Towards Identifying Networks with Internet Clients Using Public Data

Weifan Jiang[†*], Tao Luo[†*], Thomas Koch[†], Yunfan Zhang[†], Ethan Katz-Bassett[†], Matt Calder[‡†]

[†] Columbia University in the City of New York

[‡] Microsoft

*: primary authors

We measured web client activity globally

These prefixes…

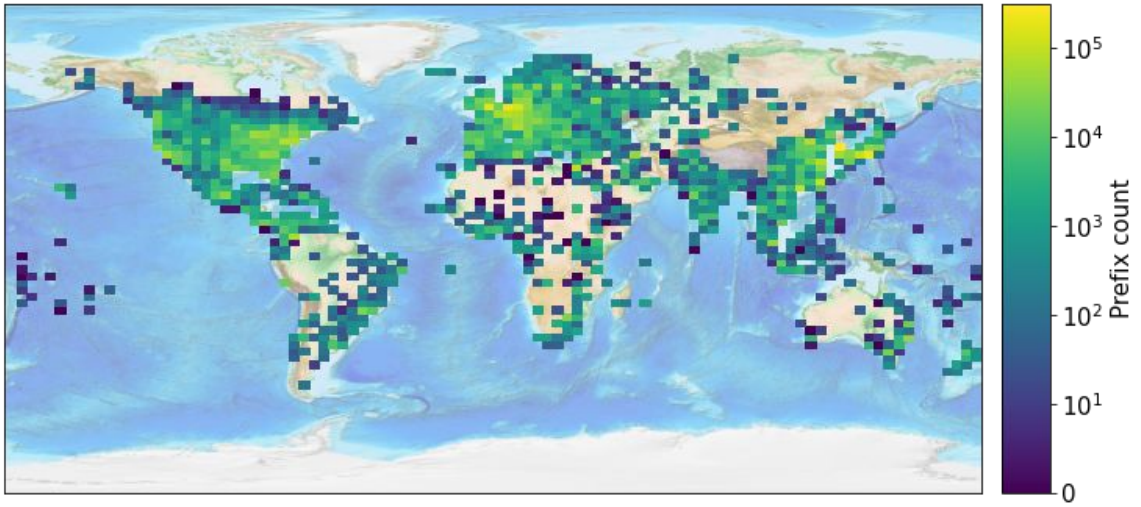

We measured web client activity globally

These prefixes…

- We built this entirely with **public** data and **replicable** methodologies

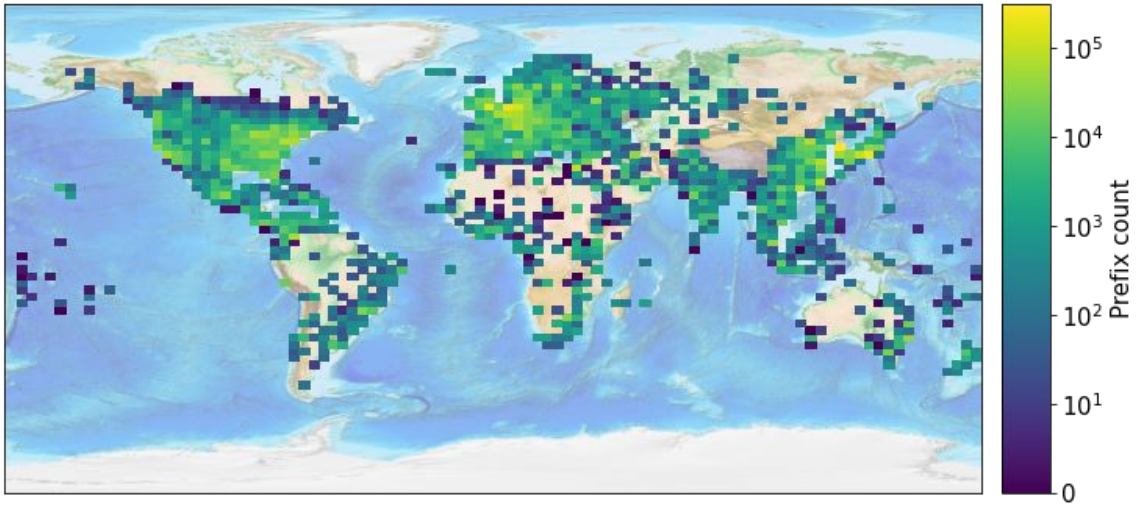

We measured web client activity globally

These prefixes…

- We built this entirely with **public** data and **replicable** methodologies
- Good coverage that rivals the set of prefixes with client activities seen by Microsoft

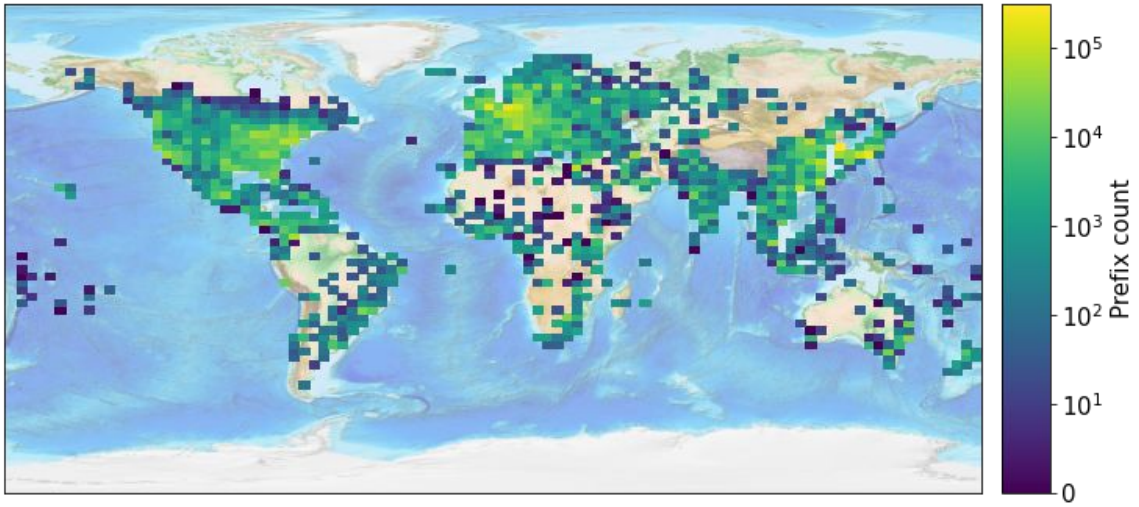

Why identifying networks with active clients?

Helps researchers to better **interpret** and **analyze** measurement results:

Why identifying networks with active clients?

Helps researchers to better **interpret** and **analyze** measurement results:

- Trinocular[1] identifies outages at prefix level. Which outages **impact clients**?

1: Quan et al. “Trinocular: Understanding Internet Reliability Through Adaptive Probing”.

Why identifying networks with active clients?

Helps researchers to better **interpret** and **analyze** measurement results:

- Trinocular[1] identifies outages at prefix level. Which outages **impact clients**?
- Geolocation databases (e.g. Maxmind[2]) are more accurate for **end-user** networks.

1: Quan et al. "Trinocular: Understanding Internet Reliability Through Adaptive Probing".
2. Maxmind GeoIP2 Databases.

Why identifying networks with active clients?

Helps researchers to better **interpret** and **analyze** measurement results:

- Trinocular[1] identifies outages at prefix level. Which outages **impact clients**?
- Geolocation databases (e.g. Maxmind[2]) are more accurate for **end-user** networks.
- ...

1: Quan et al. "Trinocular: Understanding Internet Reliability Through Adaptive Probing".
2. Maxmind GeoIP2 Databases.

Why identifying networks with active clients?

Helps researchers to better **interpret** and **analyze** measurement results:

- Trinocular[1] identifies outages at prefix level. Which outages **impact clients**?
- Geolocation databases (e.g. Maxmind[2]) are more accurate for **end-user** networks.
- ...

But we sometimes lack the data/tools to do so :(

1: Quan et al. "Trinocular: Understanding Internet Reliability Through Adaptive Probing".
2. Maxmind GeoIP2 Databases.

Limitations of existing methods

Limitations of existing methods

- Previous studies:
  - privileged data[1], out of date...
  - not client-driven[2]

1: Chiu et al. “Are We One Hop Away from a Better Internet?”.
2: Heidemann et al. “Census and Survey of the Visible Internet”.

Limitations of existing methods

- Previous studies:
  - privileged data[1], out of date...
  - not client-driven[2]
- APNIC AS Population dataset[3]: maps each AS to the number of users

1: Chiu et al. “Are We One Hop Away from a Better Internet?”.
2: Heidemann et al. “Census and Survey of the Visible Internet”.
3: APNIC AS Population dataset.

Limitations of existing methods

- Previous studies:
  - privileged data[1], out of date...
  - not client-driven[2]
- APNIC AS Population dataset[3]: maps each AS to the number of users
  - not validated (to the best of our knowledge)

1: Chiu et al. “Are We One Hop Away from a Better Internet?”.
2: Heidemann et al. “Census and Survey of the Visible Internet”.
3: APNIC AS Population dataset.

