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

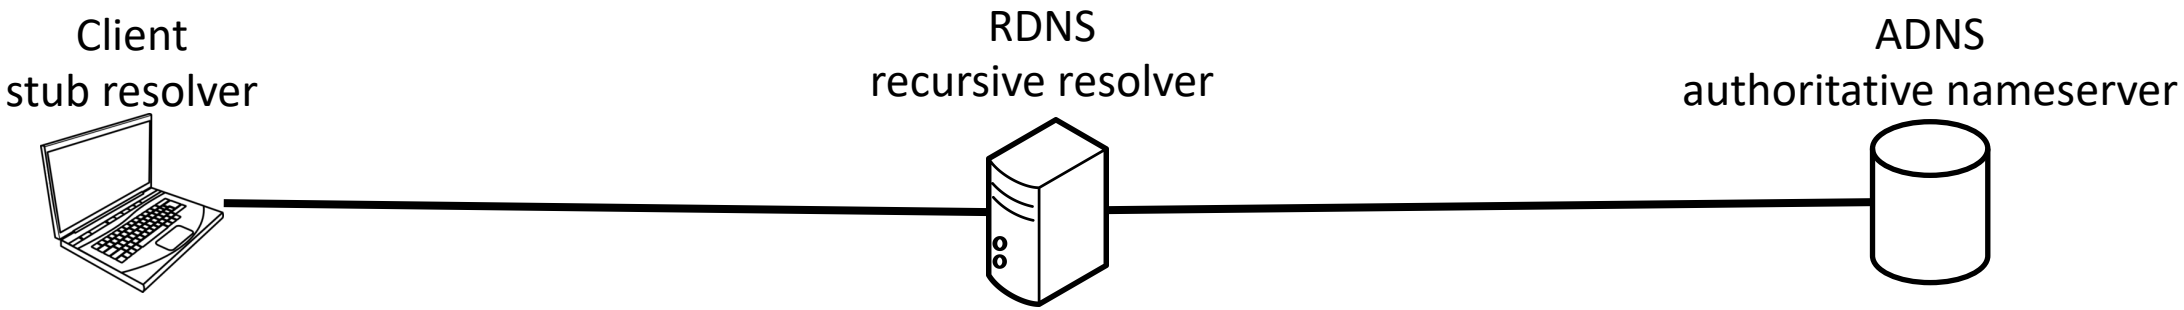

- ECS Purpose
  - Enable CDN server selection by ADNS based on client subnet
- ECS Option in DNS queries from resolvers to ADNS includes
  - Client IP address prefix
  - Source prefix length

  → 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

Goals of Study

- Exploring resolvers' ECS behavior
  - Resolver ECS support?
  - Probing behavior for ADNS support?
  - Source prefix lengths used?
  - Adherence to ECS scope's cache restrictions?
- ECS deployment pitfalls

Datasets

- Logs from Akamai's ADNS
  - Study IP addresses of recursive resolvers that send ECS option
- Internet-wide scans of recursive resolvers
  - Actively measure recursive resolver behavior

ECS Support

- Akamai's ADNS dataset
  - 3.7M total resolvers
  - 7737 sent at least one query with ECS (0.2%)
    - Majority major public DNS services
    - 3067 (40%) from Chinese AS. General skew in ECS support towards China
- Internet scan dataset
  - 2.7M open resolvers that forward to…
  - 1534 ECS-enabled recursive resolvers

ECS Probing Strategies

- Akamai's ADNS looks non-supporting to non-allowlisted resolvers.
- Unsolicited ECS queries represent probing strategies

Probing Behavior	#	% out of 4147
Always include ECS	3382	82%
Include ECS every N•30 min for one hostname	32	1%
Always include ECS for specific hostnames	258	6%
Include ECS for specific hostnames on miss	88	2%
Unable to discern	387	9%

Also, some resolvers send ECS with *all* query types, e.g., NS, where RFC 7871 suggests not to (section 7.4).

