ADD
Internet-Draft
Intended status: Standards Track
Expires: January 11, 2021

M. Boucadair
Orange
T. Reddy
McAfee
D. Wing
Citrix
N. Cook
Open-Xchange

July 10, 2020

Encrypted DNS Discovery and Deployment Considerations for Home Networks
draft-btw-add-home-07

Abstract

This document discusses DoT/DoH deployment considerations for home networks. It particularly sketches the required steps to use DoT/DoH capabilities provided by local networks.

The document specifies new DHCP and Router Advertisement Options to convey a DNS Authentication Domain Name.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 11, 2021.

Table of Contents

1. Introduction

Internet Service Providers (ISPs) traditionally provide DNS resolvers to their customers. Typically, ISPs deploy the following mechanisms to advertise a list of DNS Recursive DNS server(s) to their customers:

- Protocol Configuration Options in cellular networks [TS.24008].

- DHCP [RFC2132] (Domain Name Server Option) or DHCPv6 [RFC8415][RFC3646] (OPTION_DNS_SERVERS).
- IPv6 Router Advertisement [RFC4861][RFC8106] (Type 25 (Recursive DNS Server Option)).

The communication between a customer's device (possibly via Customer Premises Equipment (CPE)) and an ISP-supplied DNS resolver takes place by using cleartext DNS messages (Do53, [I-D.ietf-dnsop-terminology-ter]). Some examples are depicted in Figure 1. In the case of cellular networks, the cellular network will provide connectivity directly to a host (e.g., smartphone, tablet) or via a CPE. Do53 mechanisms used within the Local Area Network (LAN) are similar in both fixed and cellular CPE-based broadband service offerings.

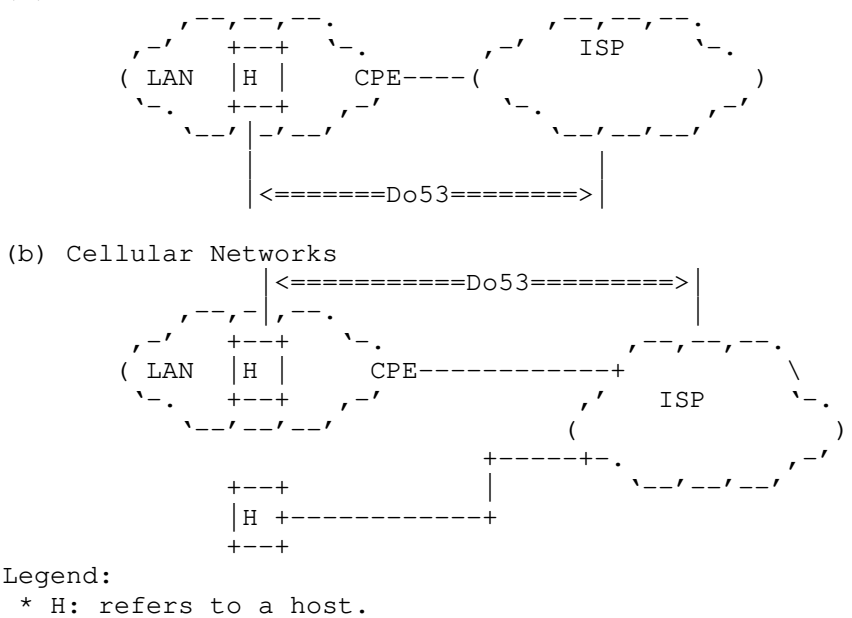

Figure 1: Sample Legacy Deployments

ISPs use DNS to provide additional services such as (but not limited to) malware filtering, parental control, or VoD (Video on Demand) optimization. DNS is also a central component for mastering the quality of experience for current latency-sensitive services, but also emerging ones (such as those services that pertain to the Ultra Reliability and Low Latency Communications (uRLLC) or Enhanced Mobile Broadband (eMBB).

> For example, the latency targets set in the context of 5G are 1ms (uRLLC) and 4ms (eMBB). An ISP will be able to address such demanding latency requirements assuming the corresponding services rely upon resources (network, compute, storage) that are located as close to the user as possible (e.g., by means of Edge Computing techniques and resources). Such latency requirements are likely to be addressed by means of optimized designs (DNS, in particular), too.

Relying upon local DNS resolvers will therefore contribute to meet the aforementioned service requirements. The use of external resolvers is likely to induce an extra service delay which exceeds by far the service target.

This document focuses on the support of DNS-over-HTTPS (DoH) [RFC8484] or DNS-over-TLS (DoT) [RFC7858] in local networks. In particular, the document describes how a local DoH/DoT server can be discovered and used by connected hosts. This document specifies options that allow DNS clients to discover local DoT/DoH servers. Section 4 describes DHCP, DHCPv6, and RA options to convey the Authentication Domain Name (ADN, defined in [RFC8310]).

Some ISPs rely upon external resolvers (e.g., outsourced service or public resolvers); these ISPs provide their customers with the IP addresses of these resolvers. These addresses are typically configured on CPEs using the same mechanisms listed above. Likewise, users can modify the default DNS configuration of their CPEs (e.g., supplied by their ISP) to configure their favorite DNS servers. This document permits such deployments.

Both managed and unmanaged CPEs are discussed in the document (Section 3). Also, considerations related to hosting a DNS forwarder in the CPE are described (Section 7).

Hosts and/or CPEs may be connected to multiple networks; each providing their own DNS configuration using the discovery mechanisms specified in this document. Nevertheless, it is out of the scope of this specification to discuss DNS selection of multi-interface devices. The reader may refer to [RFC6731] for a discussion of issues and an example of DNS server selection for multi-interfaced devices.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP

14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499] and [I-D.ietf-dnsop-terminology-ter].

Do53 refers to unencrypted DNS.

'DoH/DoT' refers to DNS-over-HTTPS and/or DNS-over-TLS.

3. Sample Deployment Scenarios

3.1. Managed CPEs

ISPs have developed an expertise in managing service-specific configuration information (e.g., CPE WAN Management Protocol [TR-069]). For example, these tools may be used to provision the authentication domain name information (ADN) to managed CPEs if DoH/DoT is supported by a local network similar to what is depicted in Figure 2.

DoH-capable (or DoT) clients establish the DoH (or DoT) session with the discovered DoH (or DoT) server.

The DNS client discovers whether the DNS server in the local network supports DoH/DoT by using a dedicated field in the discovery message: Encrypted DNS Types (Section 4).

(a) Fixed Networks

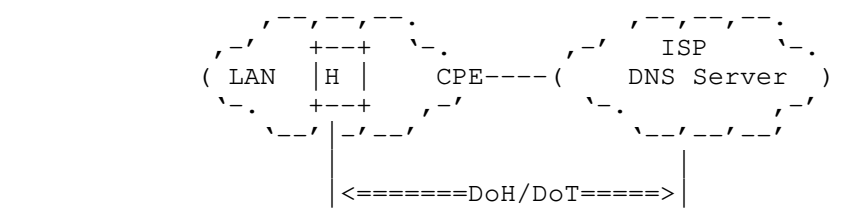

(b) Cellular Networks

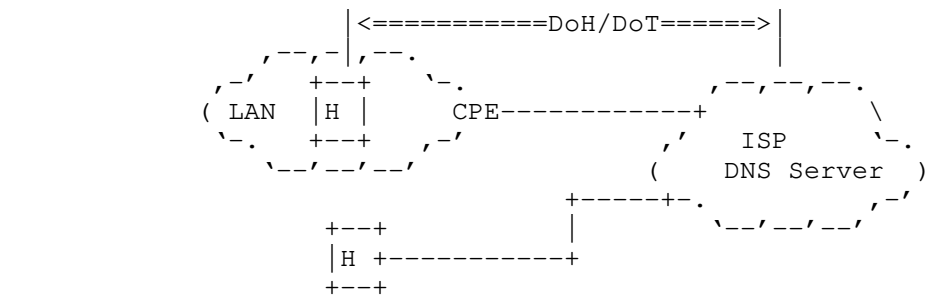

Figure 2: DoH/DoT in the WAN

Figure 2 shows the scenario where the CPE relays the list of DoT/DoH servers it learns for the network by using mechanisms like DHCP or a specific Router Advertisement message. In such context, direct DoH/DoT sessions will be established between a host serviced by a CPE and an ISP-supplied DoT/DoH server (see the example depicted in Figure 3 for a DoH/DoT-capable host).

