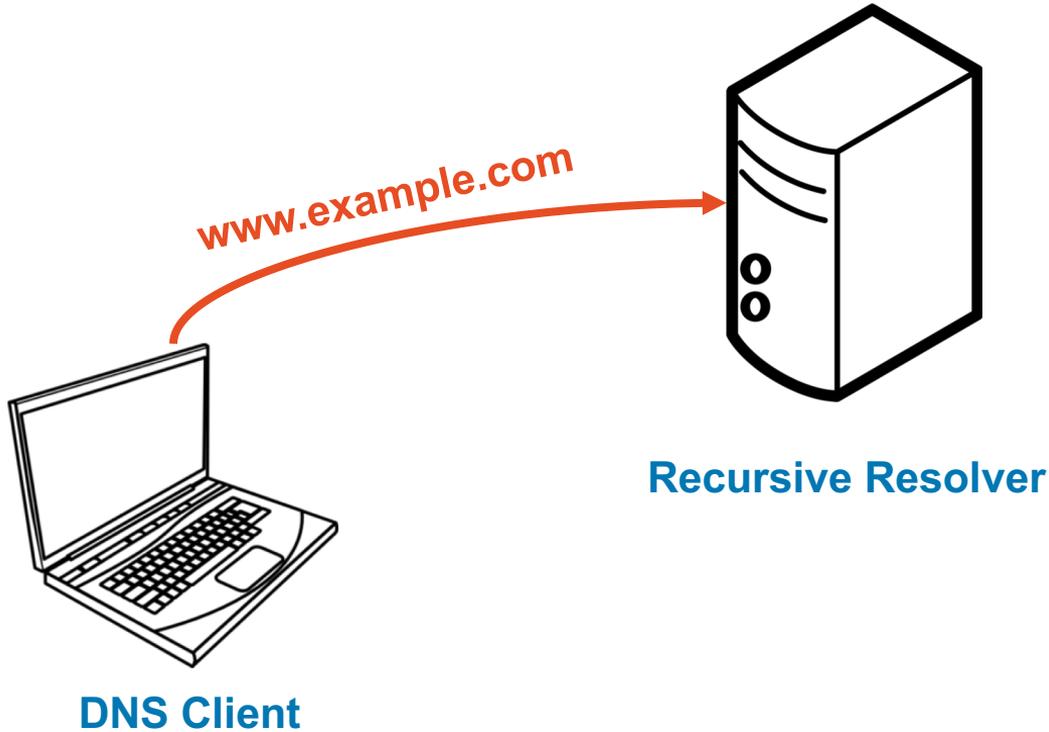


# Characterization of Collaborative Resolution in Recursive DNS Resolvers

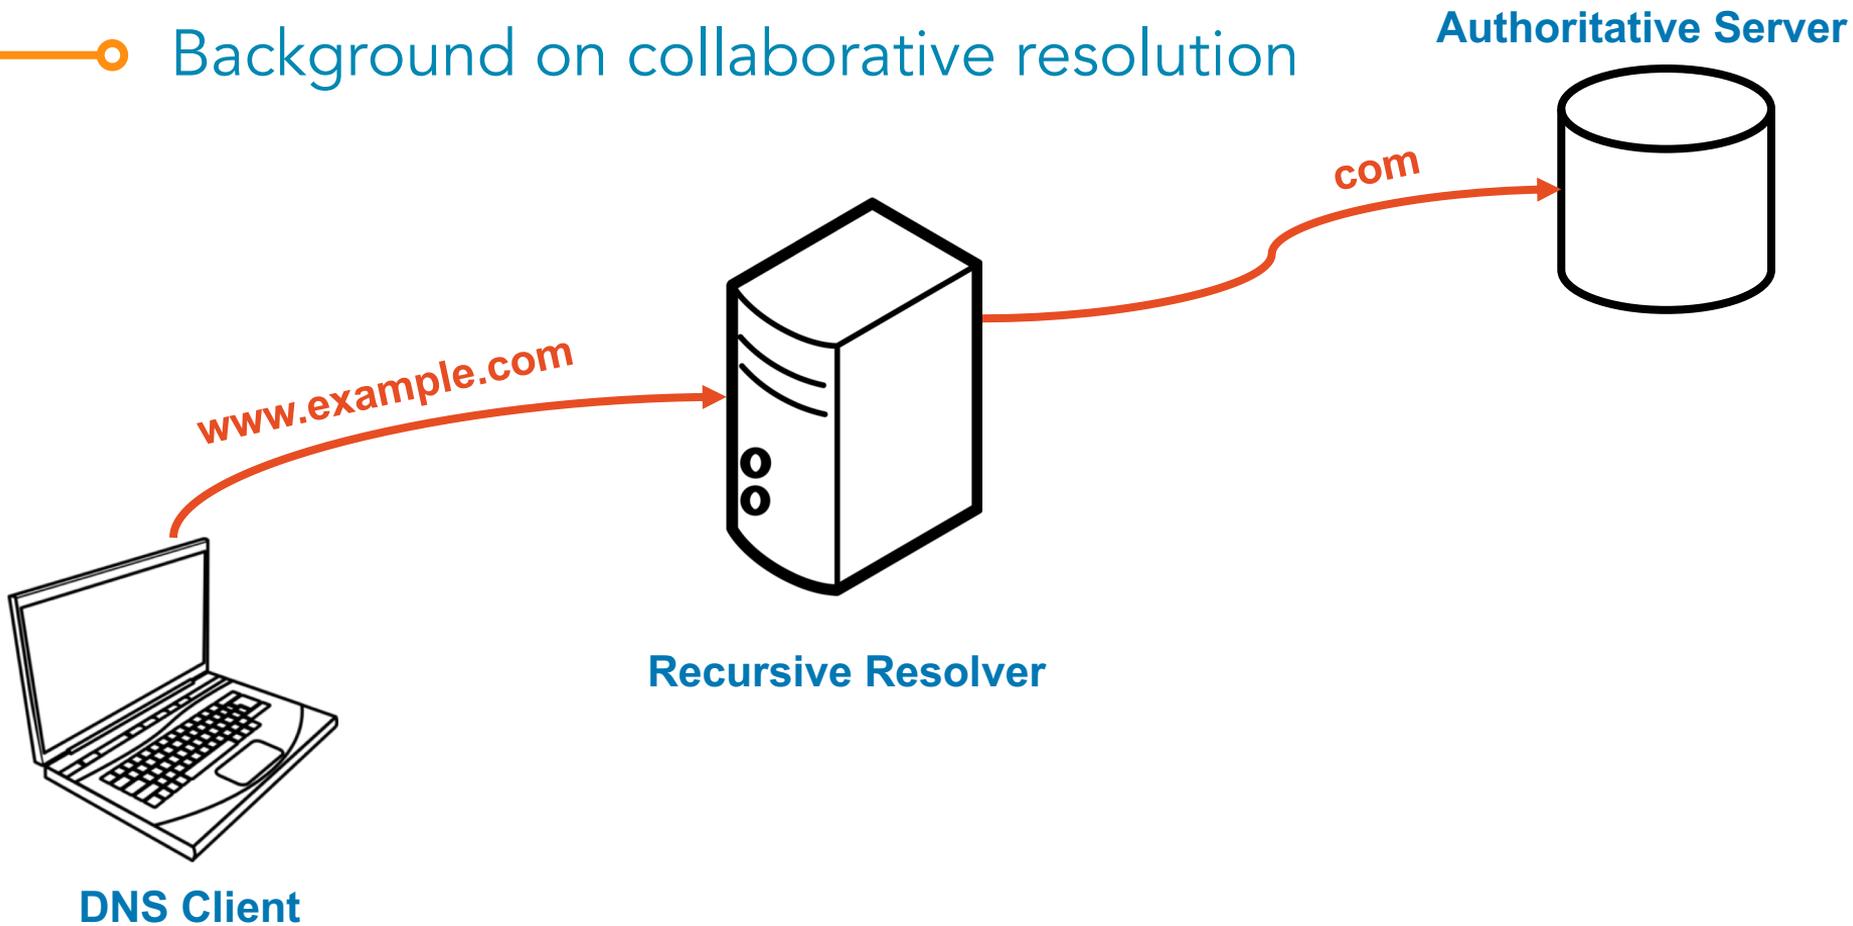
Rami Al-Dalky  
Kyle Schomp



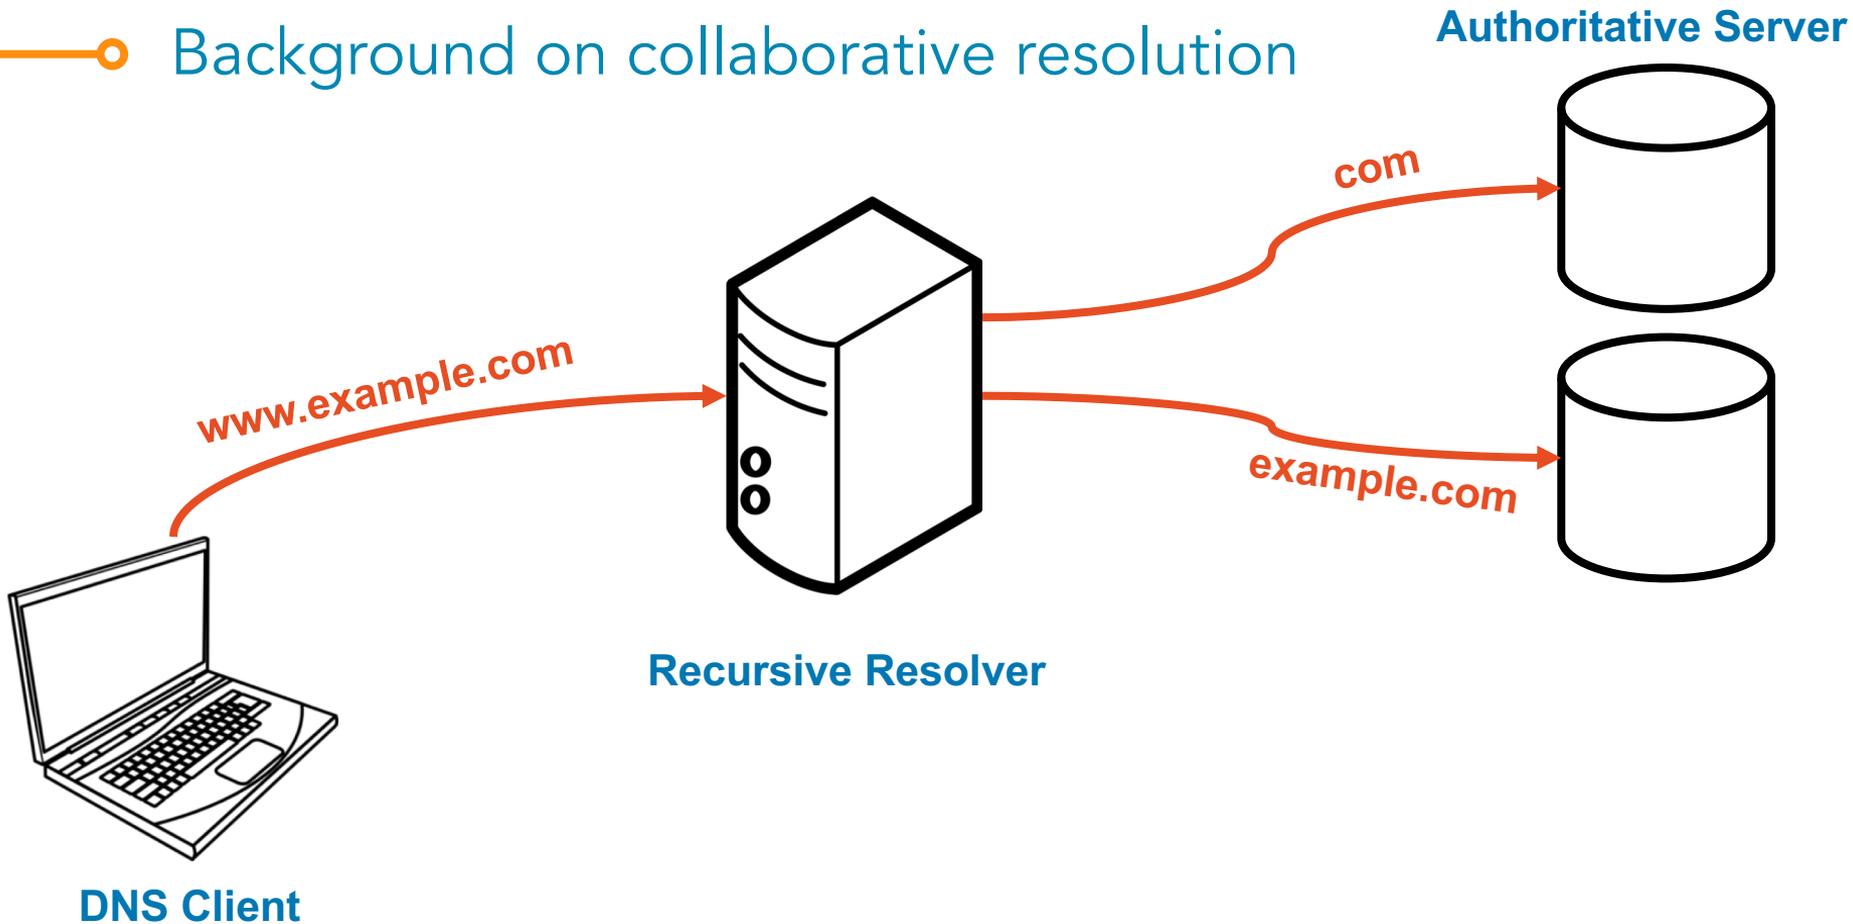
## ○ Background on collaborative resolution



