

DoH!

Geoff Huston
APNIC

Words of Caution!

I'm speaking after Paul

And I'm speaking on the DNS

This may not end well!

But here goes...

What defines "The Internet"?

- The 1990's answer was all about '**reachability**'
 - The Internet was this connected domain where every connected device could send IP packets to any other connected device
 - And the packet I sent to you is the packet you got
 - Modulo TTL and fragmentation and reassembly
- But then we invented enterprise networks, firewalls and security realms, NATs and all kinds of "value added" network services
- And the entire architecture of the Internet shifted to a client/server architecture
 - Clients could interact with servers, but not with each other

What defines "The Internet"?

- Then we turned to the server model and started playing with anycast to improve server performance by service replication
- Today:
 - Clients don't have a public IP address (NATs)
 - Servers don't have a unique public IP address (Anycast)
 - So what is the address architecture of the Internet?

What defines "The Internet"?

- If the Internet used be defined by a communications domain that shared a common name and address infrastructure then we've broken the address part and its never going to come back!
- The Internet is now defined only by a common name space

The DNS as the Internet's Glue

RFC 2826:

Effective communications between two parties requires two essential preconditions:

- The existence of a common symbol set, and
- The existence of a common semantic interpretation of these symbols.

Failure to meet the first condition implies a failure to communicate at all, while failure to meet the second implies that the meaning of the communication is lost.

The Internet's Domain Name System

What is it?

- A common set of syntax rules that defines 'valid' DNS names
- A hierarchically structured distributed database
- A common name resolution protocol that can consult this database and map a name to a value
- A collection of engines (resolvers and servers) that run a common query/response protocol that performs name resolution

The DNS as Internet Infrastructure

The Theory:

- Names are visible to all
- Names resolve consistently to the same values all the time

The Practice:

- How much of these two principles can we break and still get away with it?