Probing is not popular

Source Prefix Lengths

Privacy violations

	Source Prefix Length	# Resolvers (Scan dataset)	# Resolvers (CDN dataset)
IPv4	<24	8	82
	24	1384	757
	25	1	1
	Jammed to 32	130	3084
	32	~	221
IPv6	<48	~	88
	48	~	56
	>48	433	8

Likely misunderstanding of the RFC.
Result is misleading information

Honoring Scope Restriction on Caching

- Attempted to study 278* ECS-enabled recursive resolvers in scan
- Test scope restriction using various measurement tricks (see paper) by:
  - Send query to resolver from client in /24 X and /16 Y
  - Return answer from ADNS with scope prefix length 16
  - Send second query to resolver from client in /24 Z and /16 Y
- Results
  - 76 could not be studied
  - 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

Lie to the ADNS

*excluding a major public DNS service

ECS Deployment Pitfalls

We found several types of real-life resolver setups that interact with ECS in ways that diminish or negate its benefits.

1. Unroutable ECS Prefixes

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

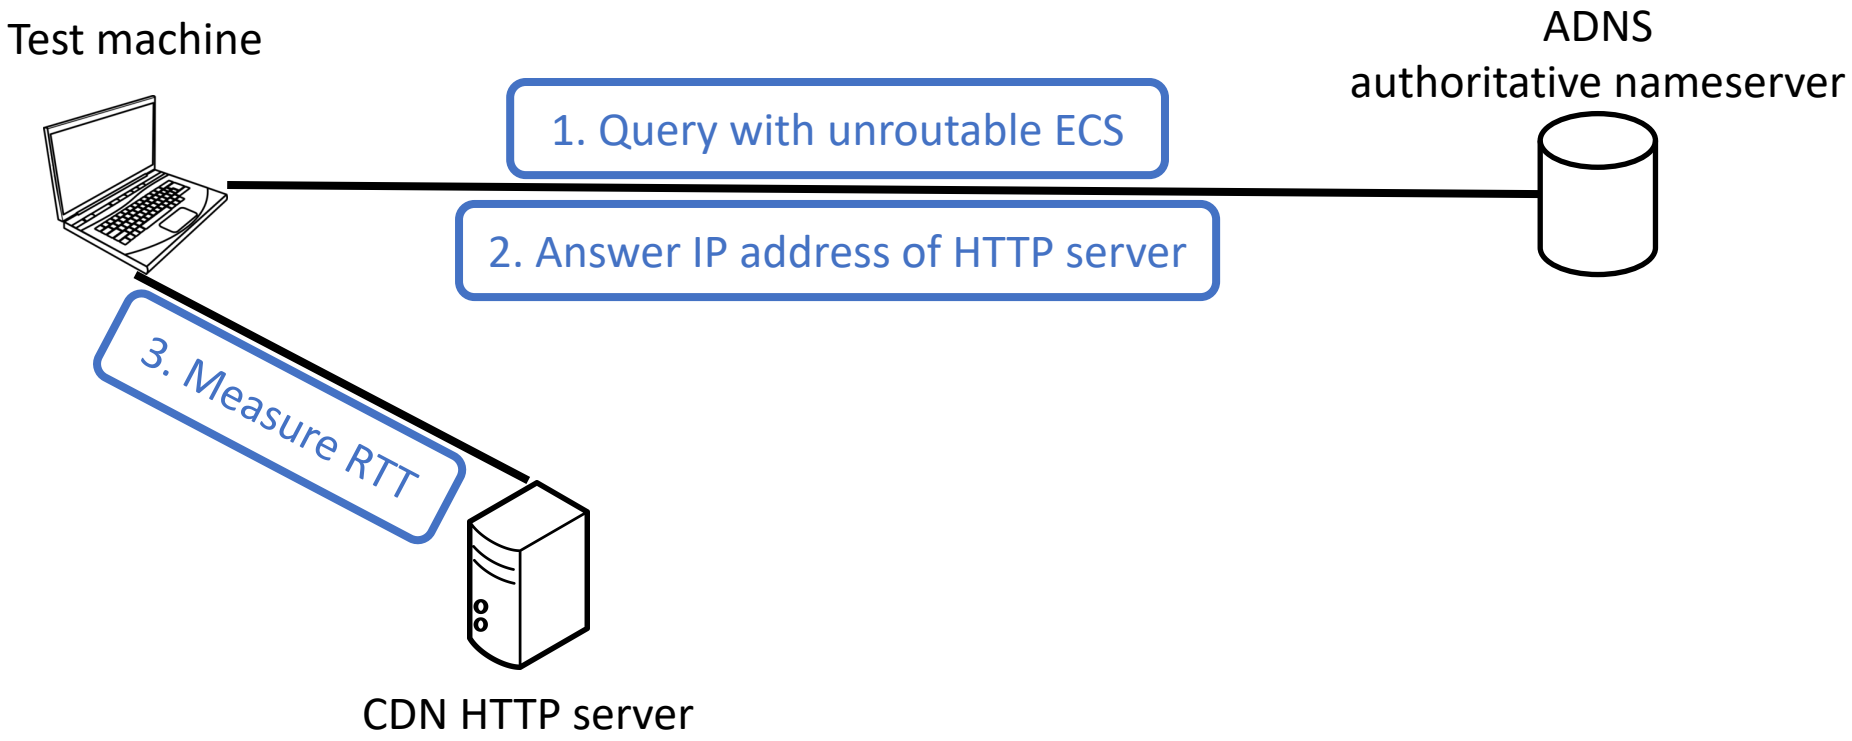

1. Unroutable ECS Prefixes

- 33 resolvers from Internet-wide scan sent queries with loopback as ECS prefix
- Could this confuse ADNSs?
- Send 5 queries from test machine in Cleveland, OH to ADNS for youtube.com:

	ECS Prefix	RTT (ms)	Location
1.	None	35	Chicago, IL
2.	/24 of test machine's IP	35	Chicago, IL
3.	127.0.0.1/32	155	Switzerland
4.	127.0.0.0/24	47	Mountain View, CA
5.	169.254.252.0/24	285	South Africa

Sending unroutable ECS prefix results in worse application performance than not sending ECS at all

1. Unroutable ECS Prefixes

- RFC 7871, section 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.”
- Further add language:
  - Recursive resolvers SHOULD send routable prefixes in the ECS option (or not send the option if that’s not possible).

