DNS Privacy in Practice - Measuring Deployment of DoT, DoH, and TFO

BY

Casey Deccio and Jacob Davis

NANOG 78

February 10, 2020, San Francisco, CA

Unless otherwise noted, images by Freepik from flaticon.com

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND No: SAND2019-14898 C

Domain Name System (DNS) Review

- DNS typically runs over UDP (original standard)
- Recursive resolver follows answers from Authoritative servers

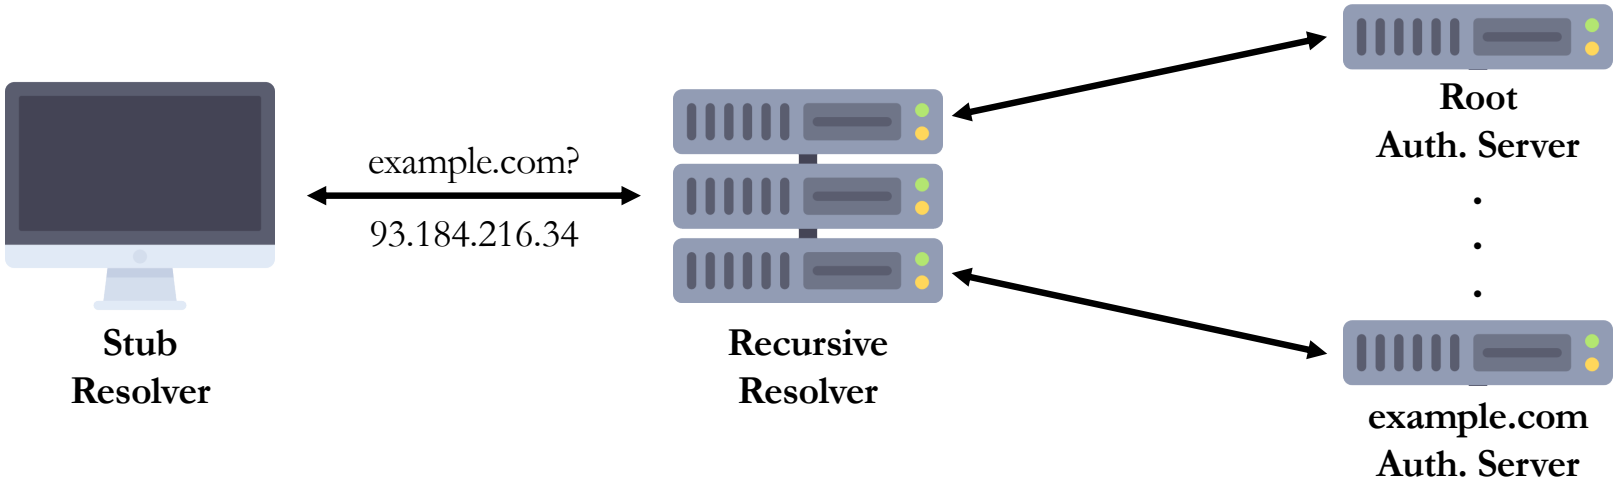

DNS Dangers

- UDP has no security measures
- Vulnerable to eavesdropping, modifications, spoofing (DDoS), etc.
- Easy to use for filtering and logging