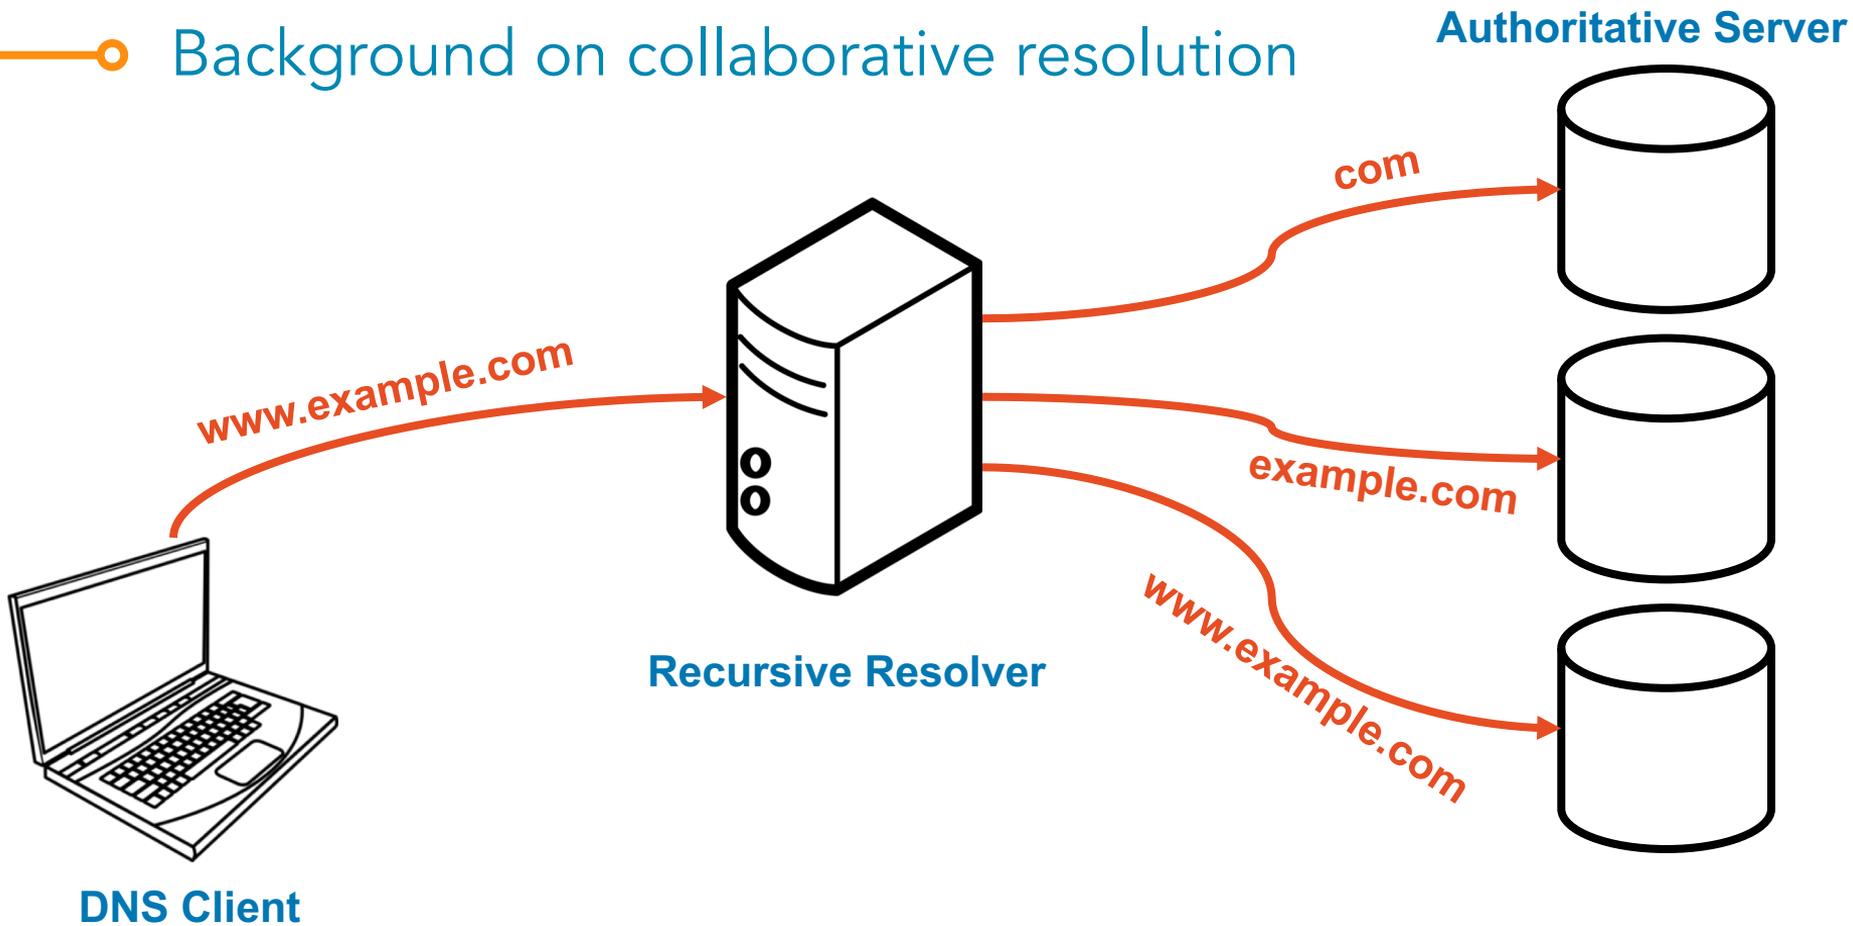
# Background on collaborative resolution



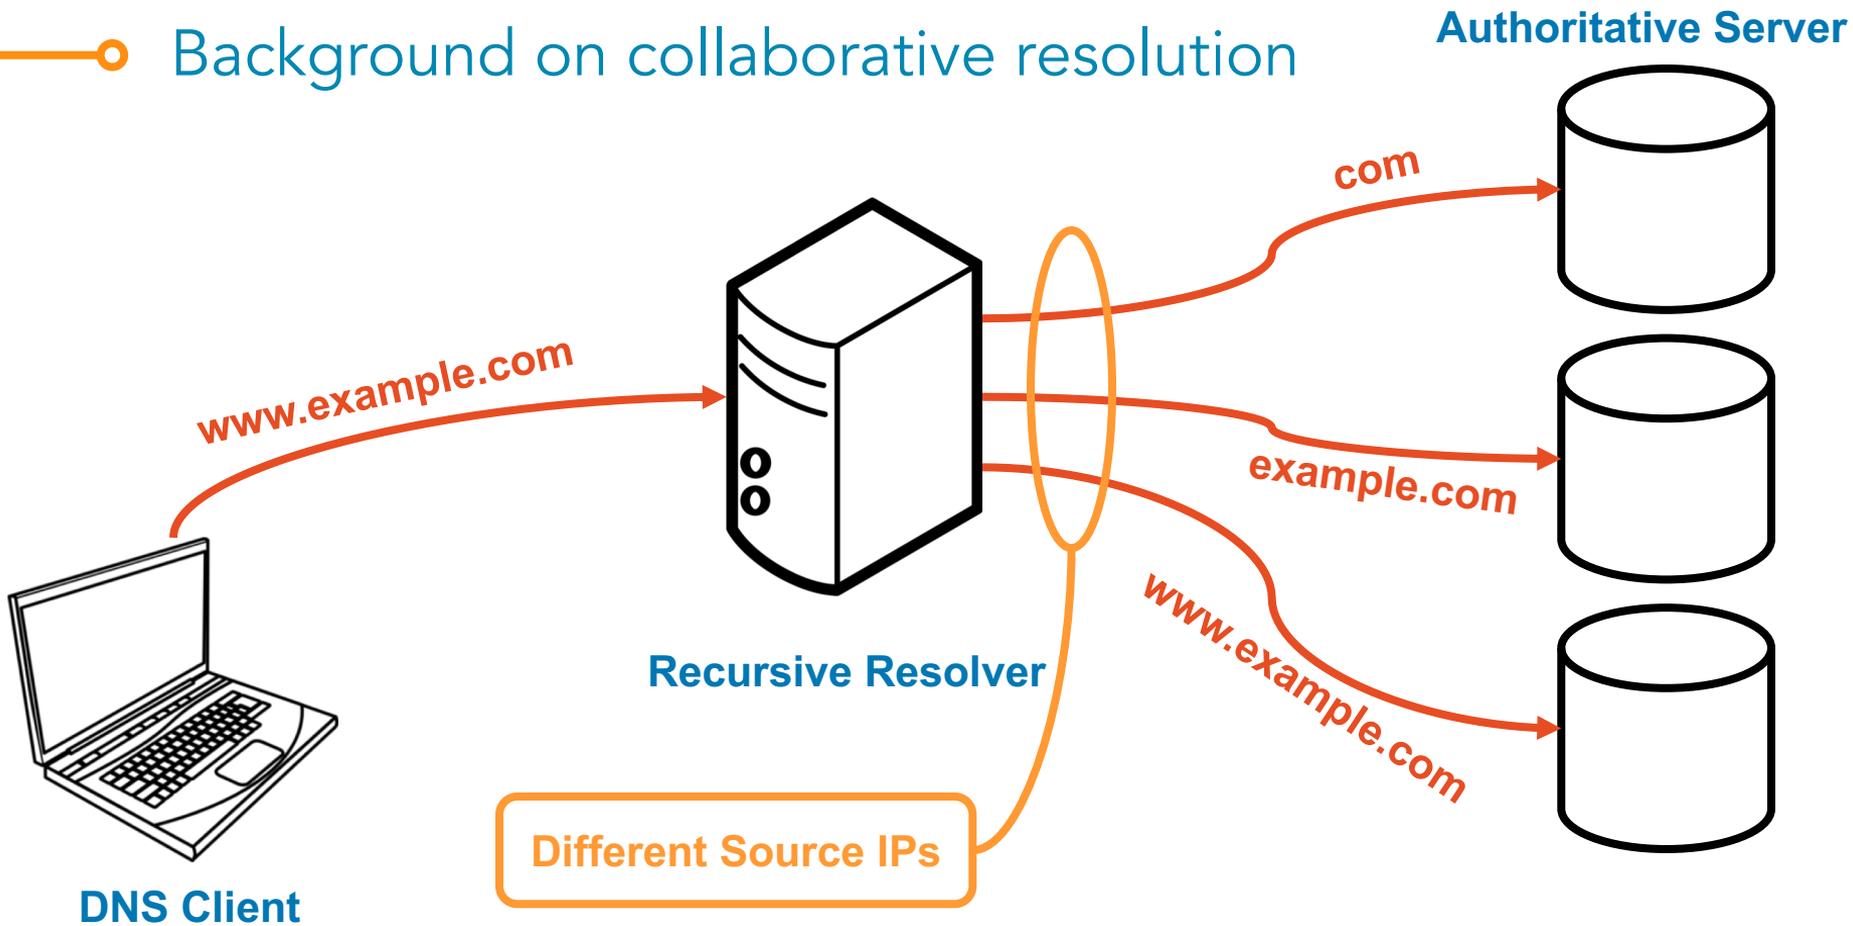
# Background on collaborative resolution



