DNS Resolvers Considered Harmful

Kyle Schomp, Mark Allman, and Michael Rabinovich

CASE WESTERN RESERVE UNIVERSITY EST. 1826

INTERNATIONAL COMPUTER SCIENCE INSTITUTE

***DNS resolvers abstract complexity and offer the possibility of improved performance and better scalability*.**

Why are they harmful?

Resolvers Are Vulnerable to Cache Injection

- Kaminsky vulnerability discovered in 2008, 16% of resolvers vulnerable to Kaminsky attack in 2012[1]
- Preplay attack discovered in 2014, millions of wifi routers acting as resolvers are vulnerable[1]
- Shulman attack discovered in 2013, 79% of resolvers vulnerable[2]

[1] Schomp, Kyle, Tom Callahan, Michael Rabinovich, and Mark Allman. "Assessing DNS Vulnerability to Record Injection." *PAM* (2014).

[2] Herzberg, Amir and Haya Shulman. “Fragmentation Considered Poisonous, or: One-domain-to-rule-them-all.org.” *CNS* (2013).

Resolvers Should Not Be Trusted

- Resolvers rewrite responses for non-existent domains, effects 24% of clients[1]
- Others intentionally participate in hijacking domains (e.g., Paxfire in 2011[3])
- Many countries use resolvers to enable censorship[4]
- ...yet we give them access to sensitive user information

[3] Weaver, Nicholas, Christian Kreibich, and Vern Paxson. "Redirecting DNS for ads and profit." *FOCI* (2011).

[4] Verkamp, John-Paul, and Minaxi Gupta. "Inferring mechanics of web censorship around the world." *2nd FOCI* (2012).

Resolvers Obscure Clients

- Client-resolver location mismatch, 7.5-15% of clients suffer reduced performance due to wrong CDN edge server[5]
- Resolvers hide client population reducing the effectiveness of DNS-based load balancing

[5] Huang, Cheng, Ivan Batanov, and Jin Li. "A practical solution to the client-LDNS mismatch problem." *SIGCOMM* (2012).

Resolvers Used In Amplification Attacks

- 24 million open resolvers on the Internet today[6]
- DNS amplification attacks are massive[7] and growing in popularity[8]

[6] http://openresolverproject.org/

[7] http://www.zdnet.com/the-largest-ddos-attack-didnt-break-the-internet-but-it-did-try-7000013225/

[8] NSFOCUS 2014 Mid-Year DDoS Threat Report. http://en.nsfocus.com/2014/SecurityReport_0922/190.html

Existing Solutions

Solutions to *some* of these issues, *e.g.*,

- Random transaction IDs and source ports mitigate guessing attacks such as Kaminsky
- Closing open resolvers thwarts amplification attacks and preplay
- EDNS-client-subnet (ECS) reveals more information about clients behind a resolver
- DNSSEC provides data integrity and protects against all fraudulent records

Existing Solutions Aren’t Working

- Resolvers are still vulnerable to Kaminsky 6 years after its publication
- Millions of open resolvers on the Internet
- Current DNSSEC standard released 10 years ago, but deployment is still low
- *Vulnerabilities still being discovered*

Looking In Another Direction

- Many security issues that are not being addressed currently
- Much of the attack surface lies on the resolvers

Why don’t we just get rid of resolvers?

