Early observations from BT and DT DoH Trials

UKNOF 45 – 15th January 2020

Andy Fidler, Principal Network Architect, BT
Nic Leymann, Senior Network Architect, DT

NOTE: ALL OBSERVATIONS ARE EARLY FINDINGS AND SUBJECT TO CHANGE

BT DoH Experimental Trial

- **Shared with industry technical community on 6th December, 2019.**
- **Available* at https://doh.bt.com/dns-query/ with test page at http://splashpage.doh.bt.com**
- **Currently testing across small base of BT employees.**
- **Built on and working with OpenXchange / PowerDNS.**
- **Supporting only IPv4 and RFC8484 implementation.**
- **For the trial providing a public / open resolver.**
- **Shortly planning to enable DNSSEC validation.**

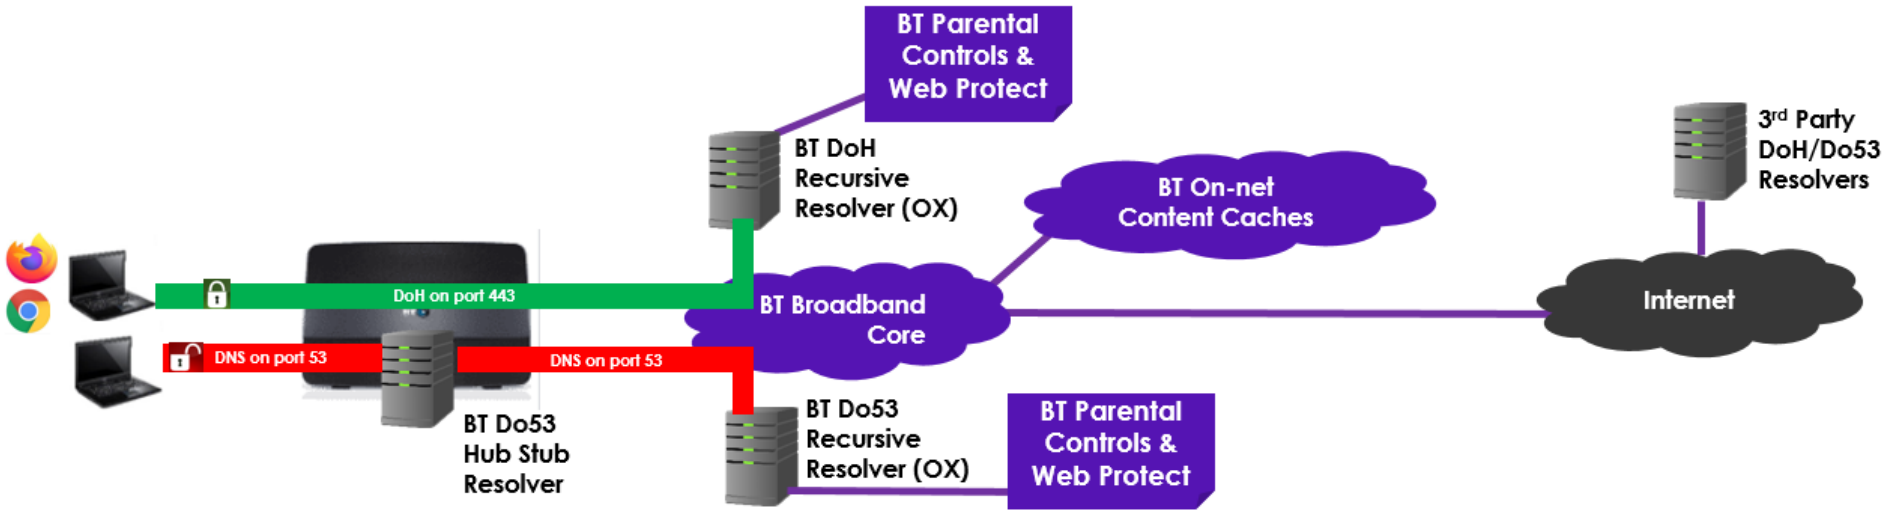

***Please note this is not an official service in any way.** It is purely experimental, may not offer similar service performance to live services and may be taken out of service without notice. The experimental capability should support any existing BT customer parental control and/or web protect settings, however if you are testing the capability on family devices we would recommend that you check that parental controls are still applied. Personal data will be processed in accordance with BT's Privacy Policy - https://www.bt.com/privacy-policy/

