



Internet Systems Consortium

DNS, DoT and DoH

DoQ Not included

Alan Clegg
December 11, 2019

RFC 1034 (DNS)

RFC 7858 (DoT)

RFC 8484 (DoH)

draft-huitema-quic-dnsquic-06 (DoQ)



© 2019 - Internet Systems Consortium





Agenda

- How this presentation is different
- Differentiate the available technologies
- Observe the current environment
- Consider an actual implementation
- Thoughts on debugging
- Future conversations
- Closing comments



Rationale



What are we talking about?



DNS

- What we know as “DNS” today
- TCP/UDP port 53
- Not encrypted
 - Easily monitored
 - Easily blocked
 - Easily redirected
 - Easily modified (unless DNSSEC implemented)



DoT

- DNS over TLS
Uses same encryption as HTTPS
- TCP port 853
- Encrypted
Easily monitored (for traffic, not content)
Easily blocked
Not easily modified
More CPU intensive (TLS setup per server contacted)



DoH

- DNS over HTTPS
- TCP port 443
- Encrypted
 - Not easily monitored (mixed in HTTPS traffic)
 - Not easily blocked
 - Not easily redirected
 - Not easily modified



Enter the Matrix

Can it Be...	Do53	DoT	DoH
Read	Yes Plain Text	No Encrypted	No Encrypted
Monitored (Traffic, not Content)	Yes	Yes	No
Blocked	Yes	Yes	No
Modified	Yes Unless using DNSSEC	No	No
Redirected	Yes	No	No

What does this look like?