2. Impact of Source Prefix Length

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

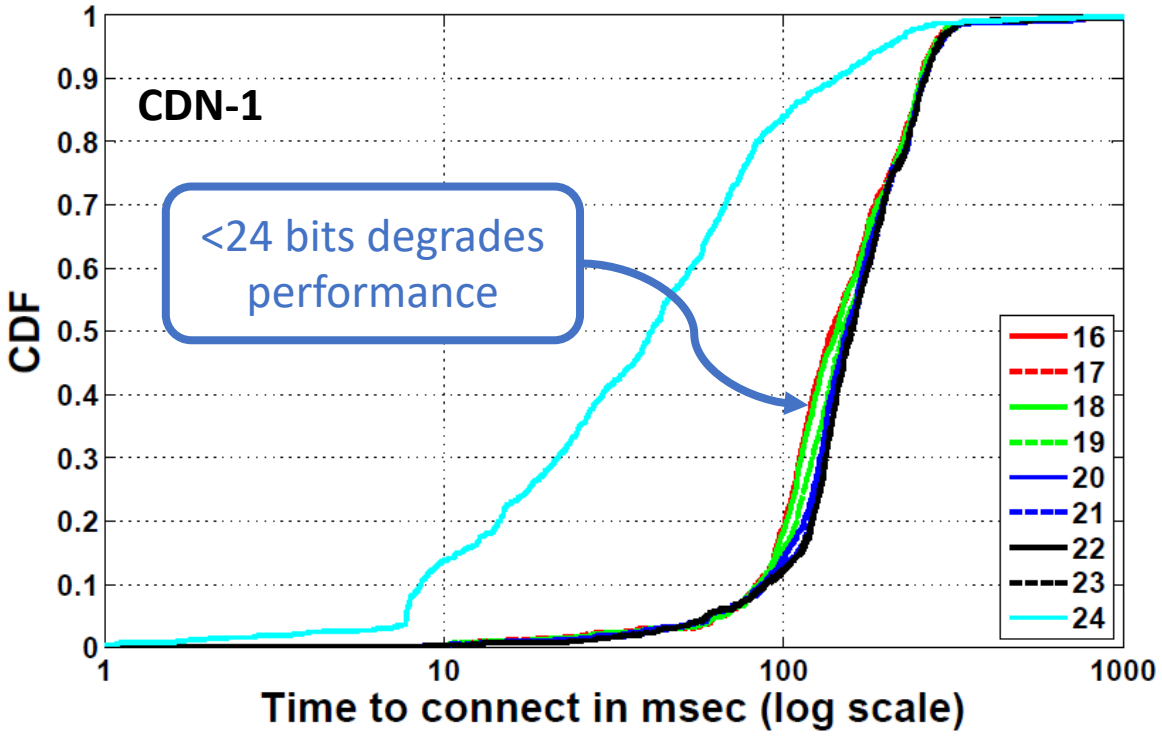

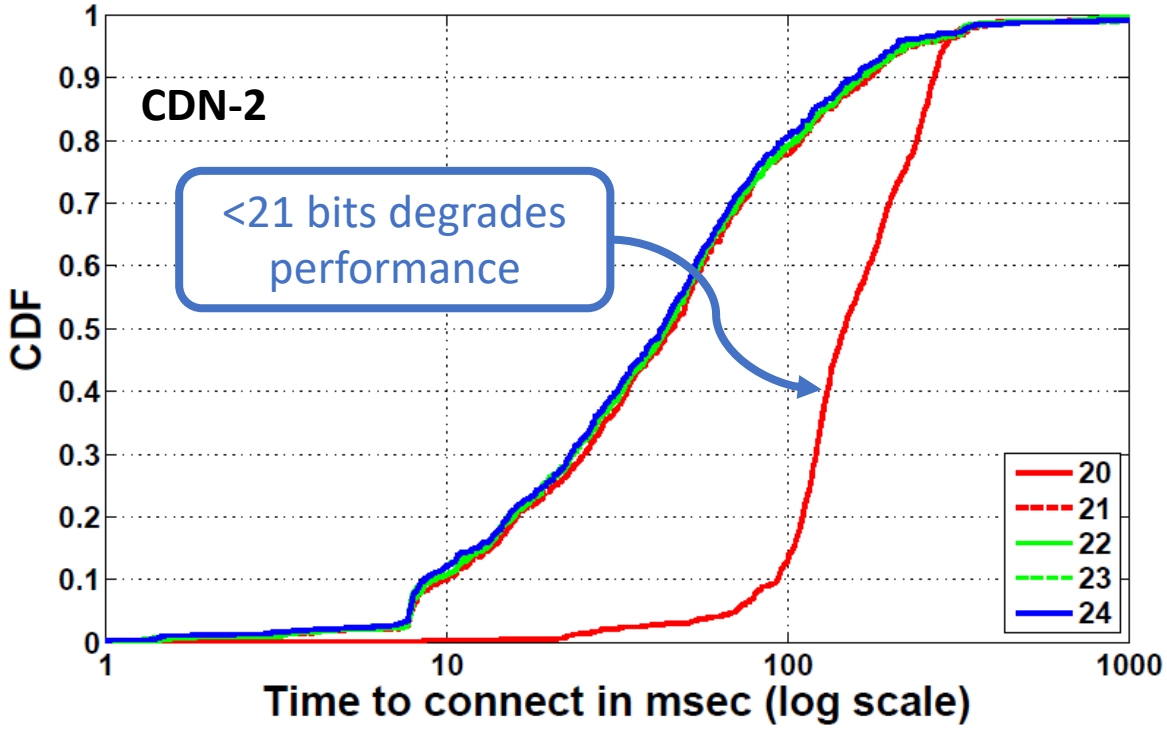

2. Impact of Source Prefix Length

- Sending > 24 bits
  - `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 < 24 bits
  - Negatively impacts some CDNs
- Resolvers can either keep source prefix length state per ADNS