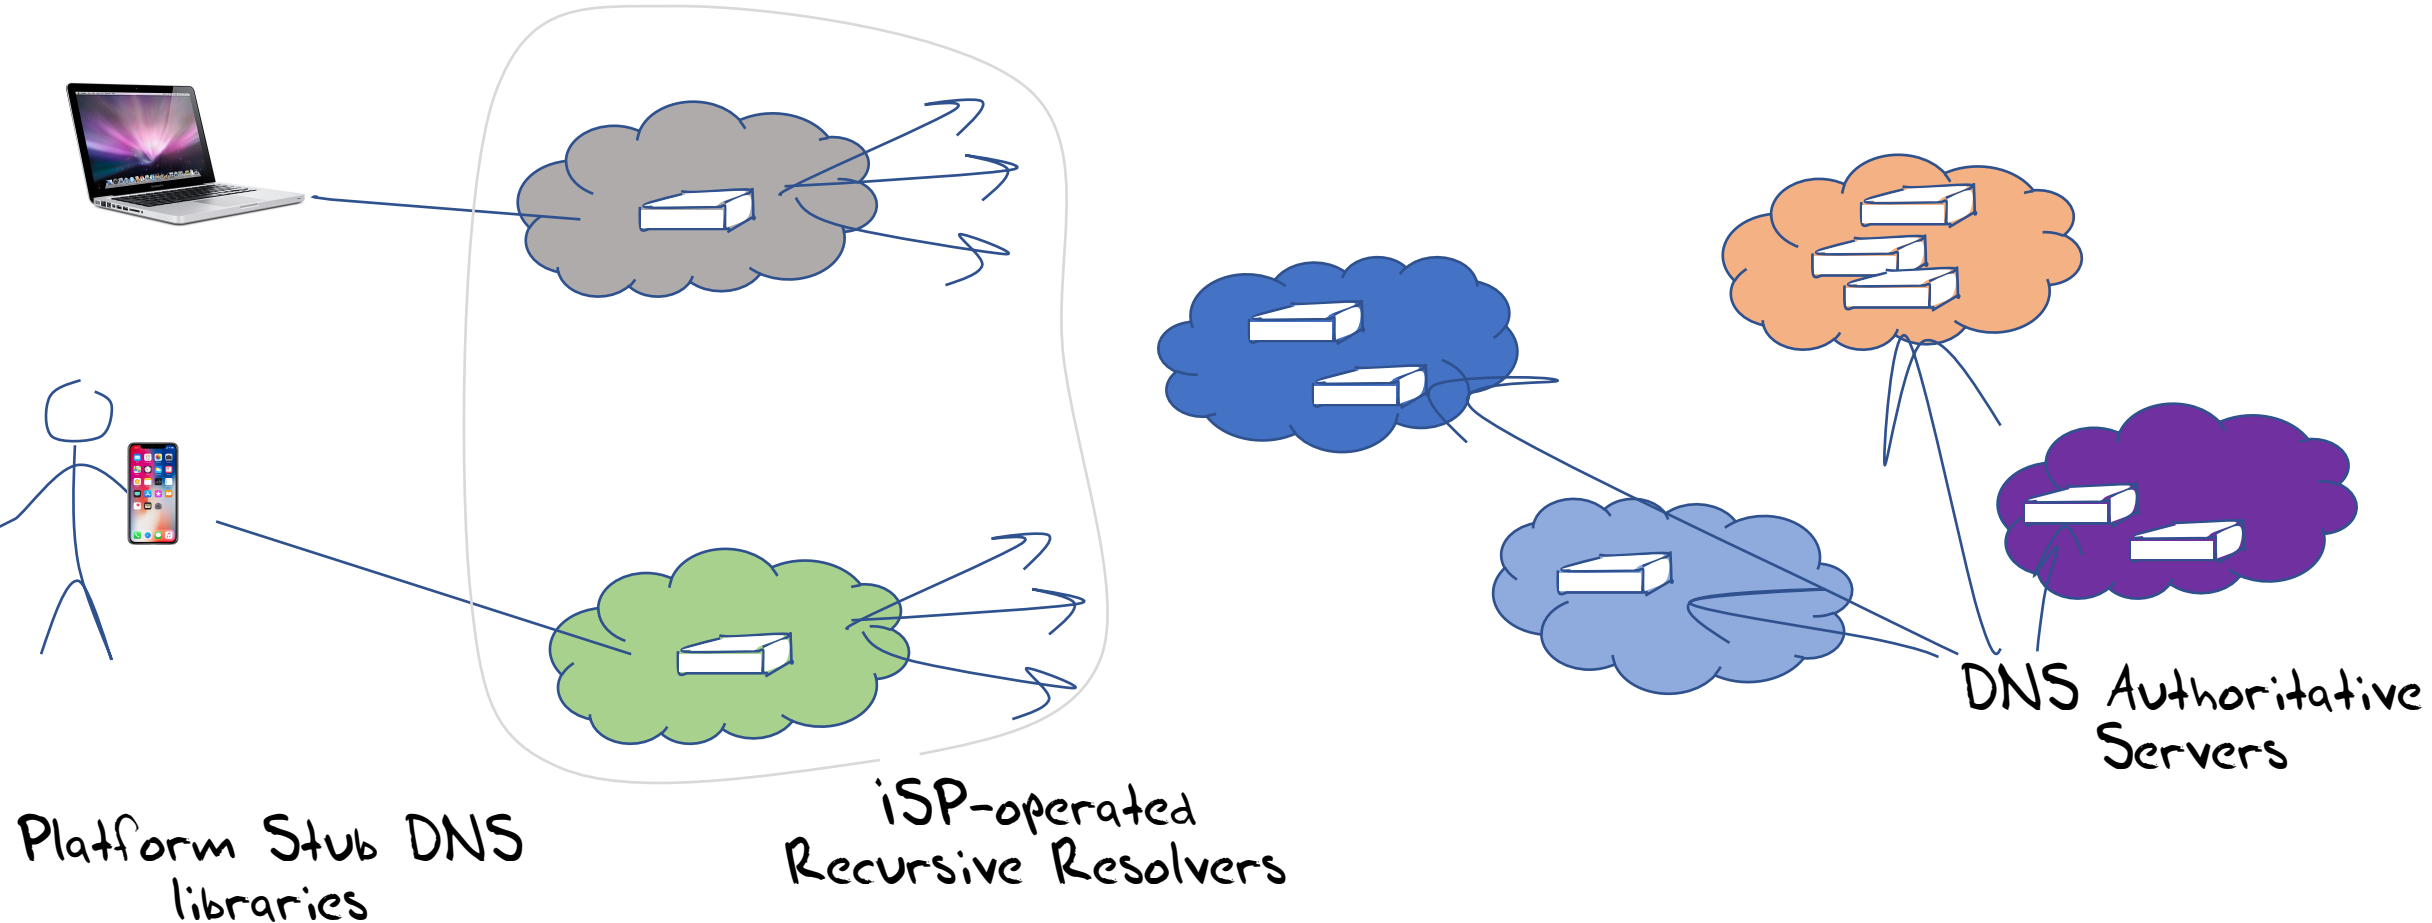
Old School DNS

- The DNS is operated as a common infrastructure (not application specific)
- The common infrastructure assumes a common and consistent name set in the DNS that is assessible to all
- If a name is defined than its definition is the same for all queriers
- If a name doesn't exist it doesn't exist for every querier

Old School DNS

- DNS resolvers configured with IP addresses as part of the connection context (DHCP)
- DNS recursive resolvers operated by the ISP as part of the ISP's service to their users
- DNS authoritative services provided in various ways (often as part of web hosting environments)
- Applications used `gethostbyname()` and tapped into the DNS common infrastructure

Old School DNS




DNS (Ab)Use

- Split Horizon DNS
- NXDOMAIN substitution
- TTL munging and Cache manipulation
- Fake Roots
- White Lies: DNS64
- DNS Geolocation
- EDNS Client Subnet

The Path to DoH

The DNS leaks information like a sieve



Why pick on the DNS?

- The DNS is very **easy to tap**
 - Its open and unencrypted
- DNS traffic is **easy to tamper with**
 - Its payload is not secured and tampering cannot be detected
 - Its predictable and false answers can be readily inserted
- The DNS is **hard for users to trace**
 - Noone knows exactly where their queries go
 - Noone can know precisely where their answers come from
- The DNS is **used by everyone**

Second-hand DNS queries are a business opportunity these days

The image shows a screenshot of the Farsight Security website. The header includes the Farsight Security logo and navigation links for Solutions, Resources, Blog, Partners, Community, and Company. The main content area features an IDC report titled "Farsight Security - Providing Real-Time DNS Data to Threat Intelligence". A central graphic reads "EVERYTHING STARTS WITH DNS". The footer contains a "LATEST NEWS" section with several article teasers.