Home/Office
User



ISP DNS
Servers



The Watcher



Alternative
(Cloud) DNS

Setup



Home/Office
User



The Watcher



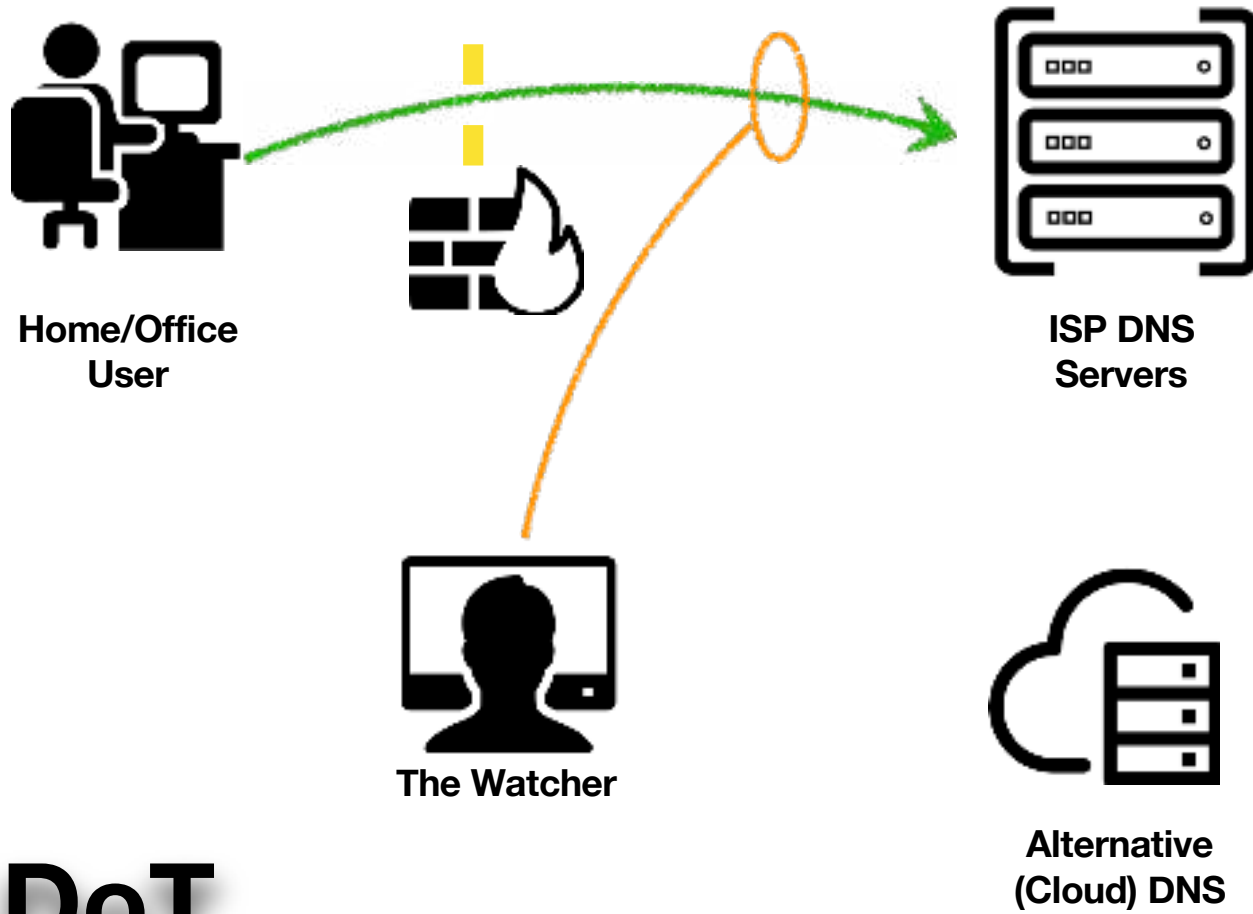
ISP DNS
Servers



Alternative
(Cloud) DNS

DNS

9



DoT



Home/Office
User



ISP DNS
Servers



The Watcher



Alternative
(Cloud) DNS

DoH



Home/Office
User



ISP DNS
Servers



The Watcher



Alternative
(Cloud) DNS

DoH Attack of
The Apps



Home/Office
User



ISP DNS
Servers




The Watcher



Alternative
(Cloud) DNS

DoH Attack of
The Malware





Home/Office
User



ISP DNS
Servers



Alternative
(Cloud) DNS

Imagine an image of a
yellow man with MG
hair/ears saying "DoH!"

Where are we now?



DoT

- There are implementations and deployments
 - Client support in Android 9+
 - Server very easy to configure as an nginx stream
 - `https://www.aaf1a1o.me/2019/03/dns-over-tls/`
 - Supported in dnsmasq



DoH

- List of known DoH servers

`https://github.com/curl/curl/wiki/DNS-over-HTTPS#publicly-available-servers`

- The problem with firewalling DoH (blocking port 443) is that if the remote server also serves web content, it is impossible to block without losing access to the content.
- What if `www.google.com` also responded to DoH queries?



DoH

- **Google**
- **Cloudflare**
- **Quad9**
- CleanBrowsing
- @chantra
- @jedisct1
- **PowerDNS**
- blahdns.com
- NekomimiRouter.com
- SecureDNS.eu
- Rubyfish.cn
- Commons Host
- dnswarden.com
- aaflalo.me
- Server EU
- Foundation for Applied Privacy



DoH

- **Google**
- **Cloudflare**
- **Quad9**
- CleanBrowsing
- @chantra
- @jedisc1
- **PowerDNS**
- blahdns.com
- NekomimiRouter.com
- SecureDNS.eu
- Rubyfish.cn
- Commons Host
- dnswarden.com
- aaflalo.me
- Server EU
- Foundation for Applied Privacy

AND ALL
THE
OTHERS



DoH

Supported in browsers and clients

Name	Version	Comments
Firefox	62	temporary docs
Bromite	67.0.3396.88	How to enable DoH
curl	7.62.0	See DOH-implementation
OkHttp	3.11	See Providers
curl-doh	n/a	basic stand-alone DoH client that uses curl
Chrome	66	https://bugs.chromium.org/p/chromium/issues/detail?id=799753

<https://github.com/curl/curl/wiki/DNS-over-HTTPS#supported-in-browsers-and-clients>



And here's news...

- Microsoft has announced the addition of DoH to Windows

Windows will improve user privacy with DNS over HTTPS

11-17-2019 09:00 PM

Brought to you by Tammy Jensen, Ivan Pashov, and Gabriel Montenegro

Here in Windows Core Networking, we're interested in keeping your traffic as private as possible, as well as fast and reliable. While there are many ways we can and do approach user privacy on the wire, today we'd like to talk about encrypted DNS. Why? Basically, because supporting encrypted DNS queries in Windows will close one of the last remaining plain-text domain name transmissions in common web traffic.

Providing encrypted DNS support without breaking existing Windows device admin configuration won't be easy. However, [at Microsoft we believe that](#) "we have to treat privacy as a human right. We have to have end-to-end cybersecurity built into technology."