Early Customer Experience observations from BT Trial

Customer Experience	Status	Observations	Industry Opportunity
Browser Manual Custom Entry Set-up		Firefox: Simple manual custom entry. Chrome: Via executable flags now, but should be addressed via options in 81.	Consideration on: 1) applying policy detection to custom entry as well as auto enablement. 2) providing visual notification to customer on DoH usage.
Future auto discovery		For customers using BT Hubs with stub resolvers presenting private IP addresses to clients, inability for applications to discover BT as ISP and DoH status.	Demonstrates clear need for a context aware DoH discovery protocol to be developed within proposed new IETF Adaptive DNS Discovery (ADD) group.
Browsing Experience		For general users a good browsing experience, however early technical measurements appear to be showing additional latency from TLS set-up and variations based on encryption settings approach.	Demonstrates benefits to be gained from creating Best Current Practices (BCP) recommendations on DoH encryption options. BCPs could be within IETF, EDDI, ISPA or GSMA.
DNS Parental Control		Verified successful co-existence of BT Parental Controls with DNS over HTTPS.	Industry standardisation of policy detection protocol and use with custom entry as well as auto enablement.
DNS Malware Protection		Verified successful co-existence of BT Web Protect with DNS over HTTPS.	Industry standardisation of policy detection protocol and use with custom entry as well as auto enablement.
Context Awareness		If custom DoH entry is unavailable (e.g. off network), then browsers may still try this first then fall-back to default Do53 settings, potentially creating a slower response.	Demonstrates need for IETF ADD group to develop a context aware DoH discovery protocol supporting broadband, mobile and 3rd party wi-fi options.
Hub / Device Set-up		Breaks simple BT hub set up GUI URL – “hub.home” link.	Future ISP hubs will need to avoid using private domains.

BT DoH Load Test Configuration

- DNS queries raised using cloud hosted BlazeMeter™.
- BT successfully ran load tests at 200 and 1000 QPS using around 830,000 different URLs, with worst case one look up per TLS session.