Current Situation

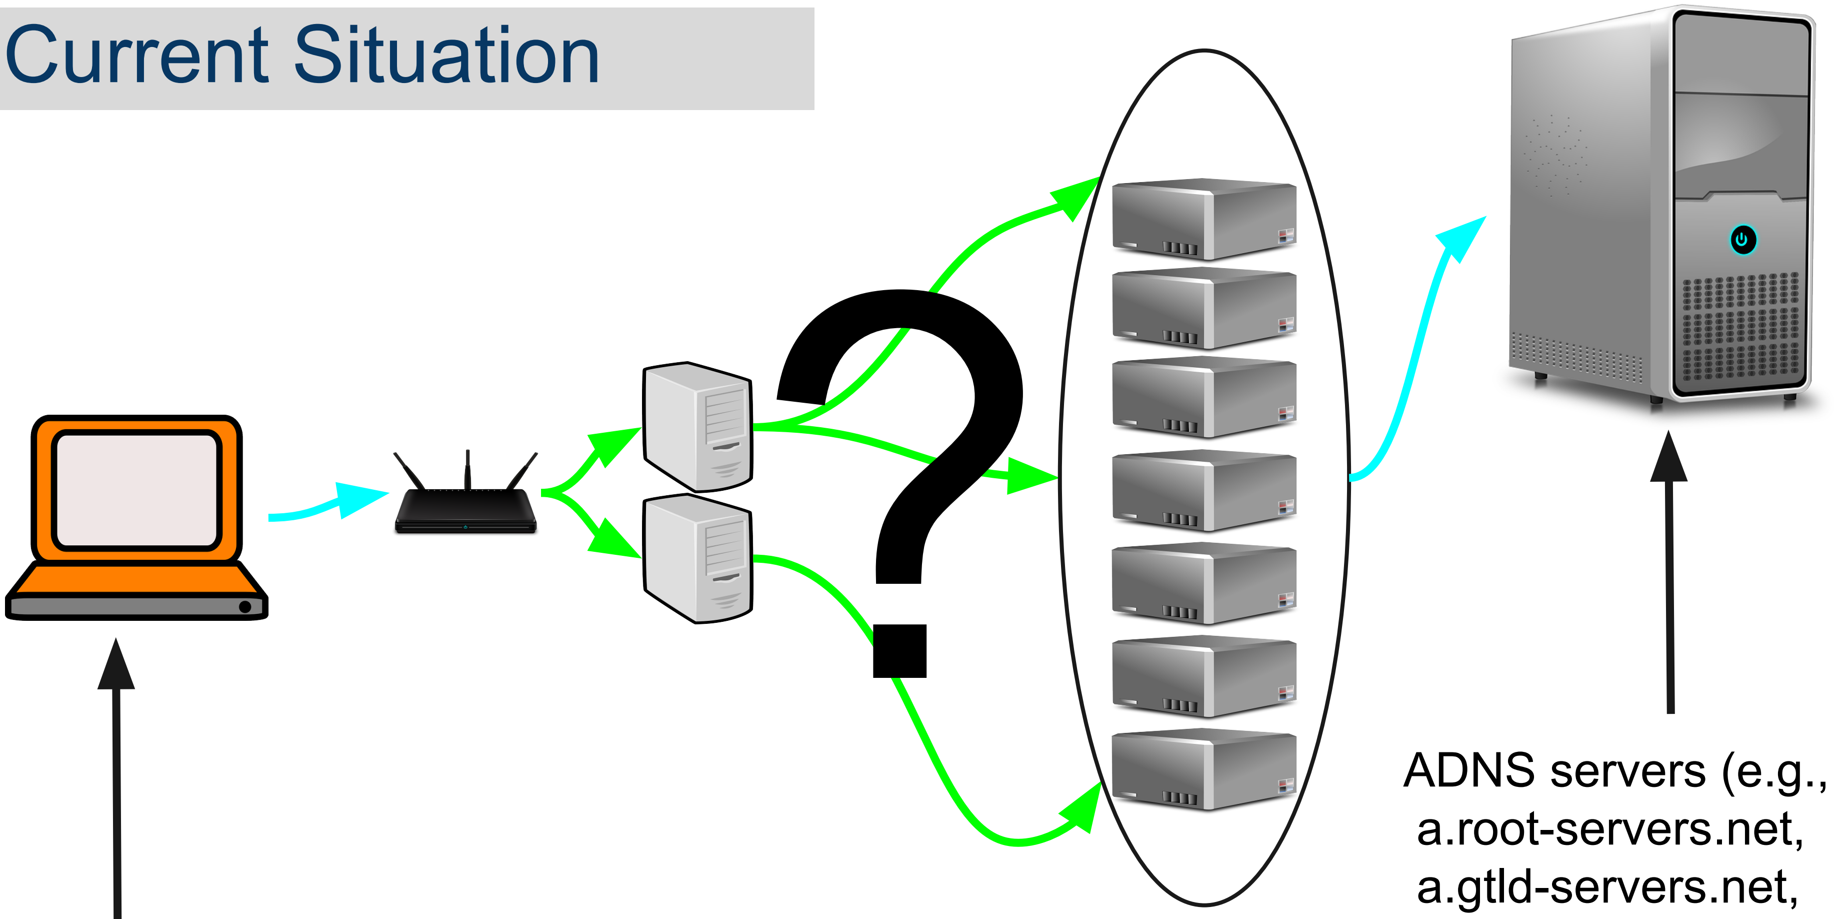

ADNS servers (e.g., a.root-servers.net, a.gtld-servers.net, ns1.google.com)