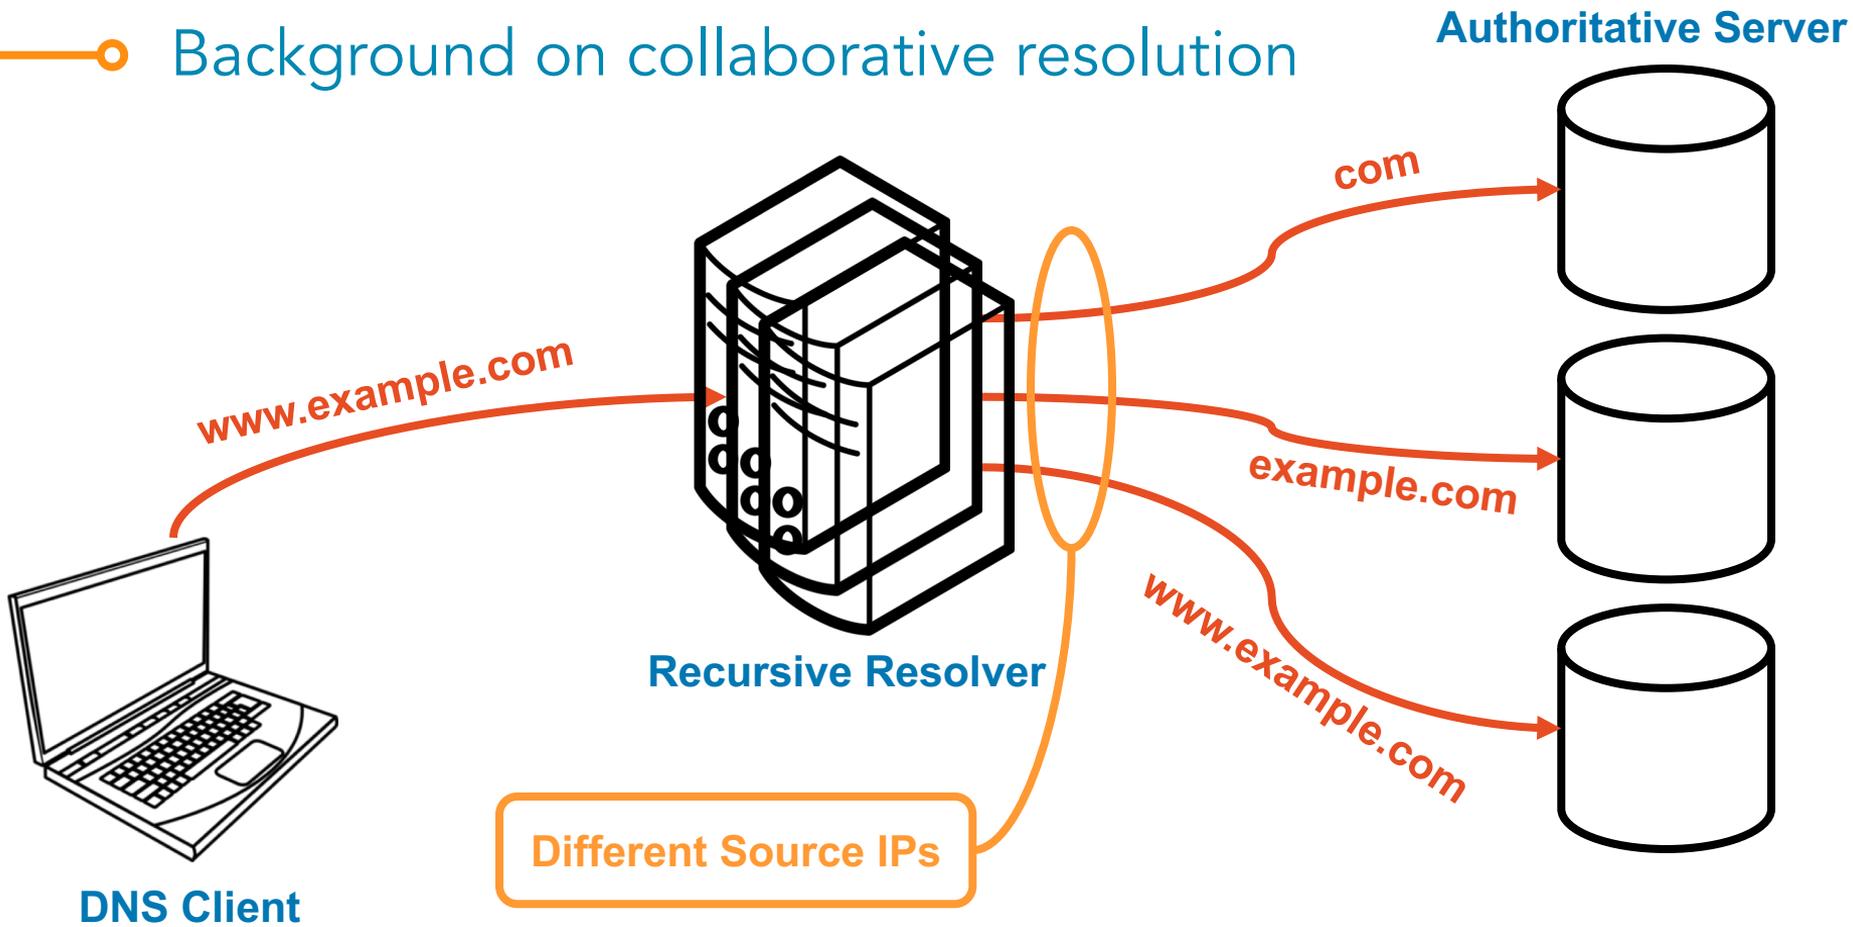
# Background on collaborative resolution



# Background on collaborative resolution



# Background on collaborative resolution



## ○ Why study recursive resolvers?

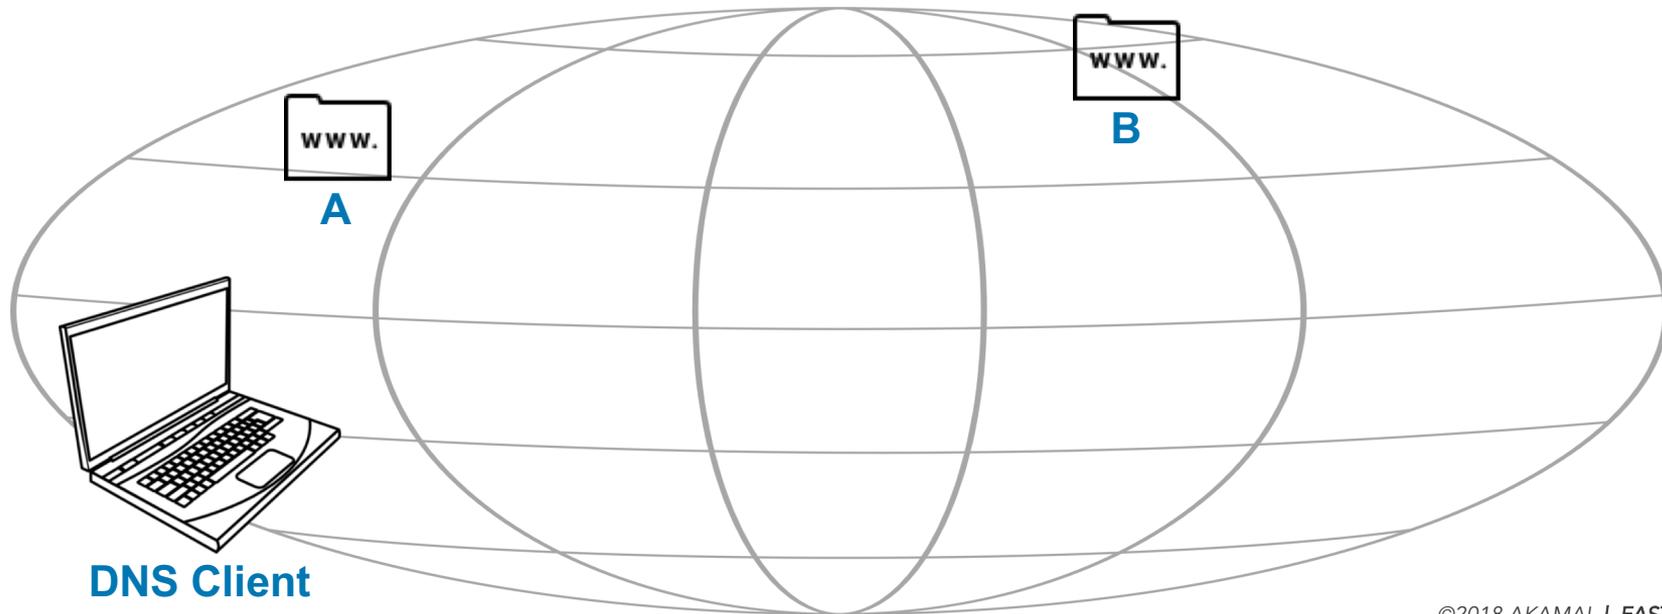
- DNS resolutions prefaces majority of web transactions
- But also, resolvers have an outsized role in performance