FARSIGHT SECURITY

Solutions ▾ Resources ▾ Blog Partners Community Company ▾

IDC ANALYTICS THE FUTURE

IDC Report:
Farsight Security - Providing Real-Time DNS Data to Threat Intelligence

Farsight Security:
Providing Real-Time DNS Data to Threat Intelligence

EVERYTHING STARTS WITH
DNS

LATEST NEWS

How ThreatConnect® Leverages DNSDB to Track...
FARSIGHT
New Research on Domain Lifetimes by Dr. Vixie at Virus...
IPs, Address Ranges, a CIDR Block Queries in...

How can we improve DNS Privacy?

And not alter the DNS architecture in fundamental ways

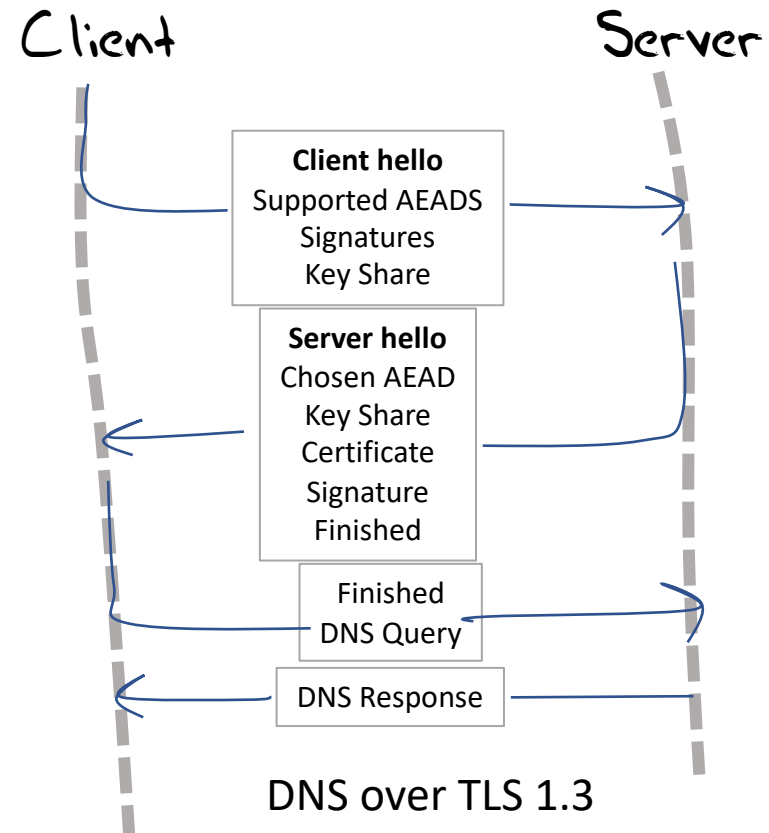
Move away from clear queries and responses and use session encryption

Encrypting the session

- Today the standard tool is TLS, which uses dynamically generated session keys to encrypt all traffic between two parties
- We could use TLS between the end client and the client's recursive resolver
 - We could probably do the same between recursive resolvers and authoritative servers, but the IETF is doing this one step at a time

DoT - DNS over TLS

- TLS is a TCP 'overlay' that adds server authentication and session encryption to TCP
- TLS uses an initial handshake to allow a client to:
 - Validate the identity of the server
 - Negotiate a session key to be used in all subsequent packets in the TCP session
- RFC 7858, RFC 8310, RFC8446



DoT - DNS over TLS

- Similar to DNS over TCP:
 - Open a **TLS** session with a recursive resolver
 - Pass the DNS query using DNS wireline format
 - Wait for the response
- Can use held DNS sessions to allow the TLS session to be used for multiple DNS queries
- The queries and the responses are hidden from intermediaries
- The client may validate the recursive resolver's identity ...

Who is at the other end of the TLS session?

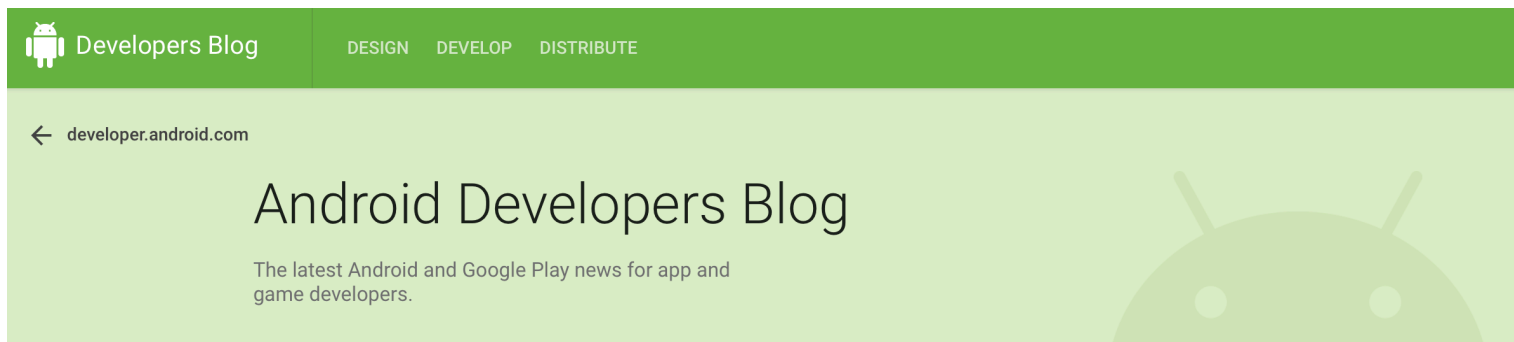
Strict Mode:

- Connect by name, and perform a TLS handshake based on authentication of the offered name certificate
(which sounds a whole lot better than it really is due to the WEB PKI CA mess!)