Limitations of existing methods

- Previous studies:
  - privileged data[1], out of date...
  - not client-driven[2]
- APNIC AS Population dataset[3]: maps each AS to the number of users
  - not validated (to the best of our knowledge)
  - only provides coarse AS-level granularities

1: Chiu et al. “Are We One Hop Away from a Better Internet?”.
2: Heidemann et al. “Census and Survey of the Visible Internet”.
3: APNIC AS Population dataset.

Goal

To identify networks with clients:

Goal

To identify networks with clients:

1. focus on client-driven activity

Goal

To identify networks with clients:

1. focus on client-driven activity
2. use replicable methods

Goal

To identify networks with clients:

1. focus on client-driven activity
2. use replicable methods
3. prefix-level granularity at global scale

Contributions

Contributions

1. Two **new**, **replicable** methodologies to identify prefixes hosting clients:
    - CACHE PROBING (will be covered in this talk)
    - DNS LOGS (please refer to our paper for details)

Contributions

1. Two **new**, **replicable** methodologies to identify prefixes hosting clients:
   - CACHE PROBING (will be covered in this talk)
   - DNS LOGS (please refer to our paper for details)
2. Cross-comparison with the public APNIC dataset and the privileged Microsoft data to show our methodologies achieve good global coverage

Inferring client activity by DNS cache snooping

Inferring client activity by DNS cache snooping

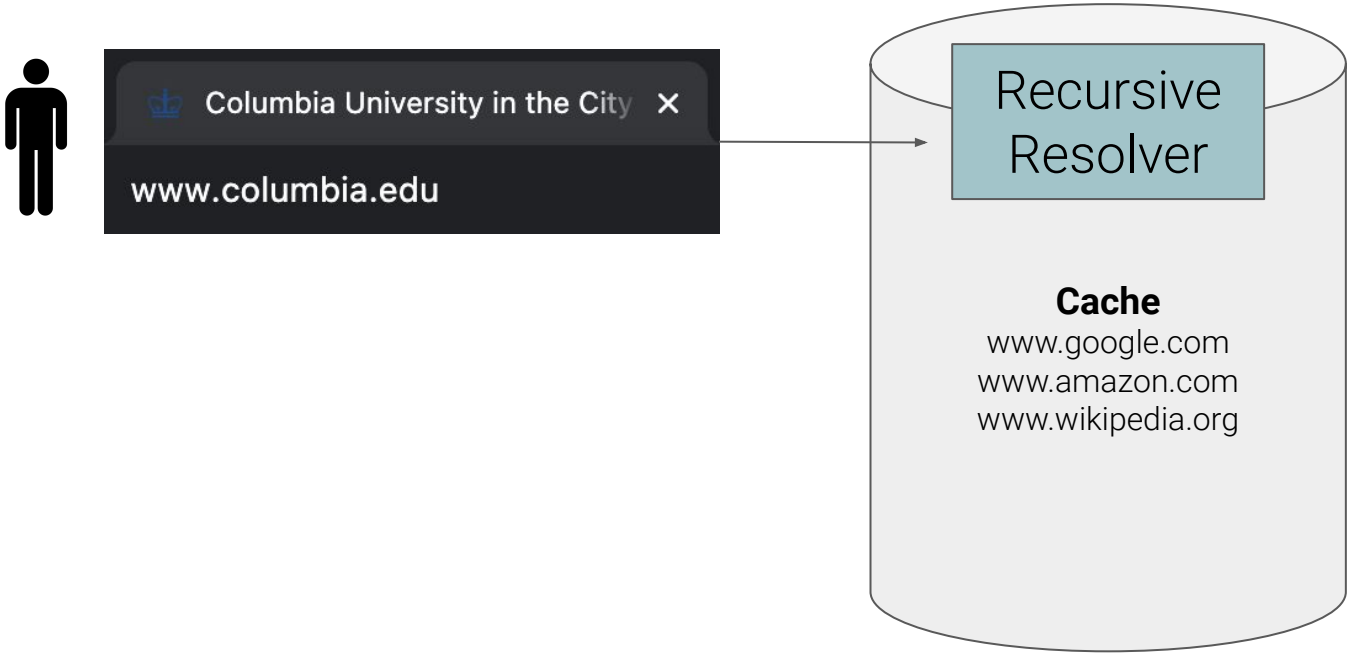

Inferring client activity by DNS cache snooping

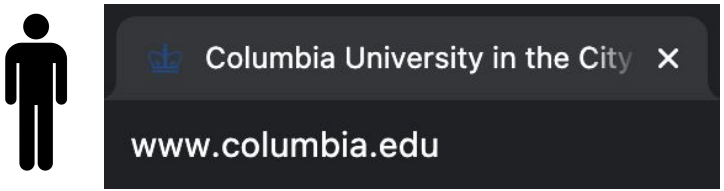

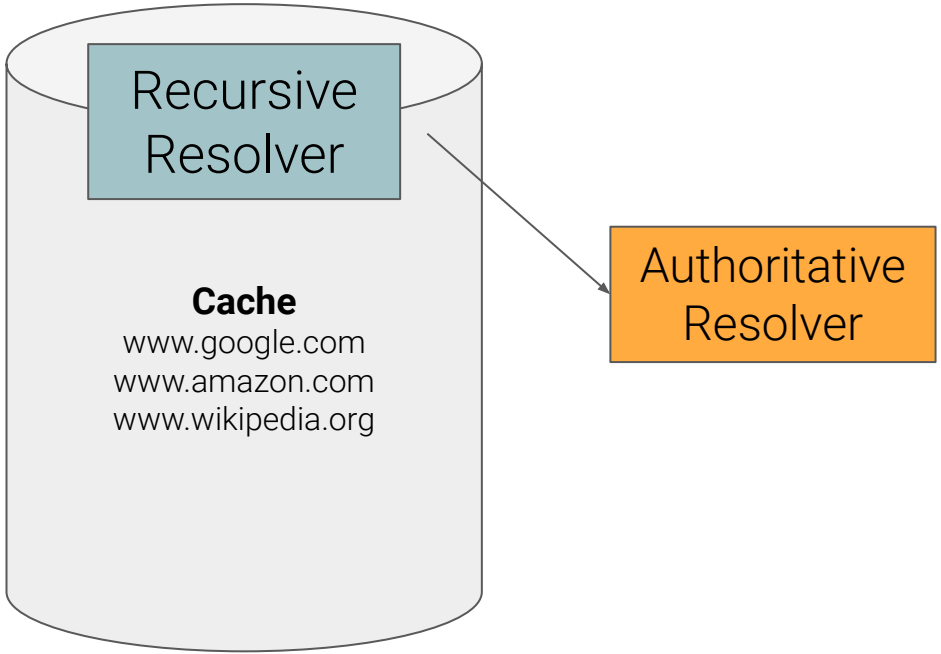

Inferring client activity by DNS cache snooping

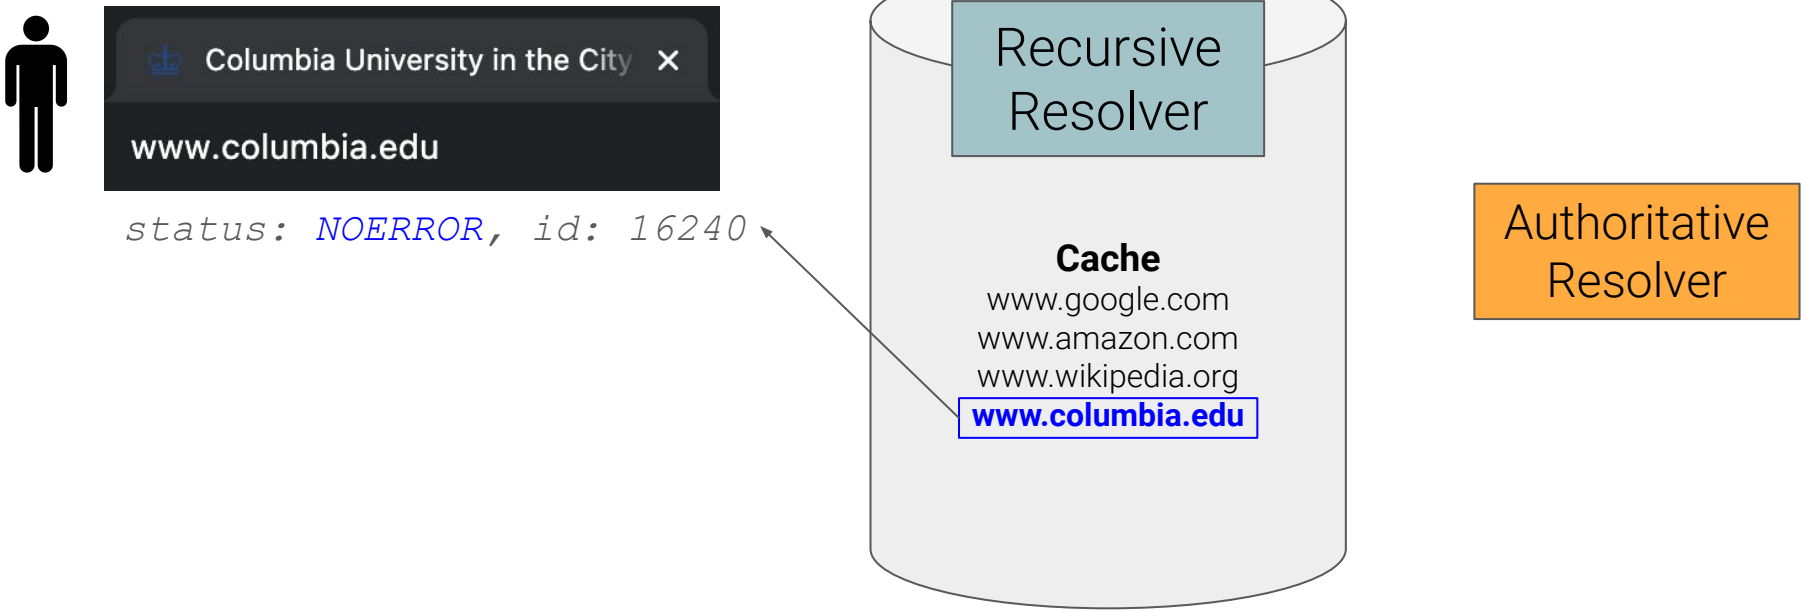

Inferring client activity by DNS cache snooping

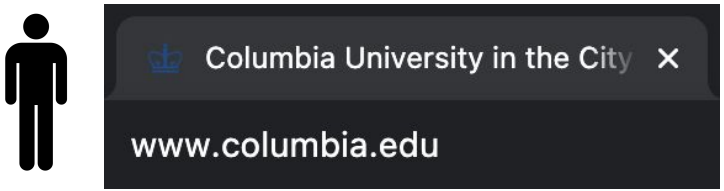

status: NOERROR, id: 16240

```
dig @8.8.8.8
www.columbia.edu +norecurse
```

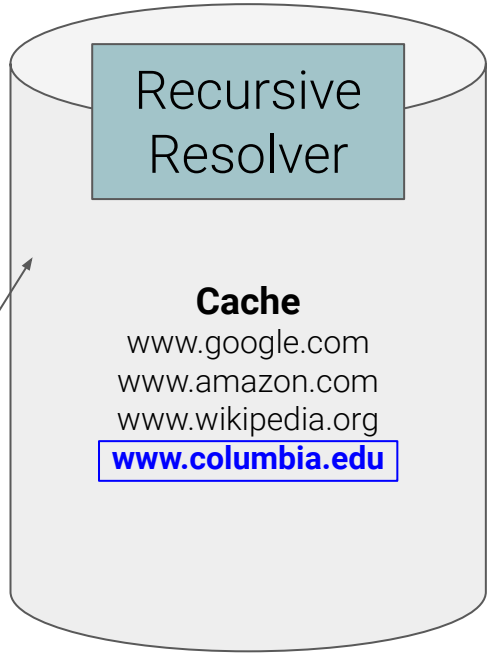

Authoritative Resolver

Inferring client activity by DNS cache snooping

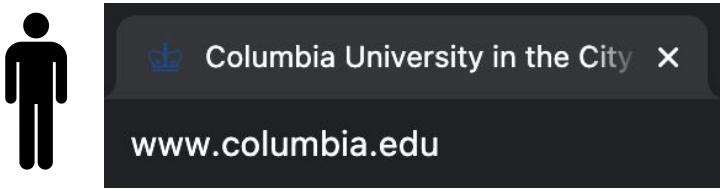

status: NOERROR, id: 16240

```
dig @8.8.8.8
www.columbia.edu +norecurse
```

status: NOERROR, id: 16625

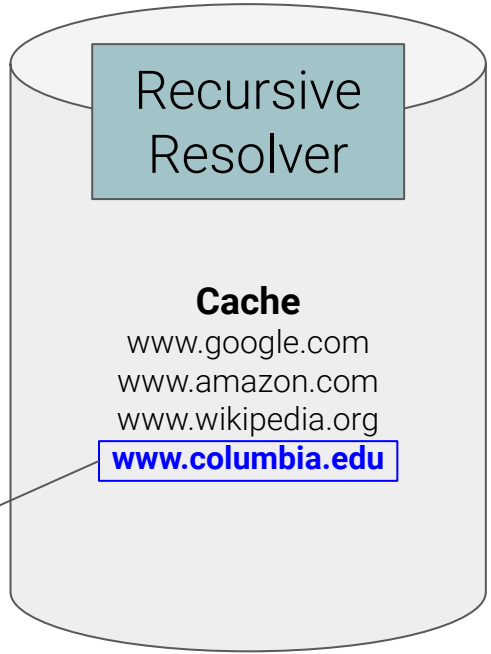

Authoritative Resolver

Inferring client activity by DNS cache snooping

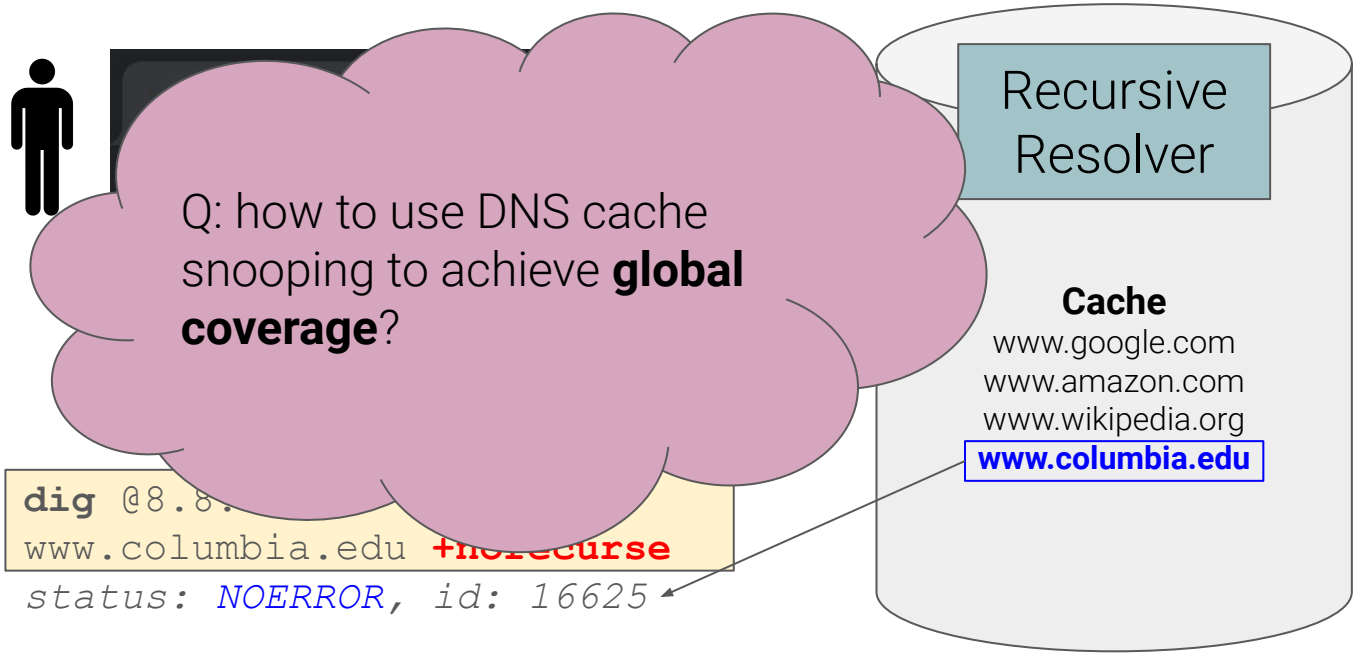