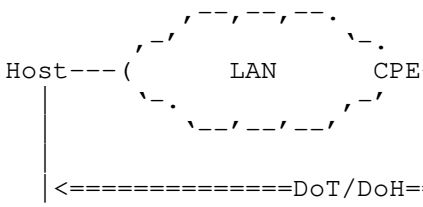

Figure 3: Direct DoH/DoT Sessions

Figure 4 shows a deployment where the CPE embeds a caching DNS forwarder. The CPE advertises itself as the default DNS server to the hosts it serves. The CPE relies upon DHCP or RA to advertise itself to internal hosts as the default DoT/DoH/Do53 server. When receiving a DNS request it cannot handle locally, the CPE forwards

the request to an upstream DoH/DoT/Do53 resolver. Such deployment is required for IPv4 service continuity purposes (e.g., [I-D.ietf-v6ops-rfc7084-bis]) or for supporting advanced services within the home (e.g., malware filtering, parental control, Manufacturer Usage Description (MUD, [RFC8520] to only allow intended communications to and from an IoT device)). When the CPE behaves as a DNS forwarder, DNS communications can be decomposed into two legs:

- The leg between an internal host and the CPE.
- The leg between the CPE and an upstream DNS resolver.

An ISP that offers DoH/DoT to its customers may enable DoH/DoT in both legs as shown in Figure 4. Additional considerations related to this deployment are discussed in Section 7.

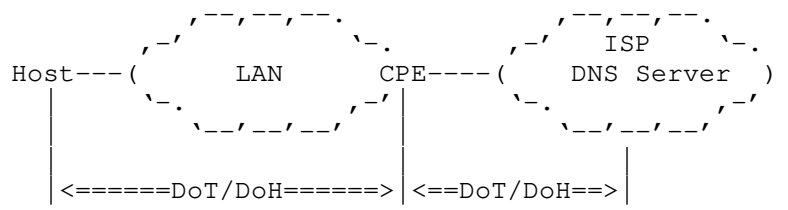

Figure 4: Proxied DoH/DoT Sessions

3.2. Unmanaged CPEs

Customers may decide to deploy unmanaged CPEs (assuming the CPE is compliant with the network access technical specification that is usually published by ISPs). Upon attachment to the network, an unmanaged CPE receives from the network its service configuration (including the DNS information) by means of, e.g., DHCP. That DNS information is shared within the LAN following the same mechanisms as those discussed in Section 3.1. A host can thus establish DoH/DoT session with a DoH/DoT server similar to what is depicted in Figure 3.

Customers may also decide to deploy internal home routers (called hereafter, Internal CPEs) for a variety of reasons that are not detailed here. Absent any explicit configuration on the internal CPE to override the DNS configuration it receives from the ISP-supplied CPE, an Internal CPE relays the DNS information it receives via DHCP/RA from the ISP-supplied CPE to connected hosts. DoH/DoT sessions can be established by a host with the DoH/DoT servers of the ISP (see Figure 5).

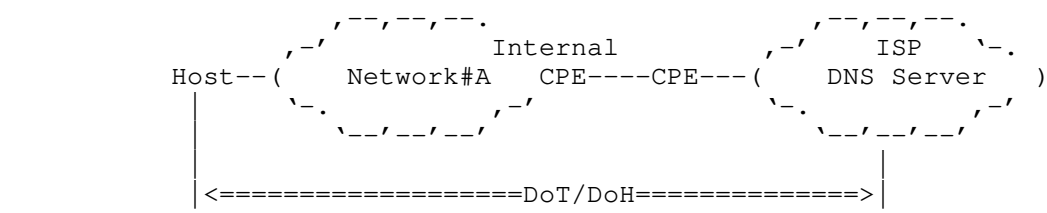

Figure 5: Direct DoH/DoT Sessions with the ISP DNS Resolver (Internal CPE)

Similar to managed CPEs, a user may modify the default DNS configuration of an unmanaged CPE to use his/her favorite DNS servers instead. DoH/DoT sessions can be established directly between a host and a 3rd Party DNS server (see Figure 6).

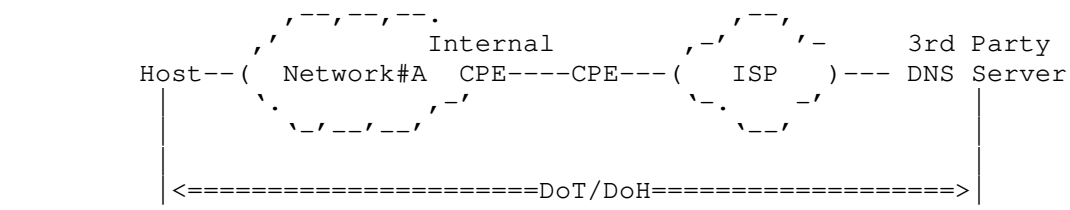

Figure 6: Direct DoH/DoT Sessions with a Third Party DNS Resolver

Section 7.2 discusses considerations related to hosting a forwarder in the Internal CPE.

4. DNS Reference Identifier Option

This section describes how a DNS client can discover the ADN of local DoH/DoT server(s) using DHCP (Sections 4.1 and 4.2) and Neighbor Discovery protocol (Section 4.3).

As reported in Section 1.7.2 of [RFC6125]:

> "few certification authorities issue server certificates based on IP addresses, but preliminary evidence indicates that such certificates are a very small percentage (less than 1%) of issued certificates".

In order to allow for PKIX-based authentication between a DNS client and a DoH/DoT server while accommodating the current best practices for issuing certificates, this document allows for configuring an authentication domain name to be presented as a reference identifier for DNS authentication purposes.

The DNS client establishes a DoH/DoT session with the discovered DNS IP address(es) (Section 6) and uses the mechanism discussed in Section 8 of [RFC8310] to authenticate the DNS server certificate using the authentication domain name conveyed in the DNS Reference Identifier.

If the DNS Reference Identifier is discovered by a host using both RA and DHCP, the rules discussed in Section 5.3.1 of [RFC8106] MUST be followed.

4.1. DHCPv6 DNS Reference Identifier Option

The DHCPv6 DNS Reference Identifier option is used to configure an authentication domain name of the DoH/DoT server. The format of this option is shown in Figure 7.

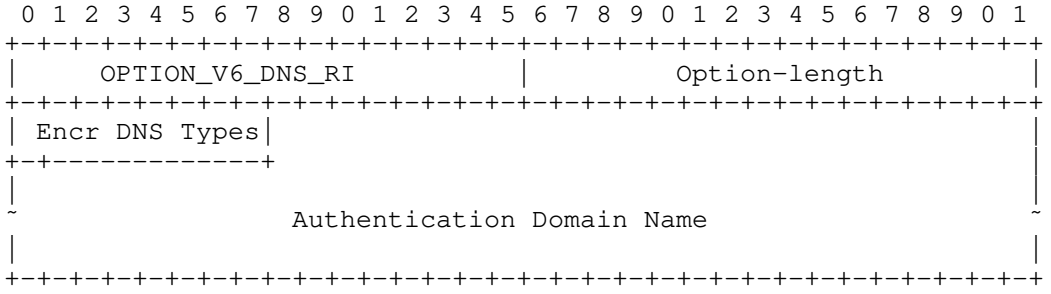

Figure 7: DHCPv6 DNS Reference Identifier Option

The fields of the option shown in Figure 7 are as follows:

- Option-code: OPTION_V6_DNS_RI (TBA1, see Section 10.1)
- Option-length: Length of the enclosed data in octets.
- Encr DNS Types (Encrypted DNS Types): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this 8-bit field is shown in Figure 8.

Figure 8: Encrypted DNS Types

T: If set, this bit indicates that the server supports DoT [RFC7858].
H: If set, this bit indicates that the server supports DoH [RFC8484].

U: Unassigned bits. These bits MUST be unset by the sender. Associating a meaning with an unassigned bit can be done via Standards Action [RFC8126].

In a request, these bits are assigned to indicate the requested encrypted DNS server type(s) by the client. In a response, these bits are set as a function of the encrypted DNS supported by the server and the requested encrypted DNS server type(s).

To keep the packet small, if more than one encrypted DNS type (e.g., both DoH and DoT) are to be returned to a requesting client and the same ADN is used for these types, the corresponding bits MUST be set in the 'Encrypted DNS Types' field of the same option instance in a response. For example, if the client requested DoH and DoTand the server supports both, then both T and H bits must be set.

- o Authentication Domain Name: A fully qualified domain name of the DoH/DoT server. This field is formatted as specified in Section 10 of [RFC8415].

An example of the Authentication Domain Name encoding is shown in Figure 9. This example conveys the FQDN "doh1.example.com.".