Opportunistic Mode:

- Use an unauthenticated encrypted session
(the client has no idea who it is talking to, but whatever is said cannot be eavesdropped in any case!)

DNS over TLS and Android

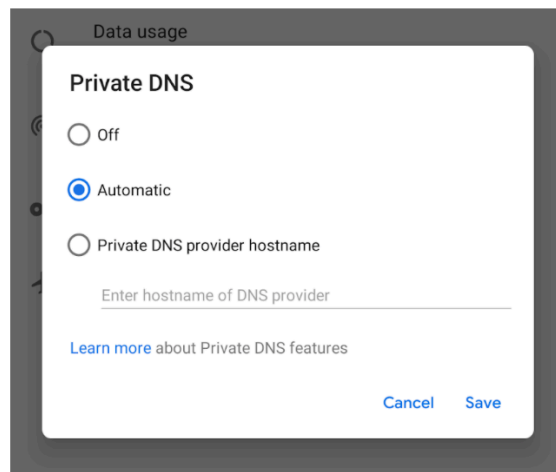


DNS over TLS in P

The Android P Developer Preview includes built-in support for DNS over TLS. We added a **Private DNS** mode to the Network & internet settings.

By default, devices automatically upgrade to DNS over TLS if a network's DNS server supports it. But users who don't want to use DNS over TLS can turn it off.

Users can enter a hostname if they want to use a private DNS provider. Android then sends all DNS queries over a secure channel to this server or marks the network as "No internet access" if it can't reach the server. (For testing purposes, see this [community-maintained list](#) of compatible servers.)



<https://android-developers.googleblog.com/2018/04/dns-over-tls-support-in-android-p.html>


DoT - DNS over TLS

- Its TCP not UDP -- May generate a higher recursive resolver memory load as each client may have a held state with one or more recursive resolvers
- The TCP session state is on port 853
 - DNS over TLS can be readily blocked by middleware
- The privacy is relative, as the recursive resolver still knows all your DNS queries
- Supported by Bind (stunnel), Unbound, KNOT, DNSDist
- Open DoT Resolvers from Google, Cloudflare (and maybe others)

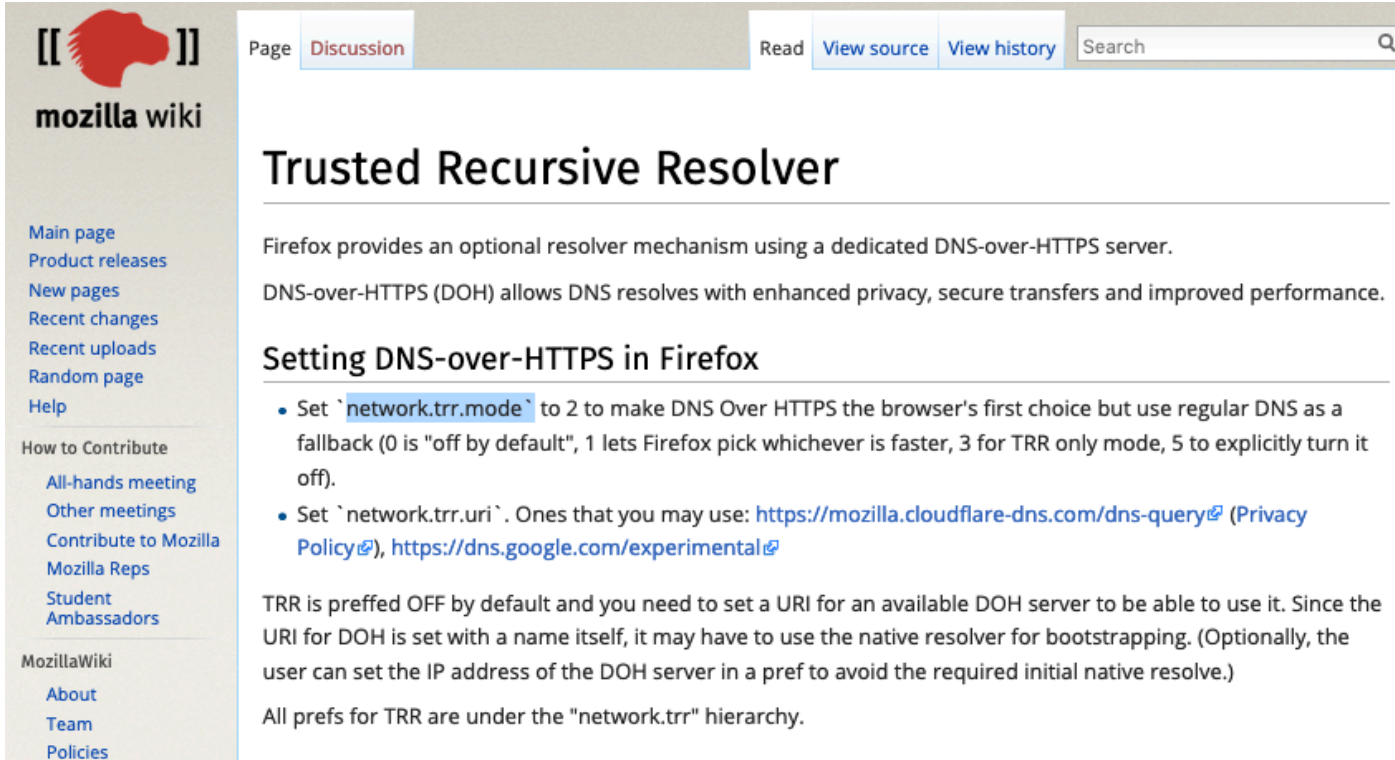
But once you are using TLS
it's a short step to...