Inferring client activity by DNS cache snooping

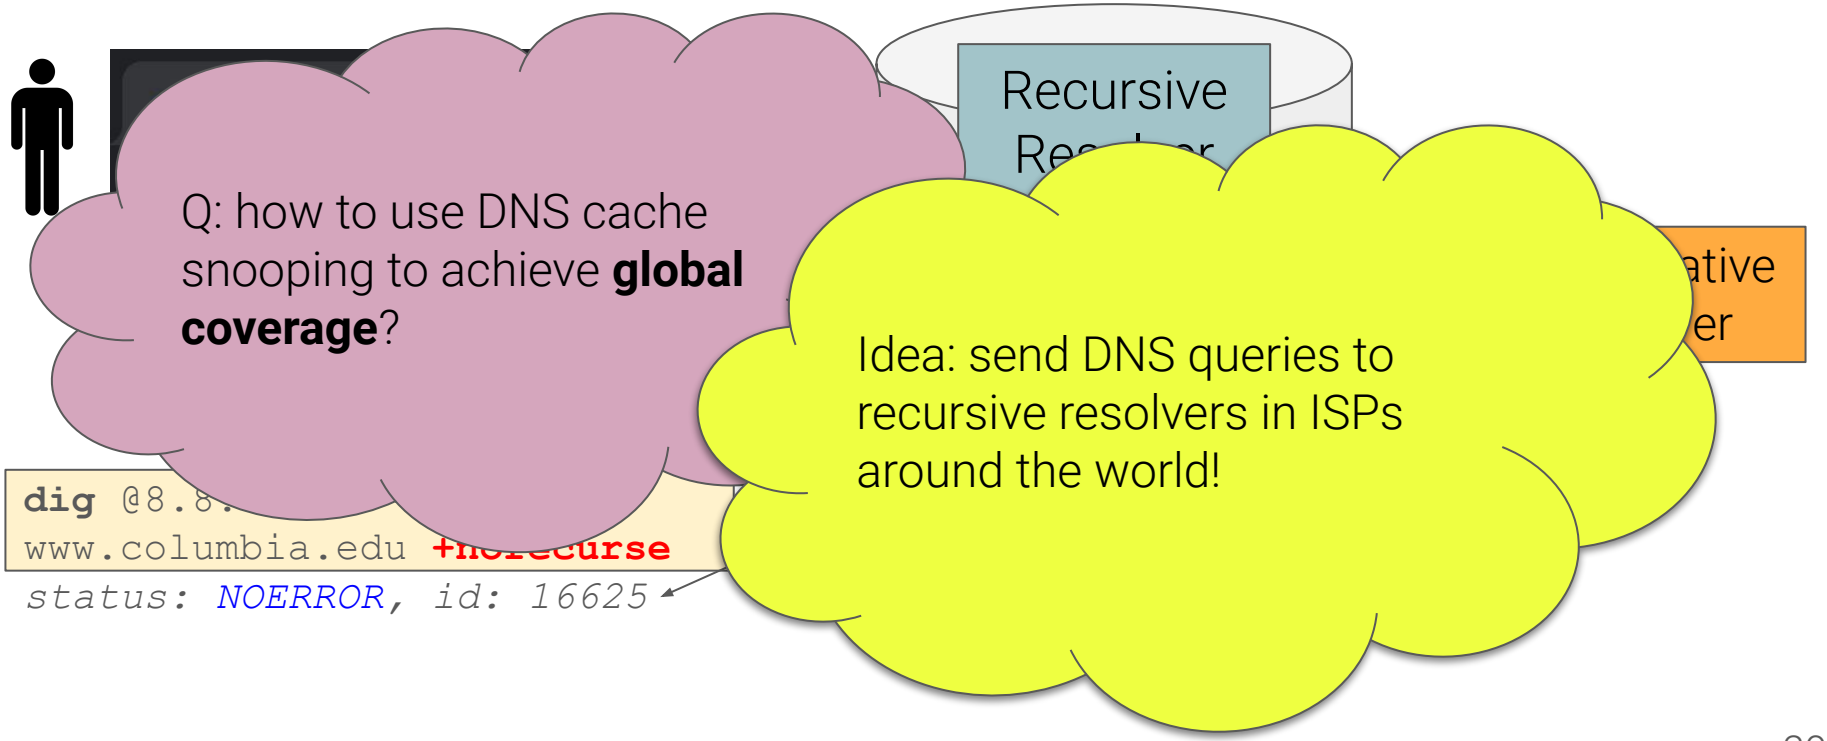

Inferring client activity by DNS cache snooping

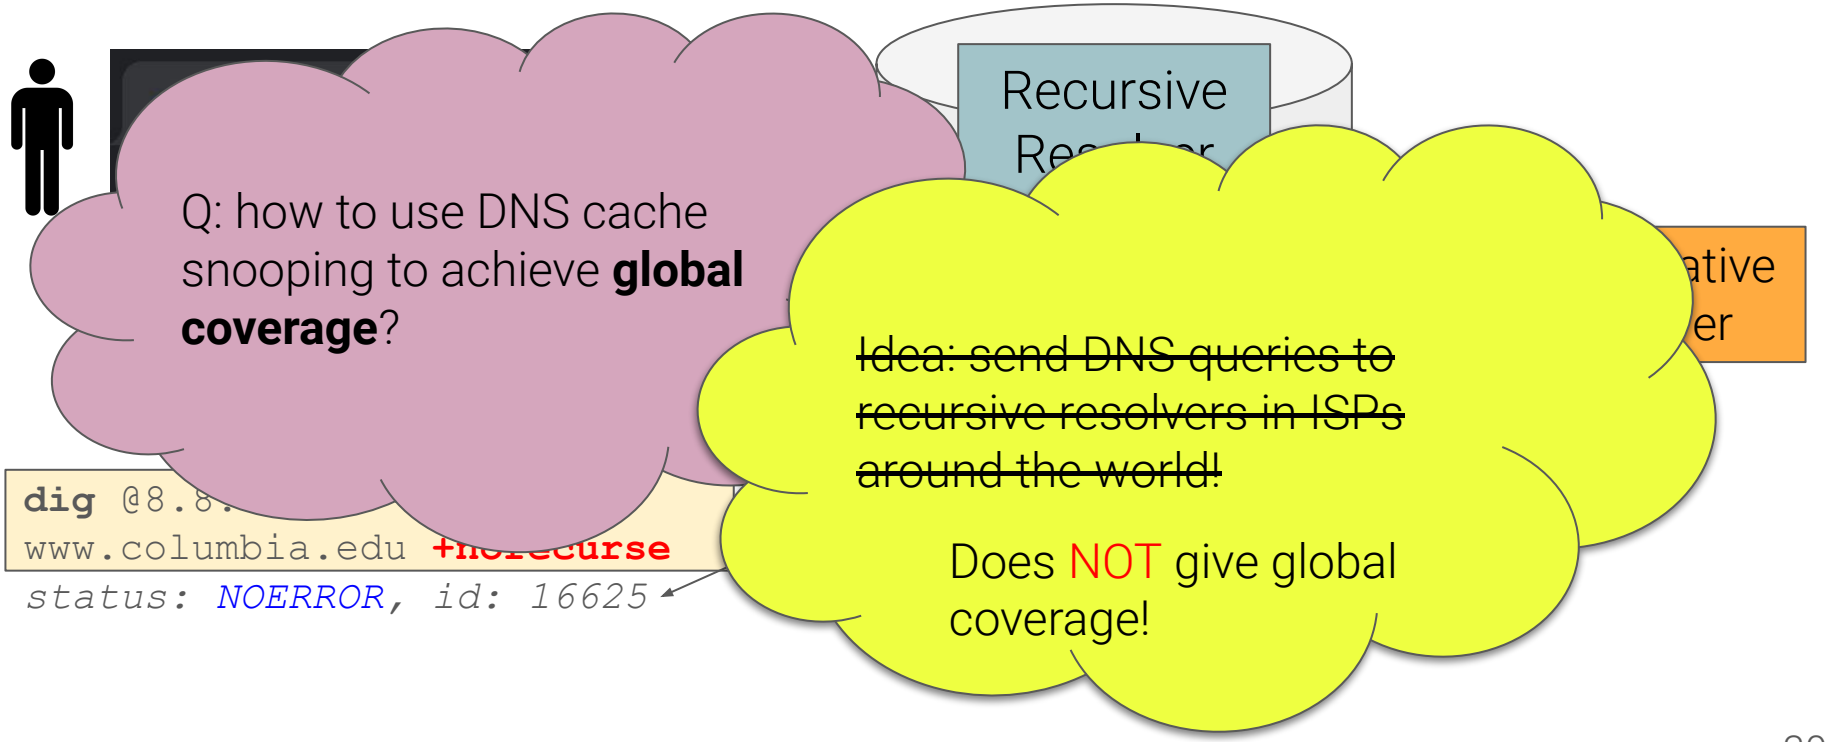


Google
Public DNS

EDNS0 Client Subnet (ECS)

EDNS0 Client Subnet (ECS)

- Allows recursive resolver to include a prefix as part of the query

EDNS0 Client Subnet (ECS)

- Allows recursive resolver to include a prefix as part of the query
  - enables authoritative to map clients based on prefixes

EDNS0 Client Subnet (ECS)

- Allows recursive resolver to include a prefix as part of the query
  - enables authoritative to map clients based on prefixes
  - if client specifies an ECS prefix, Google Public DNS will use that instead of client's own prefix

EDNS0 Client Subnet (ECS)

- Allows recursive resolver to include a prefix as part of the query
  - enables authoritative to map clients based on prefixes
  - if client specifies an ECS prefix, Google Public DNS will use that instead of client's own prefix
- Google Public DNS maintains separate cache entry per client prefix

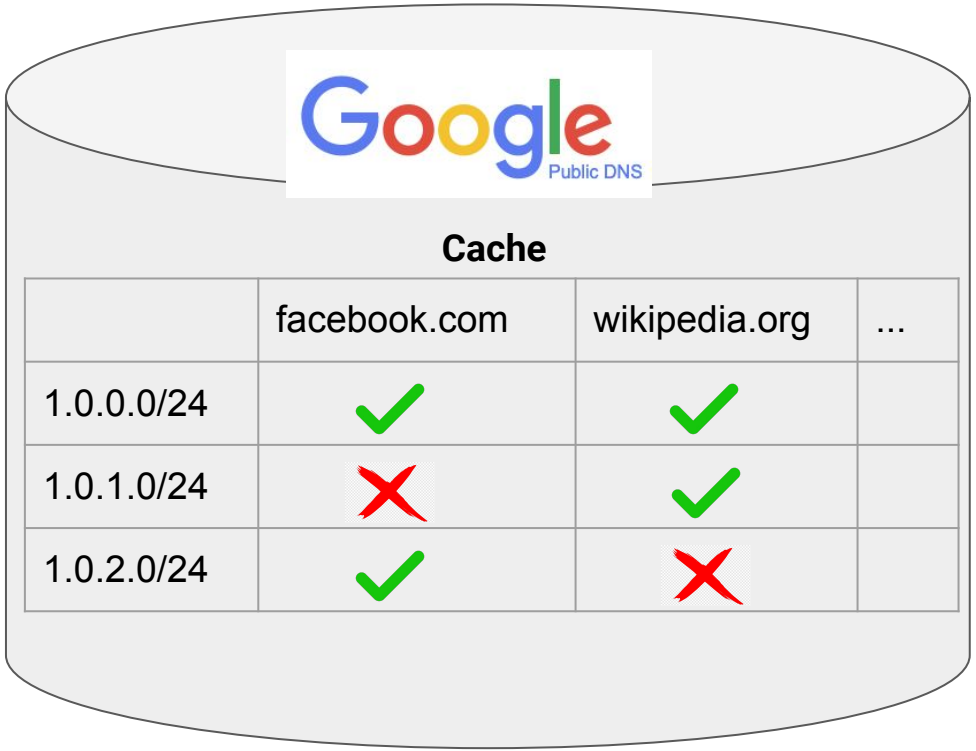

EDNS0 Client Subnet (ECS)

