A Look at the ECS Behavior of DNS Resolvers

Rami Al-Dalky*, Michael Rabinovich*, and Kyle Schomp†

* Case Western Reserve University

† Akamai Technologies

ECS: EDNS0 Client Subnet Extension

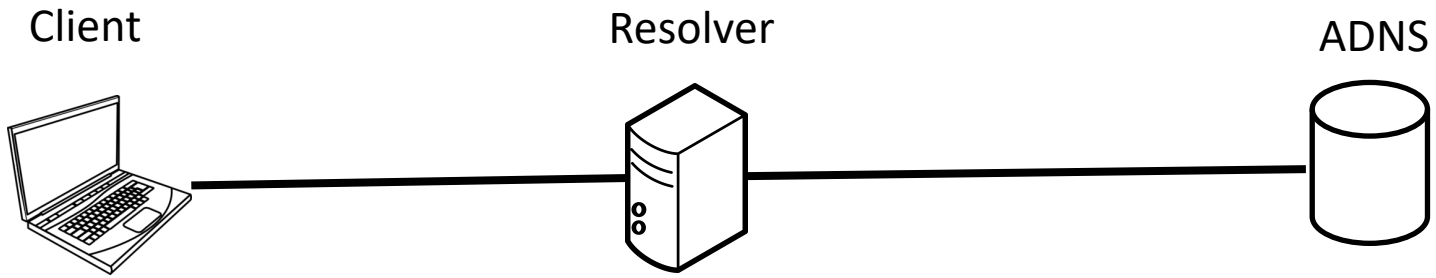

- ECS Purpose
  - Enable CDN server selection by ADNS based on client subnet
- ECS Option in DNS queries from resolvers to ADNS includes
  - Client IP address prefix
  - Source prefix length

  ADNS uses to tailor response
- ECS Option in DNS responses from ADNS to resolvers includes
  - Scope prefix length

  Resolver must only use cached response for clients covered by scope

ECS Implications

- Privacy implications: ADNS (and anyone on path) learns client IP
  - Resolvers limit client prefixes to 24 bits in IPv4 and 56 bits in IPv6
  - Resolvers probe ADNS for ECS support and don’t send ECS to non-supporting ADNS
- Security implications: easy scanning of CDN platforms by attackers
  - ADNS whitelist resolvers that are trusted via some out-of-band means

Goals of Study

- Exploring resolvers' ECS behavior
  - Probing behavior?
  - Source prefix lengths used?
  - Adherence to ECS scope's cache restrictions?
- Exploring ECS impact on resolver cache
  - Required cache size?
  - Effect on hit rate?
- ECS deployment pitfalls
- covered in this talk

Datasets

- Logs from a major CDN's ADNS
  - Look for IP addresses of recursive resolvers that send ECS option
- Internet-wide scans
  - Find open recursive resolvers that send ECS options ***or***
  - Find open forwarders that forward to recursive resolvers that send ECS option

1. Honoring Scope Restriction on Caching

- Found 278 non-Google ECS-supporting recursive resolvers via Internet-wide scanning
- Able to study 202 out of 278 via various probing strategies (check paper)
  - 76 resolvers forward /24 prefixes and honor the scope properly
  - 15 resolvers accept and forward less than /24 prefixes
  - 8 resolvers truncate incoming prefixes to at most /22
  - 1 misconfigured resolver that sends an ECS prefix from 10.0.0.0/8
  - 102 resolvers don’t obey scope caching restrictions at all
- Half of non-Google recursive resolvers using ECS are misleading ADNS!

2. Unroutable ECS Prefixes

- 33 resolvers from Internet-wide scan submitted queries with loopback ECS prefix
- Could this confuse ADNSs?

2. Unroutable ECS Prefixes

- 33 resolvers from Internet-wide scan submitted queries with loopback ECS prefix
- Could this confuse ADNSs?
- Send 5 queries from our test machine to ADNS for youtube.com:

ECS Prefix	RTT (ms)	Location
None	35	Chicago
/24 of test machine's IP	35	Chicago
127.0.0.1/32	155	Switzerland
127.0.0.0/24	47	Mountain View, CA
169.254.252.0/24	285	South Africa

2. Unroutable ECS Prefixes

- RFC 7871, 11.3:
  - "[Authoritative nameservers and recursive resolvers] SHOULD at least treat unroutable addresses, such as some of the address blocks defined in [RFC6890], as equivalent to the Recursive Resolver's own identity.”

2. Unroutable ECS Prefixes

- RFC 7871, 11.3:
  - "[Authoritative nameservers and recursive resolvers] **MUST** at least treat unroutable addresses, such as some of the address blocks defined in [RFC6890], as equivalent to the Recursive Resolver's own identity."
- Further add language:
  - Recursive resolvers MUST send routable prefixes in the ECS option (or not send the option if that's not possible).