DoH - DNS over HTTPS

- DNS over HTTPS
- Uses an HTTPS session with a resolver
- Similar to DNS over TLS, but with HTTP object semantics
- Uses TCP port 443, so can be masked within other HTTPS traffic
- Uses DNS wire format



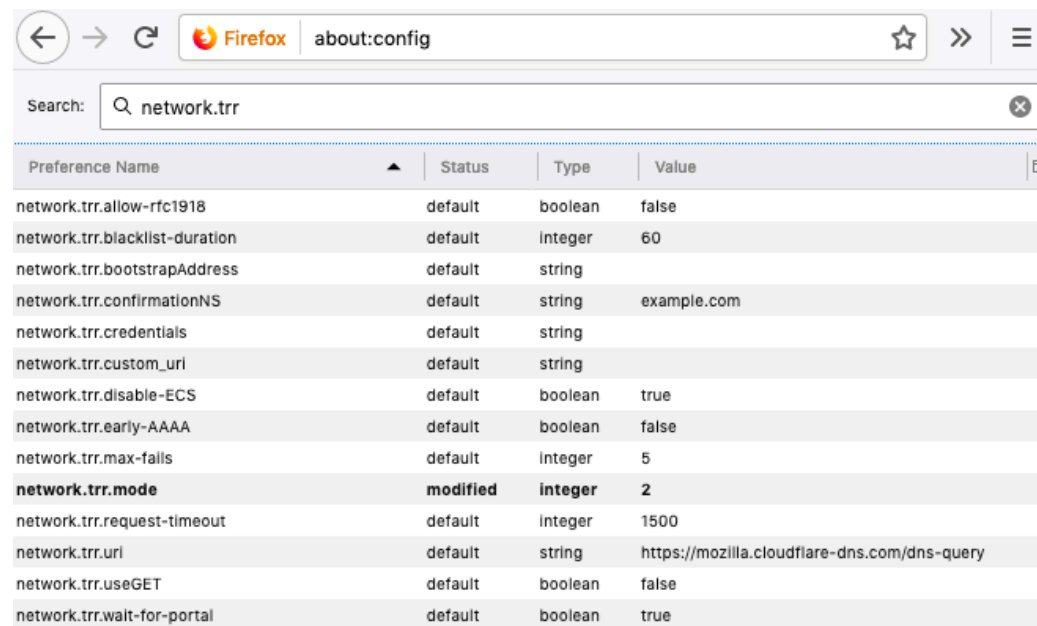
DoH - DNS within the Browser



The screenshot shows the Mozilla Wiki page for "Trusted Recursive Resolver". The page title is "Trusted Recursive Resolver". The content includes a paragraph stating that Firefox provides an optional resolver mechanism using a dedicated DNS-over-HTTPS server. It explains that DNS-over-HTTPS (DOH) allows DNS resolves with enhanced privacy, secure transfers and improved performance. Below this, there is a section titled "Setting DNS-over-HTTPS in Firefox" which contains two bullet points: "Set `network.trr.mode` to 2 to make DNS Over HTTPS the browser's first choice but use regular DNS as a fallback (0 is "off by default", 1 lets Firefox pick whichever is faster, 3 for TRR only mode, 5 to explicitly turn it off)." and "Set `network.trr.uri`. Ones that you may use: <https://mozilla.cloudflare-dns.com/dns-query> (Privacy Policy), <https://dns.google.com/experimental>".

TRR is preffed OFF by default and you need to set a URI for an available DOH server to be able to use it. Since the URI for DOH is set with a name itself, it may have to use the native resolver for bootstrapping. (Optionally, the user can set the IP address of the DOH server in a pref to avoid the required initial native resolve.)

All prefs for TRR are under the "network.trr" hierarchy.



The screenshot shows the Firefox about:config page. The search bar contains "network.trr". The table below lists various preferences, with "network.trr.mode" highlighted in bold. The value for "network.trr.mode" is 2, and its status is "modified".