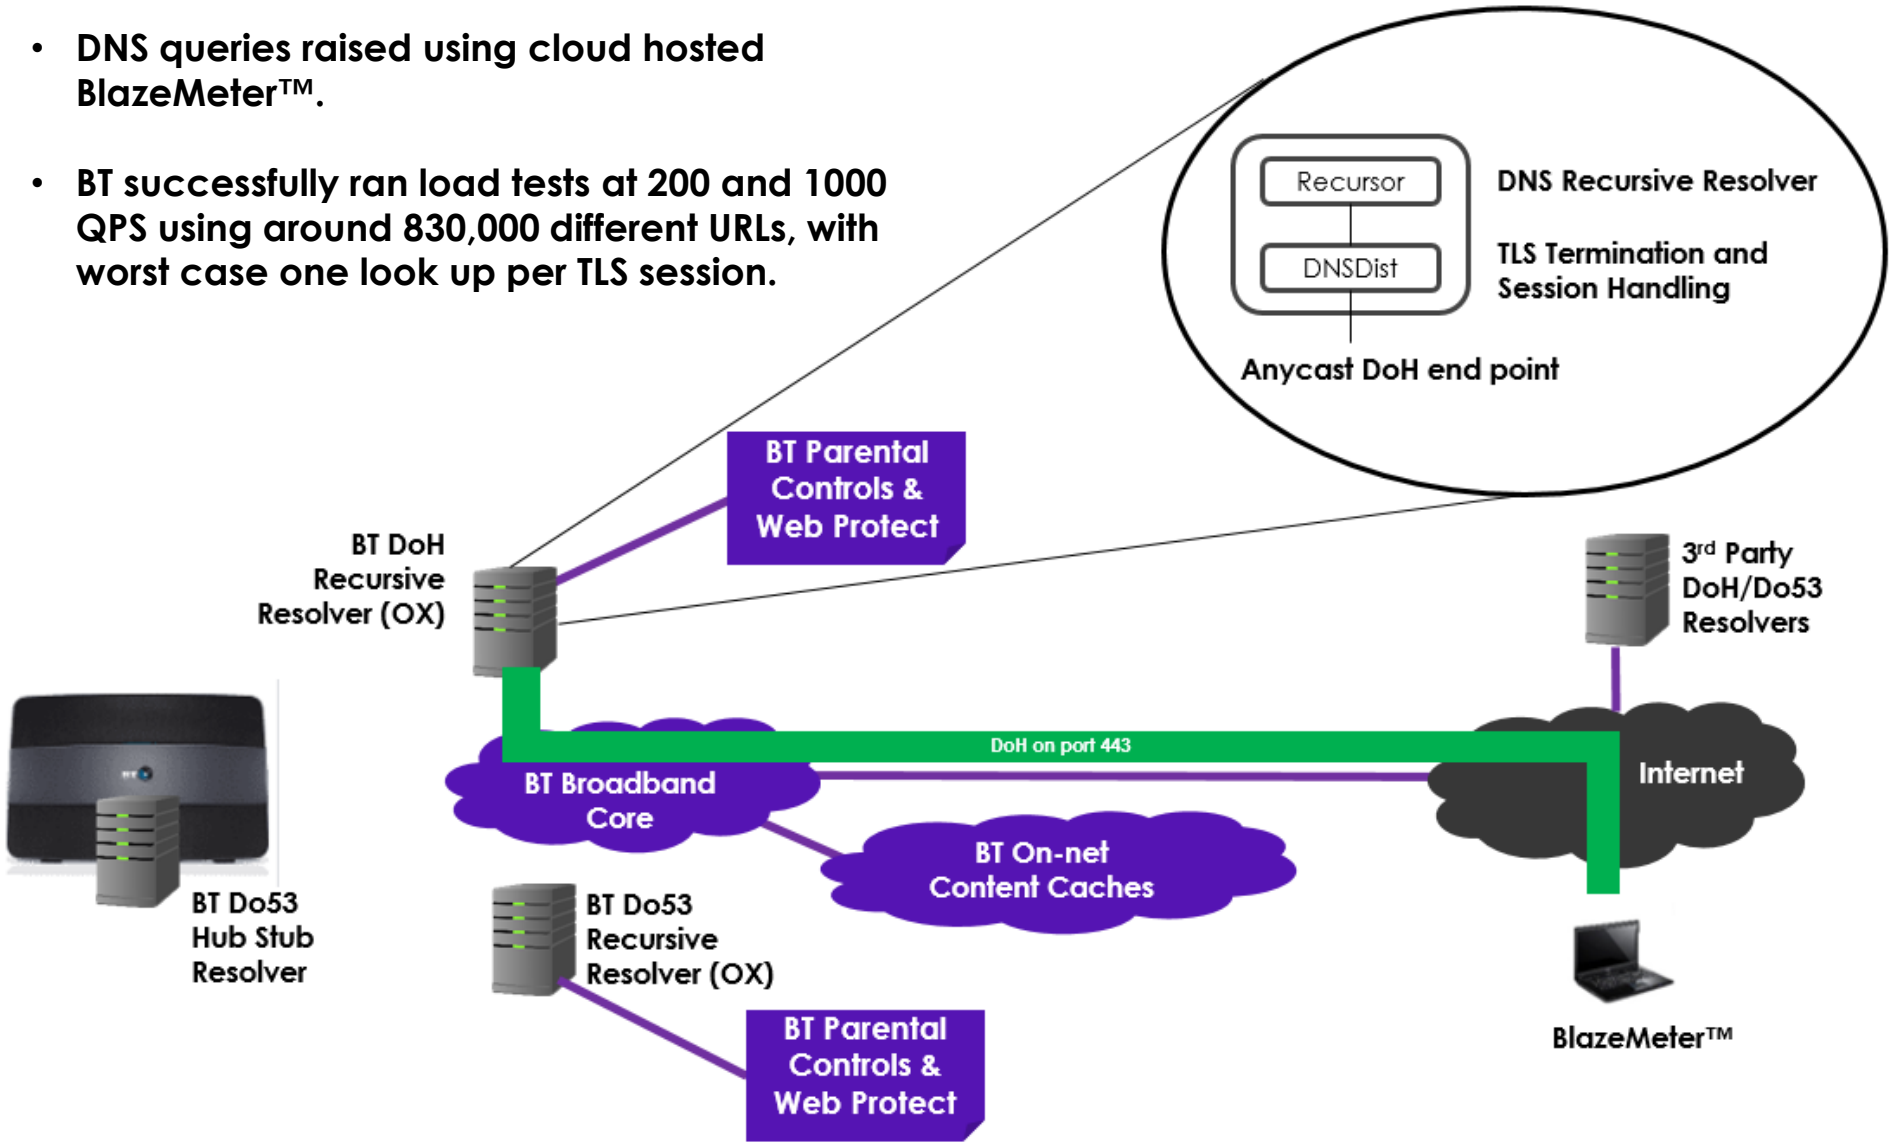

Early Performance Observations from BT DoH Trial

Full look up time in seconds from UK BT Broadband line	Cloudflare DoH	Google DoH	BT (UK) DoH	DT (Germany) DoH	Comcast (US) DoH
	TLS 1.3		TLS 1.2		
Facebook.com	0.260	0.267	0.262	0.414	0.610
a2.w10.akamai.net	0.263	0.271	0.277	0.317	0.835
google.co.uk	0.239	0.245	0.272	0.326	0.608
BT is observing that TLS 1.2 adds an overhead compared TLS 1.3					

Full look up time (s)	BT	Cloudflare	Google
DoH curl	0.34 (TLS 1.2)	0.26 (TLS 1.3)	0.20 (TLS 1.3)