<https://techcommunity.microsoft.com/t5/Networking-Blog/Windows-will-improve-user-privacy-with-DNS-over-HTTPS/ba-p/1014229>

Real Life Example



Implementation

- Branch Office:
 - Star Brilliant's "High performance DNS over HTTPS client & server" as client
 - Listens on ethernet interface, port 53
 - Internal, no protective ACLs
 - Sends "internal" qnames to `localhost:53`
 - Others: `https://doh.clegg.com/dns-query`



Implementation



- Branch Office:
 - **BIND 9.14.7**
 - Listens on loopback interface, port 53
 - Private TLD + `in-addr.arpa`



Implementation



- Cloud 1:
 - Linode - Debian 9
 - Star Brilliant's "High performance DNS over HTTPS client & server" - as Server
 - BIND 9.15.6 (development branch)
 - Nginx 1.16.1



Implementation



- Cloud 1 Configuration:
 - BIND listens on all interfaces (IPv4 & IPv6)
 - Port 53 / Recursive
 - ACLs allow queries from `localhost` and Cloud 2 only
 - Blacklisting / Adblocking via custom scripts



Implementation



- Cloud 1 Configuration:
 - Nginx
 - Existing install
 - Stream accepting **DoT** connections on **853**
 - Reverse proxy for **https://<name>/dns-query/**
 - Feeds query to **http://localhost:8053**



Implementation



- Cloud 1 Configuration:
 - **doh-server** (Star Brilliant)
 - Listening on **localhost:8053**
 - Accepts DoH queries
 - Converts raw query to DNS query
 - Passes them to **localhost:53**



Implementation



Alan's Server
In the Cloud

- Cloud 2:
 - Linode - Debian 10
 - **dnsmdist (from git)**
 - **BIND 9.14.8**
 - **Nginx**



Implementation



- Cloud 2 Configuration:
 - BIND
 - Listening on `localhost` (v4 & v6) port `5353`
 - Recursive



Implementation



- Cloud 2 Configuration:
 - dnsmist
 - Listening externally on **853 (DoT) & 443 (DoH)**
 - Load balances queries to:
 - **localhost : 5353**
 - **Cloud 1 : 53**

Debugging



Debugging this mess

- Fantastic external resources:
 - <https://getdnsapi.net/query/>
 - Allows testing of DNS over varying combinations of transport (UDP, TCP, TLS)
 - <https://github.com/dcid/>
 - DoT and DoH command line (PHP) clients



Debugging this mess

```
Dec 10 15:45:01 stargate doh-client[27418]: 2019/12/10 15:45:01 Request "ptz-cam.boat. IN A" is passed through
127.0.0.2:53.
Dec 10 15:45:01 stargate doh-client[27418]: 192.168.77.136:34059 - - [10/Dec/2019:15:45:01 +0000]
"b.info2intel.com. IN A"
Dec 10 15:45:01 stargate doh-client[27418]: 2019/12/10 15:45:01 choose upstream: upstream type: IETF, upstream
url: https://doh.clegg.com/dns-query
Dec 10 15:45:01 stargate doh-client[27418]: 192.168.77.136:34059 - - [10/Dec/2019:15:45:01 +0000]
"b.info2intel.com. IN AAAA"
Dec 10 15:45:01 stargate doh-client[27418]: 2019/12/10 15:45:01 choose upstream: upstream type: IETF, upstream
url: https://doh.clegg.com/dns-query
Dec 10 15:45:10 stargate doh-client[27418]: 192.168.77.146:53840 - - [10/Dec/2019:15:45:10 +0000]
"zimbra.isc.org. IN AAAA"
Dec 10 15:45:10 stargate doh-client[27418]: 2019/12/10 15:45:10 choose upstream: upstream type: IETF, upstream
url: https://doh.clegg.com/dns-query
Dec 10 15:45:15 stargate doh-client[27418]: 2019/12/10 15:45:15 upstream type: IETF, upstream url:
https://doh.clegg.com/dns-query, effect weight: 74
```

- Turn up logging
 - You can't see much in packet dumps
 - Oh ... privacy? Yeah, about that...
- Software is young
 - Log messages from same daemon with same data in different columns
 - Able to change without breaking everyone!