Preference Name	Status	Type	Value
network.trr.allow-rtc1918	default	boolean	false
network.trr.blacklist-duration	default	integer	60
network.trr.bootstrapAddress	default	string	
network.trr.confirmationNS	default	string	example.com
network.trr.credentials	default	string	
network.trr.custom_uri	default	string	
network.trr.disable-ECS	default	boolean	true
network.trr.early-AAAA	default	boolean	false
network.trr.max-falls	default	integer	5
network.trr.mode	modified	integer	2
network.trr.request-timeout	default	integer	1500
network.trr.uri	default	string	https://mozilla.cloudflare-dns.com/dns-query
network.trr.useGET	default	boolean	false
network.trr.wait-for-portal	default	boolean	true

DoH - DNS within the Browser

- Firefox's "Trusted Recursive Resolver"
- Avoids using the local DNS resolver library and local DNS infrastructure
- Has the browser sending its DNS queries directly to a trusted resolver over HTTPS
- Servers available from Cloudflare, Google, CleanBrowsing

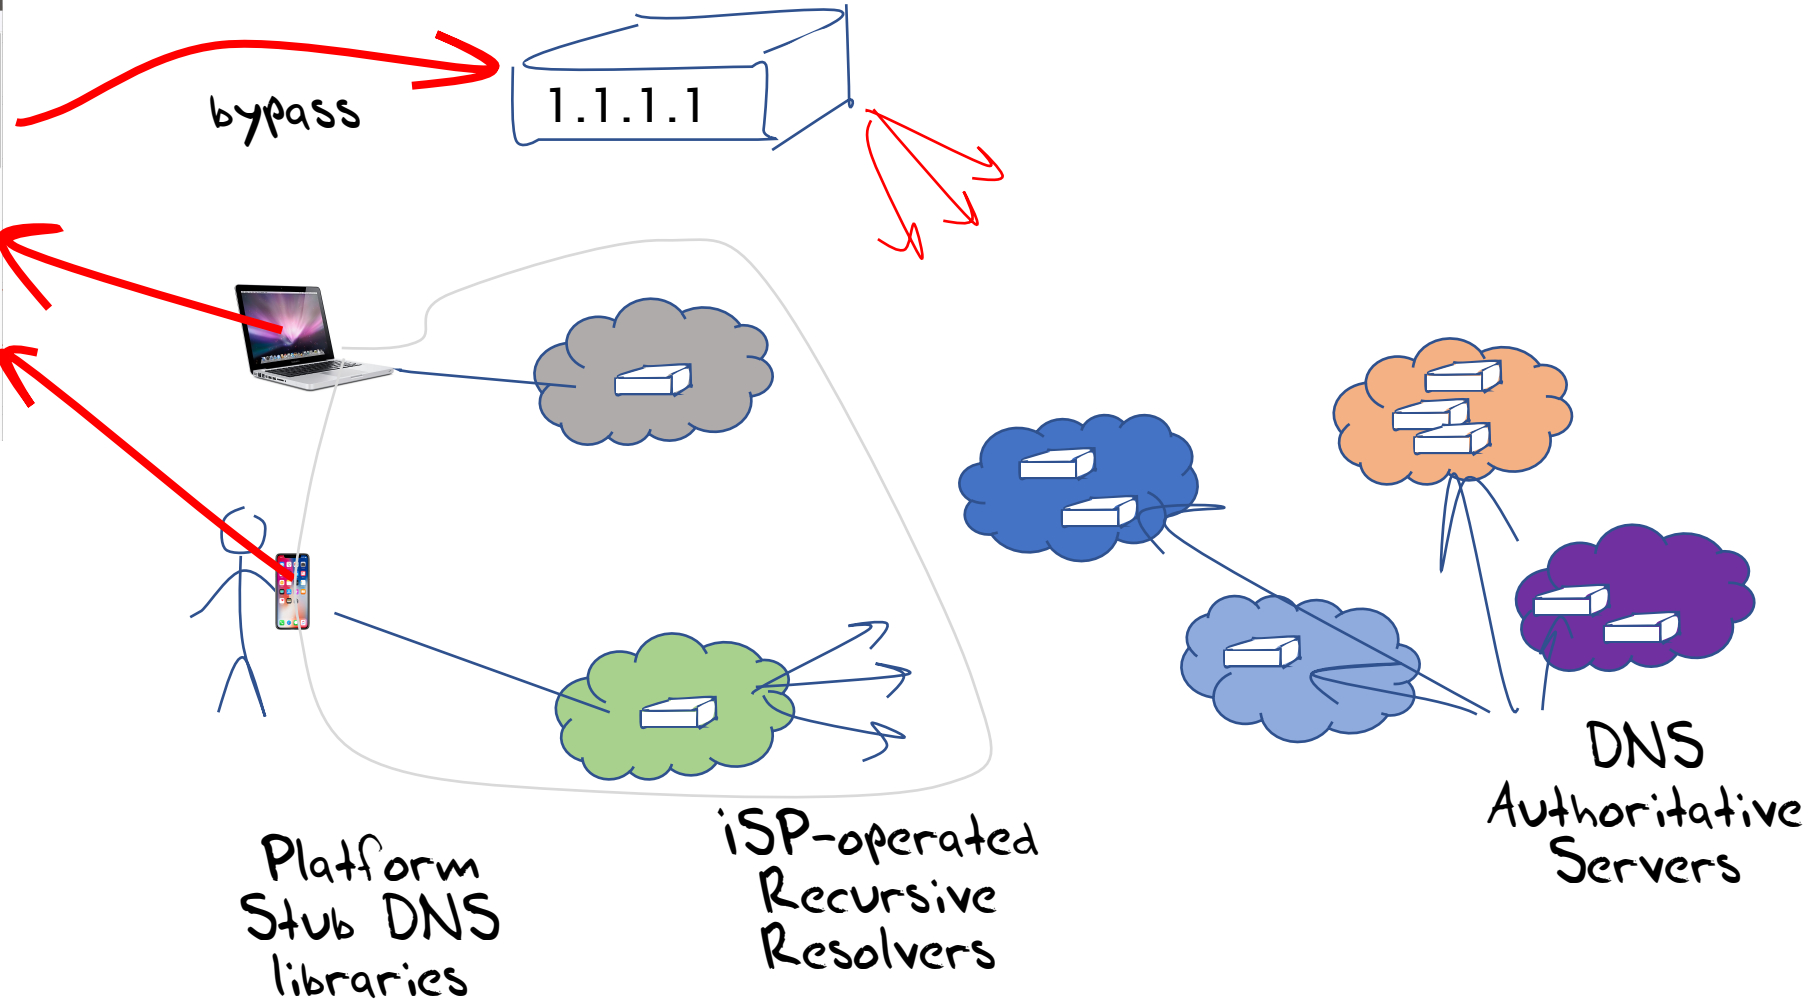
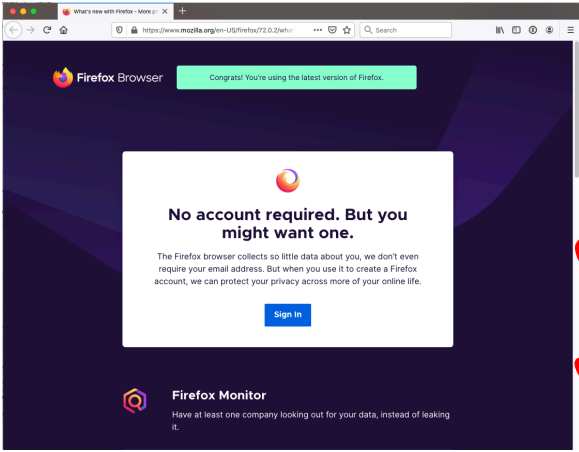
Why DoH?