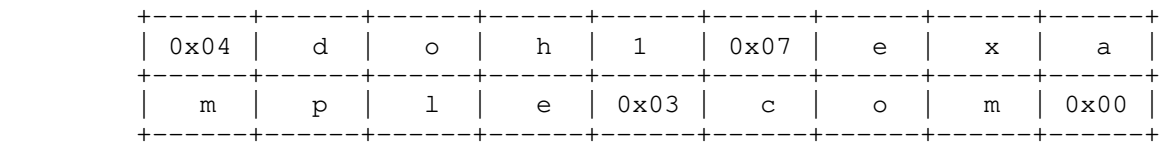

Figure 9: An example of the authentication-domain-name Encoding

Multiple instances of OPTION_V6_DNS_RI may be returned to a DHCPv6 client; each pointing to a distinct encrypted DNS server type.

To discover an encrypted DNS server, the DHCPv6 client including OPTION_V6_DNS_RI in an Option Request Option (ORO), as in Sections 18.2.1, 18.2.2, 18.2.4, 18.2.5, 18.2.6, and 21.7 of [RFC8415]. The DHCPv6 client sets the Encrypted DNS Types field to the requested encrypted DNS server type(s).

If the DHCPv6 client requested more than one encrypted DNS server type, the DHCP client MUST be prepared to receive multiple DHCP OPTION_V6_DNS_RI options; each option is to be treated as a separate encrypted DNS server.

4.2. DHCP DNS Reference Identifier Option

The DHCP DNS Reference Identifier option is used to configure an authentication domain name of the DoH/DoT server. The format of this option is illustrated in Figure 10.

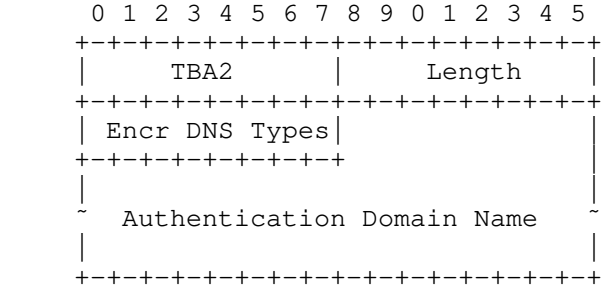

with:

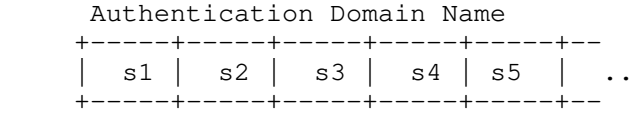

The values s1, s2, s3, etc. represent the domain name labels in the domain name encoding.

Figure 10: DHCP DNS Reference Identifier Option

The fields of the option shown in Figure 10 are as follows:

- Code: OPTION_V4_DNS_RI (TBA2, see Section 10.2).
- Length: Length of the enclosed data in octets.
- Encr DNS Types (Encrypted DNS Types): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this field is shown in Figure 8.
- Authentication Domain Name: The domain name of the DoH/DoT server. This field is formatted as specified in Section 10 of [RFC8415].

OPTION_V4_DNS_RI is a concatenation-requiring option. As such, the mechanism specified in [RFC3396] MUST be used if OPTION_V4_DNS_RI exceeds the maximum DHCP option size of 255 octets.

To discover an encrypted DNS server, the DHCP client requests the Encrypted DNS Reference Identifier by including OPTION_V4_DNS_RI in a Parameter Request List option [RFC2132]. The DHCP client sets the Encrypted DNS Types field to the requested encrypted DNS server.

If the DHCP client requested more than one encrypted DNS server type, the DHCP client MUST be prepared to receive multiple DHCP OPTION_V4_DNS_RI options; each option is to be treated as a separate encrypted DNS server.

4.3. RA DNS Reference Identifier Option

The IPv6 Router Advertisement (RA) DNS Reference Identifier option is used to configure an authentication domain name of the DoH/DoT server. The format of this option is illustrated in Figure 11.

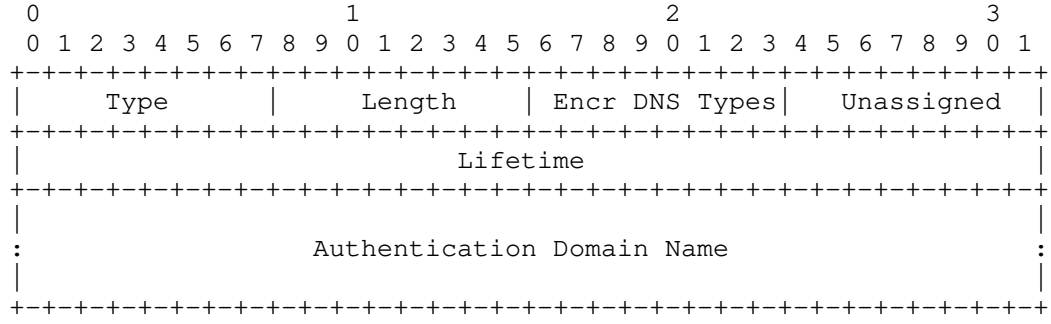

Figure 11: RA DNS Reference Identifier Option

The fields of the option shown in Figure 11 are as follows:

- Type: 8-bit identifier of the DNS Reference Identifier Option as assigned by IANA (TBA3, see Section 10.3).
- Length: 8-bit unsigned integer. The length of the option (including the Type and Length fields) is in units of 8 octets.
- Encr DNS Types (Encrypted DNS Types): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this field is shown in Figure 8.
- Unassigned: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
- Lifetime: 32-bit unsigned integer. The maximum time in seconds (relative to the time the packet is received) over which the authentication domain name MAY be used as a DNS Reference Identifier.

  The value of Lifetime SHOULD by default be at least 3 * MaxRtrAdvInterval, where MaxRtrAdvInterval is the maximum RA interval as defined in [RFC4861].

  A value of all one bits (0xffffffff) represents infinity.

A value of zero means that the DNS Reference Identifier MUST no longer be used.

- Authentication Domain Name: The domain name of the DoH/DoT server. This field is formatted as specified in Section 10 of [RFC8415].

  This field MUST be padded with zeros so that its size is a multiple of 8 octets.

5. DoH URI Templates

DoH servers may support more than one URI Template [RFC8484]. The following discusses a mechanism for a DoH client to retrieve the list of supported templates by a DoH server. Also, if the resolver hosts several DoH services (e.g., no-filtering, blocking adult content, blocking malware), these services can be discovered as templates.

Upon discovery of a DoH resolver (Section 4), the DoH client contacts that DoH resolver to retrieve the list of supported DoH services using the well-known URI defined in [I-D.btw-add-rfc8484-clarification]. DoH clients re-iterates that request regularly to retrieve an updated list of supported DoH services. Note that a "push" mode can be considered using the mechanism defined in [I-D.ietf-dnssd-push].

How a DoH client makes use of the configured DoH services is out of scope of this document.

6. Locating DoH/DoT Servers

A CPE or a host relies upon discovery mechanisms (such as PCO, DHCP, or RA) to retrieve DoH/DoT servers' reachability information. In the various scenarios sketched in Section 3, Do53, DoH, and DoT may terminate on the same IP address or distinct IP addresses. Terminating Do53/DoH/DoT on the same or distinct IP addresses is deployment-specific.

From an IP reachability standpoint, DoH/DoT servers SHOULD be located by their address literals rather than their names. This avoids adding a dependency on another server to resolve the DoH/DoT name. Concretely, if Do53/DoH/DoT terminate on same IP addresses, existing discovery mechanisms [RFC2132][RFC3646][RFC8106] can be leveraged to learn the IP addresses of DoT/DoH servers while an authentication domain name is supplied by one of the options discussed in Section 4.

The following sub-sections discusses the conditions under which discovered DoT/DoH server can be used.

6.1. DoT/DoH Auto-Upgrade

Additional considerations are discussed below for the use of DoH and DoT servers provided by local networks:

- If the DNS server's IP address discovered by using DHCP/RA is pre-configured in the OS or Browser as a verified resolver (e.g., part of an auto-upgrade program such as [Auto-upgrade]), the DNS client auto-upgrades to use the pre-configured DoH/DoT server tied to the discovered DNS server IP address. In such a case the DNS client will perform additional checks out of band, such as confirming that the Do53 IP address and the DoH server are owned and operated by the same organisation.

- Similarly, if the ADN conveyed in DHCP/RA (Section 4) is pre-configured in the OS or browser as a verified resolver, the DNS client auto-upgrades to establish a DoH/DoT session with the ADN.

  In such case, the DNS client matches the domain name in the DNS Reference Identifier DHCP/RA option with the 'DNS-ID' identifier type within subjectAltName entry in the server certificate conveyed in the TLS handshake.