Do53 pingu	0.013	0.014	0.02
Do53 curl	0.066	tbc	0.109

Early measurements are suggesting DoH has greater latency due to TLS set-up. However BT is still exploring whether existing test probes are ideal for DoH. To assist this BT will shortly be testing with whiteboxes.
It should also be noted that Curl measurements reflects worse case – TLS session per query scenario.

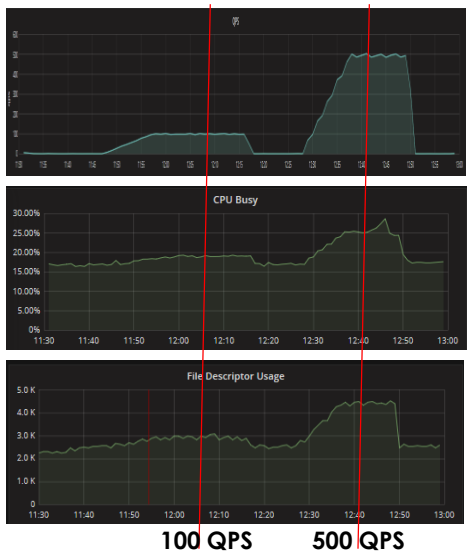

Early results from load tests seem to be indicating a higher than expected TLS overhead on server capacity.