- Lives on TCP port 443
- DNS content denoted by “application/dns- message”, allowing a server to distinguish DNS queries within the HTML stream (which is encrypted in TLS with HTTPS)
 - i.e. DNS queries and responses can be readily intertwined in other HTTPS traffic

Why DoH?

- Applications can effectively hide DNS transactions **from the network**
 - TLS 1.3 and ESNI can remove all visible indication of the DoH server name from the network
 - DoH queries and responses can use both DNS and HTML padding to disguise the payload size
- Applications can effectively hide DNS transactions **from the platform**
 - No DNS query logs on the platform
 - No cross-application spyware on the platform

DoH Bypass



DoH Futures?

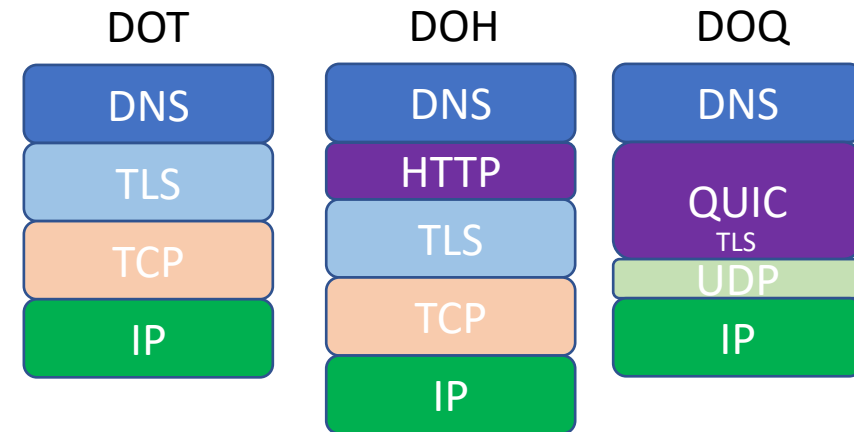
- HTML prefetch?
 - How can the client ascertain if the pushed data is genuine?
 - What is the use context of the pushed name resolution?
- DoH only names
 - Implicit client identification allowing for client customisation
- Morph the DNS into the WEB infrastructure?
 - Use HTTPS content distribution infrastructure for DoH web objects

If the web is doing it... why
not the DNS?

DoQ - DNS over QUIC

- QUIC is a transport protocol originally developed by Google and passed over to the IETF for standardised profile development
- QUIC uses a thin UDP shim and an encrypted payload
 - The payload is divided into a TCP-like transport header and a payload
- The essential difference between DOT and DOQ is the deliberate hiding of the transport protocol from network middleware with the use of QUIC
- No known implementations of DNS over QUIC exist, though IETF work continues

draft-huitema-quic-dns-quic-07



DoT, DoH, DoQ

Its not a rule, but


- It seems that applications (browsers) are looking to DoH and possibly DoQ
- Platforms are looking to use DoT as an alternative to DNS in the clear

Whose DNS is it anyway?