6.2. Other Deployment Options

Some deployment options to securely configure hosts are discussed below. These options are provided for the sake of completeness.

- If Device Provisioning Protocol (DPP) [DPP] is used, the configurator can securely configure devices in the home network with the local DoT/DoH server using DPP. If the DoT/DoH servers use raw public keys [RFC7250], the Subject Public Key Info (SPKI) pin set [RFC7250] of raw public keys may be encoded in a QR code. The configurator (e.g., mobile device) can scan the QR code and provision SPKI pin set in OS/Browser. The configurator can in-turn securely configure devices (e.g., thermostat) in the home network with the SPKI pin set using DPP.

- If a CPE is co-located with security services within the home network, the CPE can use WPA-PSK but with unique pre-shared keys for different endpoints to deal with security issues. In such networks, [I-D.reddy-add-iot-byod-bootstrap] may be used to securely bootstrap endpoint devices with the authentication domain name and DNS server certificate of the local network's DoH/DoT server.

  The OS would not know if the WPA pre-shared-key is the same for all clients or a unique pre-shared key is assigned to the host.

Hence, the user has to indicate to the system that a unique pre-shared key is assigned to trigger the bootstrapping procedure.

If the device joins a home network using a single shared password among all the attached devices, a compromised device can host a fake access point, and the device cannot be securely bootstrapped with the home network's DoH/DoT server.

7. Hosting DoH/DoT Forwarder in the CPE

7.1. Managed CPEs

The following mechanisms can be used to host a DoH/DoT forwarder in a managed CPE (Section 3.1).

7.1.1. ACME

The ISP can assign a unique FQDN (e.g., cpe1.example.com) and a domain-validated public certificate to the DoH/DoT forwarder hosted on the CPE. Automatic Certificate Management Environment (ACME) [RFC8555] can be used by the ISP to automate certificate management functions such as domain validation procedure, certificate issuance and certificate revocation.

The managed CPE should support a configuration parameter to instruct the CPE whether it has to relay the encrypted DNS server received from the ISP's network or has to announce itself as a forwarder within the local network. The default behavior of the CPE is to supply the encrypted DNS server received from the ISP's network.

7.1.2. Auto-Upgrade based on Domains and their Sub-domains

If the ADN conveyed in DHCP/RA (Section 4) is pre-configured in popular OSes or browsers as a verified resolver and the auto-upgrade (Section 6.1) is allowed for both the pre-configured ADN and its sub-domains, the DoH/DoT client will learn the local DoH/DoT forwarder using DHCP/RA and auto-upgrade because the left-most label of the pre-configured ADN would match the subjectAltName value in the server certificate. Concretely, the CPE can communicate the ADN of the local DoH forwarder (Section 7.1.1) to internal hosts using DHCP/RA (Section 4).

Let's suppose that "example.net" is pre-configured as a verified resolved in the browser or OS. If the DoH/DoT client discovers a local forwarder "cpe1-internal.example.net", the DoH/DoT client will auto-upgrade because the pre-configured ADN would match subjectAltName value "cpe1-internal.example.net" of type dNSName. As

shown in Figure 12, the auto-upgrade to a rogue server advertising "rs.example.org" will fail.

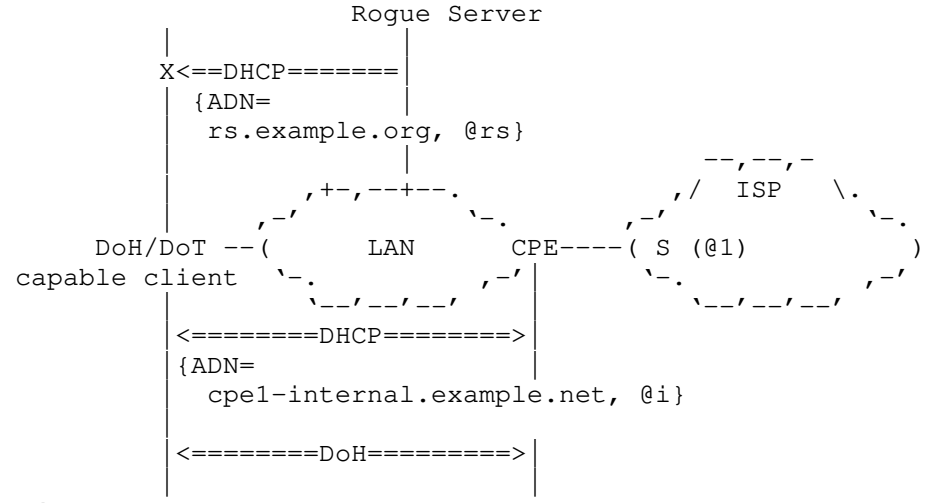

Legend:
* S: DoH/DoT server
* @1: IP address of S
* @i: internal IP address of the CPE
* @rs: IP address of a rogue server

Figure 12: A Simplified Example of Auto-upgrade based on Sub-domains

7.2. Unmanaged CPEs

The approach specified in Section 7.1 does not apply for hosting a DNS forwarder in an unmanaged CPE.

The unmanaged CPE administrator (referred to as administrator) can host a DoH/DoT forwarder on the unmanaged CPE. This assumes the following:

- The DoH/DoT server certificate is managed by the entity in-charge of hosting the DoT/DoH forwarder.

  Alternatively, a security service provider can assign a unique FQDN to the CPE. The DoH/DoT forwarder will act like a private DoT/DoH server only be accessible from within the home network.

- The DoH/DoT forwarder will either be configured to use the ISP's or a 3rd party DoH/DoT server.

o The unmanaged CPE will advertise the DoH/DoT forwarder ADN using DHCP/RA to internal hosts.

Figure 13 illustrates an example of an unmanaged CPE hosting a forwarder which connects to a 3rd party DoH/DoT server. In this example, the DNS information received from the managed CPE (and therefore from the ISP) is ignored by the Internal CPE hosting the forwarder.

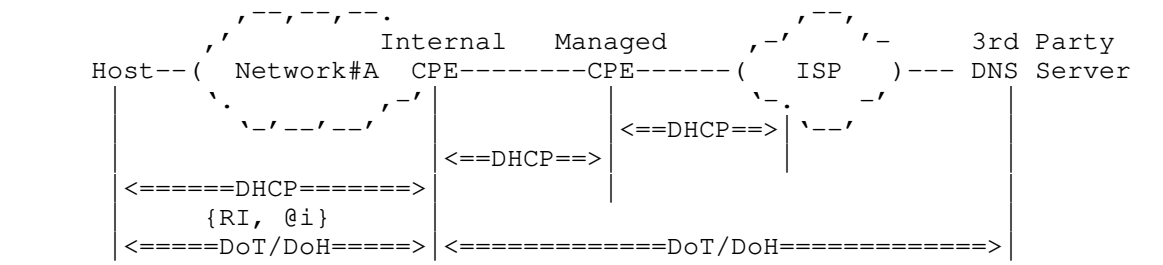

Legend:
* @i: IP address of the DNS forwarder hosted in the Internal CPE.

Figure 13: Example of an Internal CPE Hosting a Forwarder

8. Legacy CPEs

Hosts serviced by legacy CPEs that can't be upgraded to support the options defined in Section 4 won't be able to learn the DoH/DoT server hosted by the ISP, in particular. If the ADN is not discovered using DHCP/RA, such hosts will have to fallback to use the special-use domain name defined in [I-D.pp-add-resinfo] to discover the DoH/DoT server and to retrieve the list of supported DoH services using the RESINFO RRtype [I-D.pp-add-resinfo]. The DHCP/RA option to discover ADN takes precedence over special-use domain name since the special-use domain name is suseptible to both internal and external attacks whereas DHCP/RA is only vulnerable to internal attacks.

9. Security Considerations

An attacker can get a domain name, domain-validated public certificate from a CA, host a DoT/DoH server and claim the best DNS privacy preservation policy. Also, an attacker within the home network can use the public IP address, get an 'IP address'-validated public certificate from a CA, host a DoT/DoH server and claim the best DNS privacy preservation policy.

Wireless LAN as frequently deployed in home networks is vulnerable to various attacks (e.g., [Evil-Twin], [Krack], [Dragonblood]). Because of these attacks, only cryptographically authenticated communications are trusted on Wireless LAN networks. This means information provided by such networks via DHCP, DHCPv6, or RA (e.g., NTP server, DNS server, default domain) are untrusted because DHCP and RA are not authenticated.