You (on your laptop)

Client Resolution

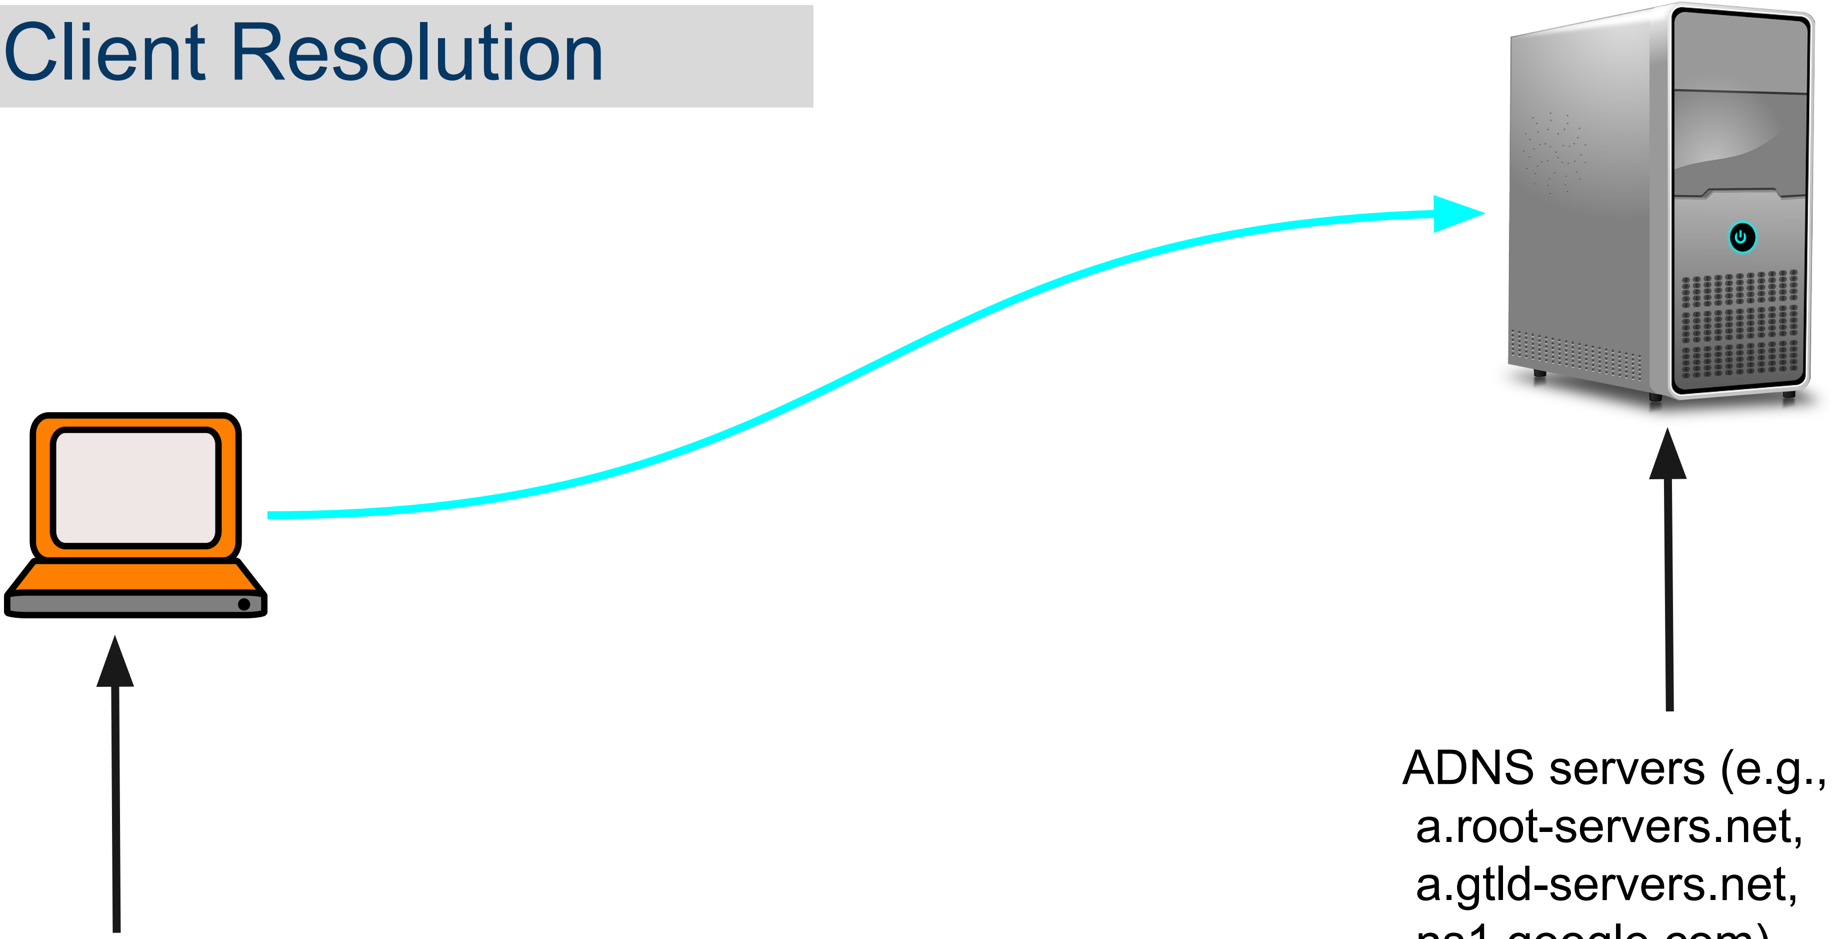

You (on your laptop)

ADNS servers (e.g.,
a.root-servers.net,
a.gtld-servers.net,
ns1.google.com)

What do we gain?

- ✓ Reduces system complexity
- ✓ Removes the target of cache injection attacks
- ✓ Client resolution not vulnerable to same attacks
- ✓ Benefits CDN load balancing and server selection

What do we lose?

- ✗ Resolver caches provide performance to the clients
- ✗ ...and scalability to the system
- ✗ Resolvers anonymize clients

Measuring The Impact