or

use 24 bits for everyone

3. Impact of Hidden Resolvers

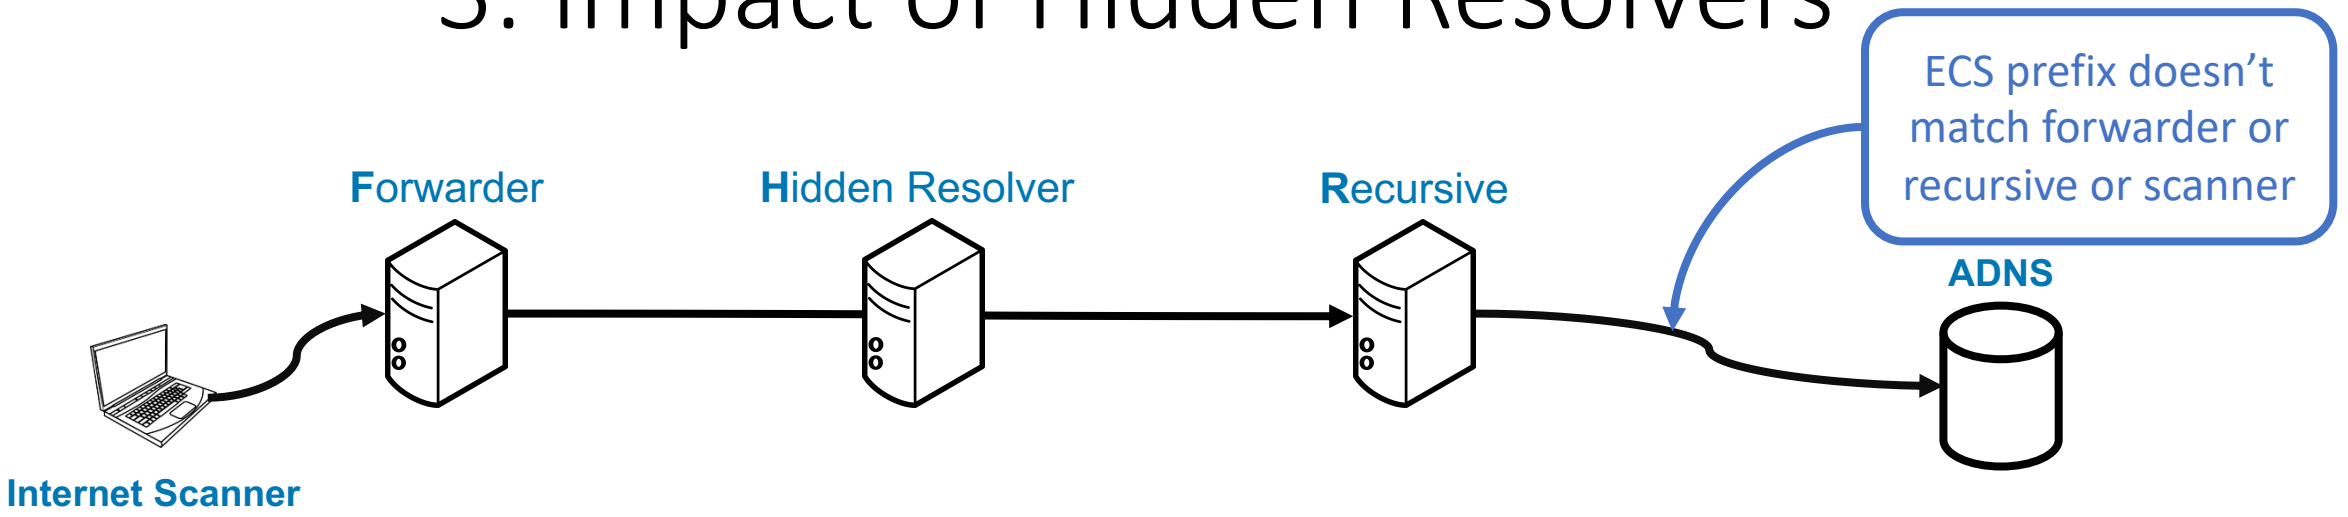

- ADNS uses ECS prefix – instead of forwarder or recursive IP address – for server selection
- Is the recursive resolver or the hidden resolver closer to the forwarder?