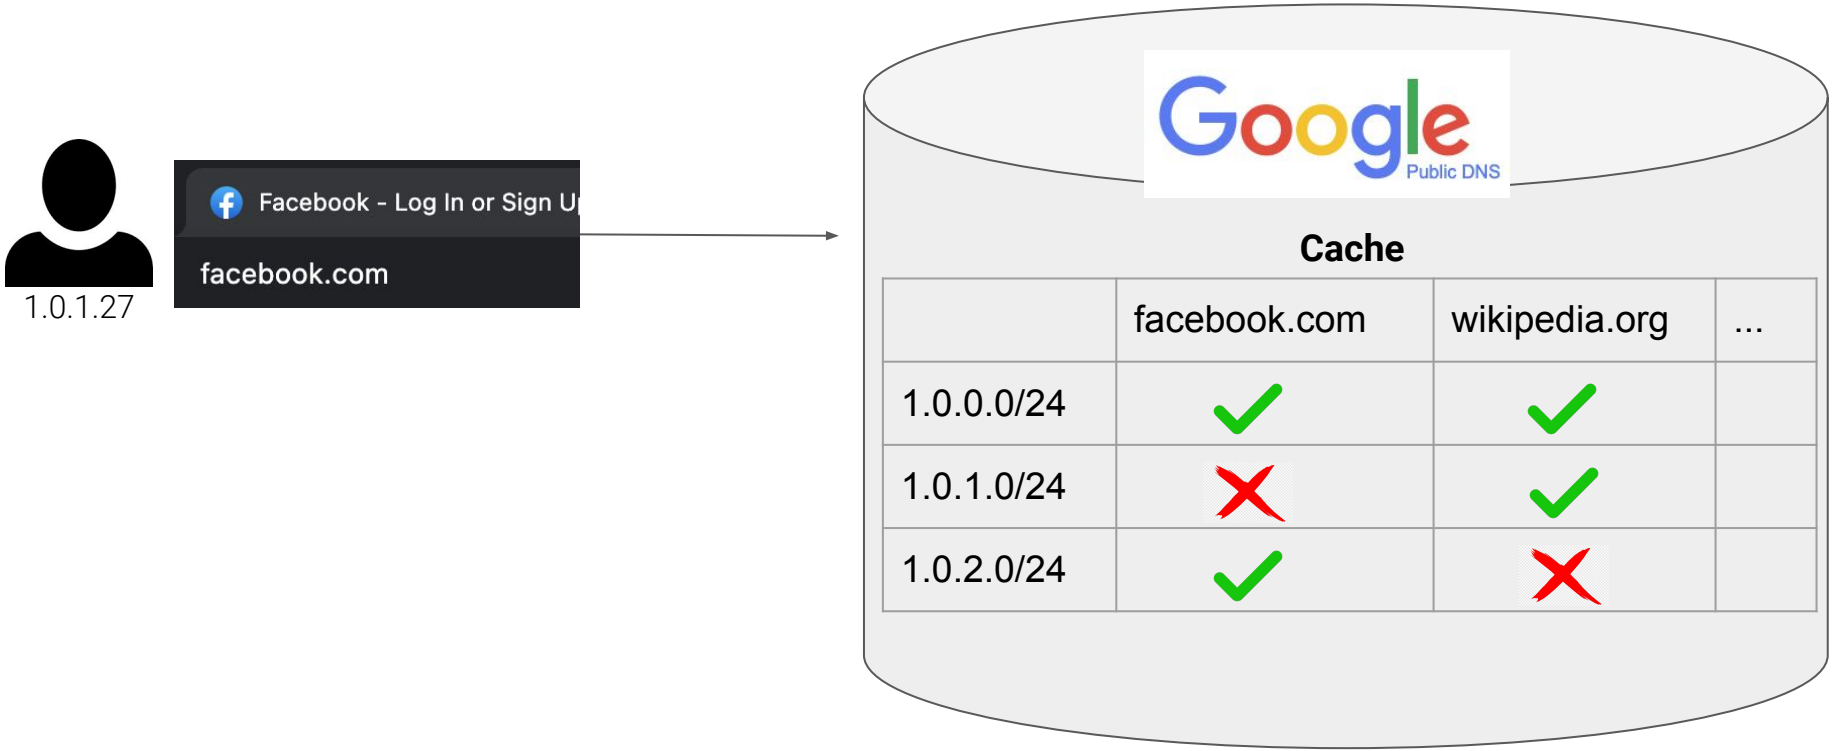

EDNS0 Client Subnet (ECS)

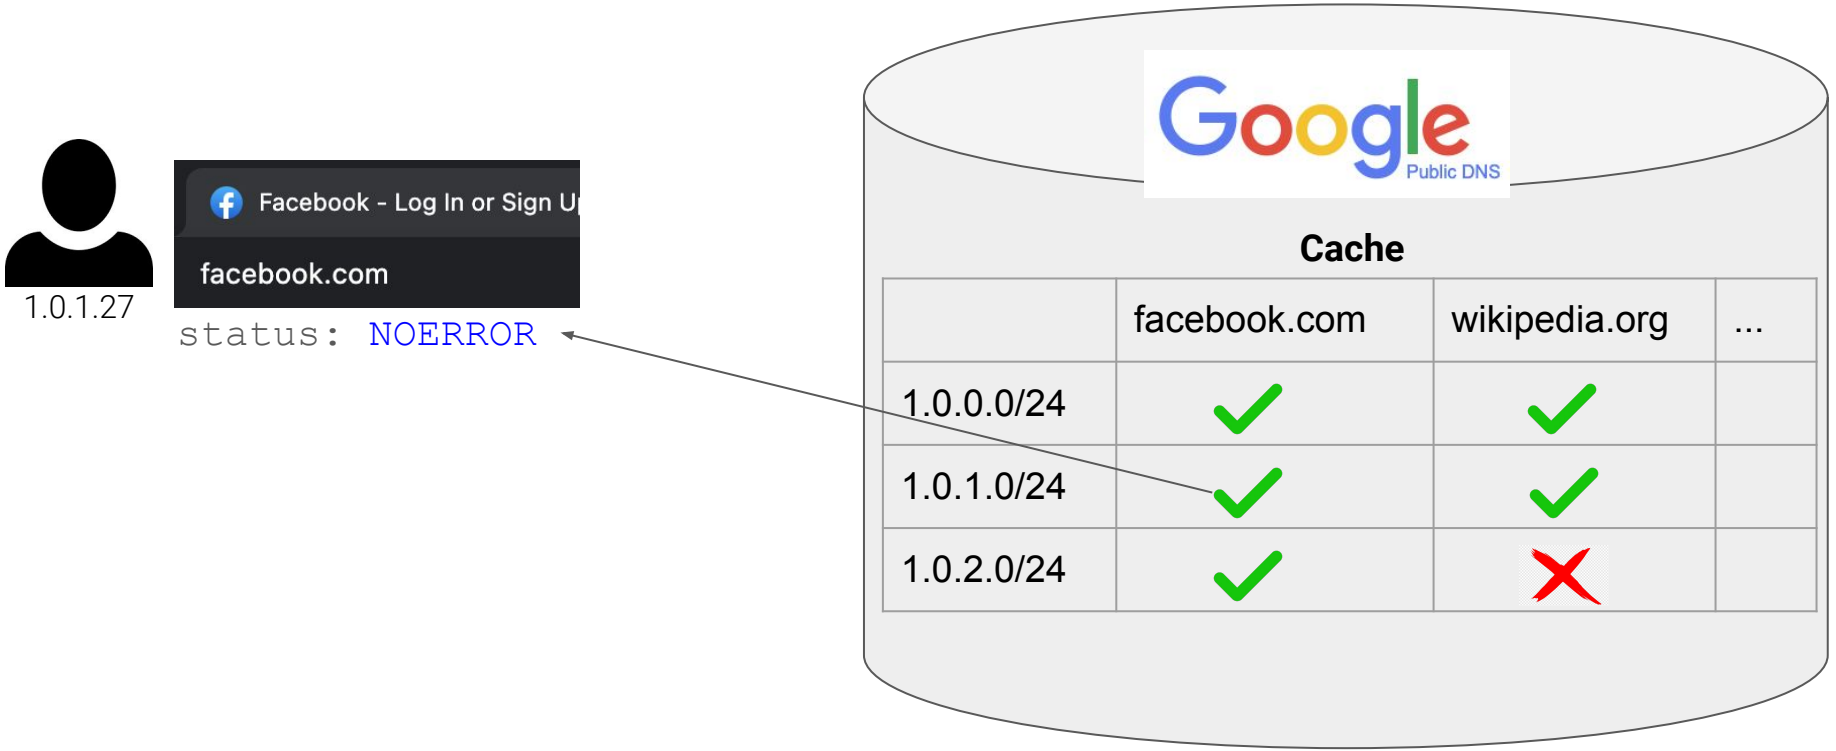

EDNS0 Client Subnet (ECS)