- Trace driven simulations to estimate client resolution’s negative impact
- The data
    - Network of approximately 100 residences
    - 2 recursive resolvers
    - 4 months of observations
    - Recursive resolutions of each domain name in the data

Effect on Performance

DNS resolvers can reduce resolution time due to shared caching.

Resolution times in trace vs. in simulated client resolution

Simulated Resolution Time

Resolutions take a bit longer.

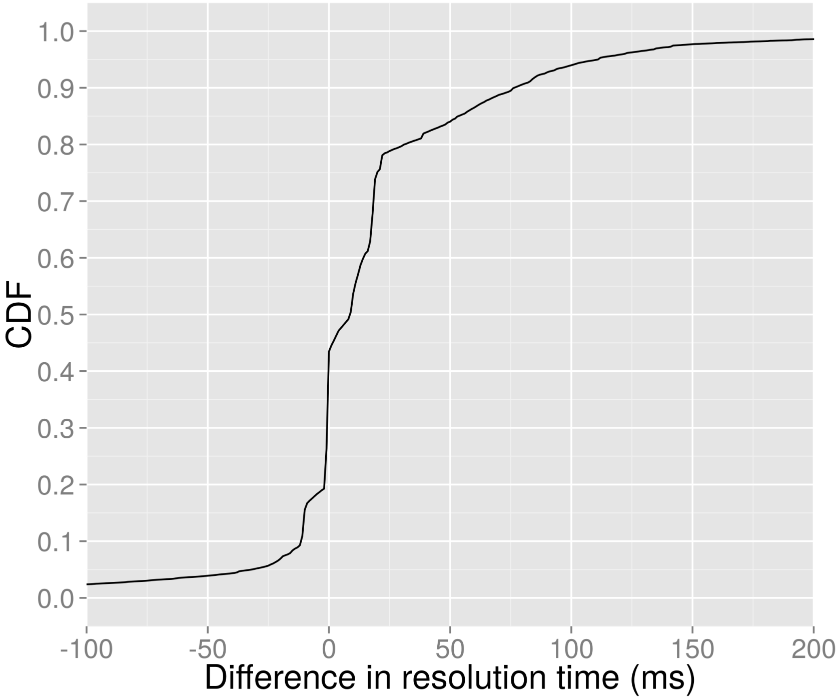

Simulated Resolution Time

Resolutions take a bit longer.