Because DHCP/RA messages are not encrypted or protected against modification in any way, their content can be spoofed or modified by active attackers (e.g., compromised devices within the home network). An active attacker (Section 3.3 of [RFC3552]) can spoof the DHCP/RA response to provide the attacker's DoT/DoH server. Note that such an attacker can launch other attacks as discussed in Section 22 of [RFC8415]. Furthermore, if the browser or the OS is pre-configured with a list of DNS servers and some of which perform malware filtering while others do not, an attacker can prevent contacting the preferred filtering DNS servers causing a downgrade attack to a non-filtering DNS server, which the attacker can leverage to deliver malware.

In this specification, DoH/DoT servers discovered using insecure discovery mechanisms (like DHCP/RA) are only used if that DoH/DoT server is pre-configured in the OS or the browser. Section 6.1 identifies a set of deployment options under which DHCP/RA RI options can be used. If the insecurely discovered DoH/DoT server is not pre-configured in the OS or browser, the client must validate the signatory (e.g., cryptographically attested by the ISP [I-D.reddy-add-server-policy-selection]). If the DHCP/RA response is dropped by the attacker, the client can fallback to use a pre-configured DoH/DoT server. However, the use of policies to select servers is out of scope of this document.

The use of DoH/DoT also depends on the user's policies. For example, the user may indicate his/her consent to use (or not) the locally-discovered DoH/DoT server or request to review human-readable privacy policy information of a selected DNS server and to assess whether that DNS server performs DNS-based content filtering (e.g., [I-D.reddy-add-server-policy-selection]). The DNS client is assumed to adhere to these policies. This document does not make any assumption about the structure of such policies nor mandates specific requirements. Such policies and their handling is out of scope.

DoT/DoH sessions with rogue servers spoofing the IP address of a DNS server will fail because the DNS client will fail to authenticate that rogue server based upon PKIX authentication [RFC6125] based upon the authentication domain name in the Reference Identifier Option. DNS clients that ignore authentication failures and accept spoofed

certificates will be subject to attacks (e.g., redirect to malicious servers, intercept sensitive data).

A passive attacker (Section 3.2 of [RFC3552]) can identify the host is using DHCP/RA to discover the DoH/DoT server and can infer the host is capable of using DoH/DoT to encrypt DNS messages. However, a passive attacker cannot spoof or modify DHCP/RA messages.

Attacks of spoofed or modified DHCP responses and RA messages by attackers in the home network can possibly be mitigated by making use of the mechanisms described in [RFC7610], [RFC7113], and [RFC7513].

TCP connections received outside the home network MUST be discarded by the DoH/DoT forwarder in the CPE. This behavior adheres to REQ#8 in [RFC6092]; it MUST apply for both IPv4 and IPv6.

10. IANA Considerations

10.1. DHCPv6 Option

IANA is requested to assign the following new DHCPv6 Option Code in the registry maintained in: https://www.iana.org/assignments/dhcpv6-parameters/dhcpv6-parameters.xhtml#dhcpv6-parameters-2.

Value	Description	Client ORO	Singleton Option	Reference
TBA1	OPTION_V6_DNS_RI	Yes	Yes	[ThisDocument]

10.2. DHCP Option

IANA is requested to assign the following new DHCP Option Code in the registry maintained in: https://www.iana.org/assignments/bootp-dhcp-parameters/bootp-dhcp-parameters.xhtml#options.

Tag	Name	Data Length	Meaning	Reference
TBA2	OPTION_V4_DNS_RI	N	DoT/DoH server authentication domain name	[ThisDocument]

10.3. RA Option

IANA is requested to assign the following new IPv6 Neighbor Discovery Option type in the "IPv6 Neighbor Discovery Option Formats" sub-registry under the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" registry maintained in http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xhtml#icmpv6-parameters-5.

Type	Description	Reference
TBA3	DNS Reference Identifier Option	[ThisDocument]

11. Acknowledgements

Many thanks to Christian Jacquenet for the review.

Thanks to Tommy Jensen, Stephen Farrell, Martin Thomson, Vittorio Bertola, and Iain Sharp for the comments.

Thanks to Mark Nottingham for the feedback on HTTP redirection.

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <https://www.rfc-editor.org/info/rfc2132>.

[RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", RFC 3396, DOI 10.17487/RFC3396, November 2002, <https://www.rfc-editor.org/info/rfc3396>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <https://www.rfc-editor.org/info/rfc8106>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

12.2. Informative References

[Auto-upgrade] The Unicode Consortium, "DoH providers: criteria, process for Chrome", <docs.google.com/document/d/128i2YTV2C7T6Gr3I-81zlQ-_Lprnsp24qzy_20Z1Psw/edit>.

[DPP] The Wi-Fi Alliance, "Device Provisioning Protocol Specification", <https://www.wi-fi.org/file/device-provisioning-protocol-specification>.

[Dragonblood]
The Unicode Consortium, "Dragonblood: Analyzing the Dragonfly Handshake of WPA3 and EAP-pwd", <https://papers.mathyvanhoef.com/dragonblood.pdf>.

[Evil-Twin]
The Unicode Consortium, "Evil twin (wireless networks)", <https://en.wikipedia.org/wiki/Evil_twin_(wireless_networks)>.

[I-D.btw-add-rfc8484-clarification]
Boucadair, M., Cook, N., Reddy.K, T., and D. Wing, "Supporting Redirection for DNS Queries over HTTPS (DoH)", draft-btw-add-rfc8484-clarification-02 (work in progress), July 2020.

[I-D.ietf-dnsop-terminology-ter]
Hoffman, P., "Terminology for DNS Transports and Location", draft-ietf-dnsop-terminology-ter-01 (work in progress), February 2020.

[I-D.ietf-dnssd-push]
Pusateri, T. and S. Cheshire, "DNS Push Notifications", draft-ietf-dnssd-push-25 (work in progress), October 2019.

[I-D.ietf-v6ops-rfc7084-bis]
Palet, J., "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-rfc7084-bis-04 (work in progress), June 2017.

[I-D.pp-add-resinfo]
Sood, P. and P. Hoffman, "DNS Resolver Information Self-publication", draft-pp-add-resinfo-02 (work in progress), June 2020.

[I-D.reddy-add-iot-byod-bootstrap]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "A Bootstrapping Procedure to Discover and Authenticate DNS-over-TLS and DNS-over-HTTPS Servers for IoT and BYOD Devices", draft-reddy-add-iot-byod-bootstrap-00 (work in progress), May 2020.

[I-D.reddy-add-server-policy-selection]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "DNS Server Selection: DNS Server Information with Assertion Token", draft-reddy-add-server-policy-selection-03 (work in progress), June 2020.

[Krack] The Unicode Consortium, "Key Reinstallation Attacks", 2017, <https://www.krackattacks.com/>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/info/rfc3552>.

[RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <https://www.rfc-editor.org/info/rfc3646>.

[RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <https://www.rfc-editor.org/info/rfc6092>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <https://www.rfc-editor.org/info/rfc6125>.

[RFC6731] Savolainen, T., Kato, J., and T. Lemon, "Improved Recursive DNS Server Selection for Multi-Interfaced Nodes", RFC 6731, DOI 10.17487/RFC6731, December 2012, <https://www.rfc-editor.org/info/rfc6731>.

[RFC7113] Gont, F., "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", RFC 7113, DOI 10.17487/RFC7113, February 2014, <https://www.rfc-editor.org/info/rfc7113>.

[RFC7250] Wouters, P., Ed., Tschofenig, H., Ed., Gilmore, J., Weiler, S., and T. Kivinen, "Using Raw Public Keys in Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", RFC 7250, DOI 10.17487/RFC7250, June 2014, <https://www.rfc-editor.org/info/rfc7250>.

[RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <https://www.rfc-editor.org/info/rfc7513>.

[RFC7610] Gont, F., Liu, W., and G. Van de Velde, "DHCPv6-Shield: Protecting against Rogue DHCPv6 Servers", BCP 199, RFC 7610, DOI 10.17487/RFC7610, August 2015, <https://www.rfc-editor.org/info/rfc7610>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <https://www.rfc-editor.org/info/rfc8520>.

[RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <https://www.rfc-editor.org/info/rfc8555>.

[TR-069] The Broadband Forum, "CPE WAN Management Protocol", December 2018, <https://www.broadband-forum.org/technical/download/TR-069.pdf>.