3. Impact of Source Prefix Length

- Try different prefix lengths against two different CDNs
- 800 random Ripe Atlas probes
- Resolve CDN-accelerated hostnames from our lab machine but with probes' ECS prefixes
- Check TCP handshake RTT from the probe to returned IP addresses for its ECS prefix

3. Impact of Source Prefix Length

- Try different prefix lengths against two different CDNs
- 800 random Ripe Atlas probes
- Resolve CDN-accelerated hostnames from our lab machine but with probes' ECS prefixes
- Check TCP handshake RTT from the probe to returned IP addresses for its ECS prefix

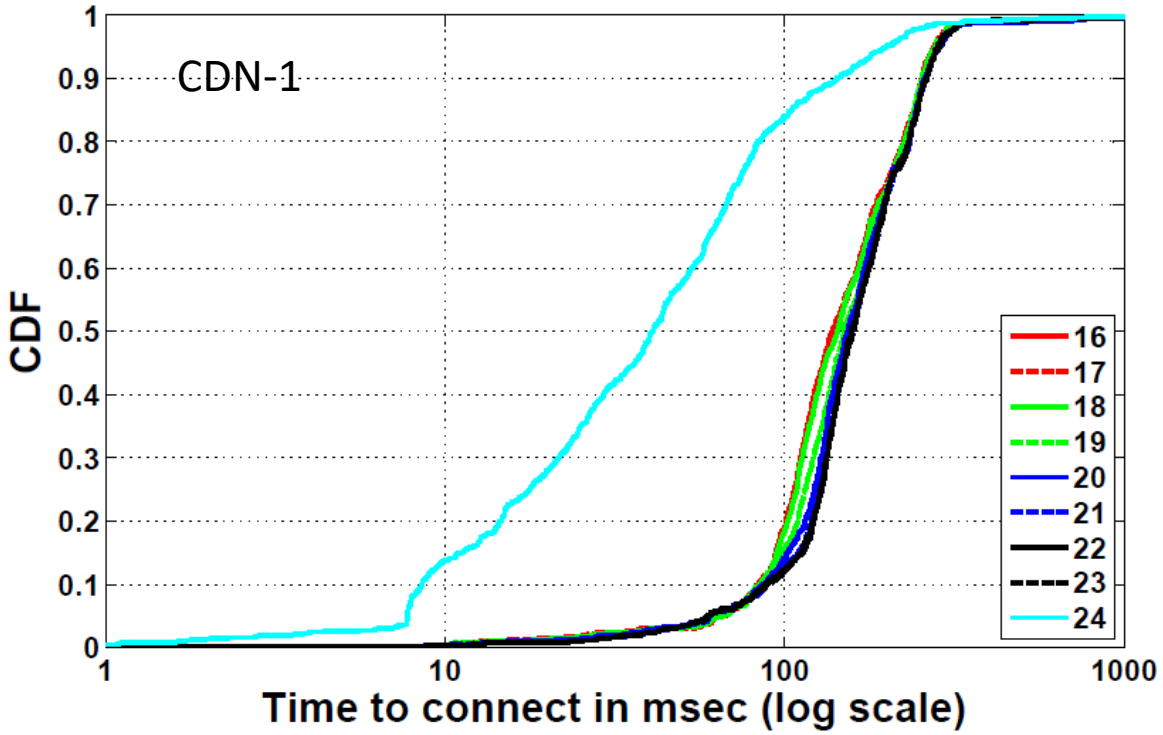

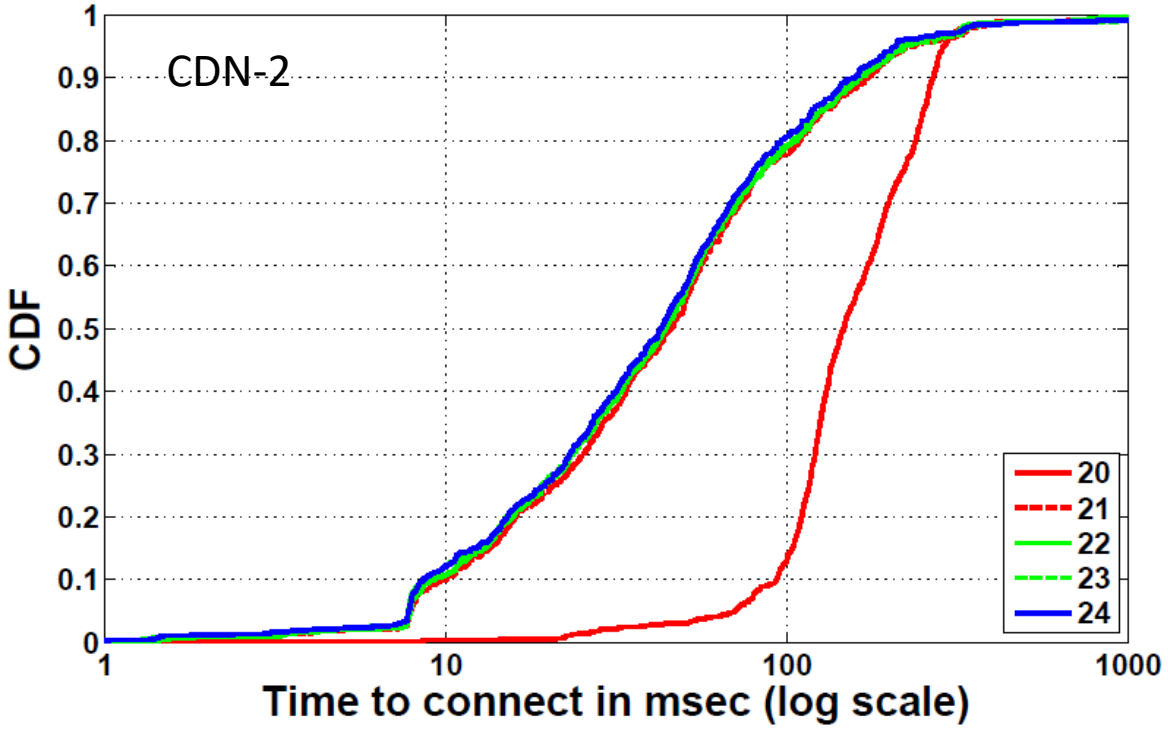

3. Impact of Source Prefix Length

- RFC 7871, 11.1:
  - "To protect users' privacy, Recursive Resolvers are strongly encouraged to conceal part of the user's IP address by truncating IPv4 addresses to 24 bits. 56 bits are recommended for IPv6, based on [RFC6177].”
- Sending less than 24 bits negatively impacts CDN mapping for some CDNs