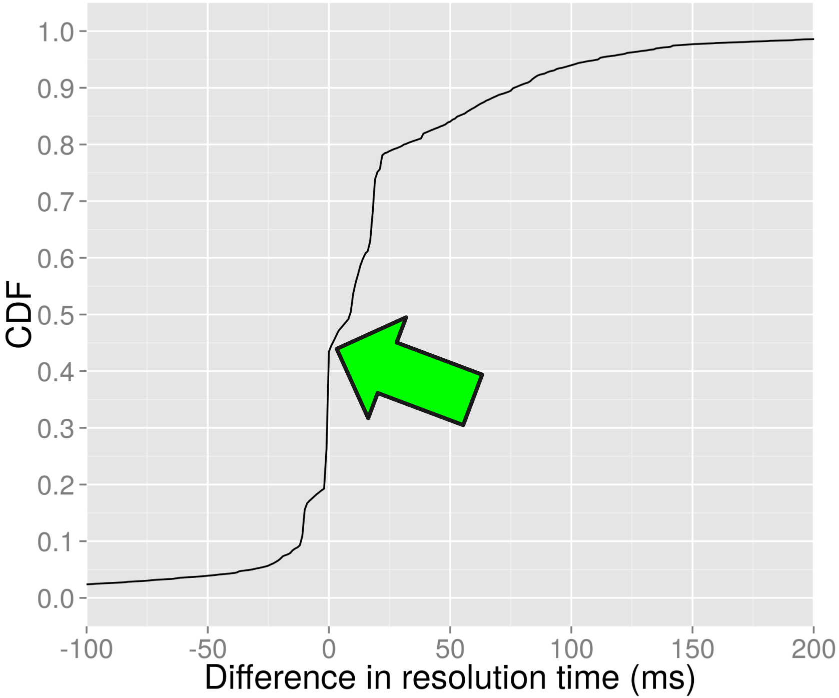

Simulated Resolution Time

Resolutions take a bit longer.

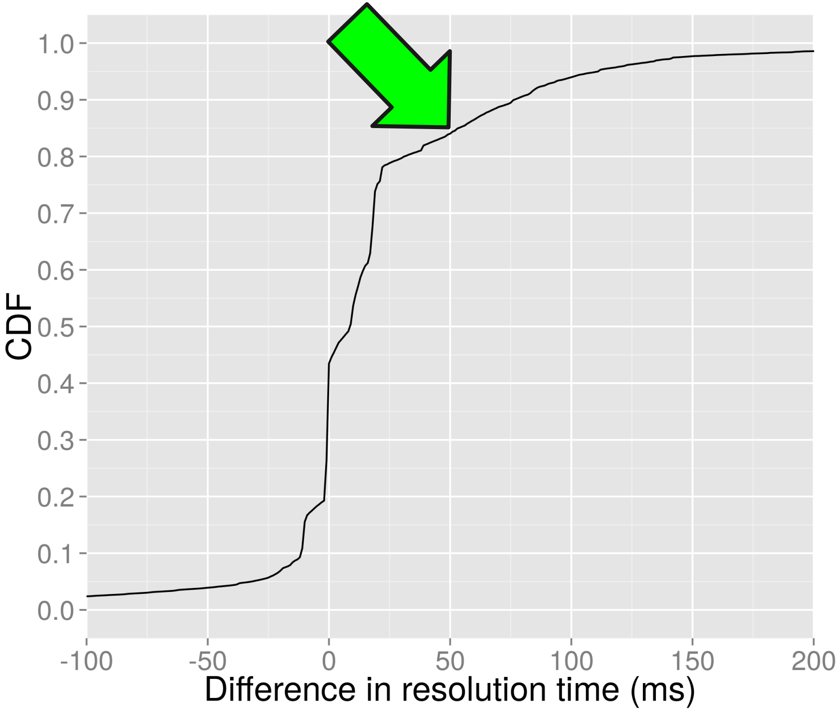

...But there's some slack

DNS responses are *not* used immediately.