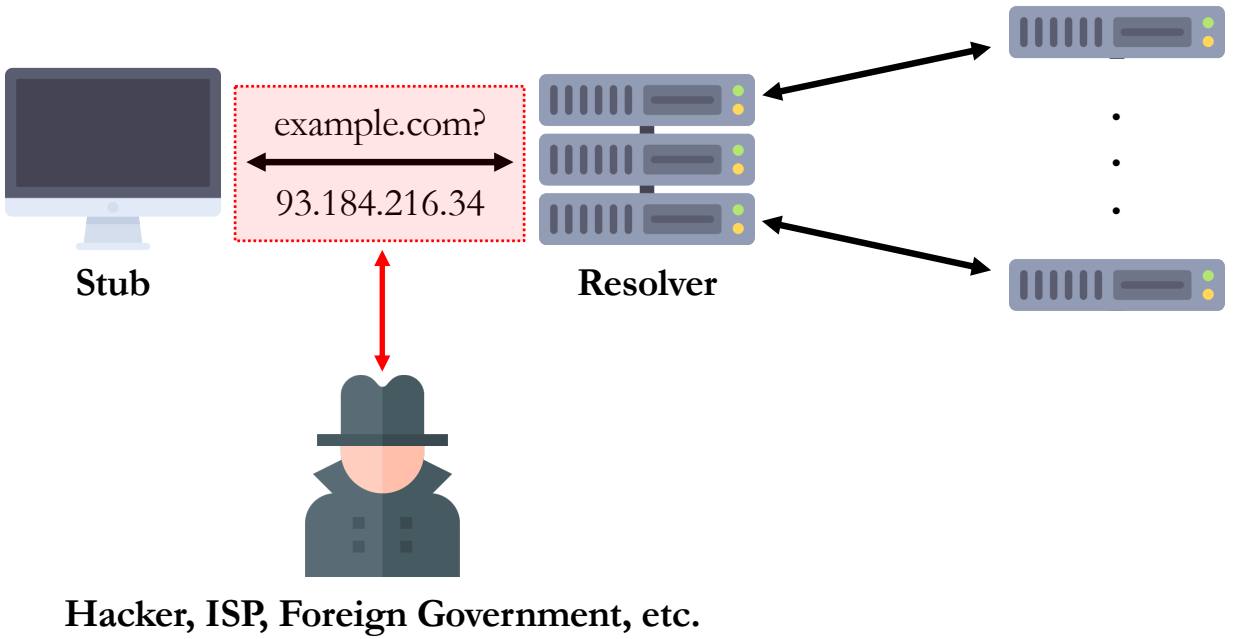

DNS Security Measures

Authenticity – Ensuring answer is correct

- DNSSEC

Confidentiality – Ensuring a connection is private

- **DNS over TLS (DoT)**
- **DNS over HTTPS (DoH)**
- DNS over DTLS
- DNS over QUIC
- DNSCrypt

DNS over TLS (2016)

- Transmit DNS queries over TLS
  - Optionally, verify server certificate is trusted
  - After handshake, everything is encrypted with shared session key
- Uses dedicated port 853
- Once handshake is complete, send queries like normal

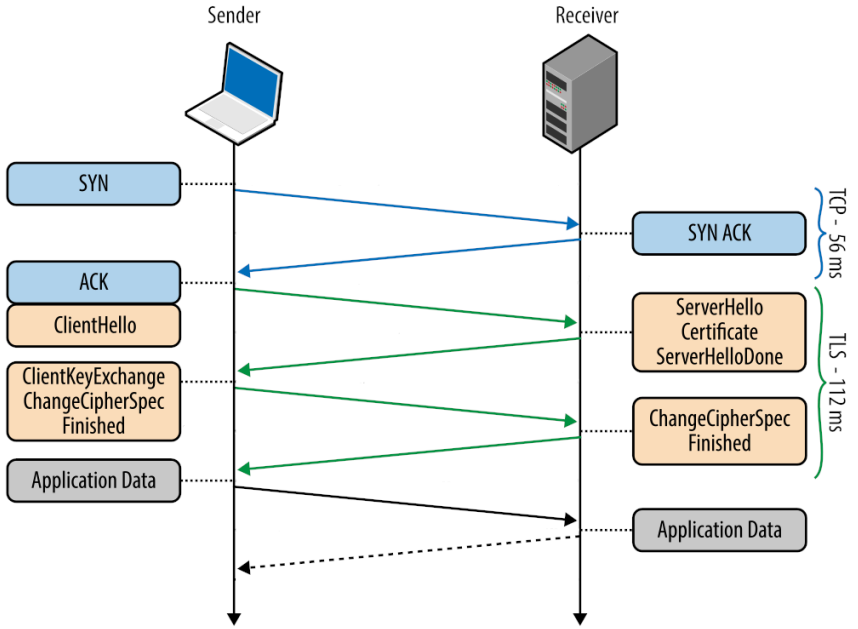

Image from hpbn.co

DNS over HTTPS (2018)

