) Netflix.

- Geo-locate IP -> country [1], Map IP -> AS [2, 3], AS -> country
- South America
  - Google: 0.02% of all on-nets, 21% of all off-nets
  - ***3 times fewer*** ASes ***peer*** with HGs

[1] https://dev.maxmind.com/geoip/geolocate-an-ip/databases
[2] https://routeviews.org/
[3] https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

Results-Access to Off-nets (Meta)

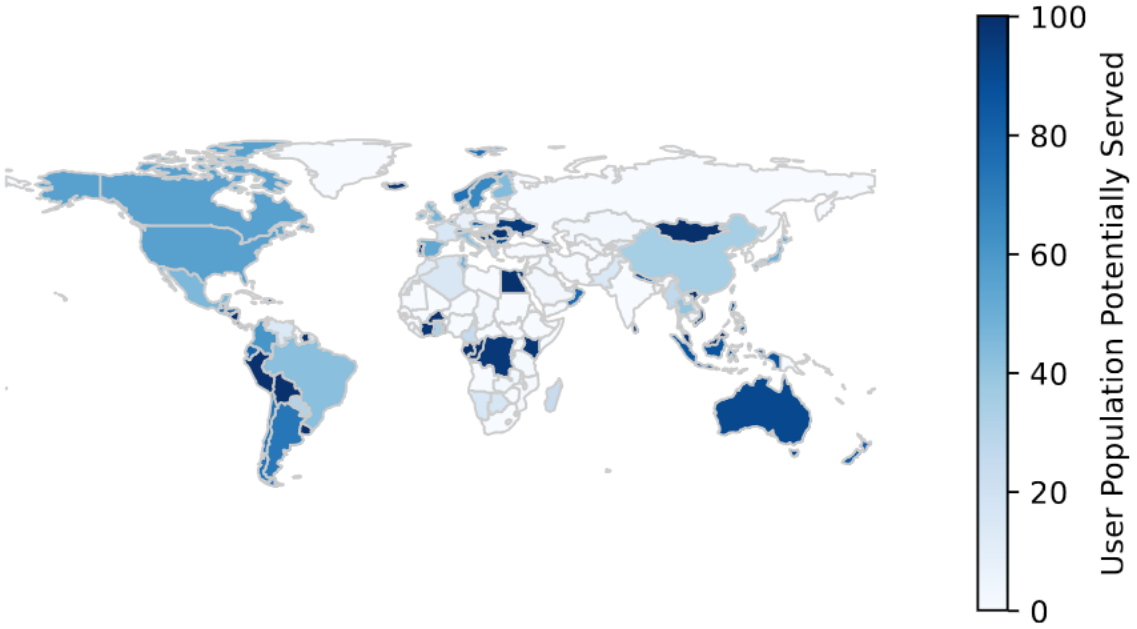

(a) IPv6.

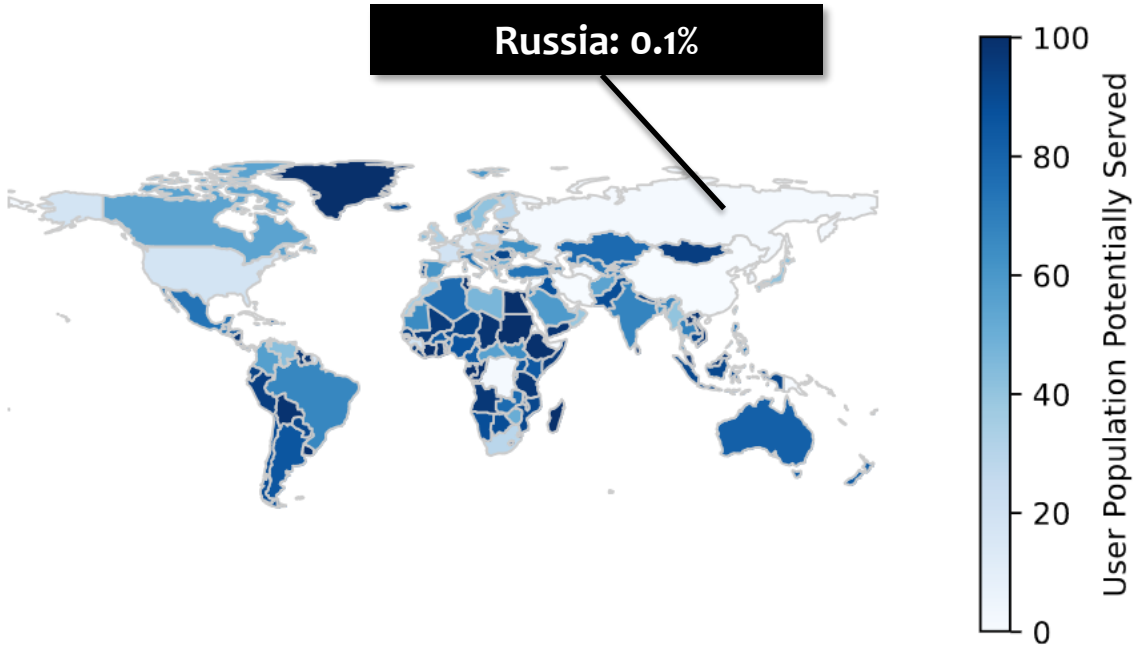

(b) IPv4.