[TS.24008] 3GPP, "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3 (Release 16)", December 2019, <http://www.3gpp.org/DynaReport/24008.htm>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: TirumaleswarReddy_Konda@McAfee.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

Neil Cook
Open-Xchange
UK

Email: neil.cook@noware.co.uk

ADD
Internet-Draft
Intended status: Standards Track
Expires: July 26, 2021

M. Boucadair
Orange
T. Reddy
McAfee
D. Wing
Citrix
N. Cook
Open-Xchange
T. Jensen
Microsoft

January 22, 2021

DHCP and Router Advertisement Options for Encrypted DNS Discovery
draft-btw-add-home-12

Abstract

The document specifies new DHCP and IPv6 Router Advertisement options to discover encrypted DNS servers (e.g., DNS-over-HTTPS, DNS-over-TLS, DNS-over-QUIC). Particularly, it allows to learn an authentication domain name together with a list of IP addresses and a port number to reach such encrypted DNS servers. The discovery of DNS-over-HTTPS URI Templates is also discussed.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 26, 2021.

Table of Contents

1. Introduction

This document focuses on the support of encrypted DNS such as DNS-over-HTTPS (DoH) [RFC8484], DNS-over-TLS (DoT) [RFC7858], or DNS-over-QUIC (DoQ) [I-D.ietf-dprive-dnsoquic] in local networks.

In particular, the document specifies how a local encrypted DNS server can be discovered and used by connected hosts by means of DHCP [RFC2132], DHCPv6 [RFC8415], and IPv6 Router Advertisement (RA) [RFC4861] options. These options are designed to convey the following information: the DNS Authentication Domain Name (ADN), a list of IP addresses, and optionally a port number. The discovery of DoH URI Templates is discussed in Section 7.

Sample target deployment scenarios are discussed in Appendix A; both managed and unmanaged Customer Premises Equipment (CPEs) are covered. It is out of the scope of this document to provide an exhaustive inventory of deployments where Encrypted DNS Options (Sections 4, 5, and 6) can be used.

Considerations related to hosting a DNS forwarder in a local network are described in Section 8.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499]. The following additional terms are used:

Do53: refers to unencrypted DNS.

Encrypted DNS: refers to a scheme where DNS exchanges are transported over an encrypted channel. Examples of encrypted DNS

are DNS-over-TLS (DoT) [RFC7858], DNS-over-HTTPS (DoH) [RFC8484], or DNS-over-QUIC (DoQ) [I-D.ietf-dprive-dnsoquic].

Managed CPE: refers to a CPE that is managed by an Internet Service Providers (ISP).

Unmanaged CPE: refers to a CPE that is not managed by an ISP.

DHCP: refers to both DHCPv4 and DHCPv6.

3. Overview and Rationale

This document describes how a DNS client can discover a local encrypted DNS server(s) using DHCP (Sections 4 and 5) and Neighbor Discovery protocol (Section 6).

As reported in Section 1.7.2 of [RFC6125]:

> Some certification authorities issue server certificates based on IP addresses, but preliminary evidence indicates that such certificates are a very small percentage (less than 1%) of issued certificates.

In order to allow for PKIX-based authentication between a DNS client and an encrypted DNS server while accommodating the current best practices for issuing certificates, this document allows for configuring an authentication domain name to be presented as a reference identifier for DNS authentication purposes.

To avoid adding a dependency on another server to resolve the ADN, the options return a list of IP addresses to locate the encrypted DNS server. In the various scenarios sketched in Appendix A, encrypted DNS servers may terminate on the same IP address or distinct IP addresses. Terminating encrypted DNS servers on the same or distinct IP addresses is deployment specific. It is RECOMMENDED to return both the ADN and a list of IP addresses to a requesting host.

Note that in order to optimize the size of discovery messages when all servers terminate on the same IP address, a host may rely upon the discovery mechanisms specified in [RFC2132][RFC3646][RFC8106] to retrieve a list of IP addresses to reach their DNS servers. Nevertheless, this approach requires a client that supports more than one encrypted DNS to probe that list of IP addresses. To avoid such probing, the options defined in the following sections associate an IP address with an encrypted DNS type. No probing is required in such design.

A list of IP addresses to reach an encrypted DNS server can be returned in the option to accommodate current deployments relying upon primary and backup servers. Whether one IP address or more are returned in an option is deployment specific. For example, a router embedding a recursive server or forwarder has to include one single IP address pointing to one of its LAN-facing interfaces. This address can be a private IPv4 address, a link-local address, a Unique Local IPv6 unicast Address (ULA), or a Global Unicast Address (GUA).

If more than one IP address are to be returned in an Encrypted DNS server option, these addresses are ordered in the preference for use by the client.

Because DoT and DoQ may make use of customized port numbers instead of default ones, the Encrypted DNS server options are designed to return alternate port numbers.

Some ISPs rely upon external resolvers (e.g., outsourced service or public resolvers); these ISPs provide their customers with the IP addresses of these resolvers. These addresses are typically configured on CPEs using dedicated management tools. Likewise, users can modify the default DNS configuration of their CPEs (e.g., supplied by their ISP) to configure their favorite DNS servers. This document permits such deployments.

If the encrypted DNS is discovered by a host using both RA and DHCP, the rules discussed in Section 5.3.1 of [RFC8106] MUST be followed.

The DNS client establishes an encrypted DNS session with the discovered DNS IP address(es) and port number, and uses the mechanism discussed in Section 8 of [RFC8310] to authenticate the DNS server certificate using the authentication domain name conveyed in the encrypted DNS options.

Devices may be connected to multiple networks; each providing their own DNS configuration using the discovery mechanisms specified in this document. Nevertheless, it is out of the scope of this specification to discuss DNS selection of multi-interface devices. The reader may refer to [RFC6731] for a discussion of issues and an example of DNS server selection for multi-interfaced devices.

DHCP/RA options to discover encrypted DNS servers (including, DoH URI Templates should the WG pursue that approach pending feedback) takes precedence over DEER [I-D.pauly-add-deer] since DEER uses unencrypted DNS to an external DNS resolver, which is susceptible to both internal and external attacks whereas DHCP/RA is only vulnerable to internal attacks.

4. DHCPv6 Encrypted DNS Options

This section defines two DHCPv6 options: DHCPv6 Encrypted DNS ADN option (Section 4.1) and DHCPv6 Encrypted DNS Address option (Section 4.2).

4.1. Encrypted DNS ADN Option

The DHCPv6 Encrypted DNS ADN option is used to configure an authentication domain name of the encrypted DNS server. The format of this option is shown in Figure 1.

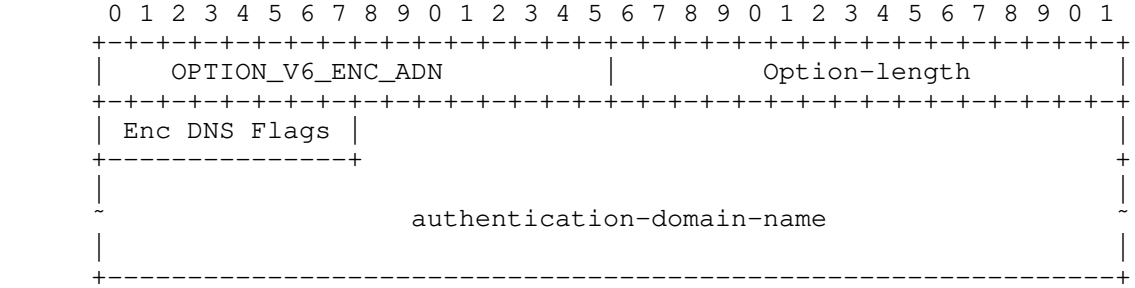

Figure 1: DHCPv6 Encrypted DNS ADN Option

The fields of the option shown in Figure 1 are as follows:

Option-code: OPTION_V6_ENC_ADN (TBA1, see Section 11.2)

Option-length: Length of the enclosed data in octets.

Enc DNS Flags (Encrypted DNS Flags): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this 8-bit field is shown in Figure 2.

Figure 2: Encrypted DNS Flags Field

T: If set, this bit indicates that the server supports DoT [RFC7858].

H: If set, this bit indicates that the server supports DoH [RFC8484].

Q: If set, this bit indicates that the server supports DoQ [I-D.ietf-dprive-dnsoquic].

U: Unassigned bits. These bits MUST be unset by the sender. Associating a meaning with an unassigned bit can be done as per Section 11.1.

In a request, these bits are assigned to indicate the requested encrypted DNS server type(s) by the client. In a response, these bits are set as a function of the encrypted DNS supported by the server and the requested encrypted DNS server type(s).

To keep the packet small, if more than one encrypted DNS type (e.g., both DoH and DoT) are to be returned to a requesting client and the same ADN is used for these types, the corresponding bits must be set in the 'Encrypted DNS Types' field of the same option instance in a response. For example, if the client requested DoH and DoT and the server supports both with the same ADN, then both T and H bits must be set.

authentication-domain-name: A fully qualified domain name of the encrypted DNS server. This field is formatted as specified in Section 10 of [RFC8415].