DNS-based Replica Selection

## Why study recursive resolvers?

- DNS resolutions prefaces majority of web transactions
- But also, resolvers have an outsized role in performance

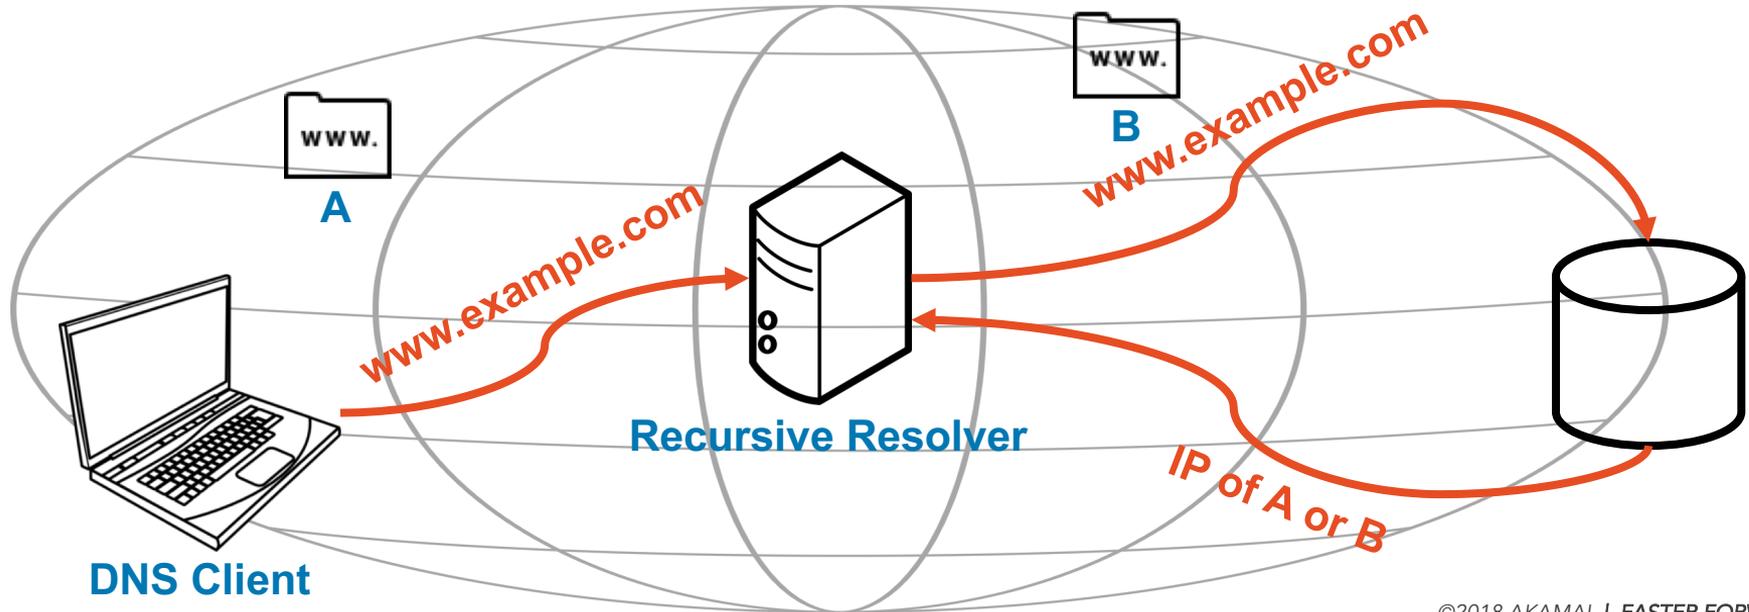
### DNS-based Replica Selection



## Why study recursive resolvers?

- DNS resolutions prefaces majority of web transactions
- But also, resolvers have an outsized role in performance

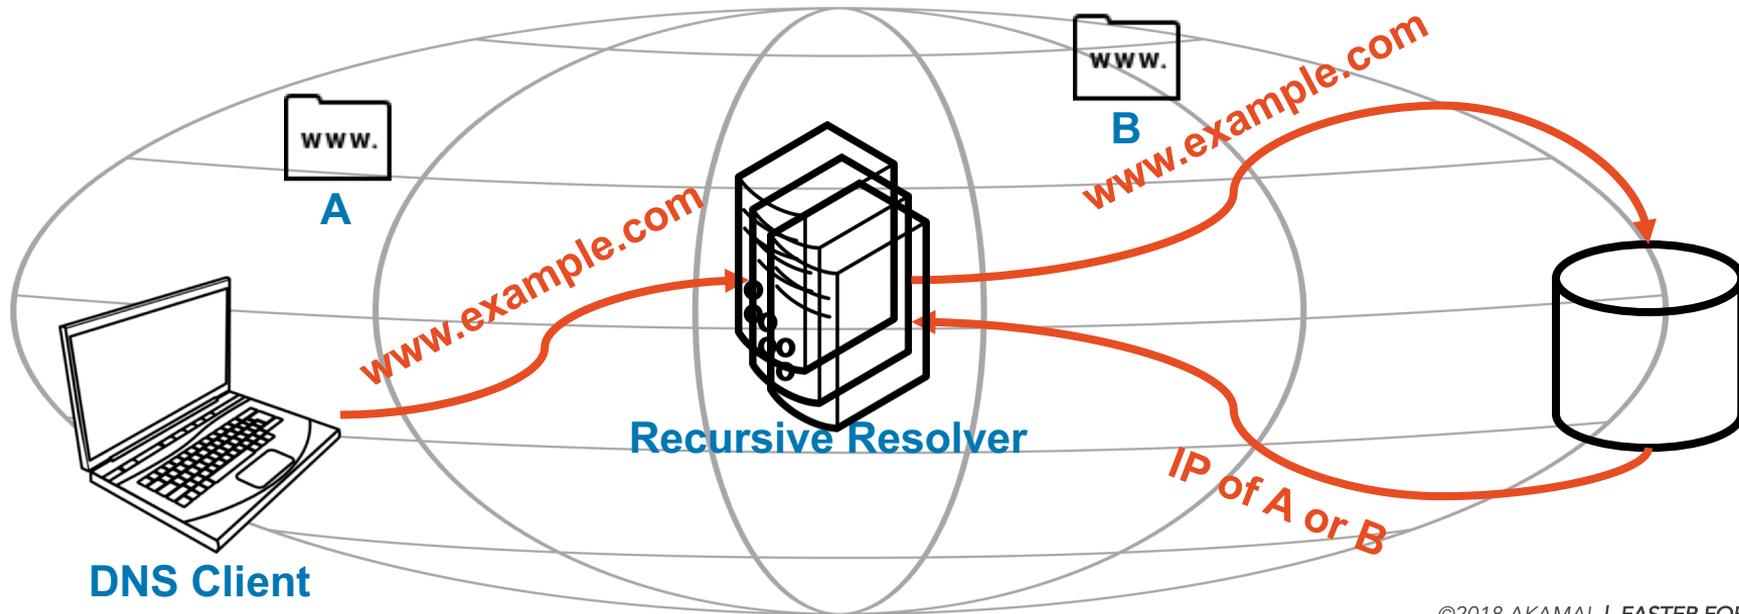
### DNS-based Replica Selection



## Why study recursive resolvers?

- DNS resolutions prefaces majority of web transactions
- But also, resolvers have an outsized role in performance

### DNS-based Replica Selection



## Data Collection

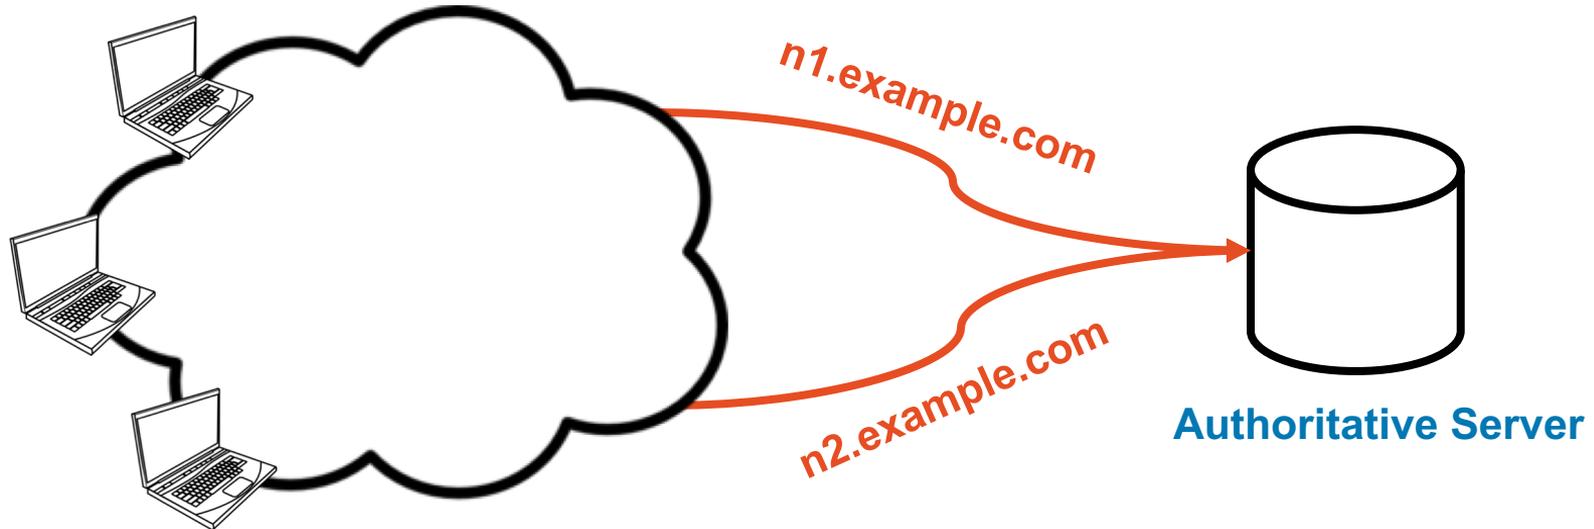
### Experimental DNS Records

|                |     |    |          |                     |
|----------------|-----|----|----------|---------------------|
| n1.example.com | 300 | IN | CNAME    | n2.example.com      |
| n2.example.com | 300 | IN | A / AAAA | 1.2.3.4 / 1:2:3:4:: |