200 & 1k QPS distributed servers.
10% CPU increase
100% file descriptor increase

NB: Background trial usage < 10 QPS

DoH via TLS 1.2 vs TLS 1.3

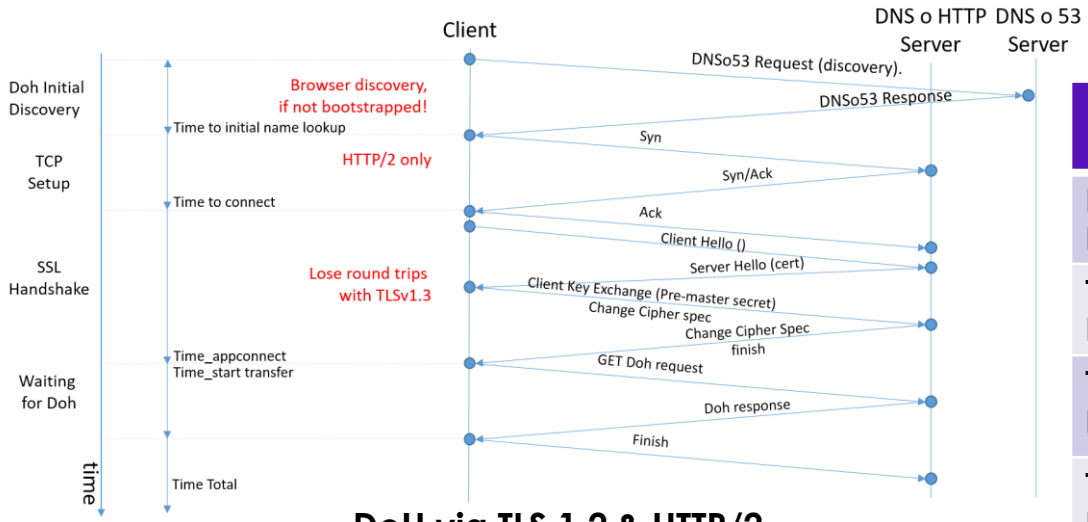

DoH via TLS 1.2 & HTTP/2

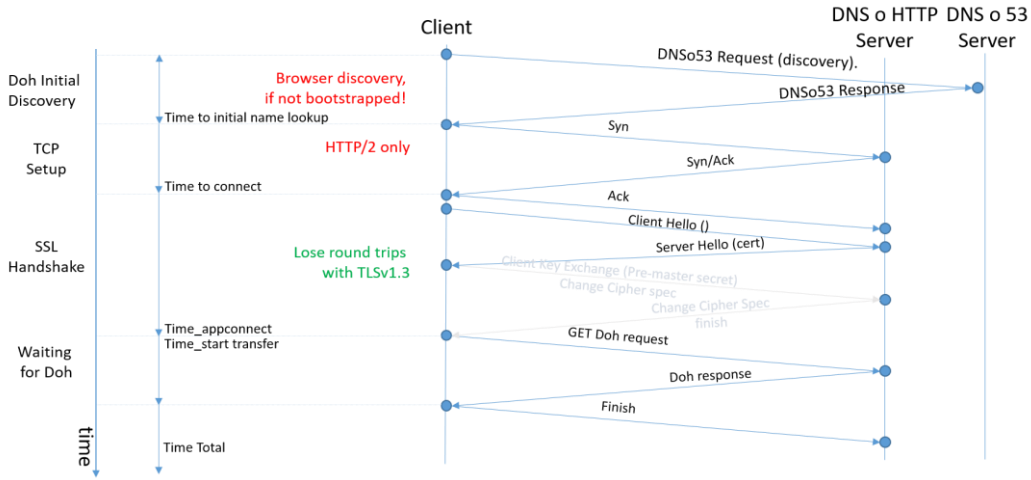

DoH via TLS 1.3 & HTTP/2

DOH Option	Round Trips	Perfect Forward Secrecy
DoH via Curl on HTTP/2 & TLS 1.2	6	Yes
Theoretical DoH on HTTP/2 & TLS 1.3	5	Yes
Theoretical DoH on HTTP/3 & TLS1.3	3.5	Yes
Theoretical DoH on HTTP/3 & TLS 1.3 Plus 0-RTT	2.5	No
Theoretical DoH on HTTP/3 & TLS 1.3 Plus 0-RTT, hardcode & no discovery	1.5	After initial

Moving to TLS 1.3 and HTTP/3 will significantly reduce TLS overheads and should ensure comparable DoH performance with Do53.

Variation in DoH resolver encryption settings

- BT has run Curl tests* against 21 DoH providers, highlighting some interesting variations and need for Best Current Practices deployment guidelines.

DoH Provider	TLS 1.3	OCSP Stapling	Session ID Duration (s)	Ticket Session (s)	Cipher Choice
Cloudflare	Yes	No	7200	172800 (2 days)	TLS_AES_256_GCM_SHA384
NextDNS	Yes	No	7200	604800 (7 days)	TLS_AES_256_GCM_SHA384
PowerDNS	Yes	No	7200	7200	TLS_AES_256_GCM_SHA384
Comcast	No (TLS 1.2)	No	7200	No	ECDHE-RSA-AES256-GCM-SHA384
Deutsche Telekom	No (TLS 1.2)	No	7200	7200	ECDHE-RSA-AES256-GCM-SHA384 ECDHE-ECDSA-AES256-GCM-SHA384
Andrews & Arnold	Yes	No	7200	7200	TLS_AES_256_GCM_SHA384
Google	Yes	No	7200	172800 (2 days)	TLS_AES_256_GCM_SHA384
BT Plc	No (TLS 1.2)	Yes (7 days)	7200	300	ECDHE-RSA-AES256-GCM-SHA384 ECDHE-ECDSA-AES256-GCM-SHA384

*Results based on tests run on 27/12/19

Saves client having to check status with CA.

Plus what about in-band authentication?

Clients & servers need to hold session resumption artefacts. Will 7200s take-up too much memory as DoH scales, should it be lower?

Why so varied and some so long?