1.0.1.27

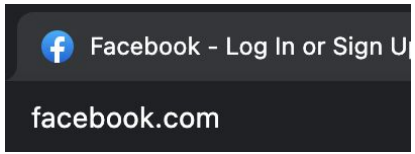

`status: NOERROR`

```
dig @8.8.8.8 facebook.com
+norecurse
+subnet=1.0.1.0/24
```

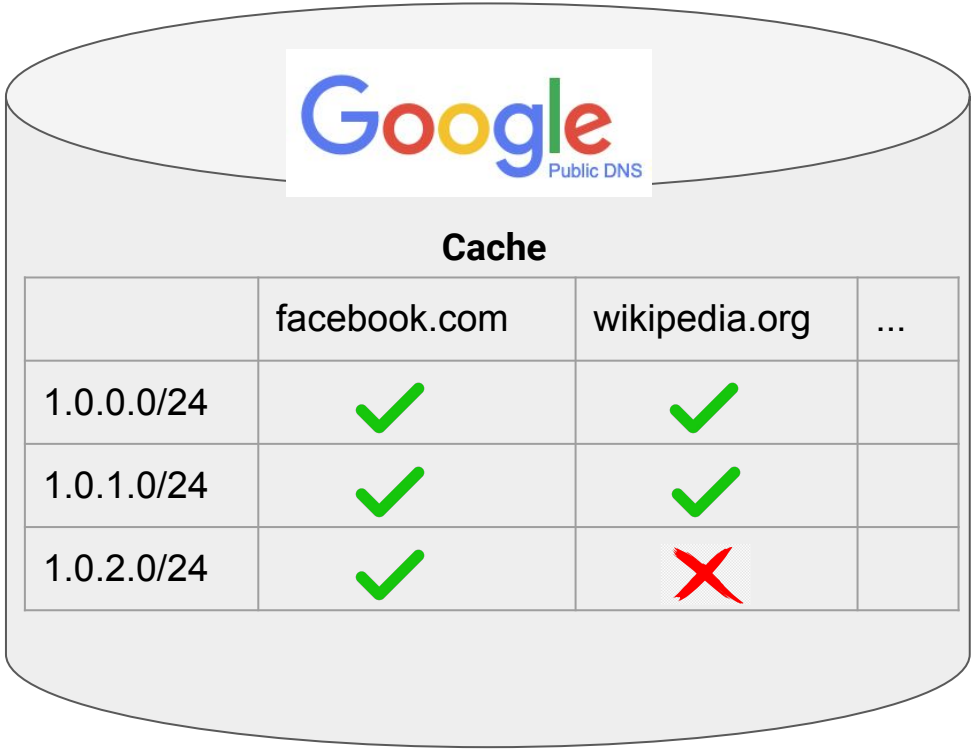

EDNS0 Client Subnet (ECS)

1.0.1.27

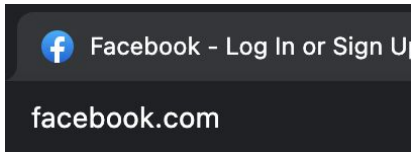


```
status: NOERROR
```

```
dig @8.8.8.8 facebook.com
+norecurse
+subnet=1.0.1.0/24
```

```
status: NOERROR
```

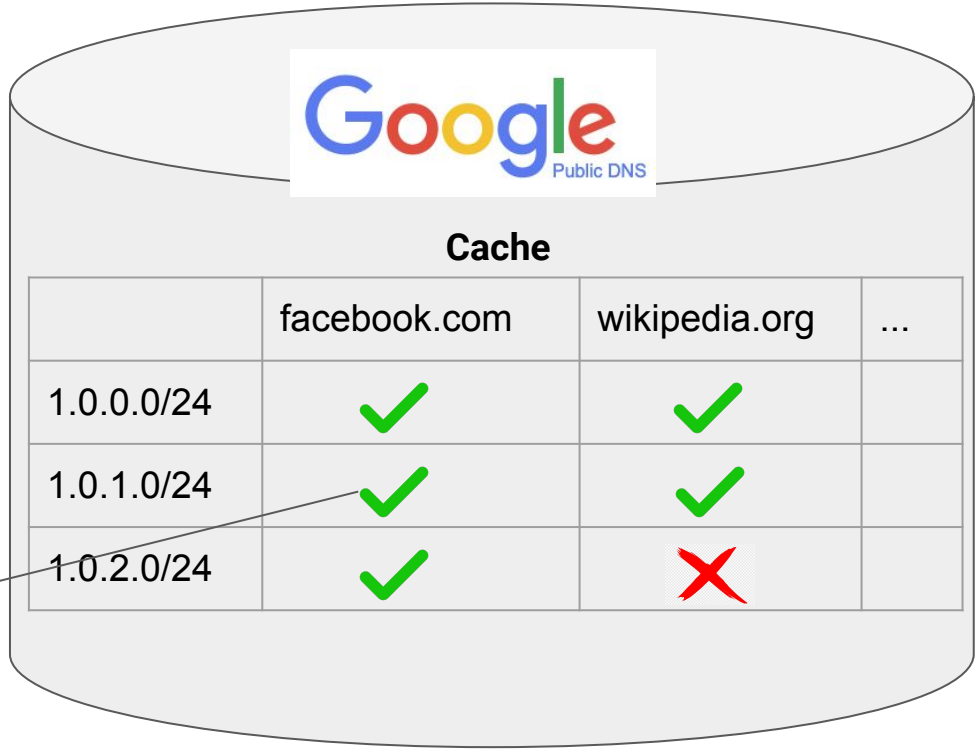

EDNS0 Client Subnet (ECS)