  - Keep source prefix length state per ADNS or use 24 bits for everyone

4. Impact of Hidden Resolvers on ECS

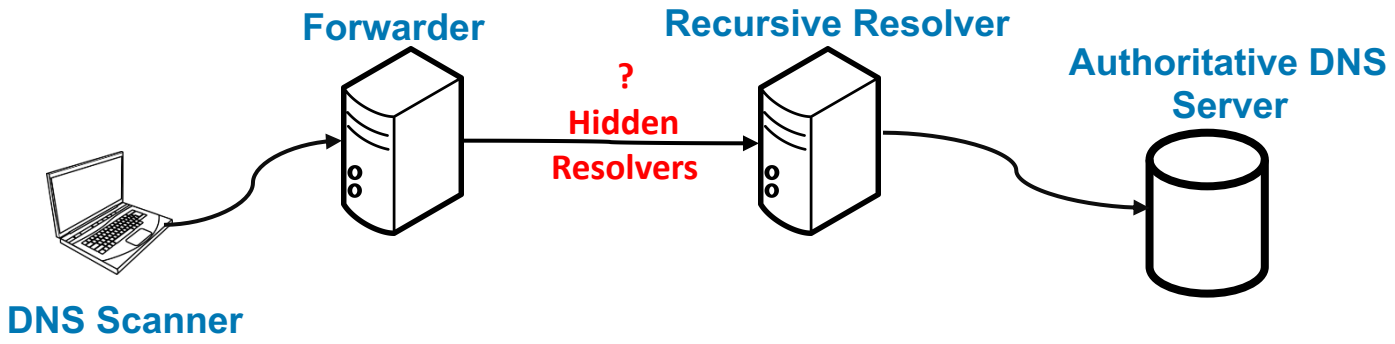

- Internet-wide scan had a number of queries with ECS source prefix different from both the probed address and egress address
- ADNS uses these prefixes – instead of forwarder or resolver IP address -- for server selection
- How close is the hidden resolver to the forwarder?

4. Impact of Hidden Resolvers on ECS

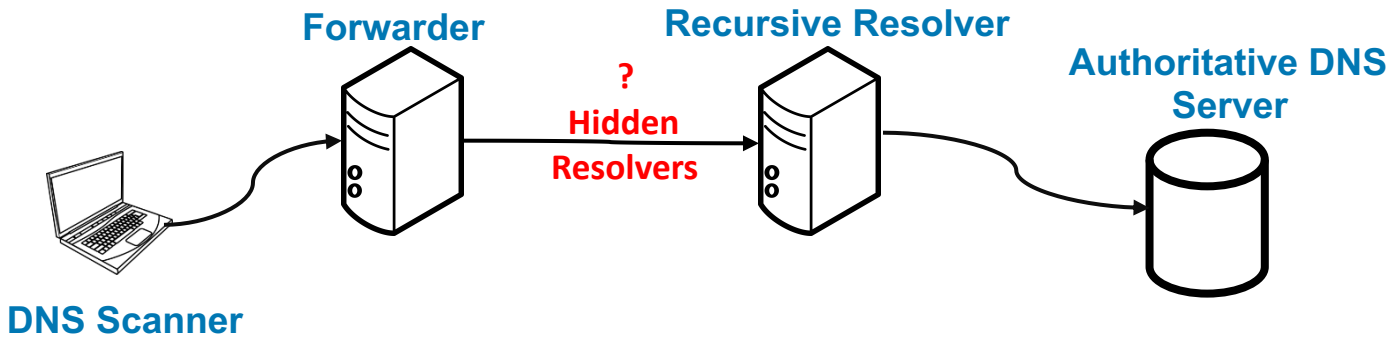

Hidden and Resolver same distance from Forwarder

Hidden further from Forwarder than Resolver

Hidden resolver impact	F/H/R Combinations (Non-Major Public Resolvers)
ECS hurts	7.8%
ECS does not help	19.5%
ECS still helps	72.7%

Conclusions and Take-Aways

- Many resolvers disregard scope caching restrictions
  - Disrupts ADNS's server selection policy & violates RFC
- Pitfalls can make ECS useless or harmful to user mapping
  - Unroutable ECS source prefixes
  - Short ECS source prefixes
  - Hidden resolvers
- Check the paper for our other observations

Questions?