- Russia
  - Reduction since 2021

Results-Access to Off-nets (Meta)

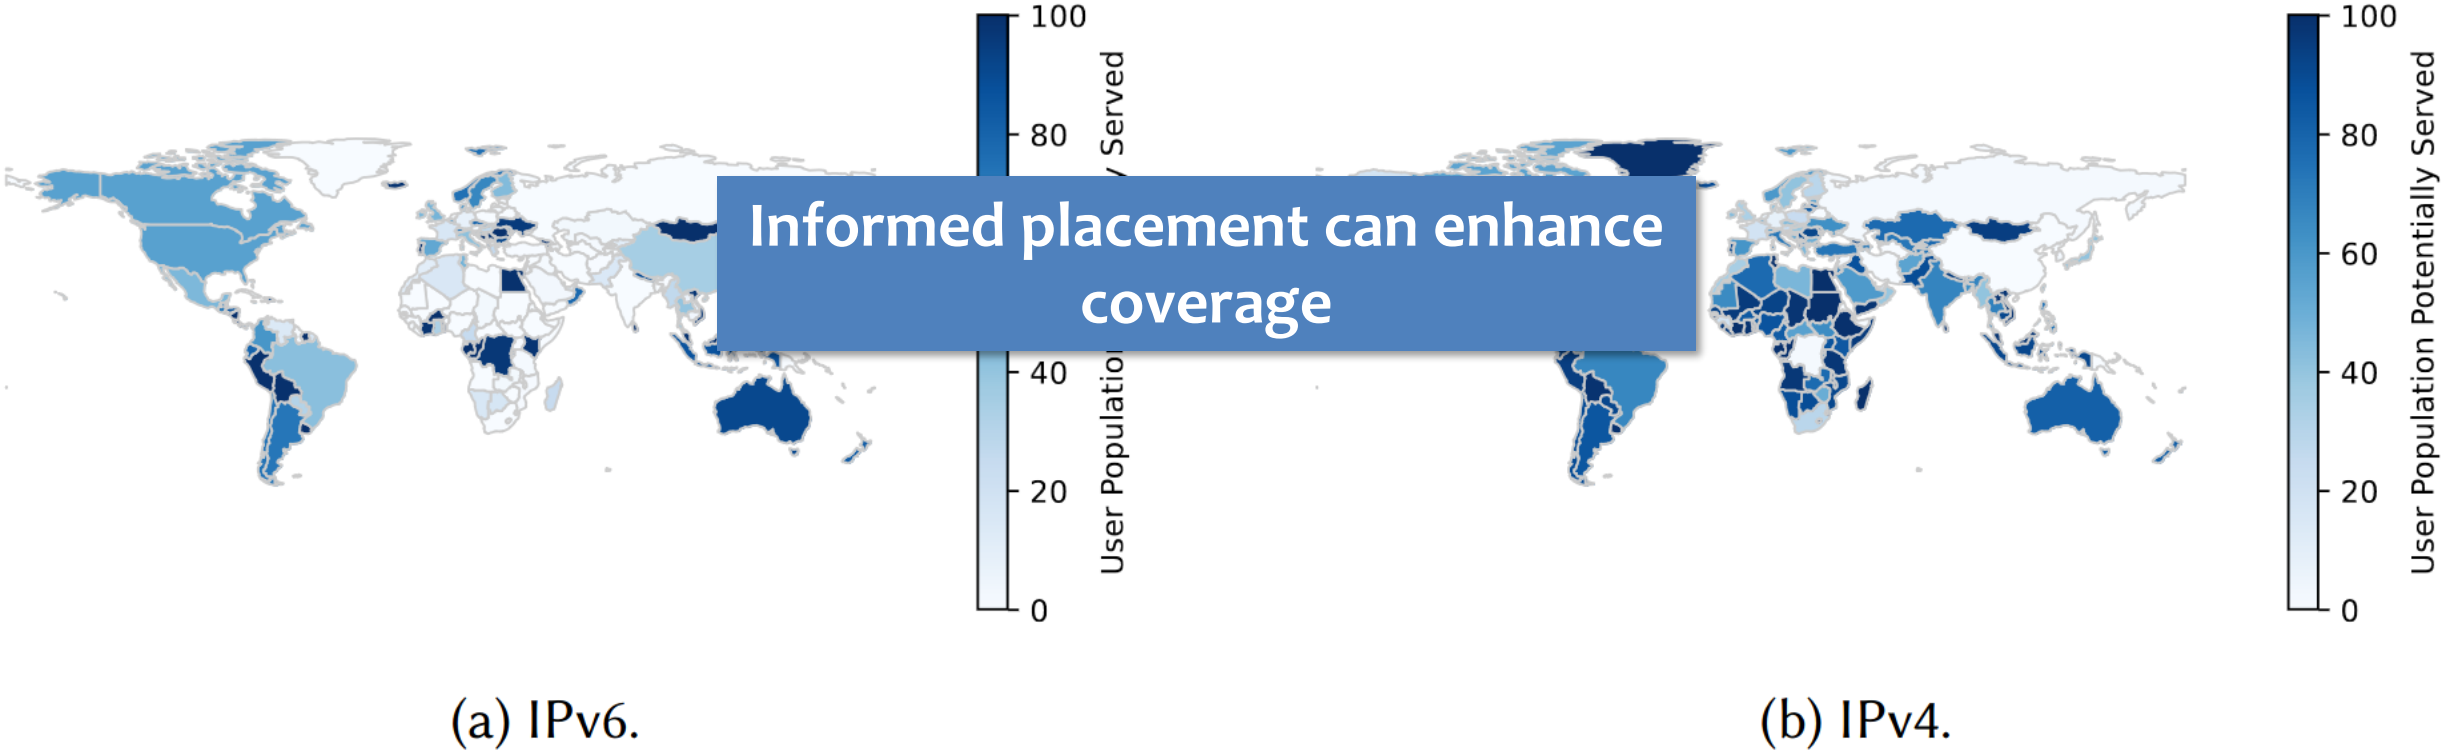

(a) IPv6.

(b) IPv4.

- Russia