# Data Collection

## Experimental DNS Records

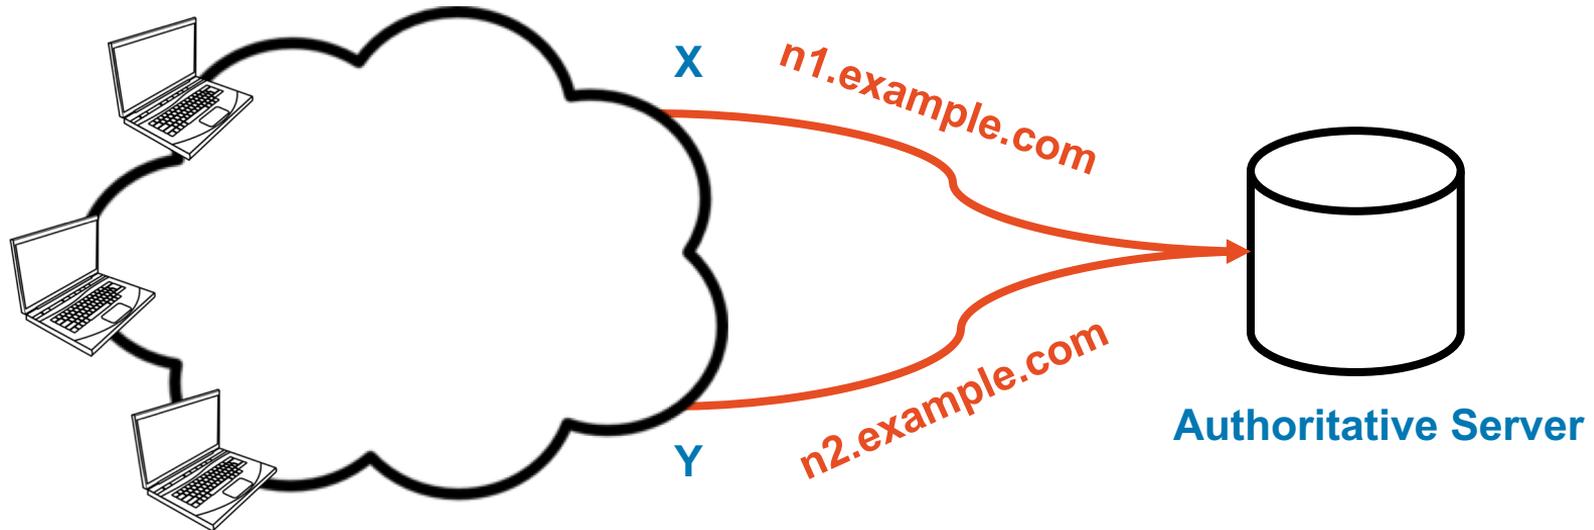
|                |     |    |          |                     |
|----------------|-----|----|----------|---------------------|
| n1.example.com | 300 | IN | CNAME    | n2.example.com      |
| n2.example.com | 300 | IN | A / AAAA | 1.2.3.4 / 1:2:3:4:: |



# Data Collection

## Experimental DNS Records

|                |     |    |          |                     |
|----------------|-----|----|----------|---------------------|
| n1.example.com | 300 | IN | CNAME    | n2.example.com      |
| n2.example.com | 300 | IN | A / AAAA | 1.2.3.4 / 1:2:3:4:: |



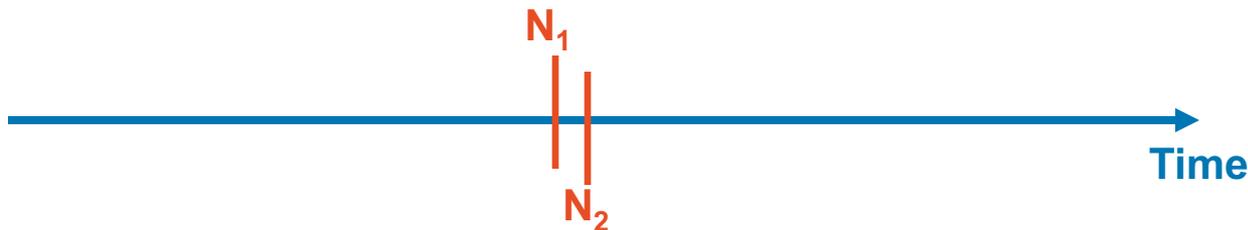
Directed Edge (X → Y) is a Pair

## ○ Removing False Positive Edges

1. Happy eyeballs
  - Include query type (A or AAAA) in the tuple when pairing queries
2. Multiple clients resolving the experimental name at the same time
  - Hostnames encode the client subnet
3. “Racing” recursive resolvers for fastest response
4. Re-resolution of one, but not both, hostnames
  - Window-based noise elimination

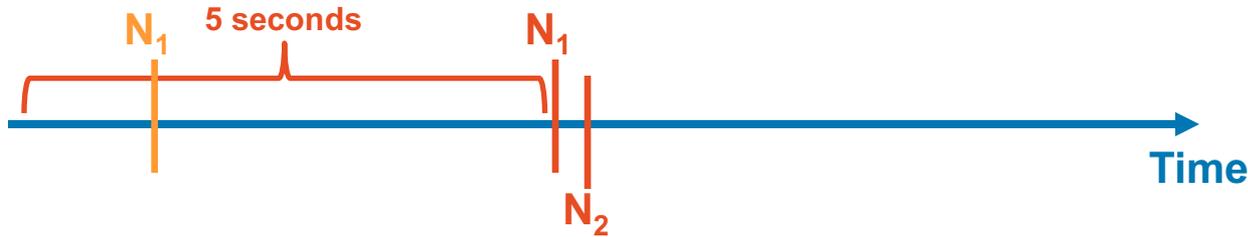
## Removing False Positive Edges

1. Happy eyeballs
  - Include query type (A or AAAA) in the tuple when pairing queries
2. Multiple clients resolving the experimental name at the same time
  - Hostnames encode the client subnet
3. "Racing" recursive resolvers for fastest response
4. Re-resolution of one, but not both, hostnames
  - Window-based noise elimination



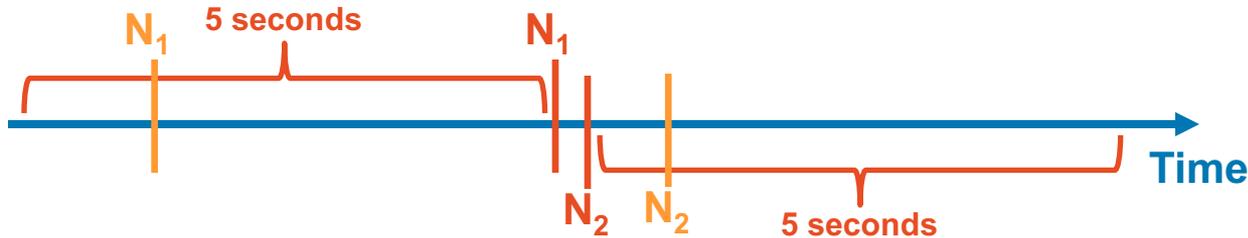
## Removing False Positive Edges

1. Happy eyeballs
  - Include query type (A or AAAA) in the tuple when pairing queries
2. Multiple clients resolving the experimental name at the same time
  - Hostnames encode the client subnet
3. "Racing" recursive resolvers for fastest response
4. Re-resolution of one, but not both, hostnames
  - Window-based noise elimination



## Removing False Positive Edges

1. Happy eyeballs
  - Include query type (A or AAAA) in the tuple when pairing queries
2. Multiple clients resolving the experimental name at the same time
  - Hostnames encode the client subnet
3. "Racing" recursive resolvers for fastest response
4. Re-resolution of one, but not both, hostnames
  - Window-based noise elimination



## Data

- 1 week of authoritative DNS query logs

**DNS Queries**

**820M**

## Data

- 1 week of authoritative DNS query logs
- Filter down to pairs from the queries

|                    |             |
|--------------------|-------------|
| <b>DNS Queries</b> | <b>820M</b> |
| <b>Pairs</b>       | <b>109M</b> |

## Data

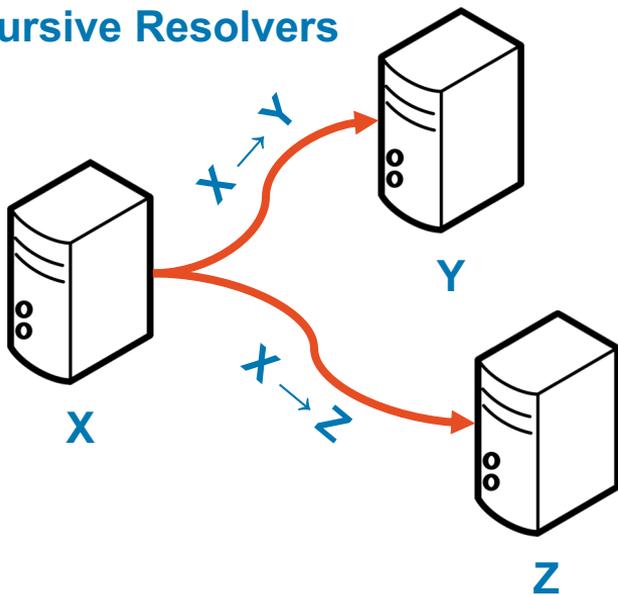
- 1 week of authoritative DNS query logs
- Filter down to pairs from the queries
- Cluster *same* initiator into pools

|                    |             |
|--------------------|-------------|
| <b>DNS Queries</b> | <b>820M</b> |
| <b>Pairs</b>       | <b>109M</b> |
| <b>Clusters</b>    | <b>421K</b> |

## Data

- 1 week of authoritative DNS query logs
- Filter down to pairs from the queries
- Cluster *same* initiator into pools

### Recursive Resolvers



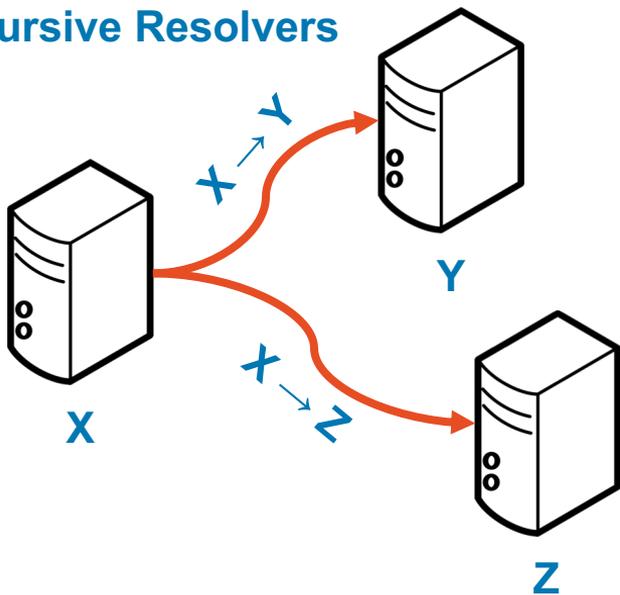
**Pool: X → Y, Z**

|                    |             |
|--------------------|-------------|
| <b>DNS Queries</b> | <b>820M</b> |
| <b>Pairs</b>       | <b>109M</b> |
| <b>Clusters</b>    | <b>421K</b> |

## Data

- 1 week of authoritative DNS query logs
- Filter down to pairs from the queries
- Cluster *same* initiator into pools