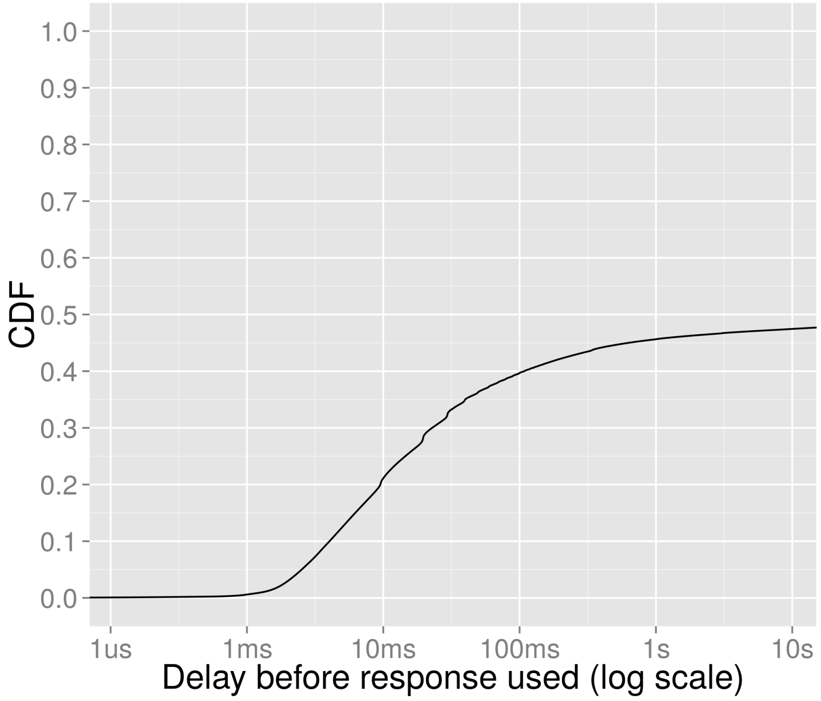

...But there’s some slack

DNS responses are *not* used immediately.

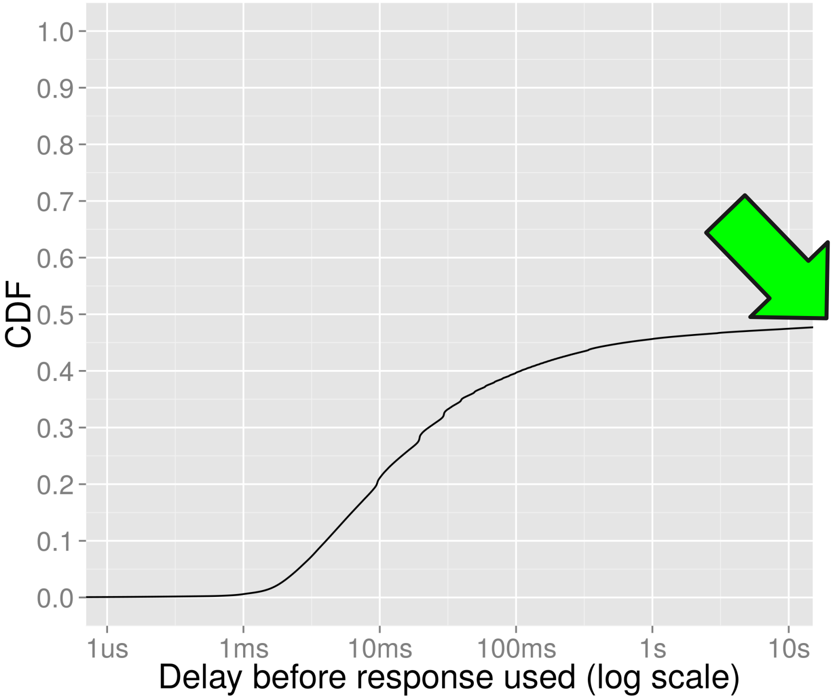

...But there's some slack

DNS responses are *not* used immediately.