3. Impact of Hidden Resolvers

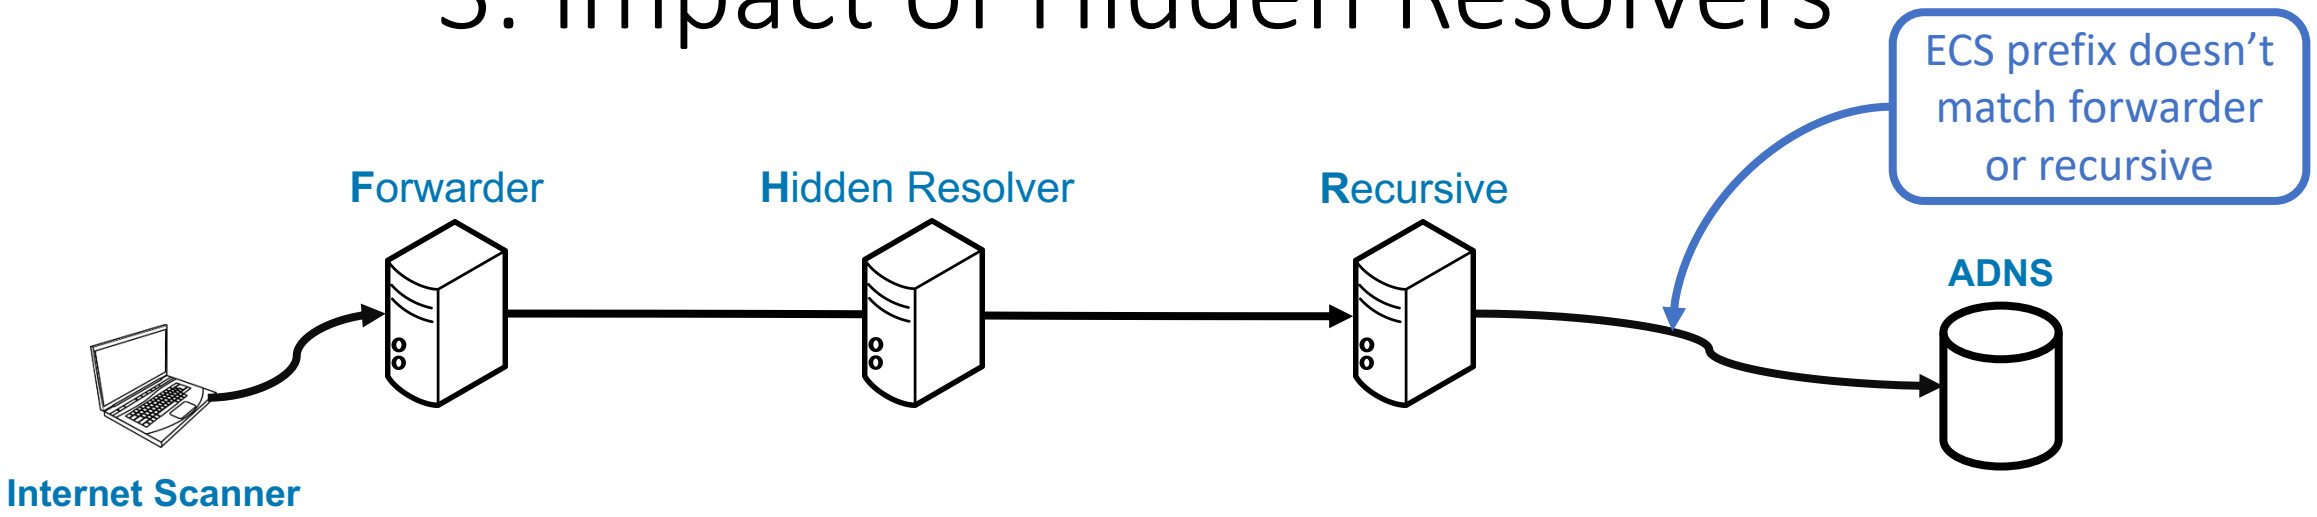

If not applied carefully, ECS may offer no benefit or actively harm performance

Scenario	% Measurements
Distance (F,H) > Distance (F,R)	7.8%
Distance (F,H) = Distance (F,R)	19.5%
Distance (F,H) < Distance (F,R)	72.7%

Summary & Takeaways

- Deployment is not widespread
  - Existing deployments concentrated to a few public DNS providers / China
- Probing for ECS support is virtually unused
  - Resolvers either send ECS in all queries or none
- Evidence of privacy erosion via long ECS prefixes
  - Research needed to understand the significance
- 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 prefixes
  - Too short ECS prefixes
  - Hidden resolvers

Questions?

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

* Case Western Reserve University

✢ Akamai Technologies