What's the best balance here between privacy and user experience?

Variation in Cipher Choice.

Variation in DoH Protocol Support & HTTP Response Status Codes

DoH Provider	DoH RFC8484	DoH-JSON (response code)	Support HTTP/1.0 Head Request (Response Code)	Support HTTP/1.1 Head Request (Response Code)	Support HTTP/2 Head Request (Response Code)	HTTP/3 Head Request (Response Code)
Cloudflare	Yes	Yes	No (200)	Yes (200)	Yes (200)	No (200)
NextDNS	Yes	Yes	Yes (405)	Yes (405)	Yes (405)	No
PowerDNS	Yes	No (400)	No (400)	Yes (400)	Yes (400)	No
Comcast	Yes	No (400)	No (400)	Yes (400)	Yes (400)	No
Deutsche Telekom	Yes	No (400)	No (404)	Yes (404)	Yes (404)	No
Andrews & Arnold	Yes	No (400)	?	Yes (302)	Yes (302)	No (302)
Google	Yes	No (400)	Yes (200)	Yes (200)	Yes (200)	Yes
BT Plc	Yes	No (400)	No (400)	Yes (400)	Yes (400)	No

Only Cloudflare & NextDNS supporting non-standard JSON

Noticed different listeners and variation in HTTP response status codes return for head requests, how will clients handle this variation?

Does DoH HTTP status response codes approach need to be covered in BCPs and thoughts on test tools?

DoH Cookie Observations from BT Trial

- User interfaces and policies may not be clear on how cookies are handled across browser and DoH databases. We appear to be seeing the browser side mention cookies for DoH domains.
  - We assume this is due to visiting the domain itself, but would welcome user interface clarity on which cookies are present in which database, and confirmation that browsers and DoH servers are not sending / accepting cookies in DoH messages.