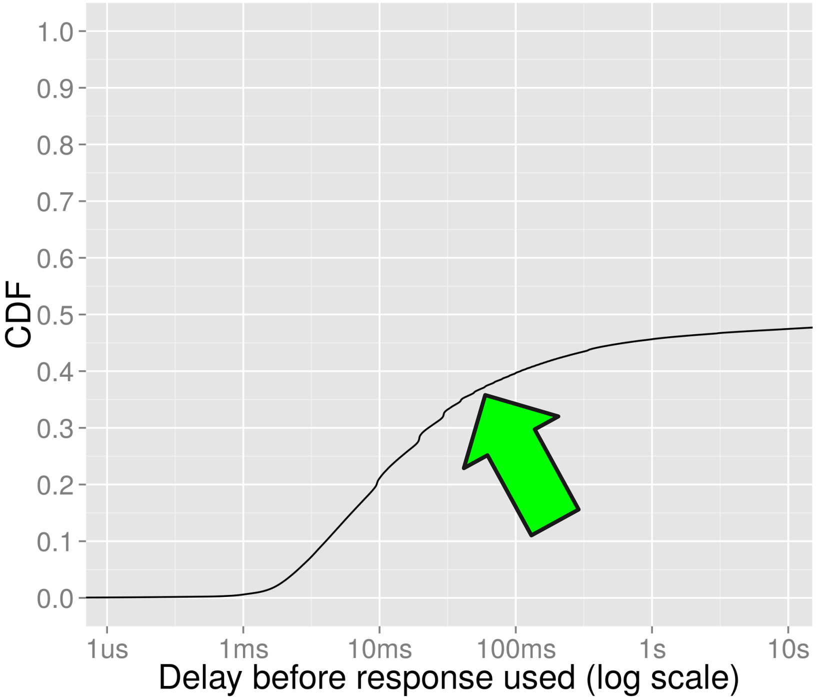

Delay On Connections

Only a small % of connections are delayed at all!

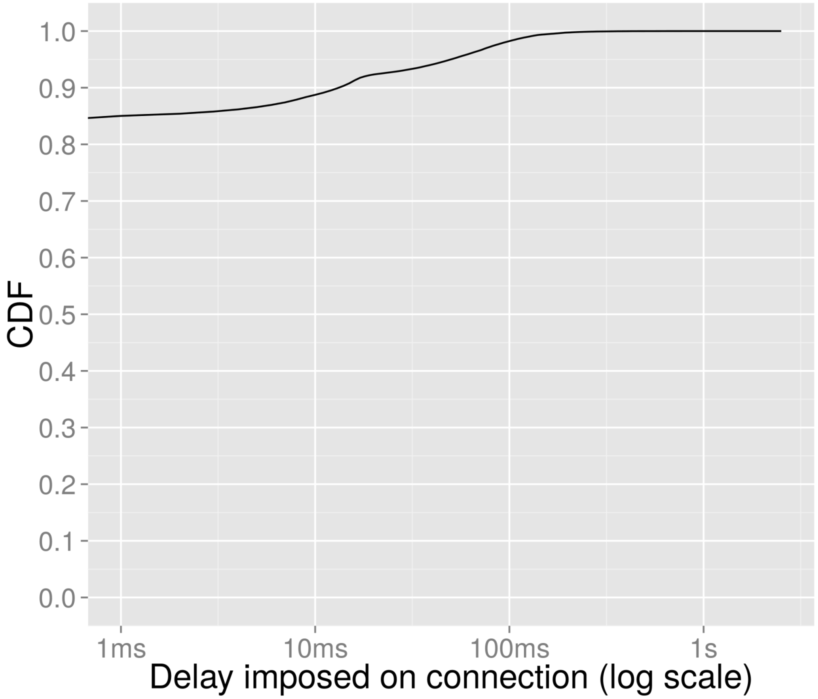

(more details in the paper)

Delay On Connections

Only a small % of connections are delayed at all!