Rami Al-Dalky*, Michael Rabinovich*, and Kyle Schomp✢

* Case Western Reserve University

✢ Akamai Technologies

EXTRA SLIDES

Datasets

- CDN Dataset: ECS transactions from all ADNS servers of a major CDN
  - 1-day log: 3.7M total resolvers, 7737 ECS resolvers (3590 whitelisted and 4147 non-whitelisted)
  - Non-whitelisted:
    - 1.5B total queries
    - 847M queries with an ECS option, **including 3067 from one Chinese AS**
    - 4002 IPv4 and 145 IPv6 resolvers
- Public Resolver/CDN Dataset: ECS transactions from a major public resolution service to the CDN
  - 3-hour log
  - 3.8B queries
  - 2370 IP addresses
- All-Names Resolver Dataset: ECS transactions between anycasted resolver front-end and egress
  - Contains full client IP addresses and ECS scope of responses
  - 11.1M A/AAAA transactions
  - 76.2K client IP addresses (37.4K IPv4 and 38.8K IPv6)
  - 12.3K /24 IPv4 client subnets and 2.8K /48 IPv6 client subnets
- Active: Scan Dataset
  - 2.7M open resolvers
  - 1.5M produced ECS queries using 1534 egress resolvers
    - Including both open and closed egress resolvers
  - 1256 egress IP addresses are from Google Public DNS

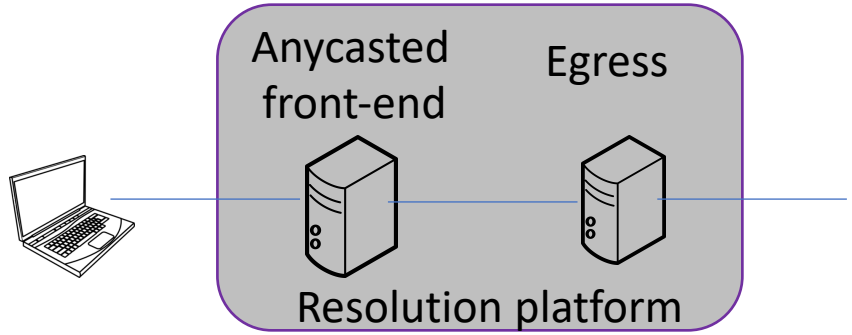

ECS Probing Strategies

- Using CDN dataset
  - CDN's ADNS looks like non-ECS to non-whitelisted resolvers.
  - Dataset represents probing of non-ECS ADNS

Probing Behavior	# resolvers
Always include ECS	3382
Include ECS every N*30 min for one hostname	32
Always include ECS for specific hostnames	258
Include ECS for specific hostnames on a miss	88
Unable to discern	387
Total	**4147**

Source Prefix Lengths

Source Prefix Length	# of Resolvers (Scan dataset)	# of Resolvers (CDN dataset)
18		3
21		60
22	8	19
24	1384	757
24,25,32/jammed last byte		1
24,32/jammed last byte		3
25	1	1
25,32/jammed last byte		78
32/jammed last byte	130	3002
32		221
32 (IPv6)		28
44 (IPv6)		60
48 (IPv6)		56
56 (IPv6)	433	4
64 (IPv6)		1
64,96,128 (IPv6)		3

- A number of resolvers violate 24-v4/56-v6 bit prefix restriction
- Even with jammed last byte, 32-bit prefixes are misleading and incorrect

Honoring Scope Restriction on Caching

- Deliver two successive queries with common /16 and distinct /24 client addr
- Respond to first query with ECS scope of /24. Will the second query be a miss?
- Use different tricks for different resolvers
  - For open egress resolvers that accept ECS option in queries, send appropriate ECS source prefixes
  - Else use open forwarders with suitable IP addresses to probe egress ECS resolvers
  - Else try send queries through hidden resolvers with suitable IP addresses (see paper)
- Able to study 202 out of 278 non Google egress ECS resolvers plus one host from Google IP range
  - 103 resolvers don’t obey scope caching restrictions
  - 76 resolvers forward /24 prefixes and honor the scope properly
  - 15 resolvers accept and forward >/24 prefixes
  - 8 resolvers truncate incoming prefixes to at most /22
  - 1 misconfigured resolver that sends an ECS prefix from 10.0.0.0/8

Impact on cache size – Public DNS service

- Simulation of the resolvers' cache with and without ECS
- Reflects cache size for the CDN-destined transactions
- Assumes isolated caches at resolvers
- No premature evictions