### Recursive Resolvers



DNS Queries

820M

Pairs

109M

Clusters

421K

Pool:  $X \rightarrow Y, Z$

$N_{X \rightarrow Y}, M_{X \rightarrow Z}$

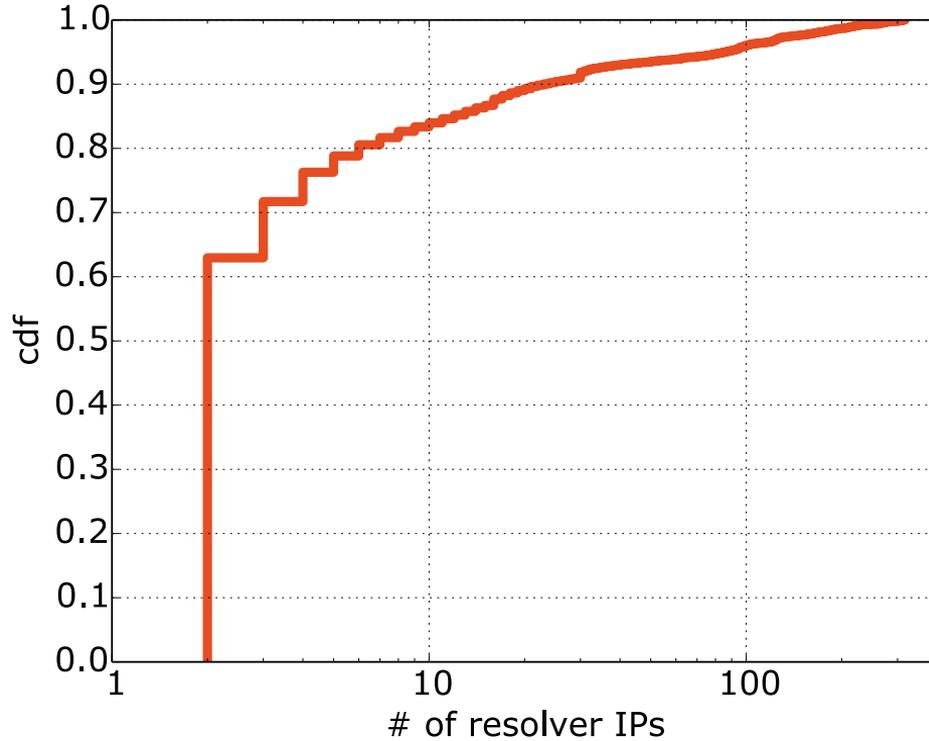
## Data

- 1 week of authoritative DNS query logs
- Filter down to pairs from the queries
- Cluster *same* initiator into pools
- 14% of clusters contain more than 1 IP

|                     |             |
|---------------------|-------------|
| <b>DNS Queries</b>  | <b>820M</b> |
| <b>Pairs</b>        | <b>109M</b> |
| <b>Clusters</b>     | <b>421K</b> |
| → <b>Singletons</b> | <b>360K</b> |
| → <b>Pools</b>      | <b>61K</b>  |

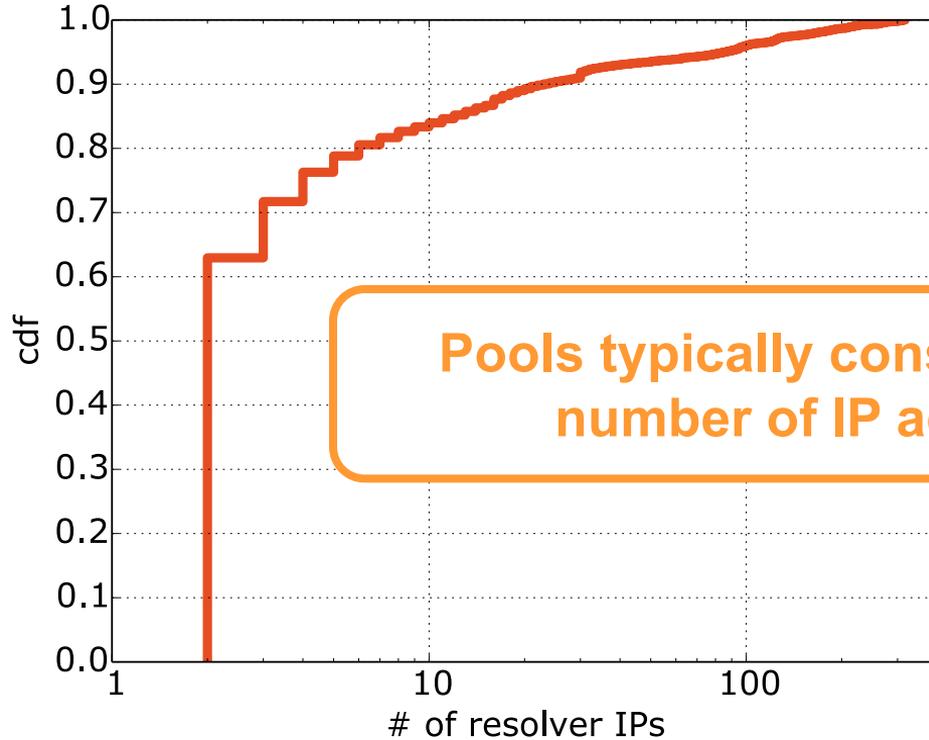
# What do the pools look like?

## Number of recursive IPs per pool



# ○ What do the pools look like?

## Number of recursive IPs per pool



**Pools typically consist of a small number of IP addresses**

# What do the pools look like?

## Distribution in IP-Space

| IPv4 representation | Binary representation                   |
|---------------------|---|
| 1.2.3.1             | 0000 0001 0000 0010 0000 0011 0000 0001 |
| 1.2.3.128           | 0000 0001 0000 0010 0000 0011 1000 0000 |
| 1.2.4.1             | 0000 0001 0000 0010 0000 0100 0000 0001 |

# What do the pools look like?

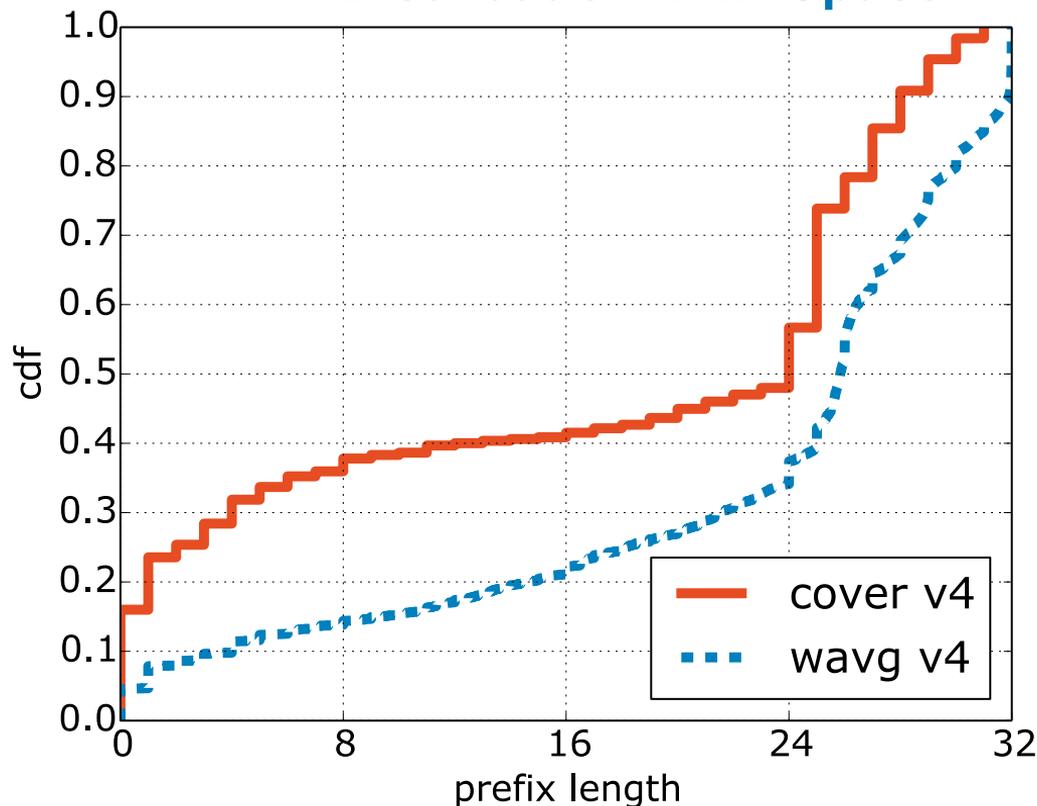
## Distribution in IP-Space

| IPv4 representation | Binary representation                   |
|---------------------|---|
| 1.2.3.1             | 0000 0001 0000 0010 0000 0011 0000 0001 |
| 1.2.3.128           | 0000 0001 0000 0010 0000 0011 1000 0000 |
| 1.2.4.1             | 0000 0001 0000 0010 0000 0100 0000 0001 |

Length of covering prefix is 21

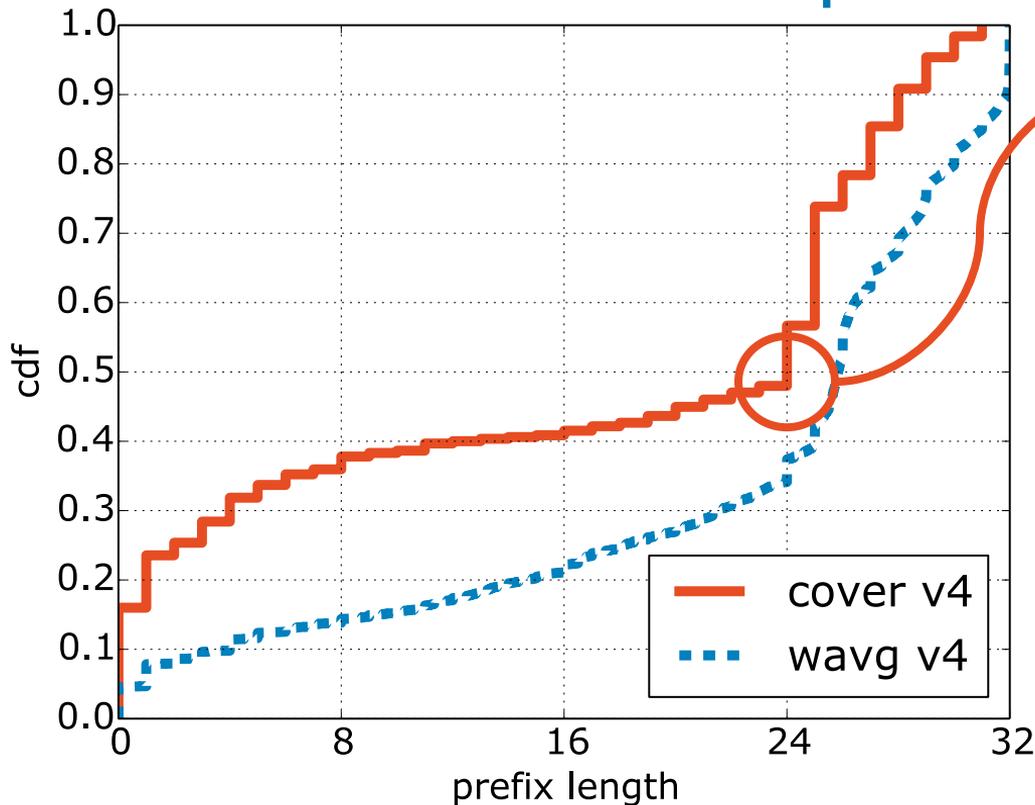
# What do the pools look like?