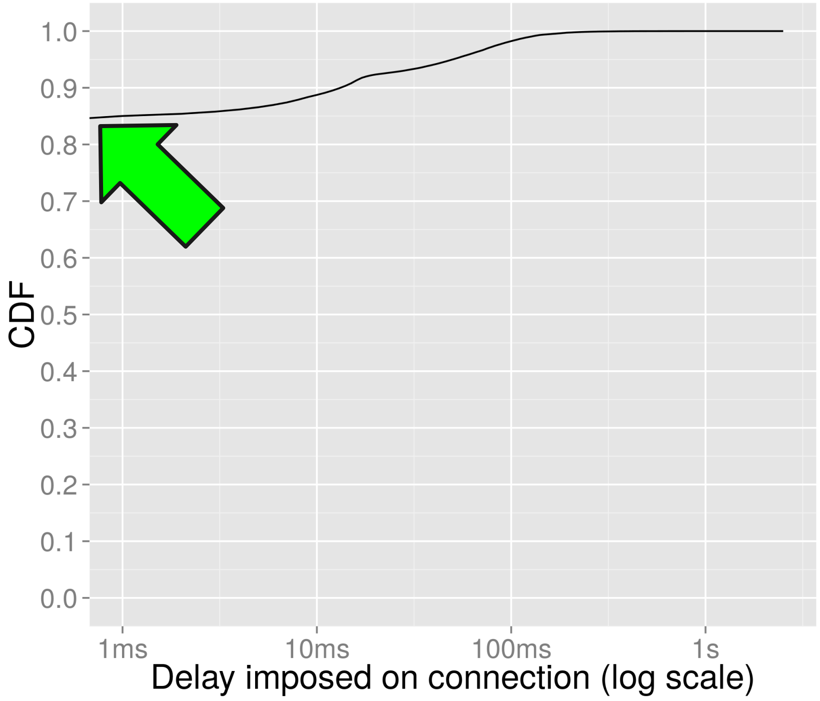

(more details in the paper)

Delay On Connections

Only a small % of connections are delayed at all!

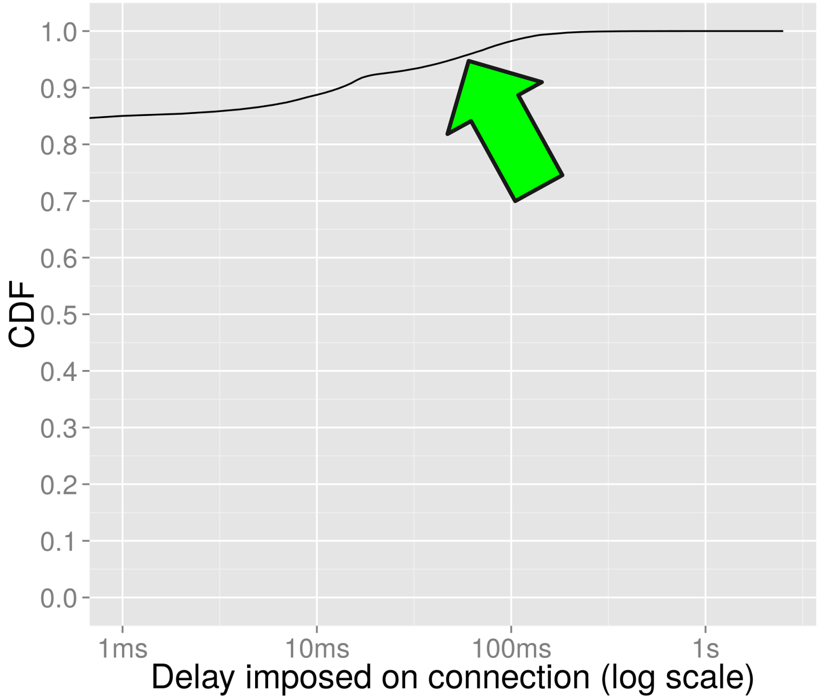

(more details in the paper)

Finding #1:

Performance impact of client resolution is small

Effect on Scalability

DNS resolvers reduce number of resolutions reaching authoritative servers

Resolutions per authoritative domain in trace vs. in client resolution

Load on Auth. Domains

93% of authoritative domains will not see an increase in load

~but~

popular domains *(e.g., com, google.com)* will

- ○ use *com* as exemplar

com Domain Load

- Average load increases by 3.41 times!
- Peak load only increases by 1.14 times
- Which is more representative of impact on *com* domain?
  - Uncertain, let’s make both manageable

Increase Record Time-to-Live

- SLD records normally have 2 days TTLs
- Roughly 1.1% of those records change during a week
- What happens when the TTL is 1 week?

Average Load	3.41 => 2.13 times trace load
Peak Load	1.14 => 1.03 times trace load

Increase Questions Per Query

- Currently 1 question per DNS query
- Protocol can support multiple questions
- What happens when we ask 2 questions per query?

Average Load	3.41 => 1.61 times trace load
Peak Load	1.14 => 1.06 times trace load

(reduces number of packets, not number of queries)

Increase TTL And Questions

- What happens when TTL is 1 week and we ask 2 questions per query?

Average Load	3.41 => 1.33 times trace load
Peak Load	1.14 => 1.06 times trace load

Finding #2:

Scalability impact of client resolution is manageable

Final Thoughts

- Removing resolvers offers many advantages
- ...and small loses
  - Loss of anonymity in queries
  - Increase in authoritative domain load
- DNS prefetching has reduced reliance upon shared caches

?

Thank you!

email me at kgs7@case.edu