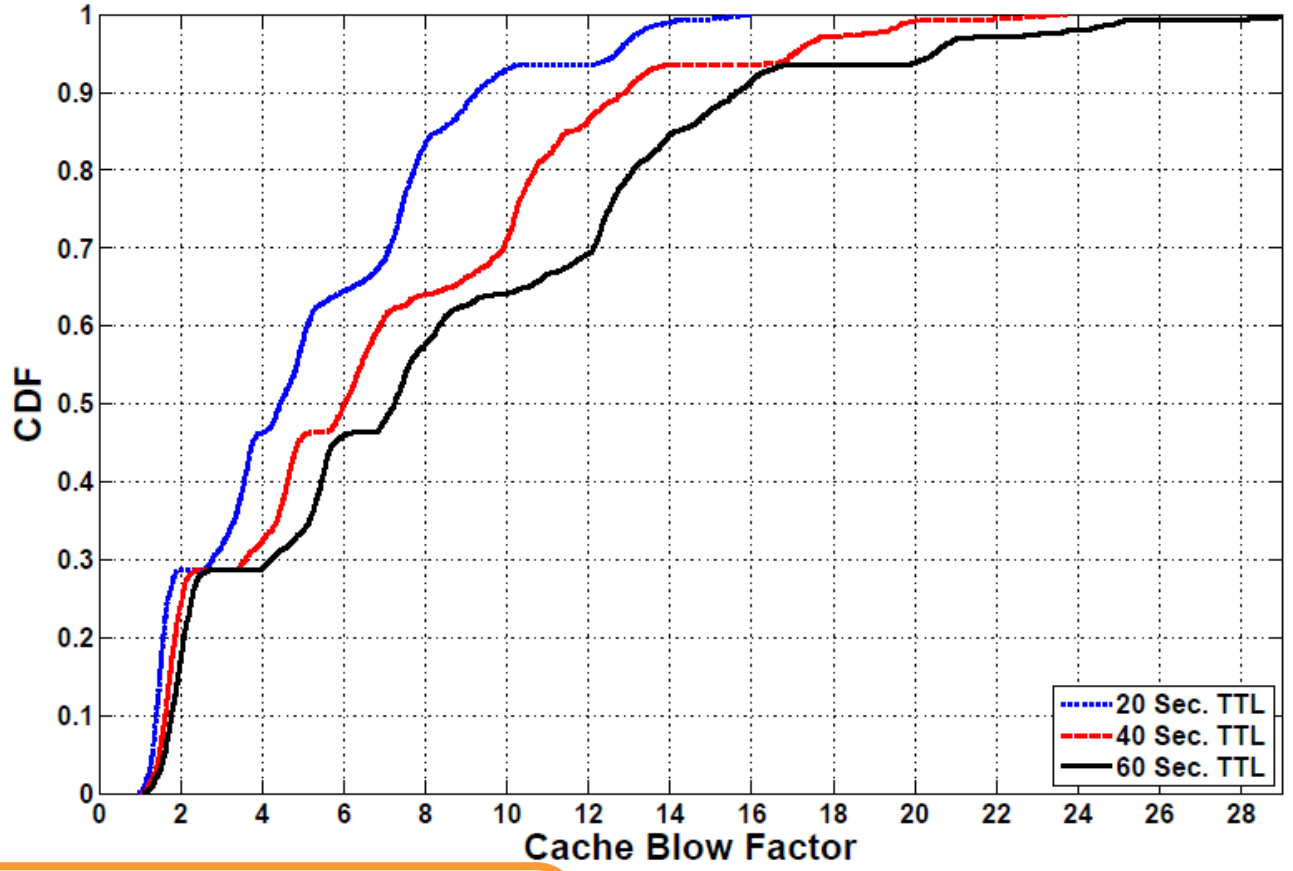

Large TTL values have a significant impact on the recursive resolvers' cache size

Impact on cache size – Busy resolver

- Simulation of the resolver' cache with and without ECS
- Reflects cache size for all ECS transactions to all destinations
- No premature evictions