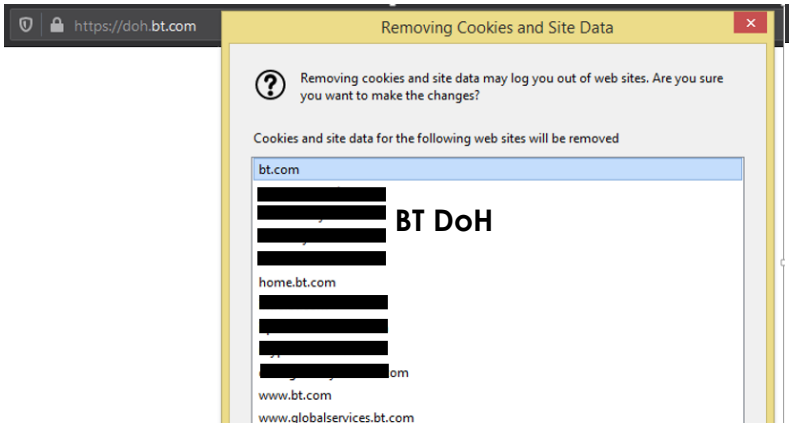

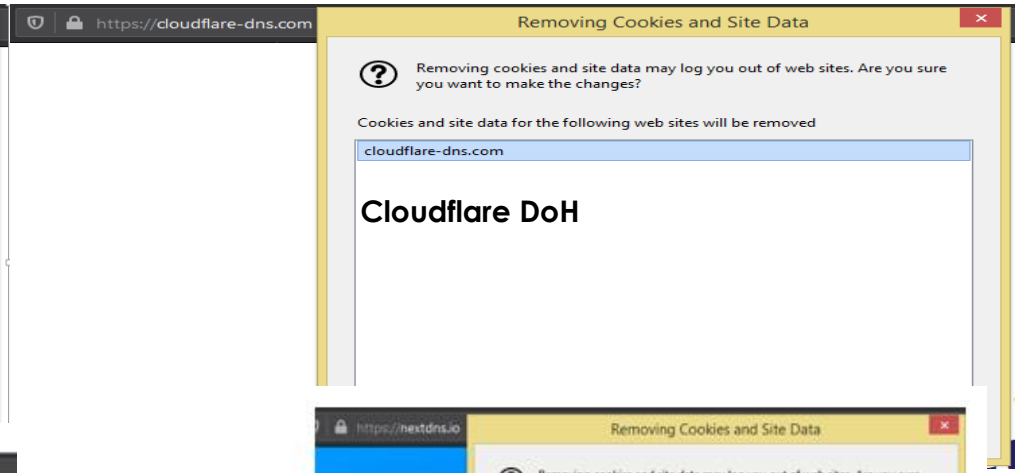

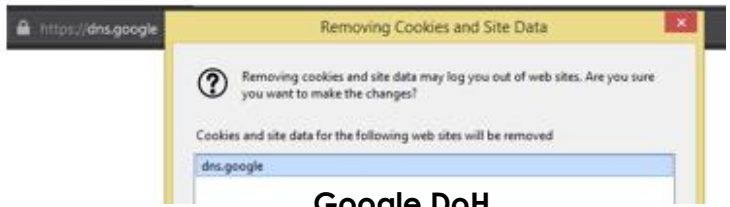

- Further clarification may needed in DoH BCPs and subsequent I-D's / RFCs to state that:
  - Clients should not accept "Set-Cookie" as part of a DoH response.
  - Clients should not send "Cookie" headers they have previously learned for the relevant domain.
  - DoH servers should disregard Cookies.
  - Guidance on DoH namespace.

The Evolution of DNS at DT
DT DNS Platform

DT runs a huge, high performance DNS infrastructure, fully redundant IPv4/IPv6 enabled.

About 2 Million DNS requests per second are handled

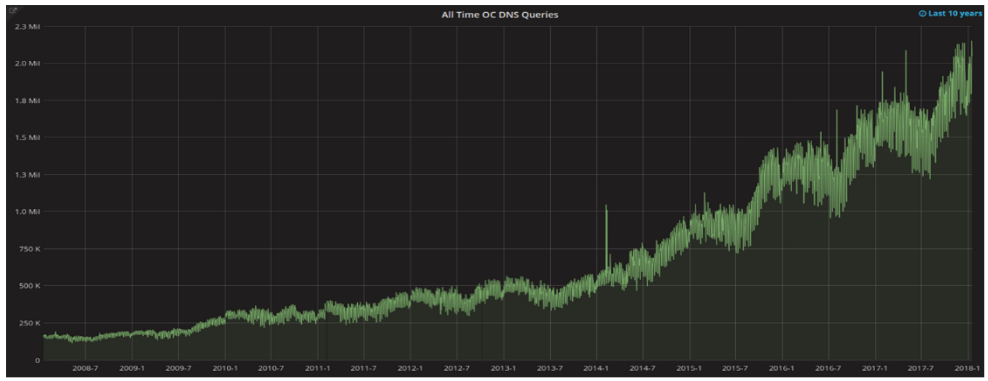

The DT DNS platform is the foundation for implementing a wide variety of user services. Those services REQUIRE that end users are using the DT DNS infrastructure. This includes security features, NAT64/DNS64 in Mobile Networks, Load Balancing for CDNs, …