  - AS 51659 (Baxet), AS 8359 (MTS) improve by over ***50%***

Results-Off-net Performance

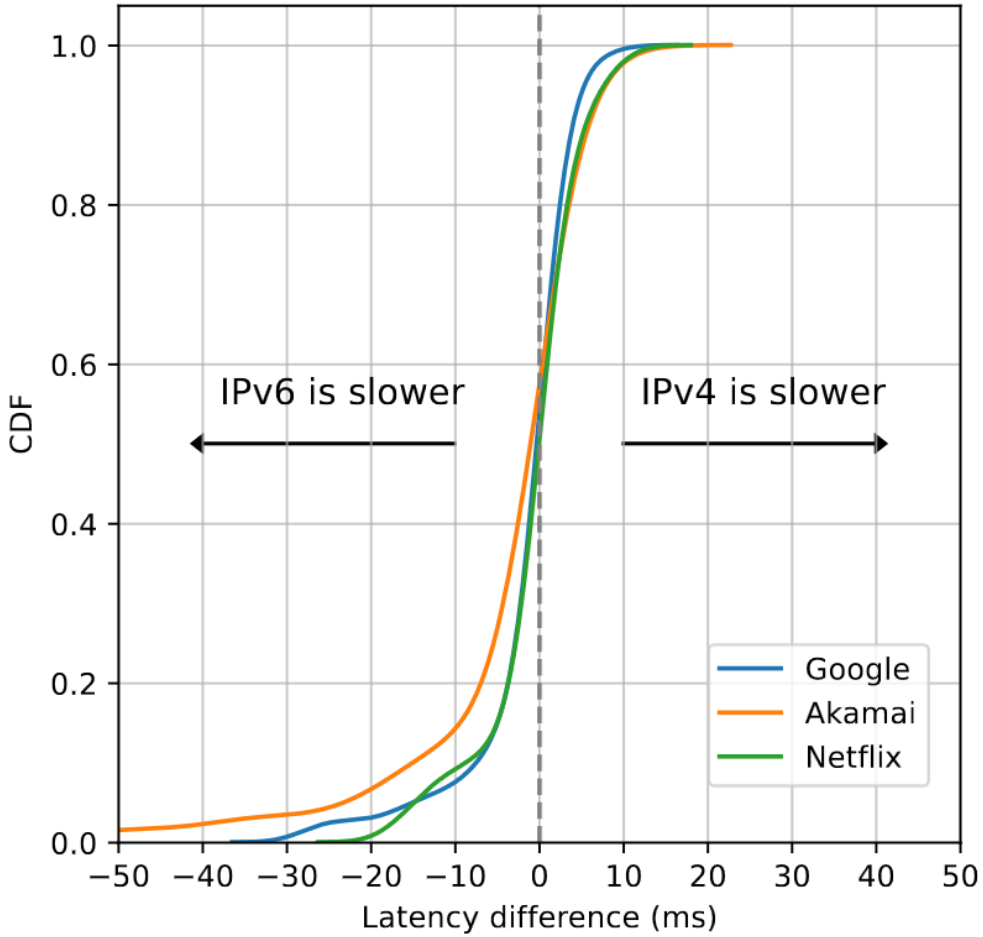

- RIPE Atlas probes ASes to IPv6, IPv4 off-nets