## Distribution in IP-Space



# What do the pools look like?

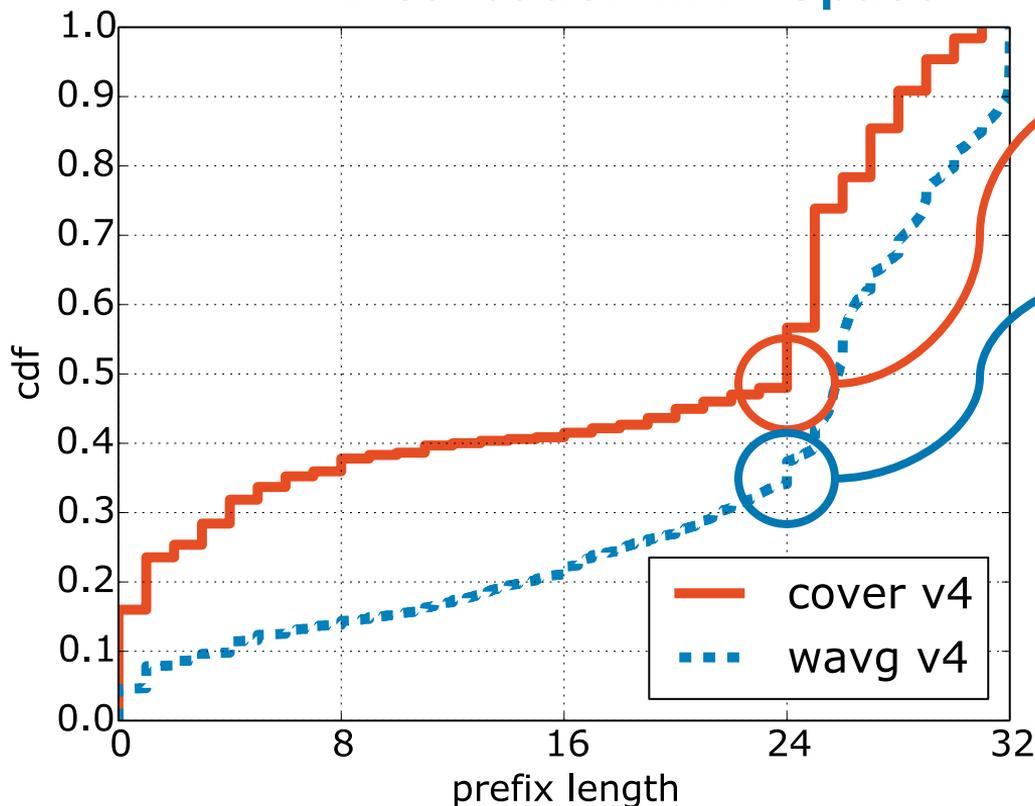
## Distribution in IP-Space



Only 50% of pools are in a single /24

# What do the pools look like?

## Distribution in IP-Space

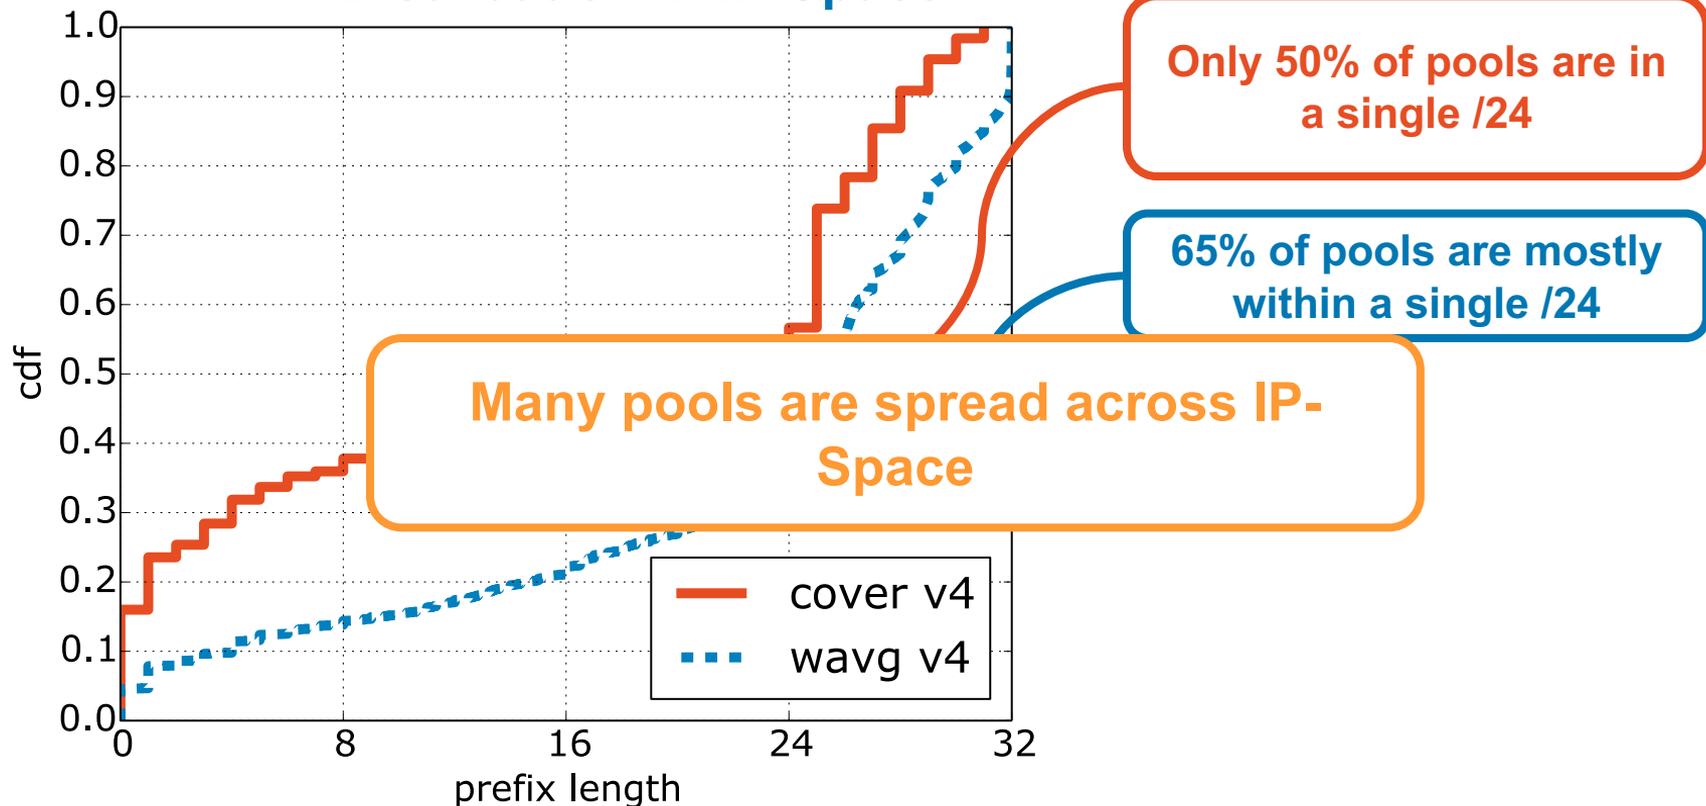


Only 50% of pools are in a single /24

65% of pools are mostly within a single /24

# What do the pools look like?

## Distribution in IP-Space



## ○ What do the pools look like?

### **Autonomous Systems per Pool**

- 85% of pools are within a single AS
- 14% of pools are in 2 ASs
- Comparing WHOIS entries of most frequently occurring AS pairs shows same organization

## ○ What do the pools look like?

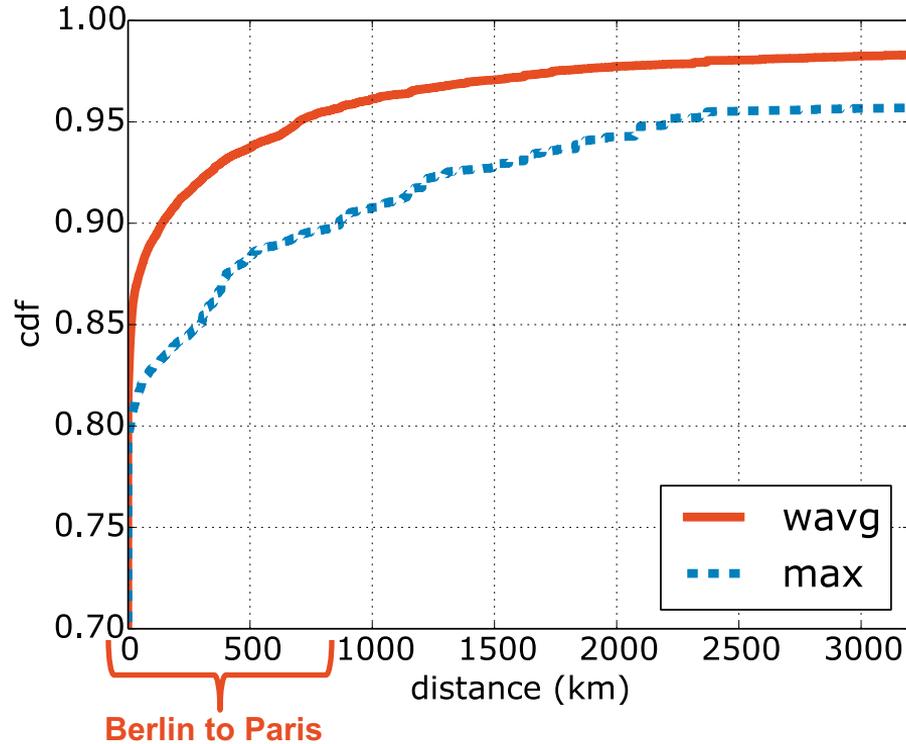
### **Autonomous Systems per Pool**

- 85% of pools are within a single AS
- 14% of pools are in 2 ASs
- Comparing WHOIS entries of most frequently occurring AS pairs shows same organization

**Pools very rarely cross organizational boundaries**

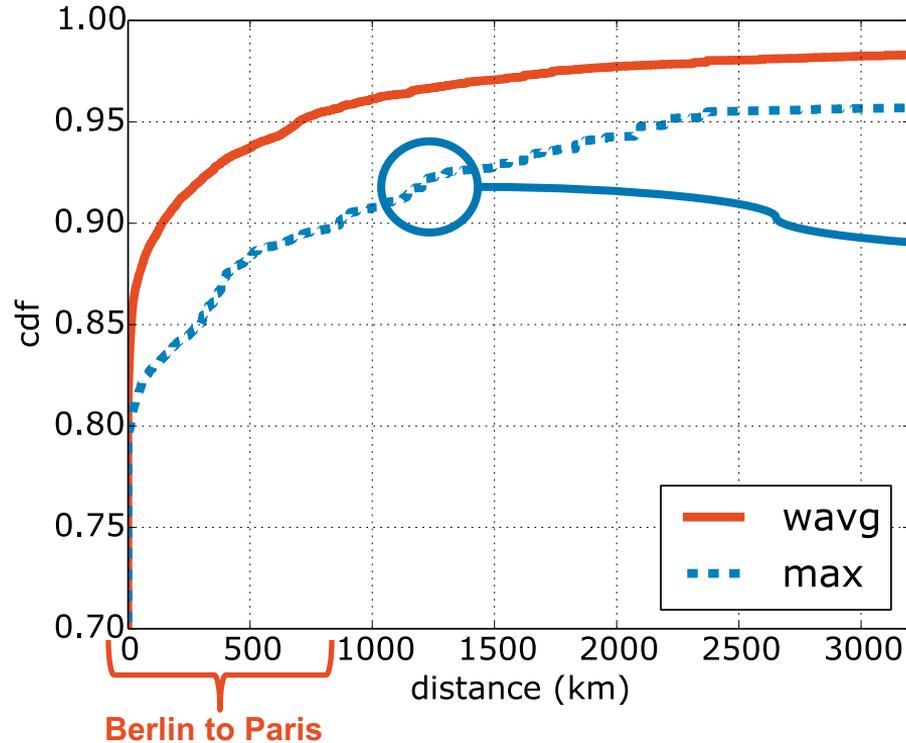
# What do the pools look like?