- ISP-provided DNS infrastructure
- User configured DNS resolvers can override ISP defaults
 - Although open DNS and DoT can be blocked at the ISP level by port level blocking and interception
 - It's unclear whether DoH and DoQ can be blocked so readily
- Application selected resolvers can override ISP and platform configured defaults
 - It's unclear whether an applications use of DoH can even be detected by the platform, let alone by the ISP

DNS use in the Internet

Top 10



Provider's DNS


Same country
(Provider's DNS)

Google's Public
DNS

All the rest (Open
DNS platforms)

DNS use in New Zealand

Top 10



Provider's DNS

Same country (Provider's DNS)

Google's Public DNS

All the rest (Open DNS platforms)

DNS use in New Zealand

ASN	AS Name	sameas	samecc	googlepdns	diffcc	cloudflare	opendns	quad9	level3	neustar	114dns	Samples
AS4771	SPARKNZ Spark New Zealand Trading Ltd.	97.042%	0.827%	6.044%	0.475%	0.646%	0.420%	0.053%	0.011%	0.020%	0.000%	72,245
AS9500	VODAFONE-TRANSIT-AS Vodafone NZ Ltd.	96.367%	0.682%	6.470%	0.498%	0.684%	1.045%	0.072%	0.022%	0.000%	0.000%	22,777
AS9790	VOCUSGROUPNZ VocusGroup	91.952%	1.144%	13.462%	0.812%	1.050%	0.941%	0.062%	0.073%	0.021%	0.000%	15,855
AS38793	NZCOMMS-AS-AP Two Degrees Mobile Limited	99.732%	0.412%	5.462%	0.000%	0.078%	0.000%	0.000%	0.000%	0.000%	0.000%	6,015
AS55850	TRUSTPOWERLTD-AS-AP TrustPower Ltd	97.155%	0.661%	7.085%	0.800%	0.313%	0.578%	0.156%	0.031%	0.061%	0.059%	5,197
AS23655	SNAP-NZ-AS Snap Internet Limited	85.159%	2.613%	18.705%	3.957%	1.559%	1.803%	0.120%	0.183%	0.118%	0.000%	2,819
AS133579	MYREPNZ-AS-AP MYREPUBLIC LIMITED	93.490%	0.709%	10.680%	0.505%	1.873%	2.359%	0.054%	0.000%	0.000%	0.000%	2,658
AS56030	VOYAGERNET-AS-AP Voyager Internet Ltd.	71.703%	4.503%	33.218%	2.446%	9.889%	4.214%	0.000%	0.000%	0.000%	0.000%	1,549
AS20473	AS-CHOOPA	52.743%	0.000%	49.783%	14.259%	8.271%	3.208%	0.104%	0.114%	0.000%	2.314%	1,504
AS9876	NOWNEW-AS-AP NOW New Zealand Ltd.	94.010%	0.711%	11.869%	0.117%	1.298%	0.804%	0.000%	0.000%	0.000%	0.000%	1,415
AS4768	VFNZ-INET-AS Vodafone NZ Ltd	27.866%	68.589%	19.769%	4.622%	2.697%	5.110%	5.227%	0.000%	0.364%	0.000%	1,327
AS45177	DEVOLI-AS-AP Devoli	88.439%	1.184%	24.739%	1.433%	3.496%	2.180%	0.764%	0.582%	1.412%	0.000%	1,053
AS136442	OCEANWAVE-AS-AP Ocean Wave Communication Co., Ltd	48.835%	0.000%	100.000%	4.005%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	1,051
AS58600	FLIP-AS-AP Flip Services Limited	0.000%	94.785%	0.000%	5.200%	0.637%	0.000%	0.000%	0.000%	0.000%	0.000%	1,048
AS17705	INSPIRENET-AS-AP InSPire Net Ltd	80.326%	0.980%	24.265%	0.696%	1.007%	1.736%	0.000%	0.000%	0.000%	0.000%	984
AS55853	MEGATEL-AS-AP Megatel	5.023%	0.503%	95.850%	0.353%	0.862%	0.540%	0.353%	0.000%	0.000%	0.000%	972
AS4648	SPARK-NZ Global-Gateway Internet	71.779%	8.954%	22.279%	9.033%	5.633%	5.059%	3.990%	0.365%	0.000%	0.000%	967
AS55872	BAYCITY-AS-AP BayCity Communications Limited	99.112%	19.910%	15.352%	1.256%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	889
AS45267	LIGHTWIRE-AS-AP Lightwire LTD	92.457%	7.883%	19.591%	8.017%	2.940%	1.897%	0.857%	0.000%	0.000%	0.000%	785
AS18199	LINKTELECOM-NZ-AP Link Telecom (NZ) Limited	96.556%	0.222%	5.294%	2.279%	1.110%	0.874%	1.346%	0.000%	0.000%	0.000%	683
AS4764	WIDEBAND-AS-AP Aussie Broadband	82.824%	0.000%	15.429%	0.312%	5.772%	1.771%	0.000%	0.000%	0.000%	0.000%	520
AS9245	COMPASS-NZ-AP COMPASS	95.182%	0.969%	10.410%	0.340%	0.977%	0.000%	0.000%	0.000%	0.000%	0.000%	497

Thanks!