Debugging this mess

```
aclegg@stargate:~/local/bin $ doh-client \  
> --domain dns.dnsoverhttps.net \  
> --qname sigfail.verteiltssysteme.net \  
> --dnssec  
flag provided but not defined: -domain  
Usage of doh-client:  
  -conf string  
        Configuration file (default "doh-client.conf")  
  -pid-file string  
        PID file for legacy supervision systems lacking support for reliable cgroup-based process tracking  
  -verbose  
        Enable logging  
  -version  
        Show software version and exit  
aclegg@stargate:~/local/bin $ ./doh-client --domain dns.dnsoverhttps.net \  
> --qname sigfail.verteiltssysteme.net --dnssec  
Traceback (most recent call last):  
  File "./doh-client", line 6, in <module>  
    from dohproxy.client import main  
  File "/home/aclegg/.local/lib/python3.7/site-packages/aiohttp2/__init__.py", line 2, in <module>  
    from .helper import *  
  File "/home/aclegg/.local/lib/python3.7/site-packages/aiohttp2/helper.py", line 89  
    async_task = asyncio.async  
                    ^  
SyntaxError: invalid syntax
```

- Everything is named the same
 - And it explodes...



Debugging this mess

```
aclegg@stargate:~/bin $ cat getchain
echo | openssl s_client -connect $1:853 |grep -B 2 -A 5 "Certificate chain"

aclegg@stargate:~/bin $ getchain dns.google
depth=2 OU = GlobalSign Root CA - R2, O = GlobalSign, CN = GlobalSign
verify return:1
depth=1 C = US, O = Google Trust Services, CN = GTS CA 101
verify return:1
depth=0 C = US, ST = California, L = Mountain View, O = Google LLC, CN = dns.google
verify return:1
DONE
CONNECTED(00000003)
---
Certificate chain
 0 s:C = US, ST = California, L = Mountain View, O = Google LLC, CN = dns.google
   i:C = US, O = Google Trust Services, CN = GTS CA 101
 1 s:C = US, O = Google Trust Services, CN = GTS CA 101
   i:OU = GlobalSign Root CA - R2, O = GlobalSign, CN = GlobalSign
---
```

- Learn about OpenSSL!
 - Paid certificates from CA
 - Free certificates from Let's Encrypt
 - No more self-signed certificates if possible



Debugging this mess

```
aclegg@stargate:~ $ dnstls alan.clegg.com doh.clegg.com
alan.clegg.com has address 45.33.100.174
aclegg@stargate:~ $ dnstls some-bad-site.com doh.clegg.com
pornhub.com has address 66.254.114.41
aclegg@stargate:~ $ dnstls some-bad-site.com cleanbrowsing
Host pornhub.com not found: 3(NXDOMAIN)
aclegg@stargate:~ $ dnstls 00author.com doh.clegg.com
00author.com has address 0.0.0.0
aclegg@stargate:~ $ dnstls 00author.com cleanbrowsing
00author.com has address 64.136.20.41
```

- Everything that looks simple ... Isn't.
 - Wrapping DNS queries in TLS
 - Generating a DNS query over HTTPS
- Which client is talking to which server?
 - Things blocked (firewalled) on one server are not blocked on another
 - Blocking methods differ so results will differ
 - Applications on the same client may be talking to different servers with different policies
- Caches are now all over the place
 - In the DoH/DoT code, in the recursive server, in the client...





Debugging this mess

- Learn your data paths
 - No longer directly from client to resolver to auth
 - Where did that query go / vanish?
- Keep software up-to-date
 - Rapidly changing
 - Not yet packaged - dependency hell



DNS/53
"local"



DoH/443



Corporate Server
In the Cloud

- Learn your data paths
- No-longer directly from client to resolver to auth
- Where did that query go / vanish?

DNS/53
Recursion



Auth Servers


Conversations We Are Going To Have



Recent Conversation Starters

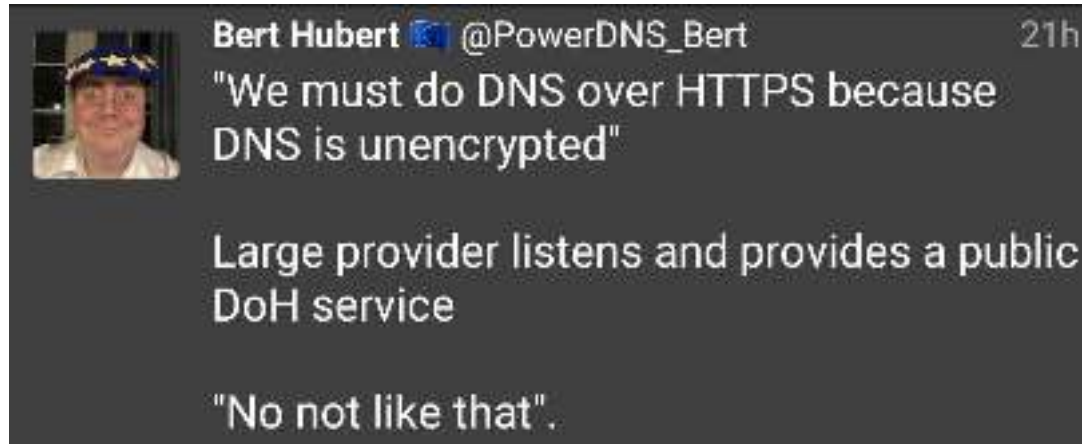
 **David Ulevitch** @davidu 6 Dec
The right answer is that everyone should be running a feature-complete caching + forwarding resolver on localhost. All the rest of these discussions are noise from companies that want eyeballs. [twitter.com/CarolineGreer/...](https://twitter.com/CarolineGreer/)
via Twitter Web App