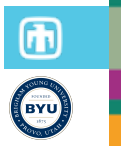

- Send queries like normal web traffic (port 443)
  - Harder to block/detect as a result
  - Easier to implement for applications
- Use either GET or POST requests
  - POST: include wire format message in body
  - GET: include wire format message encoded in Base64url as a URL parameter

```
:method = POST
   :scheme = https
   :authority = dnsserver.example.net
   :path = /dns-query
   accept = application/dns-message
   content-type = application/dns-message
   content-length = 33

   <33 bytes represented by the following hex encoding>
   00 00 01 00 00 01 00 00  00 00 00 00 03 77 77 77
   07 65 78 61 6d 70 6c 65  03 63 6f 6d 00 00 01 00
   01
```

```
:method = GET
   :scheme = https
   :authority = dnsserver.example.net
   :path = /dns-query?
     dns=AAABAAABAAAAAAAAA3d3dwdleGFtcGxlA2NvbQAAAQAB
   accept = application/dns-message
```

Examples from RFC in HTTP2 format

Comparison of DoT and DoH

DoT

- System-wide: can be used for all applications, i.e., with `getaddrinfo`
- Follows same paradigm as DNS over TCP
- Uses a dedicated port– 853

DoH

- Easy to implement at application level
  - Use standard OS HTTPS libraries
  - Websites via Javascript
- Typically uses port 443
- Harder to block because it looks like normal Web/HTTPS traffic (i.e., same port)

Measuring Support at Resolvers

- Get a list of open recursive resolvers
  - Query every IPv4 address and see if it behaves like a resolver (RA flag + RCODE of `NOERROR` or `NXDOMAIN`)
  - **1,197,794** discovered
- Test discovered resolvers with TCP (pre-req for DoT/DoH)
- Send a query using DoT and DoH to those who responded

- Method excludes IPv6 and DoT or DoH resolvers that do not response over UDP

DoT and DoH Resolver Results

- **1,197,794** open resolvers
- **1,747** (0.15%) IPs responded to DoT
  - 1,529 of those from a single entity, CleanBrowsing
  - 87 unique autonomous systems
- **9** IPs responded over DoH
  - All owned by Quad9 or Cloudflare
  - More up-to-date sources list 35 public DoH resolvers

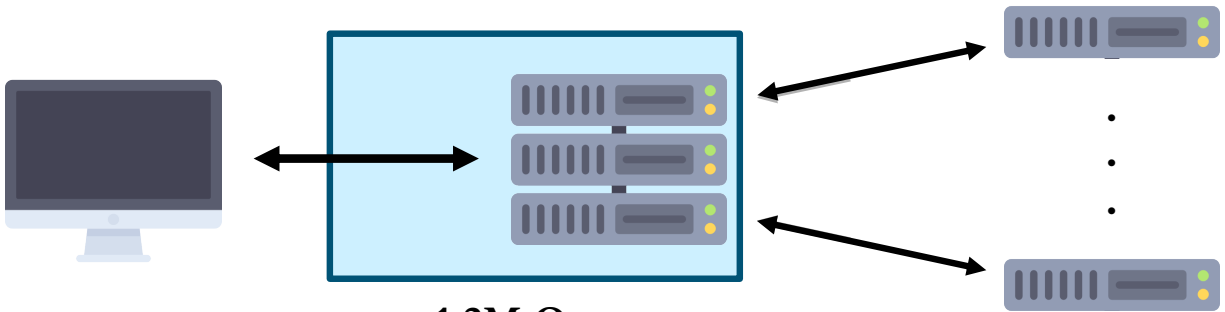

1.2M Open Resolvers

TLS Related Results

- **22** unique certificate signers were observed
  - GoDaddy and Let's Encrypt were most popular
- **11** certificates were self-signed (Issuer matched Subject)

- **79** (4.5%) IPs supported TLS 1.3
  - Important for reduced RTT (2→1) and potential for 0–RTT
- **1,701** (97%) IPs supported TLS 1.2
- **80** IPs did not support TLS 1 or TLS 1.1

Measuring Support at Auth Servers

- Limit scope to nameservers for top 5,000 Alexa sites and all TLDs
- For each Alexa site, query for nameservers, then IPv4 and IPv6 addresses of those nameservers
- Repeat steps done for resolvers

Image from linkody.com

DoT Authoritative Results

- Limited scope to nameservers for top 5K Alexa sites and all TLDS (1,530)
  - **6,817** unique IP addresses for TLDS
  - **10,214** unique IP addresses for Alexa Sites

- No TLD responded over DoT
- **12** Alexa IPs responded over DoT
  - All IPs that responded were owned by Facebook
  - Corroborates with Cloudflare blog experimenting with DoT to Facebook