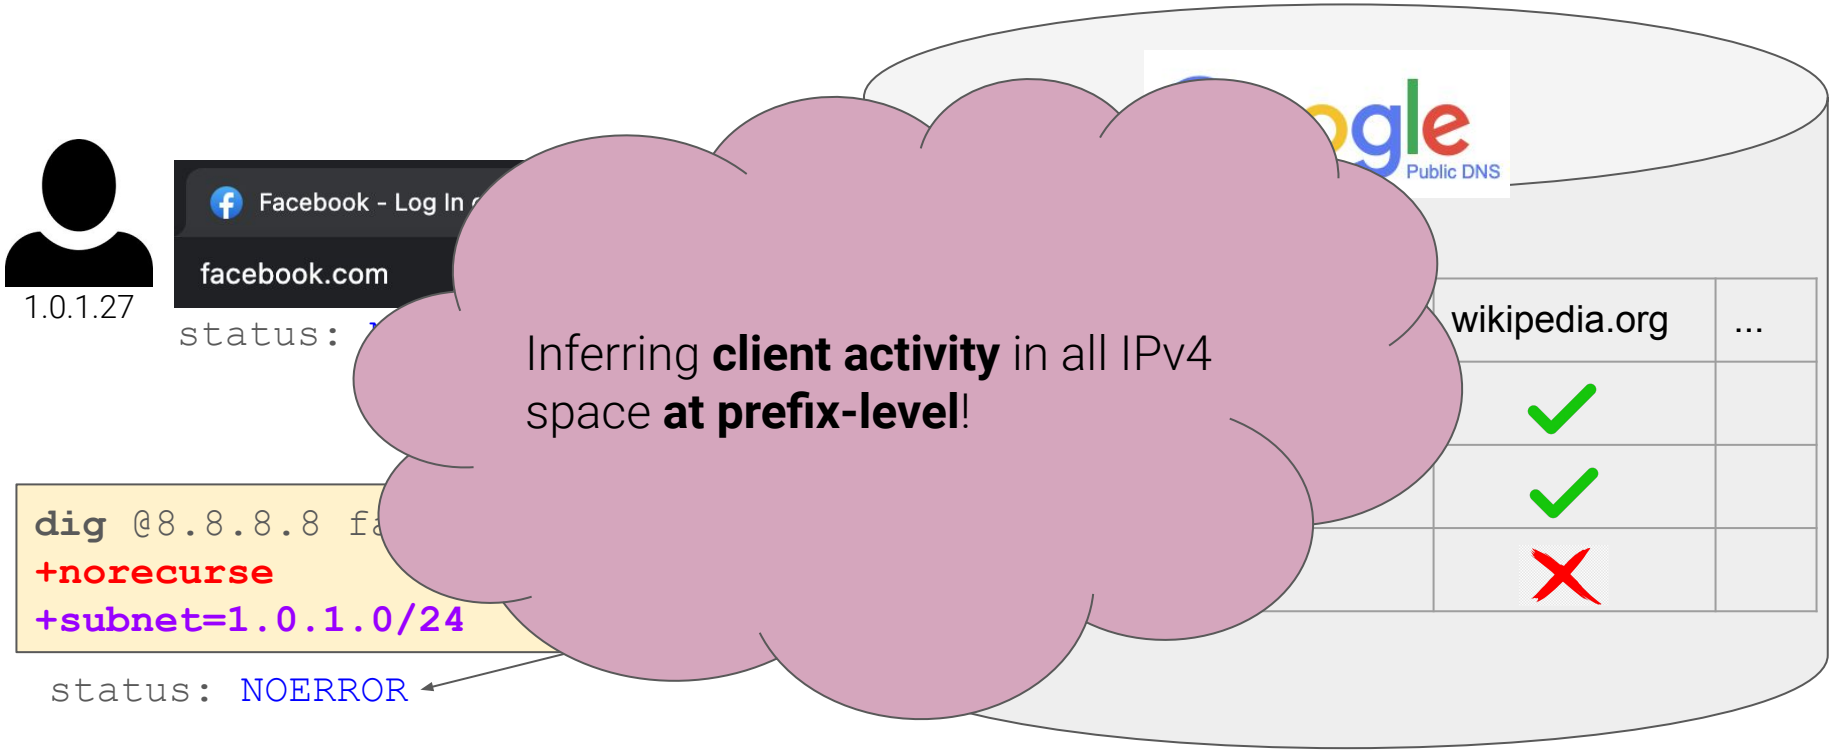

Geo-distributed Measurements

- Our goal: global coverage

Geo-distributed Measurements

- Our goal: global coverage
- Google Public DNS structure:
  - multiple PoPs with independent caches

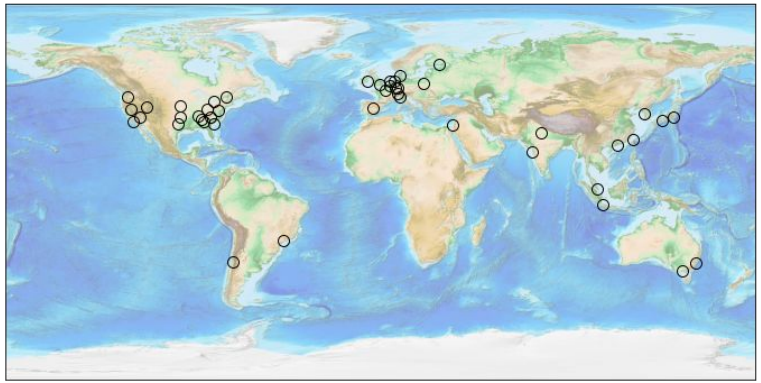

Geo-distributed Measurements

- Our goal: global coverage
- Google Public DNS structure:
  - multiple PoPs with independent caches
  - use anycast to direct queries to PoPs

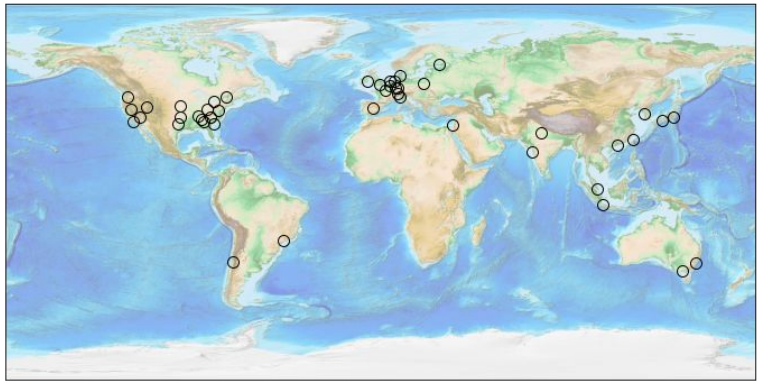

Geo-distributed Measurements

- Our goal: global coverage
- Google Public DNS structure:
  - multiple PoPs with independent caches
  - use anycast to direct queries to PoPs
- Strategy: query from vantage points around globe

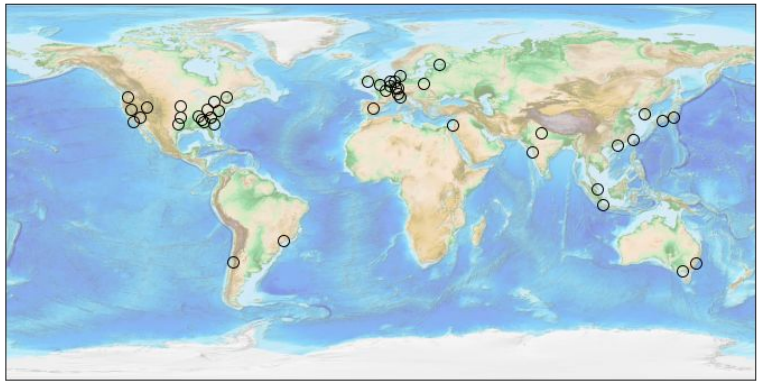

Geo-distributed Measurements

- Our goal: global coverage
- Google Public DNS structure:
  - multiple PoPs with independent caches
  - use anycast to direct queries to PoPs
- Strategy: query from vantage points around globe
- We hit 22 PoPs = 95% of all Google public DNS traffic to Microsoft

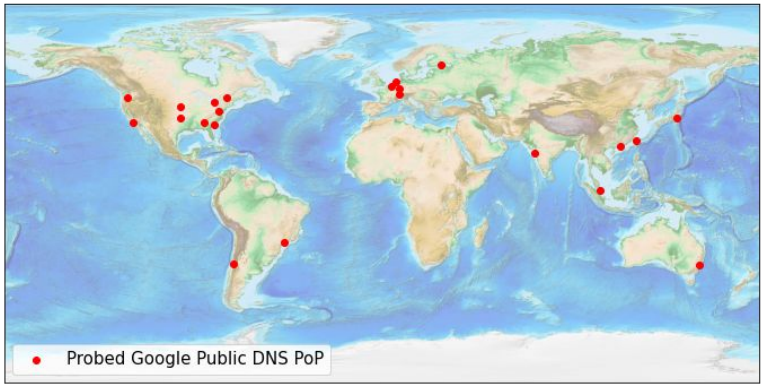

Geo-distributed Measurements

- Our goal: global coverage
- Google Public DNS structure:
  - multiple PoPs with independent caches
  - use anycast to direct queries to PoPs
- Strategy: query from vantage points around globe
- We hit 22 PoPs = 95% of all Google public DNS traffic to Microsoft

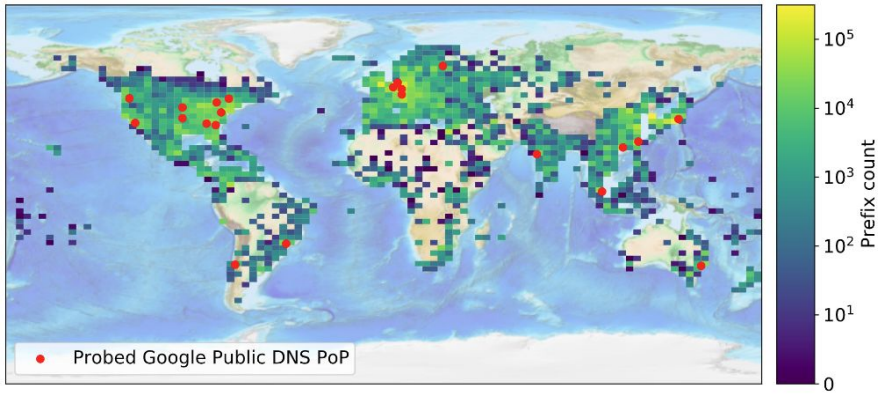

Global coverage of IPv4 prefixes!

Validation of methodologies

Validation of methodologies

- % of Microsoft traffic volume from identified networks with client activity?

Validation of methodologies

- % of Microsoft traffic volume from identified networks with client activity?

Granularity	CACHE PROBING and DNS LOGS	APNIC AS Population Dataset
AS	**98.8%**	92%

Validation of methodologies

- % of Microsoft traffic volume from identified networks with client activity?

Granularity	CACHE PROBING and DNS LOGS	APNIC AS Population Dataset
AS	**98.8%**	92%
Prefix	**95.2%**	N/A

Validation of methodologies

- % of Microsoft traffic volume from identified networks with client activity?

Granularity	CACHE PROBING and DNS LOGS	APNIC AS Population Dataset
AS	**98.8%**	92%
Prefix	**95.2%**	N/A

- **Less than 1%** of active prefixes identified by us do not contact Microsoft

Validation of methodologies

- % of Microsoft traffic volume from identified networks with client activity?

Granularity	CACHE PROBING and DNS LOGS	APNIC AS Population Dataset
AS	**98.8%**	92%
Prefix	**95.2%**	N/A

- **Less than 1%** of active prefixes identified by us do not contact Microsoft

Our methodologies produce a **good approximation** of privileged dataset

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients**

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies,

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities,

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities, and with **global-scale** coverage

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities, and with **global-scale** coverage
- **Good approximation** of Microsoft's privileged data

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities, and with **global-scale** coverage.
- **Good approximation** of Microsoft's privileged data
- We are happy to share the data/code upon request

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities, and with **global-scale** coverage
- **Good approximation** of Microsoft's privileged data
- We are happy to share the data/code upon request

Future work: what is the relative activity levels among different prefixes?