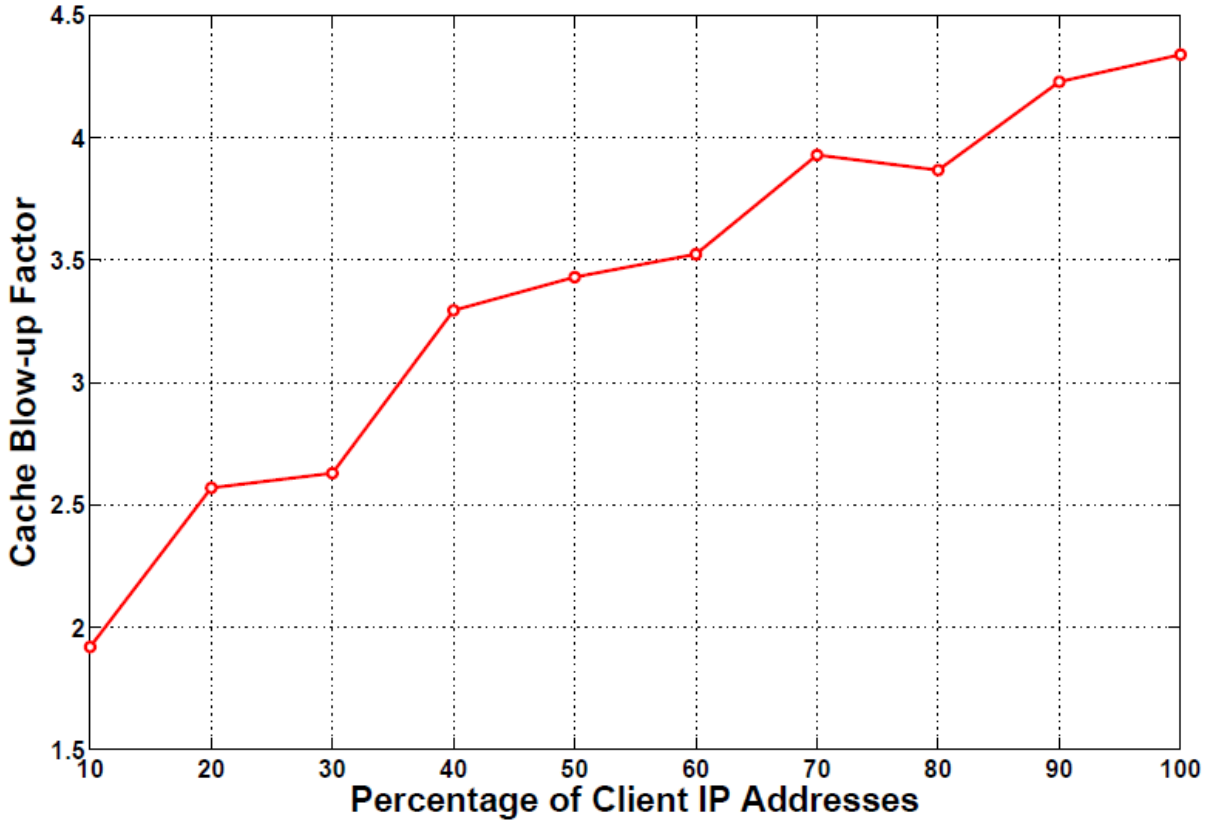

Diverse client population would result in a large increase of the cache size of recursive resolvers

Impact on cache hit rate – Busy resolver

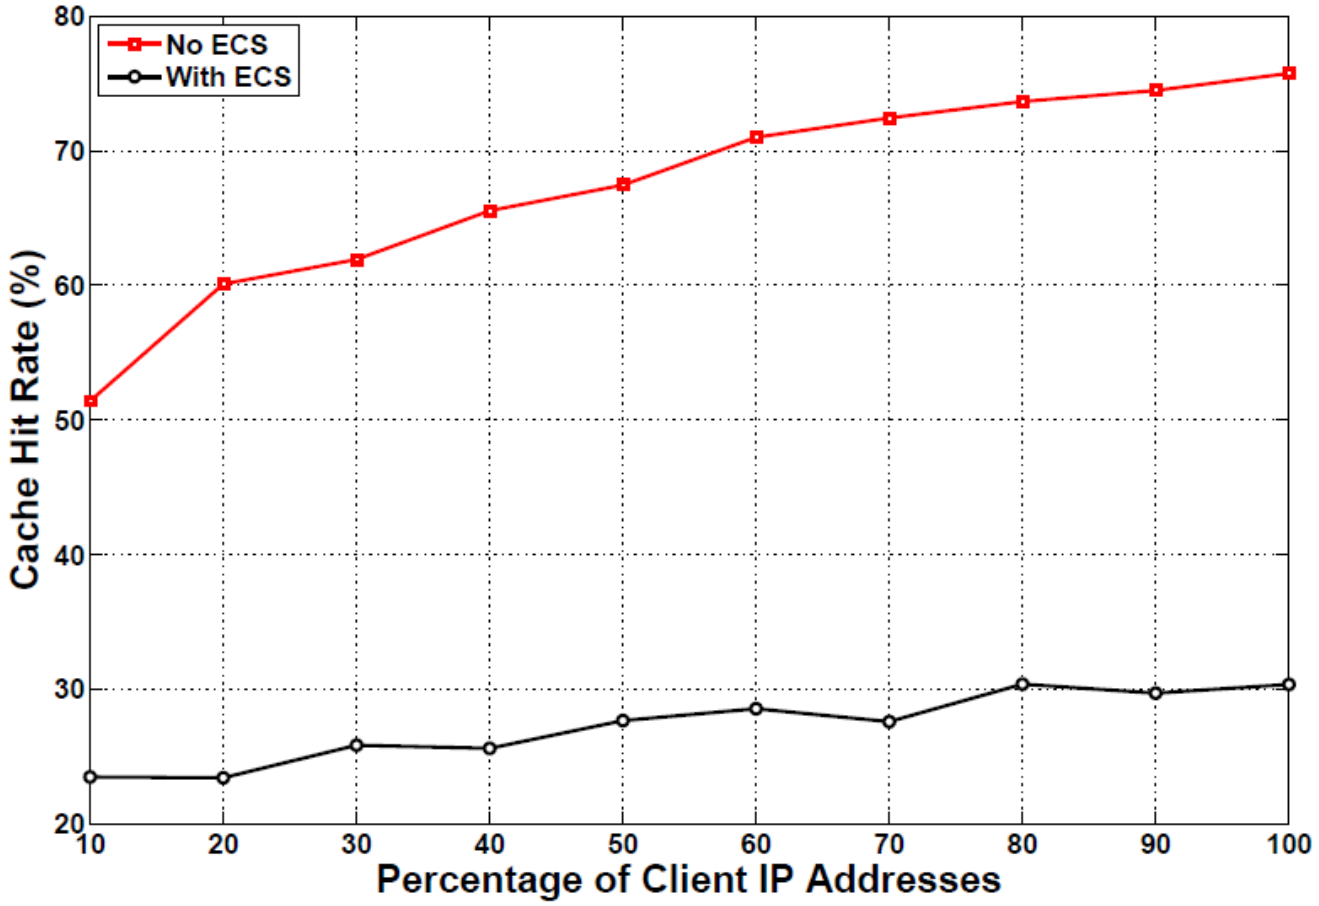

Large ECS impact on hit rate

CNAME Flattening

- CNAME flattening allows authoritative nameserver to define a CNAME at the domain apex.