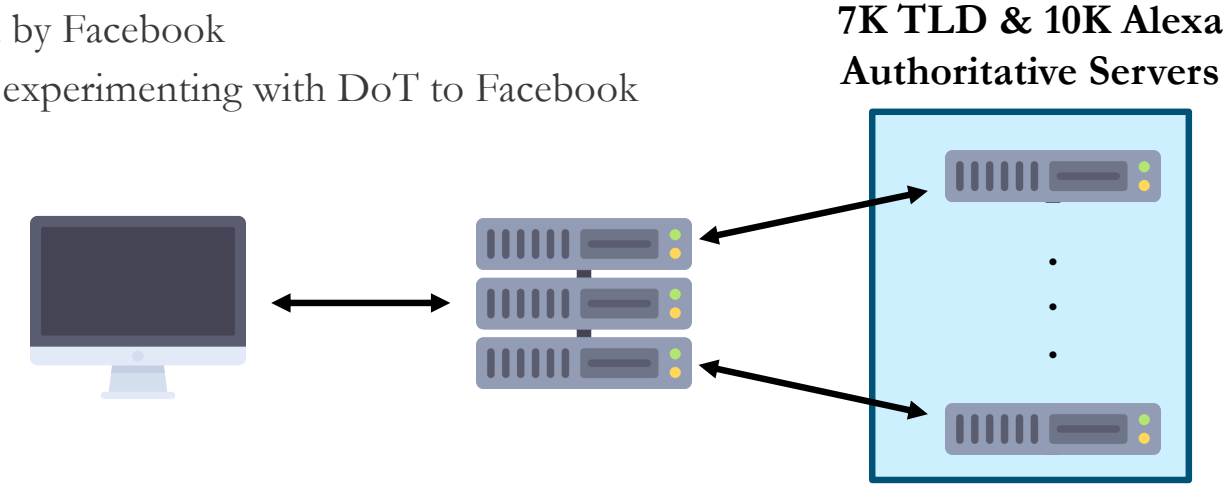

TCP Fast Open Overview (2014)

- A major drawback of security is increased delay
- TFO fixes this in subsequent connections
  - Server gives client cookie in first connection
  - Client can reconnect with cookie + data in SYN

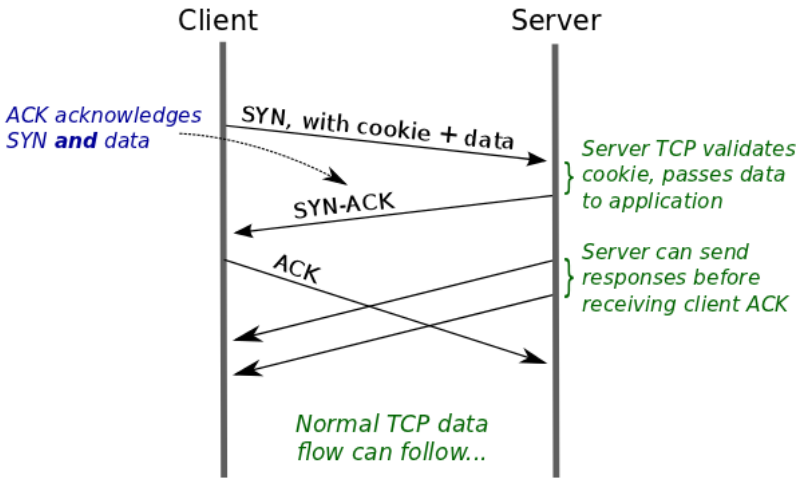

Image from lwn.net

Measuring TFO support

- Enable client TFO support
- Send TCP SYN with TFO
  - Check if the SYN-ACK has TFO option
- Send two queries with TFO back-to-back
  - Second query should have valid cookie
  - Check that data in second SYN was

TFO Results for Authoritative Servers

- Like DoT work, used nameservers for top 5K Alexa sites and all TLDS (1,530)
  - **6,743** unique IP addresses for TLDS
  - **9,558** unique IP addresses for Alexa Sites

- **11** TLD IPs included TFO option
  - 10 of these were Google's
- **5** ACKed data