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities, and with **global-scale** coverage.
- **Good approximation** of Microsoft's privileged data
- We are happy to share the data/code upon request

Future work: what is the relative activity levels among different prefixes?
Starting point: how often do prefixes get cache hits

Conclusion and future work

- Identifying networks with clients helps researchers interpret and analyze data
- We **identify networks with clients** using **only public** data and **replicable** methodologies, at **prefix-level** granularities, and with **global-scale** coverage.
- **Good approximation** of Microsoft's privileged data
- We are happy to share the data/code upon request

Future work: what is the relative activity levels among different prefixes?
Starting point: how often do prefixes get cache hits

Towards a traffic map of the Internet (HotNets'21)

Thomas Koch, Weifan Jiang, Tao Luo, Petros Gigis, Yunfan Zhang, Kevin Vermeulen, Emile Aben, Matt Calder, Ethan Katz-Bassett, Lefteris Manassakis, Georgios Smaragdakis, Narseo Vallina-Rodriguez

What Google Public DNS PoPs do we hit?

- `dig @8.8.8.8 o-o.myaddr.l.google.com -t TXT` tells us the IP address of the particular Google Public DNS PoP reached.
- Google publishes the IP range and the closest airport code for each PoP[1].

```
$ dig @8.8.8.8 o-o.myaddr.l.google.com -t TXT
…
;; ANSWER SECTION:
o-o.myaddr.l.google.com. 60    IN    TXT  "74.125.18.67"
```

1: Google Public DNS: Frequently Asked Questions.

Assigning prefixes to vantage points

- We **randomly selected** 78,637 prefixes and **queried them at all vantage points**.
- For each vantage point, we compute the **geographical radius** that would include **90%** of all cache-hit prefixes in the sample.
- We **use the 90% radius** to assign all 15,527,909 public /24 prefixes to vantage points to reduce probing overhead.
- For prefixes not assigned to any PoPs with above heuristics, we **assign it to the closest 2 PoPs**.

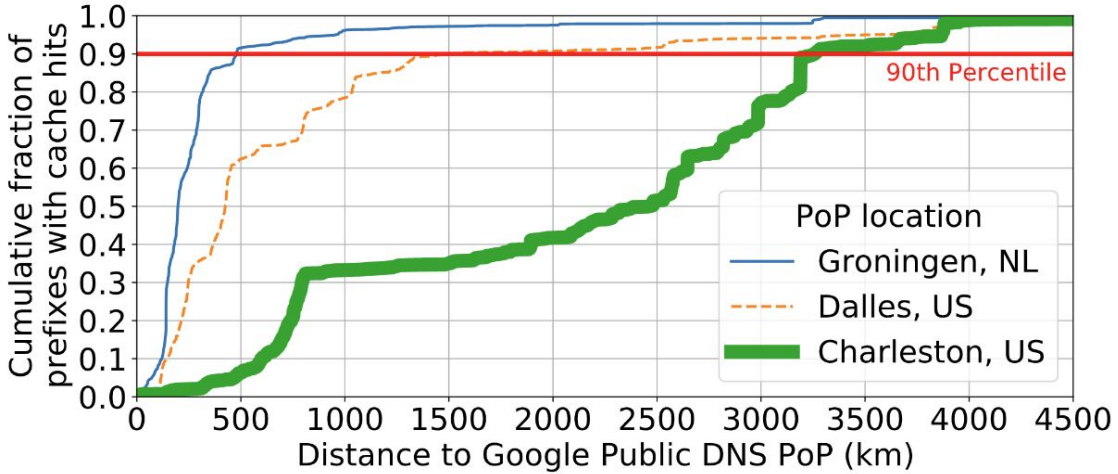

Domain names

Domain	Alexa Topsite Rank - Global (as of Sep 2021)
Google	1
YouTube	2
facebook	7
Wikipedia	13

Methodology 2: DNS LOGS

- Chromium detects DNS interception by **querying for random strings** of 7-15 lowercase letters:
  - when browser starts, and
  - when device's IP address or DNS configuration changes
- These queries **should not result in cache hit** at recursive resolvers (due to lacking a valid TLD such as ".com"), so the queries **go to a DNS root server**.
- We identify chromium queries using a heuristic that **randomly generated strings should have few collisions** across all roots in one day (based on empirical study, 7 is a good threshold).
- We look for queries matching this pattern **in the DITL traces**. Those queries contain the **IP address of the querier**, which is generally the recursive resolver used by the Chromium client.

Coverage Analysis

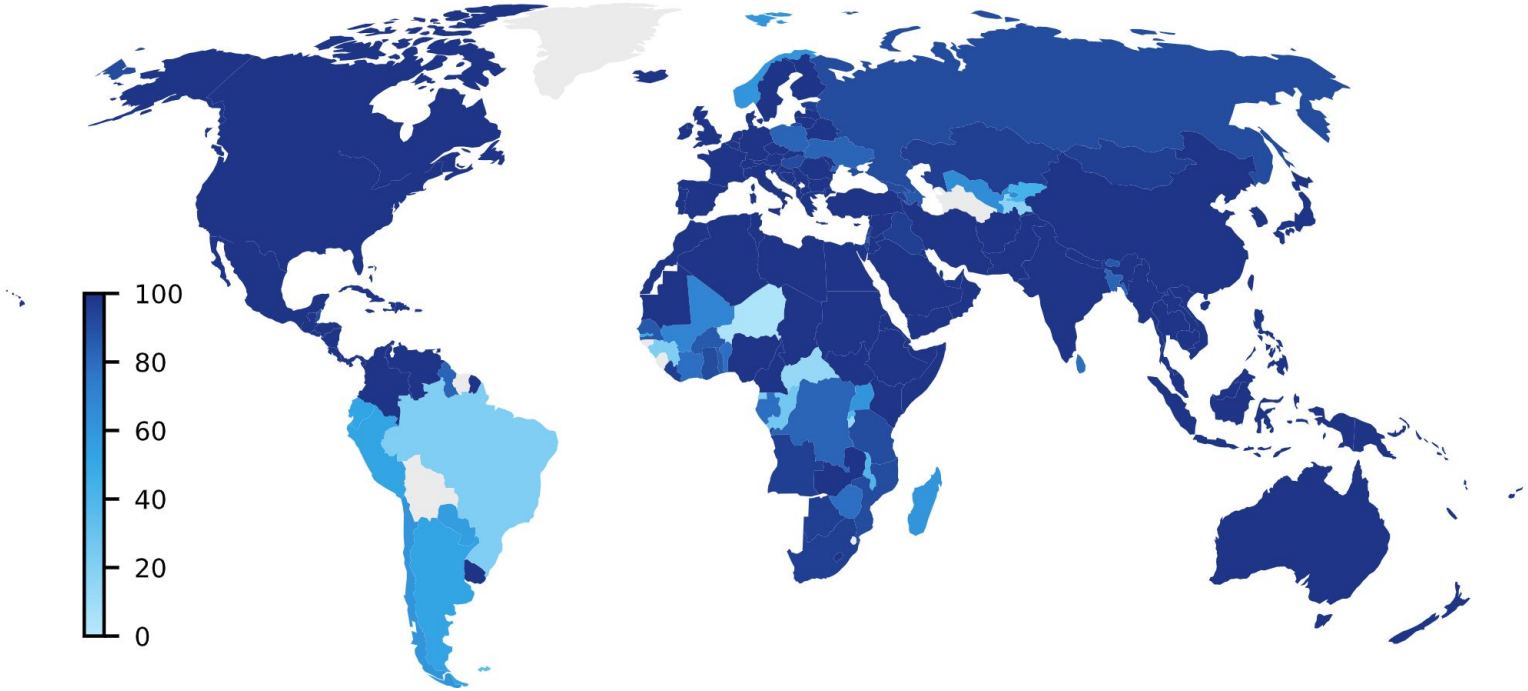

Colorscale: percentage of country's APNIC Internet user population seen by CACHE PROBING.

Validate with Microsoft data

What our methodologies saw:

- The ASes found by us are responsible for most of the Microsoft traffic.
- Implication: the ASes missed by us are very small.

	# of ASes seen by MSFT	Volume of traffic to MSFT
CACHE PROBING	55.5%	**94.9%**
DNS LOGS	59.9%	**97.4%**
CACHE PROBING U DNS LOGS	77.2%	**98.8%**

Result analysis

What activity did APNIC miss, but seen by us?

- ASDB[1] (from IMC'21!) tells us what categories an AS belongs to.
- Out of the ASes detected by our methods but missed by APNIC:
  - 10,998 (**39.5%**) are **Internet Service Providers (ISPs)**
  - Outside of ISPs, 4,823 (**17.4%**) are **hosting/cloud providers** → may reflect **non-human web clients**
  - Outside of ISPs, 1,723 (**6.2%**) are **schools** → likely **host human users**

1: Ziv et al. "ASdb: A System for Classifying Owners of Autonomous Systems".