 **Bert Hubert** @PowerDNS_Bert 7 Dec
@davidu @jpmens I struggle with this - what should they forward to? A single-user resolver has terrible performance, but sending your traffic some a random cache also has downsides. I'd love to somehow square the privacy, performance & reliability circle.
via Twitter Web App in reply to @davidu



Recent Conversation Starters

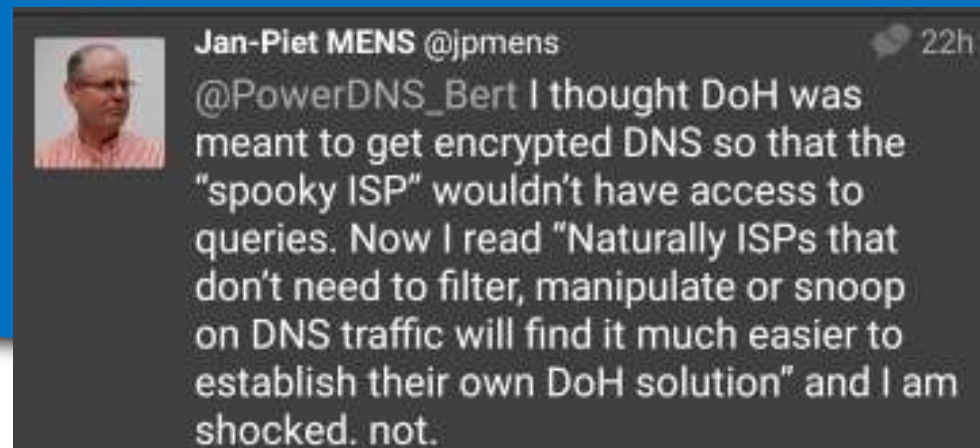


Bert Hubert  @PowerDNS_Bert 21h

"We must do DNS over HTTPS because DNS is unencrypted"

Large provider listens and provides a public DoH service

"No not like that".





Jan-Piet MENS @jpmens 22h

@PowerDNS_Bert I thought DoH was meant to get encrypted DNS so that the "spooky ISP" wouldn't have access to queries. Now I read "Naturally ISPs that don't need to filter, manipulate or snoop on DNS traffic will find it much easier to establish their own DoH solution" and I am shocked. not.




Recent Conversation Starters



 **Paul_IPv6** @Paul_IPv6 16h
@kolkman i am encouraged that several very large consumer ISPs are already doing trials of DoT and DoH, so we should have a better idea of at least what more TCP might cost in terms of hardware/network support. and i would still love to see stub DNSSEC validating in end devices as a start

by Olaf Kolkman via Twitter for iPhone

 **Paul_IPv6** @Paul_IPv6 16h
@kolkman another piece we need to consider/test is that current DoT/DoH is designed for last mile, not to auth. ADoT is going to need different tuning and experience. when do persistant connections, multiple response per connection, etc make sense operationally? lots to research yet.

by Olaf Kolkman via Twitter for iPhone



Moving forward..

- We have the technology available to deploy
 - Will it scale?
 - Is it supportable?
 - My stuff broke, who's going to fix it?
- Timeframe for support in BIND



Software Links

- BIND installed from source code (`ftp.isc.org`)
- `https://github.com/m13253/dns-over-https`
- `https://github.com/PowerDNS/pdns`



Contact Information

- Alan Clegg
 - E-Mail: aclegg@isc.org
 - Twitter: [@AlanAtISC](https://twitter.com/AlanAtISC)

Questions?

Comments?



<https://www.isc.org>