- **726** (7.1%) Alexa IPs sent TFO option
- **18** (0.19%) ACKed data

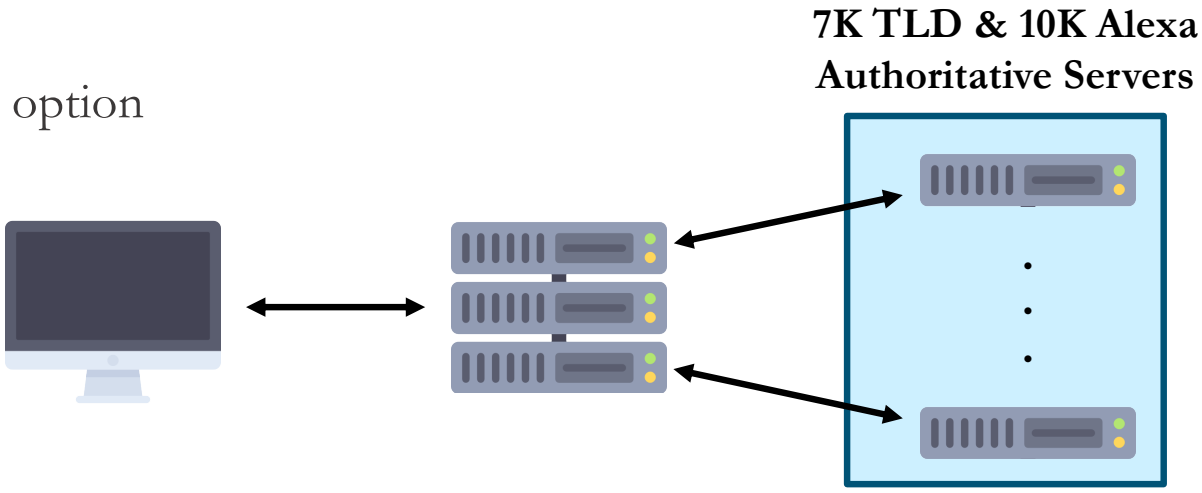

TFO Client Results at Root Servers

- Analyzed 48 hours of queries sent to root server (minus g-root)
- **3,769,471** unique IPs queried roots
- **89** IPs included TFO option
- **32** included cookie, but didn't send data in SYN

- Needs to be studied further
- Does not appear the root servers supported TFO

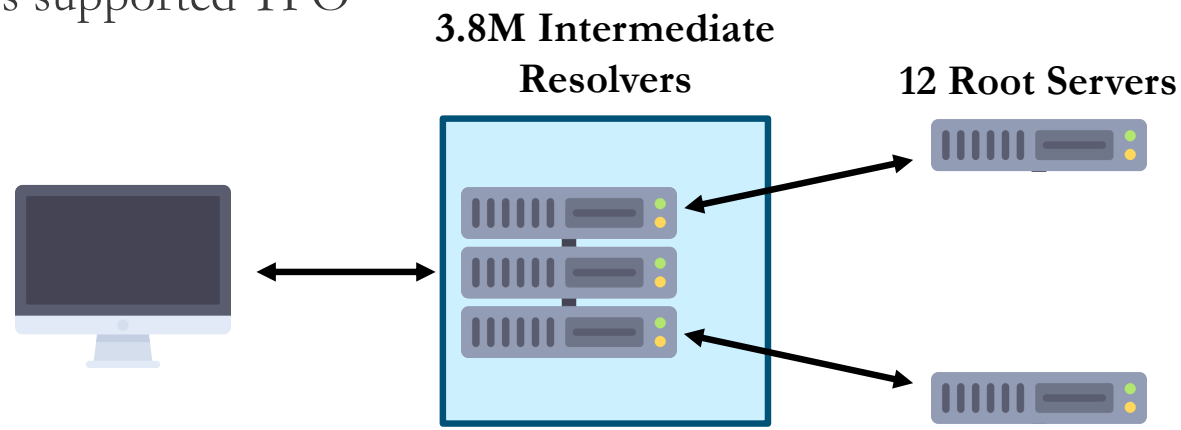

Conclusion

- Both DoT and DoH offer security to the DNS
- DoT adoption is limited, but includes most well-known resolvers
- DoH is newer, but will likely surpass DoT in adoption

- TFO can help reduce delay of DoT and DoH but support is very limited
- Many IPs are sending TFO option, but not ACKing data

Questions

Contact:
casey@byu.edu
jacdavi@sandia.gov