## Geographic Distance within Pools



# What do the pools look like?

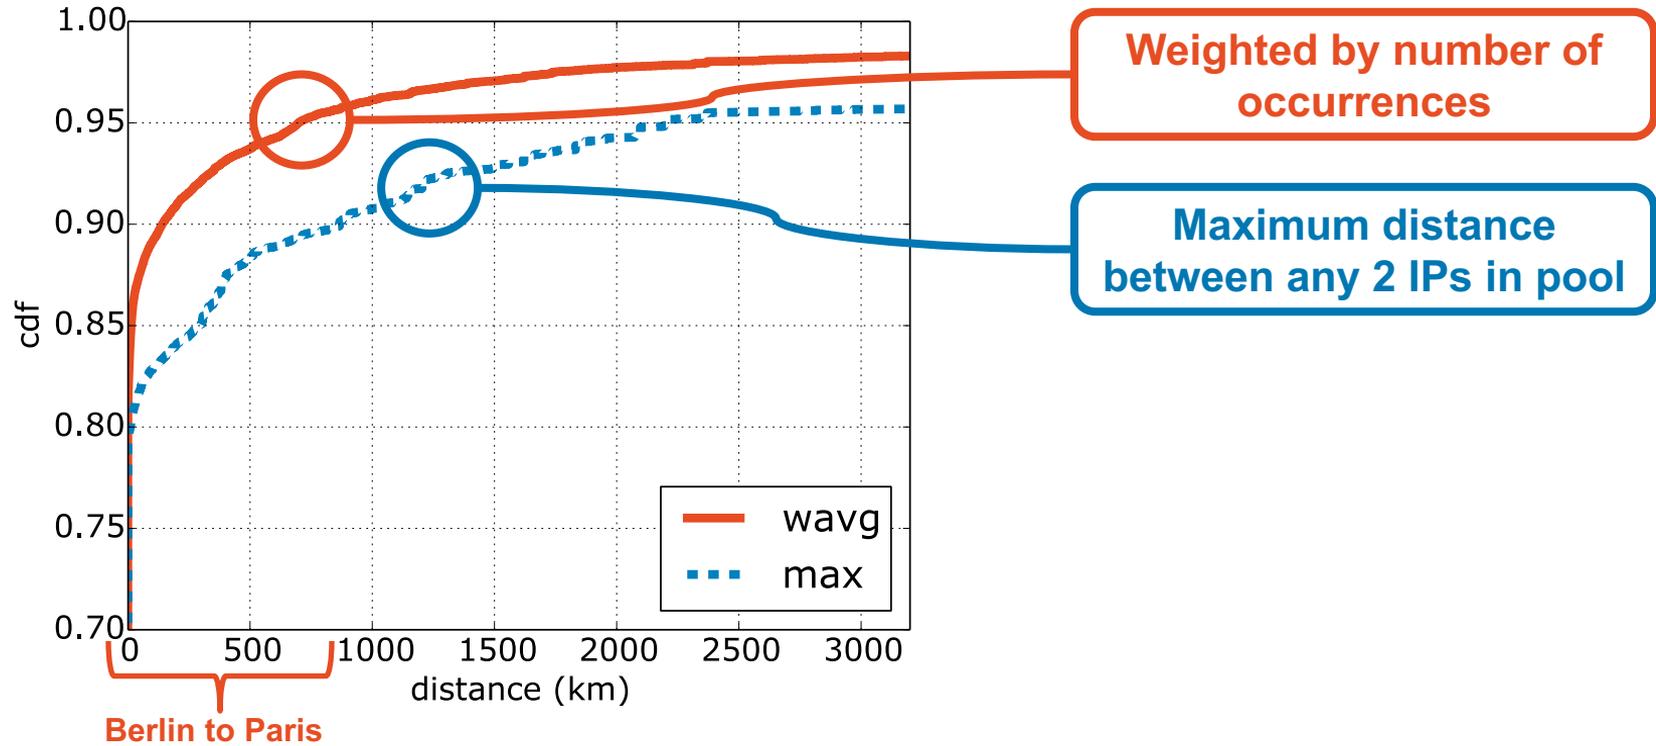
## Geographic Distance within Pools



Maximum distance  
between any 2 IPs in pool

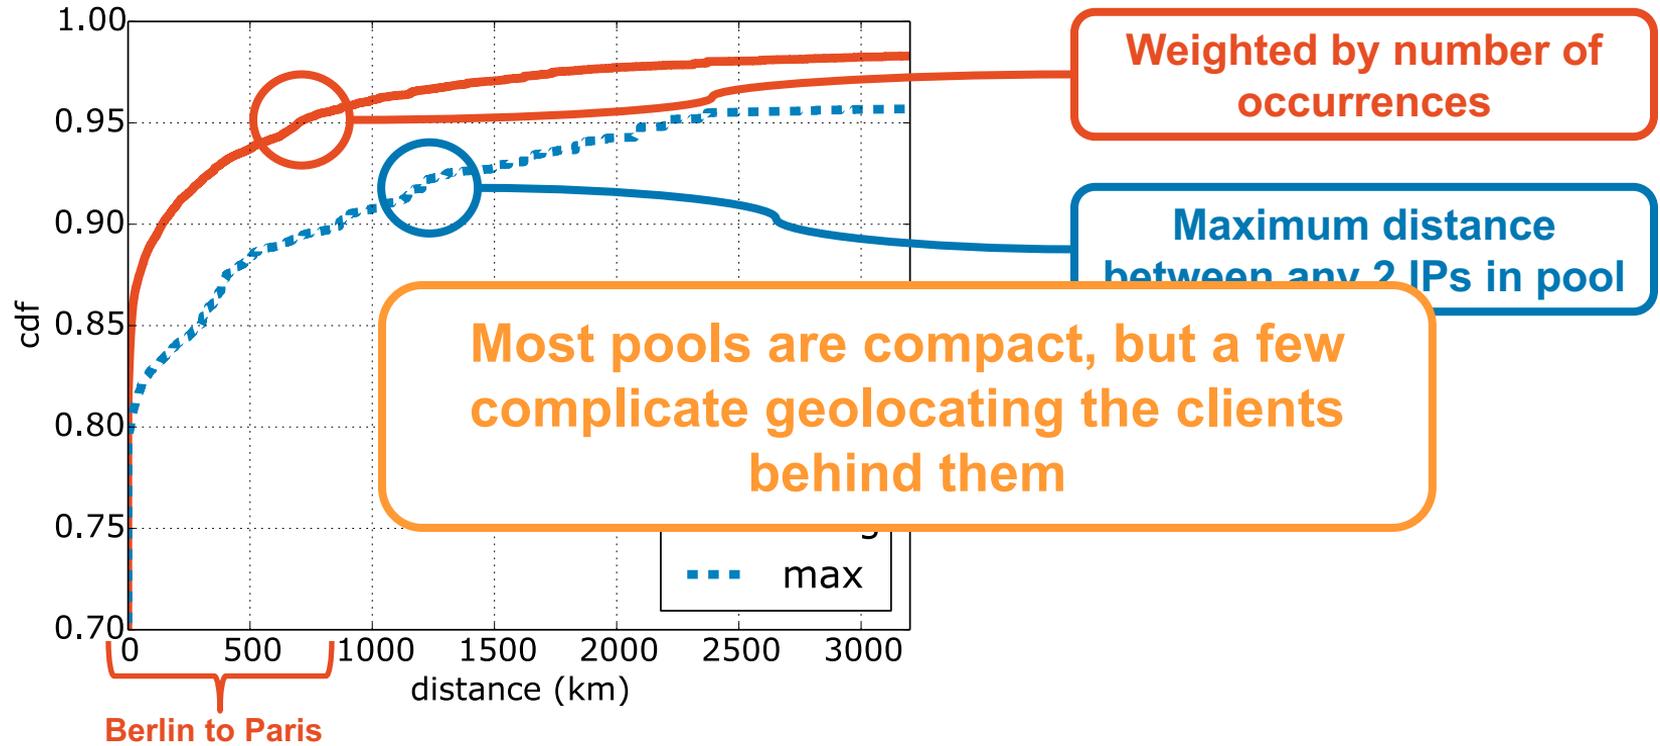
# What do the pools look like?

## Geographic Distance within Pools



# What do the pools look like?

## Geographic Distance within Pools



# How are queries distributed within the pool?

Category

Pools

100%

## How are queries distributed within the pool?

### Category

- Pools with 1 IPv4 / 1 IPv6 address likely dual-stack resolvers

29% have identifiable patterns:

|                                  |            |
|----------------------------------|------------|
| <b>1.2.3.4 and 89ab::1:2:3:4</b> | <b>10%</b> |
|----------------------------------|------------|

|                            |            |
|----------------------------|------------|
| <b>1.2.3.4 and 89ab::4</b> | <b>19%</b> |
|----------------------------|------------|

|                      |             |
|----------------------|-------------|
| <b>Pools</b>         | <b>100%</b> |
| → <b>Dual-Stacks</b> | <b>36%</b>  |

## How are queries distributed within the pool?

### Category

- Pools with 1 IPv4 / 1 IPv6 address likely dual-stack resolvers
- Use  $\chi^2$  to test for uniformly distributing queries within pool

|                           |             |
|---------------------------|-------------|
| <b>Pools</b>              | <b>100%</b> |
| → <b>Dual-Stacks</b>      | <b>36%</b>  |
| → <b>Uniform Load Bal</b> | <b>14%</b>  |

## How are queries distributed within the pool?

### Category

- Pools with 1 IPv4 / 1 IPv6 address likely dual-stack resolvers
- Use  $\chi^2$  to test for uniformly distributing queries within pool
- Pools where the initiator rarely offloads queries to other resolvers

|                           |             |
|---------------------------|-------------|
| <b>Pools</b>              | <b>100%</b> |
| → <b>Dual-Stacks</b>      | <b>36%</b>  |
| → <b>Uniform Load Bal</b> | <b>14%</b>  |
| → <b>Rare Offloading</b>  | <b>27%</b>  |

## How are queries distributed within the pool?

### Category

- Pools with 1 IPv4 / 1 IPv6 address likely dual-stack resolvers
- Use  $\chi^2$  to test for uniformly distributing queries within pool
- Pools where the initiator rarely offloads queries to other resolvers
- Other / unknown

|                           |             |
|---------------------------|-------------|
| <b>Pools</b>              | <b>100%</b> |
| → <b>Dual-Stacks</b>      | <b>36%</b>  |
| → <b>Uniform Load Bal</b> | <b>14%</b>  |
| → <b>Rare Offloading</b>  | <b>27%</b>  |
| → <b>Other/Unknown</b>    | <b>23%</b>  |

## How are queries distributed within the pool?

### Category

- Pools with 1 IPv4 / 1 IPv6 address likely dual-stack resolvers
- Use  $\chi^2$  to test for uniformly distributing queries within pool
- Pools where queries to other pools
- Other / unknown

|                           |             |
|---------------------------|-------------|
| <b>Pools</b>              | <b>100%</b> |
| → <b>Dual-Stacks</b>      | <b>36%</b>  |
| → <b>Uniform Load Bal</b> | <b>14%</b>  |
| → <b>Rare Offloading</b>  | <b>27%</b>  |
| → <b>Other/Unknown</b>    | <b>23%</b>  |

**Wide range of behaviors suggest that recursive resolver pools are used for a variety of purposes**

# Questions?



Kyle Schomp  
kschomp@akamai.com