An example of the authentication-domain-name encoding is shown in Figure 3. This example conveys the FQDN "doh1.example.com.", and the resulting Option-length field is 18.

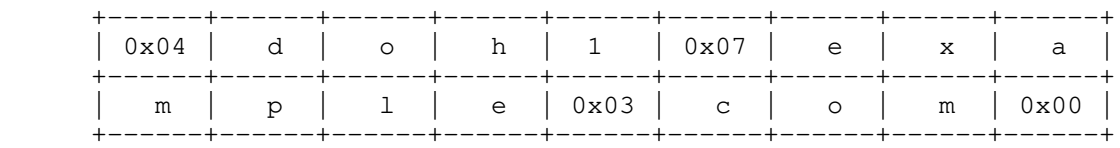

Figure 3: An Example of the DNS authentication-domain-name Encoding

4.2. Encrypted DNS Address Option

The DHCPv6 Encrypted DNS Address option is used to configure a list of IP addresses and a port number of the encrypted DNS server. The format of this option is shown in Figure 4.

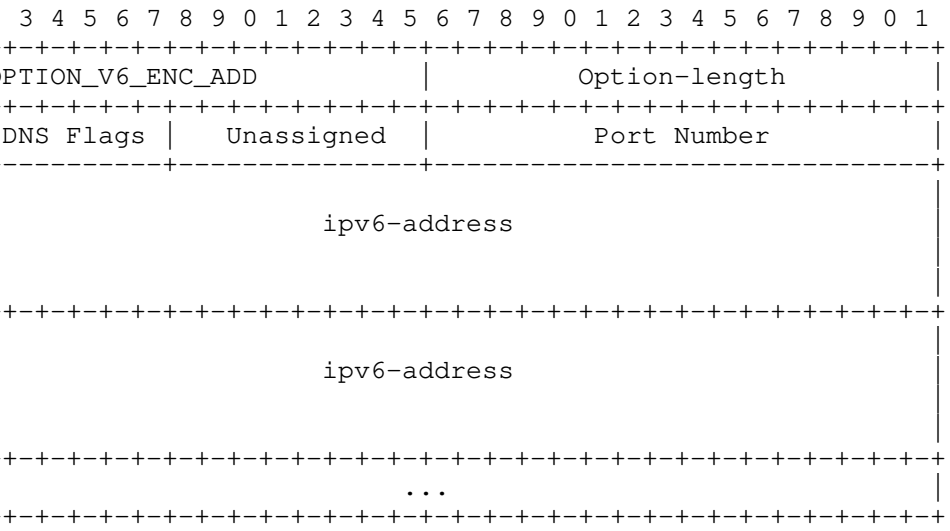

Figure 4: DHCPv6 Encrypted DNS Address Option

The fields of the option shown in Figure 4 are as follows:

Option-code: OPTION_V6_ENC_ADD (TBA2, see Section 11.2)

Option-length: Length of the enclosed data in octets.

Enc DNS Flags (Encrypted DNS Flags): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this 8-bit field is shown in Figure 2. In a request, these bits are set to indicate the requested encrypted DNS server type(s) by the client. In a response, these bits are set as a function of the encrypted DNS supported by the server and the requested encrypted DNS server type(s).

Unassigned: These bits MUST be unset by the sender. Associating a meaning with an unassigned bit can be done via Standards Action [RFC8126].

Port Number: If not null, it indicates the port number to be used for the encrypted DNS. If this field is set to zero, this indicates that default port numbers should be used. As a reminder, the default port number is 853 for DoT and 443 for DoH.

ipv6-address(es): Indicates one or more IPv6 addresses to reach the encrypted DNS server. An address can be link-local, ULA, or GUA.

Multiple instances of OPTION_V6_ENC_ADN (or OPTION_V6_ENC_ADD) may be returned to a DHCPv6 client; each pointing to a distinct encrypted DNS server type.

If more than one encrypted DNS server types is supported on the same IP address and default port numbers are used, one instance of OPTION_V6_ENC_ADD option with the appropriate bits set in "Encr DNS Types" field should be returned by the DHCP server.

4.3. DHCPv6 Client Behavior

To discover an encrypted DNS server, the DHCPv6 client MUST include OPTION_V6_ENC_ADN and OPTION_V6_ENC_ADD in an Option Request Option (ORO), as in Sections 18.2.1, 18.2.2, 18.2.4, 18.2.5, 18.2.6, and 21.7 of [RFC8415]. The DHCPv6 client sets the Encrypted DNS Types field to the requested encrypted DNS server type(s).

If the DHCPv6 client requested more than one encrypted DNS server type, the DHCP client MUST be prepared to receive multiple OPTION_V6_ENC_ADN (or OPTION_V6_ENC_ADD) options; each option is to be treated as a separate encrypted DNS server.

The DHCPv6 client MUST silently discard multicast and host loopback addresses conveyed in OPTION_V6_ENC_ADD.

5. DHCPv4 Encrypted DNS Option

5.1. Encrypted DNS Option

The DHCPv4 Encrypted DNS option is used to configure an authentication domain name, a list of IP addresses, and a port number of the encrypted DNS server. The format of this option is illustrated in Figure 5.

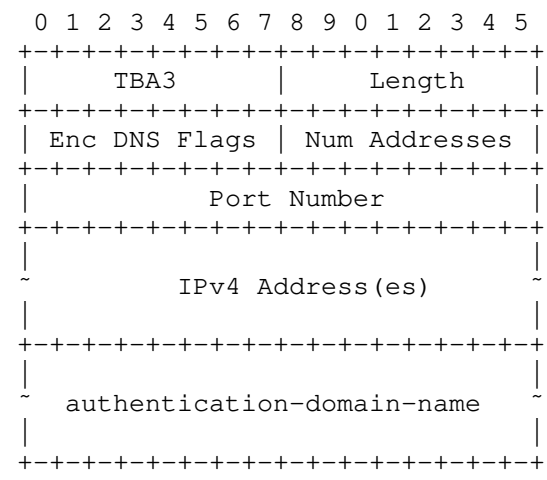

Figure 5: DHCPv4 Encrypted DNS Option

The fields of the option shown in Figure 5 are as follows:

Code: OPTION_V4_ENC_DNS (TBA3, see Section 11.3).

Length: Length of the enclosed data in octets.

Enc DNS Flags (Encrypted DNS Flags): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this field is shown in Figure 2.

Num Addresses: Indicates the number of included IPv4 addresses.

Port Number: If not null, it indicates the port number to be used for the encrypted DNS. A null value indicates that default port numbers are used. As a reminder, the default port number is 853 for DoT and 443 for DoH.

IPv4 Address(es): Indicates one or more IPv4 addresses to reach the encrypted DNS server. Both private and public IPv4 addresses can be included in this field. The format of this field is shown in Figure 6. This format assumes that an IPv4 address is encoded as a1.a2.a3.a4.

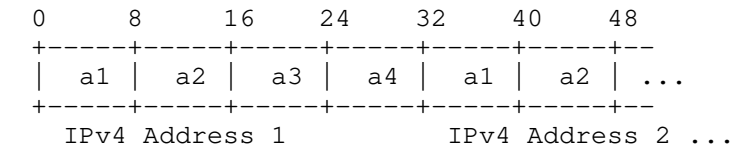

Figure 6: Format of the IPv4 Addresses Field

authentication-domain-name: The domain name of the encrypted DNS server. This field is formatted as specified in Section 10 of [RFC8415]. The format of this field is shown in Figure 7. The values s1, s2, s3, etc. represent the domain name labels in the domain name encoding.

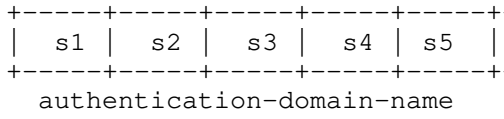

Figure 7: Format of the Authentication Domain Name Field

OPTION_V4_ENC_DNS is a concatenation-requiring option. As such, the mechanism specified in [RFC3396] MUST be used if OPTION_V4_ENC_DNS exceeds the maximum DHCPv4 option size of 255 octets.

5.2. DHCPv4 Client Behavior

To discover an encrypted DNS server, the DHCPv4 client requests the Encrypted DNS server by including OPTION_V4_ENC_DNS in a Parameter Request List option [RFC2132]. The DHCPv4 client sets the Encrypted DNS Types field to the requested encrypted DNS server.

If the DHCPv4 client requested more than one encrypted DNS server type, the DHCPv4 client MUST be prepared to receive multiple DHCP OPTION_V4_ENC_DNS options; each option is to be treated as a separate encrypted DNS server.