The Evolution of DNS at DT
Standard DNS Deployment in DT Fixed Network

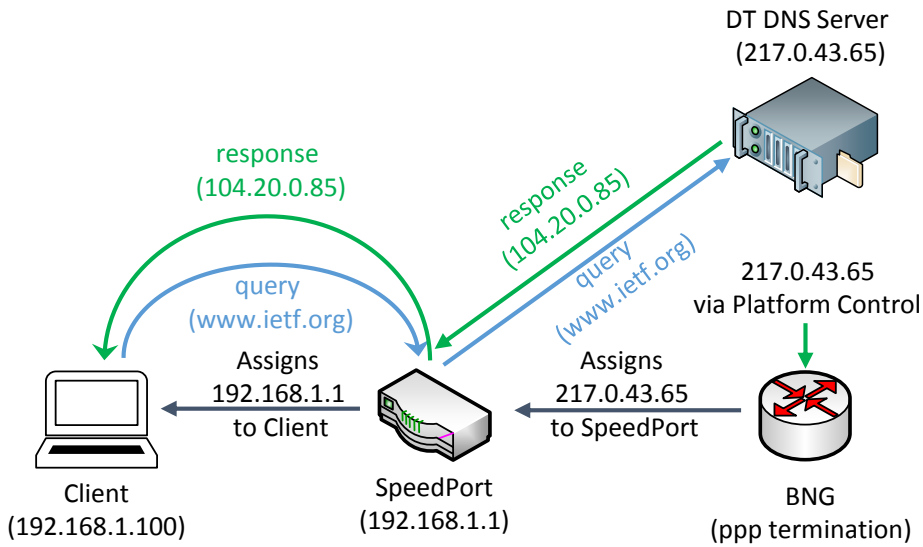

(*) Customer is able to overwrite/change DNS settings in Home Gateway

Remarks

- BNG obtains addresses of DT DNS servers via Platform Control
- BNG assigns addresses of DT DNS servers to Home Gateway (e.g. SpeedPort) during PPPoE session setup.
- Home Gateway acts as a DNS Proxy on behalf of home network devices
- Address for DNS queries is assigned to all local clients by Home Gateway (usually via DHCP) (*)
- All end devices in home network are using the IP address of the Home Router as DNS server address.
- Home Router forwards requests to DT DNS servers

DNS Server addresses are under control of the service provider.

The Evolution of DNS at DT
Status of DNS Implementation

- **Several different DNS platforms, Mobile and Fixed Line DNS largest implementations (ongoing consolidation of DNS at DT)**
  - Same software platform
- **DNS under control of DT, platform provides name resolution and DNS based services**
  - Guarantees privacy, reliability and high performance DNS implementation
  - DNS problems can be tracked and solved by DNS operations (only if customer uses DT platform)
  - Note: Customers are not mandated to use DT DNS platform (but more than 92% of customers are using the DT DNS as default), about 6% of DNS traffic to google, 2% other DNS providers)
- **Currently only DNS53 implemented (due to lack of client implementation in home gateways)**
  - DoT seen as evolution path from DNS53 towards encrypted DNS, not changing the deployment model and responsibilities
- **DT evaluating DoT and DoH implementation in existing DNS platform (for mobile and fixed network customers)**

The Evolution of DNS at DT
DT DoH and DoT Experimental Implementation

- **DT started test phase for DoT/DoH in an experimental setup in one of the existing DNS locations**
- **Both protocols are offered in parallel on the existing resolvers …**
  - … but DoH being limited to a subset
- **Based on PowerDNS solution**
- **Supports IPv4 and IPv6, IPv6 preferred protocol**
  - Side node: DT enables Dual Stack for all residential customers by default
- **Internal testing only, with a small number of employees**

The Evolution of DNS at DT
Comparison DoT and DoH

DoT – DNS over TLS

- Preferred solution for encrypting DNS traffic
- Implementation in Home Gateway (straight forward approach to encrypt DNS traffic)
- Same DNS IP addresses as standard DNS53, as transparent as possible for end customers (no customer impact)
- Home Gateway (or client) should probe for DoT support on existing DNS addresses and choose based on customers settings
  - Additional configurations (e.g. certificates) need to be provided

DoH – DNS over HTTPS

- Implemented to address move towards OTT DNS/DoH
  - Not our preferred solution for encrypted DNS
- More complex operational model, different (additional) DNS server infrastructure (to protect DNS53/DoT servers)
  - Due to the lack of dynamic discovery, DoH infrastructure open to all users (not only DT customers)
  - Larger target of attacks
- Without discovery, complex model for providing DoH addresses to end customers

The Evolution of DNS at DT
Open Issues

- **Discovery mechanism for DoH necessary**
  - Server information need to be provided by the network (no static configuration without asking customer in application)
  - Same operational model as today, DoH only used if provider supports DoH (or if end customer changes manually the DNS configuration)
  - If DoH is not supported by provider, fallback to either DoT or standard DNS53
  - Necessary to minimize customer impact
- **Set of uses cases for DoH necessary, including clear policies how OTT DNS is handled**
- **DoH can cause a lot of operational impact, if OTT services are used**
  - Performance, debugging, non working services, …
- **How does the user know if DoH is used (especially in the opportunistic scenarios)?**
  - Browser / APP should show some status information

Conclusions

- Good direction with more DoH resolvers and trials
- However many open issues still exist
- Many of which would benefit from the creation of the proposed new IETF working group
- A standardised DoH discovery protocol is required
  - and this needs to be context aware
  - and support scenario where DNS stub / proxy resolvers are used in hubs with private IP addresses
- Best Current Practice guidelines are needed to address potential variations in DoH and TLS settings.
- Further work needed on comparison of DoH vs Do53 performance and additional server capacity overhead.