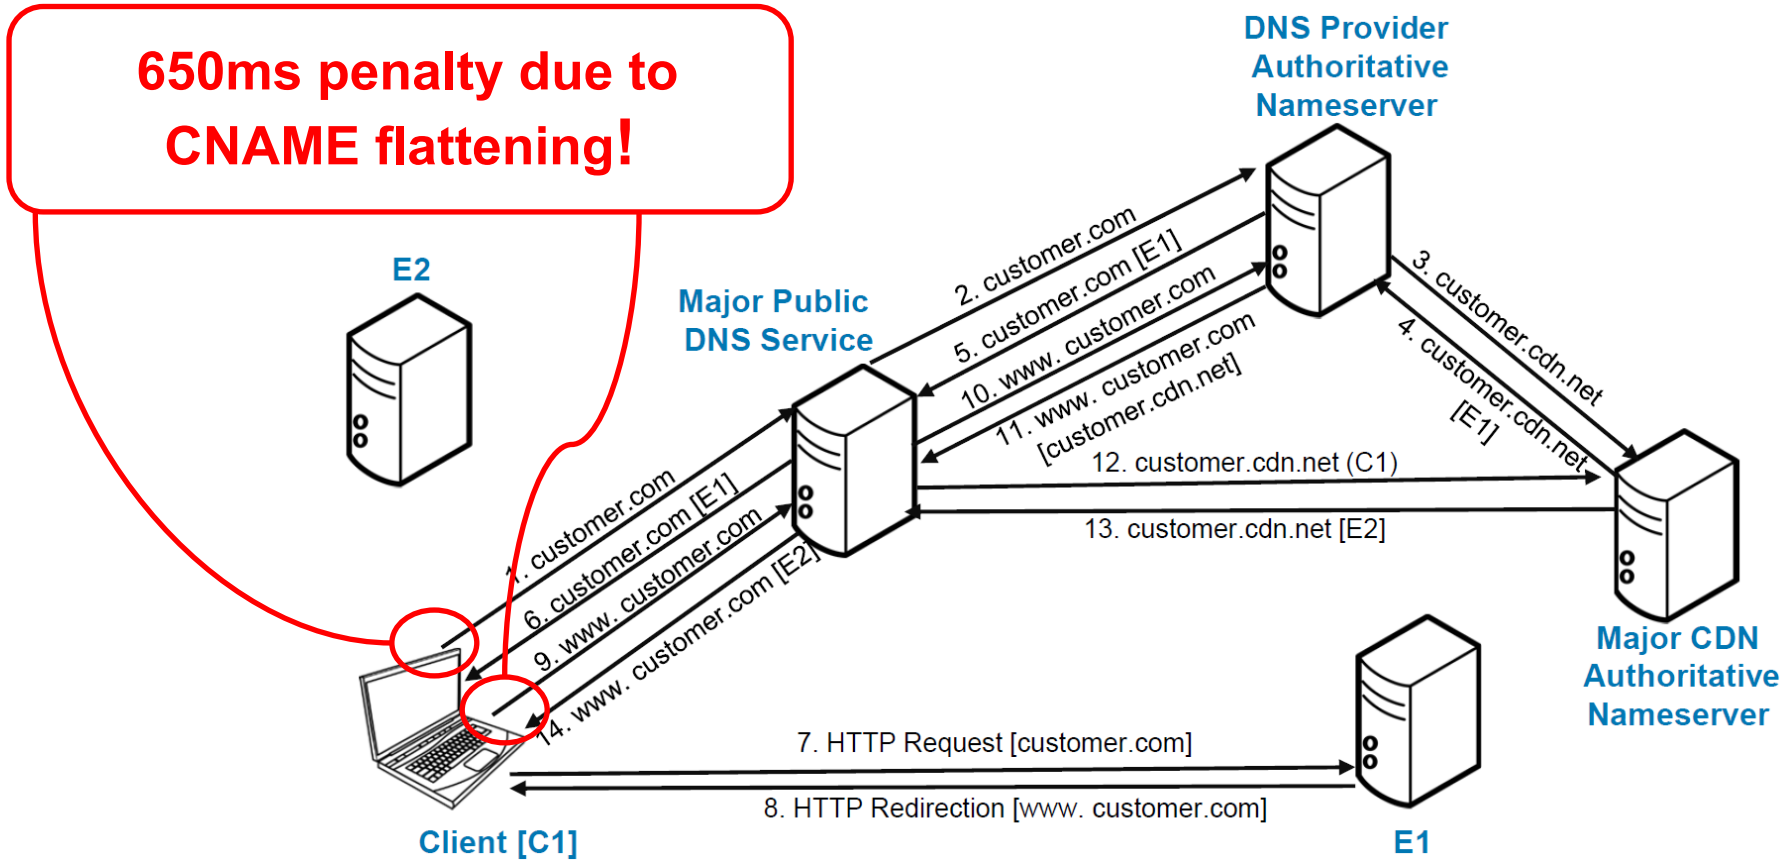

Impact of Hidden Resolvers on ECS

Major public resolver

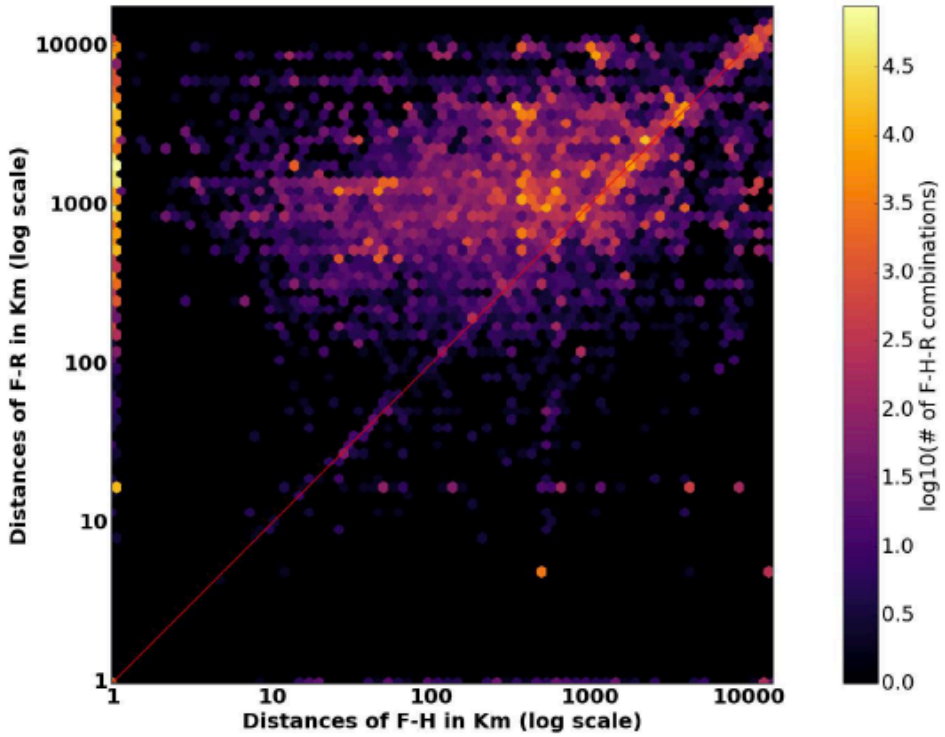

Other resolvers

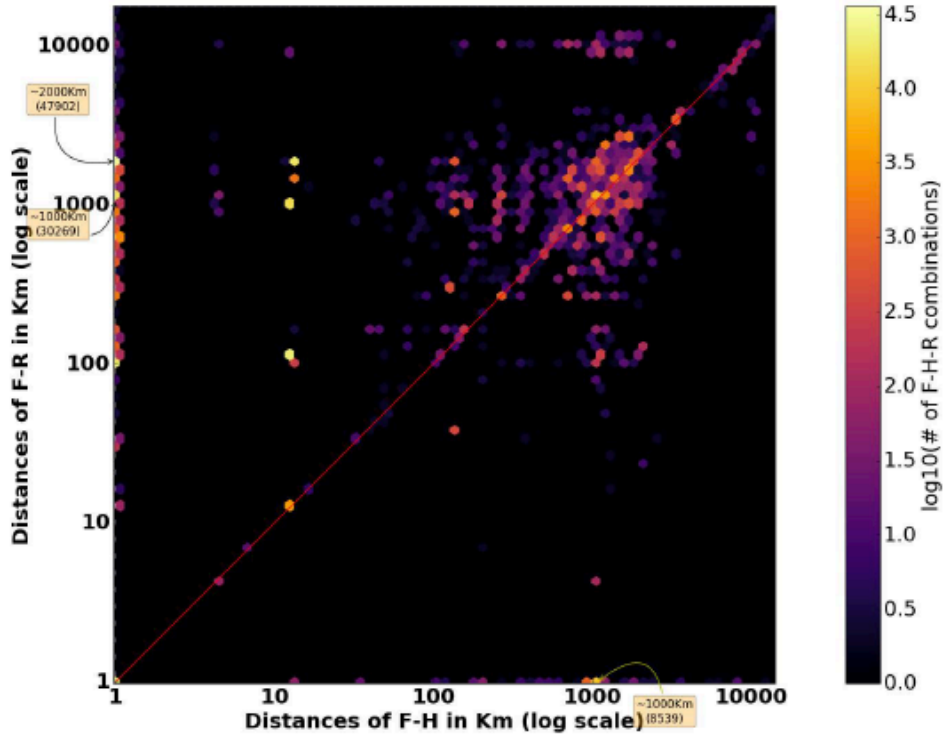

Hidden resolver impact	F/H/R Combinations (Major Public Resolver)	F/H/R Combinations (Other Resolvers)
ECS hurts	8%	7.8%
ECS does not help	1.3%	19.5%
ECS still helps	90.7%	72.7%