The DHCPv4 client MUST silently discard multicast and host loopback addresses conveyed in OPTION_V4_ENC_DNS.

6. IPv6 RA Encrypted DNS Options

This section defines two Neighbor Discovery [RFC4861]: IPv6 Router Advertisement (RA) Encrypted DNS ADN option (Section 6.1) and IPv6 RA Encrypted DNS Address option (Section 6.2). These options are useful in contexts similar to those discussed in Section 1.1 of [RFC8106].

6.1. Encrypted DNS ADN Option

The IPv6 RA Encrypted DNS ADN option is used to configure an authentication domain name of the encrypted DNS server. The format of this option is illustrated in Figure 8.

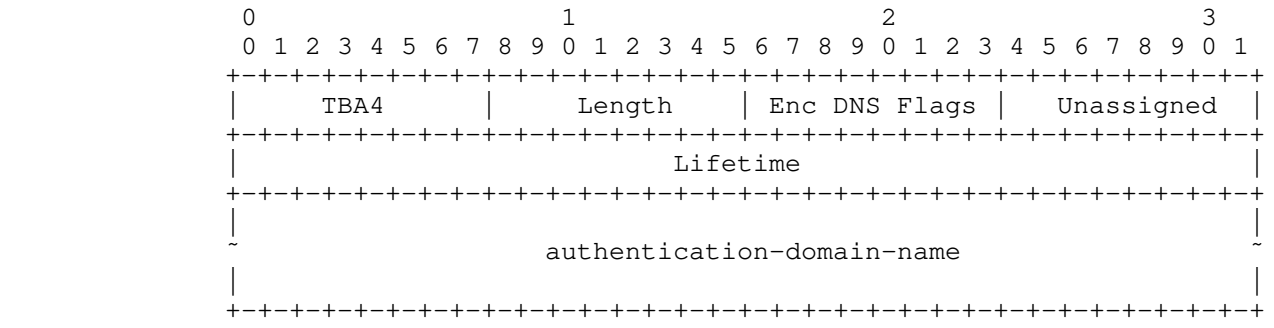

Figure 8: RA Encrypted DNS ADN Option

The fields of the option shown in Figure 8 are as follows:

Type: 8-bit identifier of the Encrypted DNS Option as assigned by IANA (TBA4, see Section 11.4).

Length: 8-bit unsigned integer. The length of the option (including the Type and Length fields) is in units of 8 octets.

Enc DNS Flags (Encrypted DNS Flags): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this field is shown in Figure 2.

Unassigned: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Lifetime: 32-bit unsigned integer. The maximum time in seconds (relative to the time the packet is received) over which the discovered Authentication Domain Name is valid.

The value of Lifetime SHOULD by default be at least 3 * MaxRtrAdvInterval, where MaxRtrAdvInterval is the maximum RA interval as defined in [RFC4861].

A value of all one bits (0xffffffff) represents infinity.

A value of zero means that this Authentication Domain Name MUST no longer be used.

authentication Domain Name: The domain name of the encrypted DNS server. This field is formatted as specified in Section 10 of [RFC8415].

This field MUST be padded with zeros so that its size is a multiple of 8 octets.

6.2. Encrypted DNS Address Option

The IPv6 RA Encrypted DNS Address option is used to configure a port number and a list of IPv6 addresses of the encrypted DNS server. The format of this option is illustrated in Figure 9. All of the addresses share the same Lifetime value. Similar to [RFC8106], if it is desirable to have different Lifetime values per IP address, multiple Encrypted DNS Address options may be used.

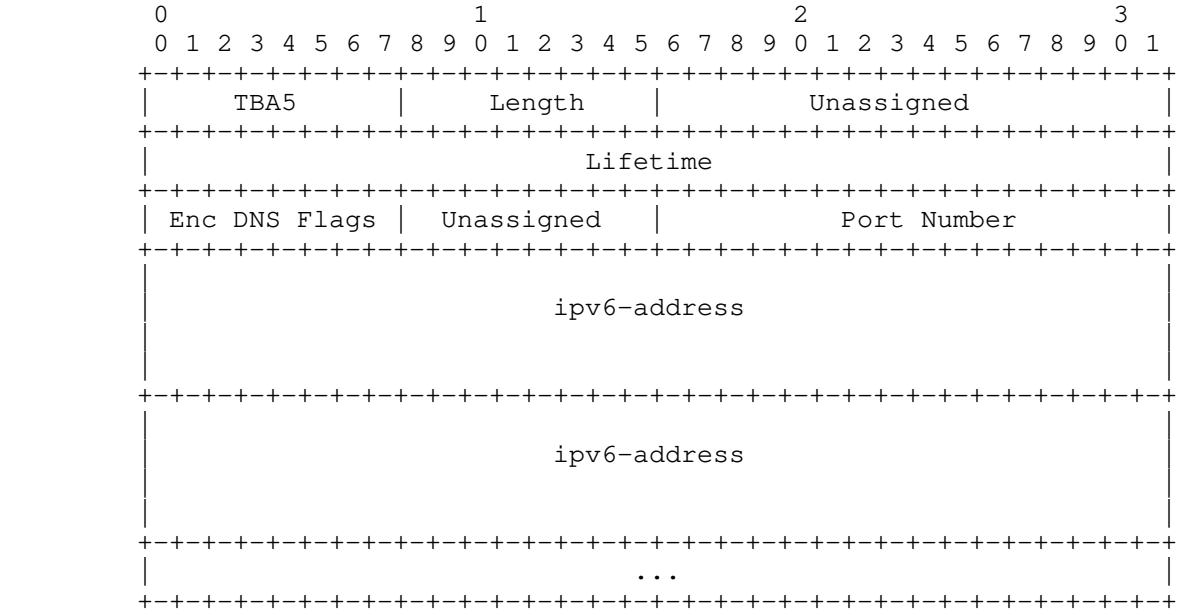

Figure 9: RA Encrypted DNS Address Option

The fields of the RA Encrypted DNS Address option shown in Figure 9 are as follows:

Type: 8-bit identifier of the Encrypted DNS Address Option as assigned by IANA (TBA5, see Section 11.4).

Length: 8-bit unsigned integer. The length of the option (including the Type and Length fields) is in units of 8 octets.

Unassigned: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Lifetime: 32-bit unsigned integer. The maximum time in seconds (relative to the time the packet is received) over which the discovered encrypted DNS IPv6 addresses are valid.

The value of Lifetime SHOULD by default be at least 3 * MaxRtrAdvInterval, where MaxRtrAdvInterval is the maximum RA interval as defined in [RFC4861].

A value of all one bits (0xffffffff) represents infinity.

A value of zero means that these IPv6 addresses MUST no longer be used.

Enc DNS Flags (Encrypted DNS Flags): Indicates the type(s) of the encrypted DNS server conveyed in this attribute. The format of this field is shown in Figure 2.

Port Number: If not null, it indicates the port number to be used for the encrypted DNS. A null value indicates that default port numbers must be used. As a reminder, the default port number is 853 for DoT and 443 for DoH.

ipv6-address(es): One or more IPv6 addresses of the encrypted DNS server. An address can be link-local, ULA, or GUA.

7. DoH URI Templates

DoH servers may support more than one URI Template [RFC8484]. Also, if the resolver hosts several DoH services (e.g., no-filtering, blocking adult content, blocking malware), these services can be discovered as templates. The following discusses a mechanism for a DoH client to retrieve the list of supported templates by a DoH server.

Upon discovery of a DoH resolver (Sections 4, 5, and 6), the DoH client may contact that DoH resolver to retrieve the list of supported DoH services using DEER [I-D.pauly-add-deer]. This will allow the client to discover the resolver's supported DoH templates or DoH resolvers that the discovered resolver designates using DNS SVCB queries [I-D.schwartz-svcb-dns]. The designated DoH resolvers and DoH resolver discovered using DHCP/RA may be hosted on the same or distinct IP addresses.

Let's suppose that a host has discovered an encrypted DNS server that is DoH-capable. The host has also discovered the following information:

- ADN: doh.example.com
- Locator: 2001:db8:1::1

The client will use DEER [I-D.pauly-add-deer] to discover the DoH templates supported by the DNS server at the Locator (2001:db8:1::1). In addition to the checks included in DEER, clients should verify the ADN (doh.example.com) is valid for the certificate provided by the DoH resolver. However, the IP address of the DEER-discovered resolver may differ from the Locator field value. This will allow the ISP to offer different DoH services to the endpoints attached to local networks.

Alternatively, dedicated DHCP/RA options may be defined to convey an URI template in order to avoid additional network traffic to bootstrap DoH configuration. An example of the format of such an option is depicted in Figure 10.

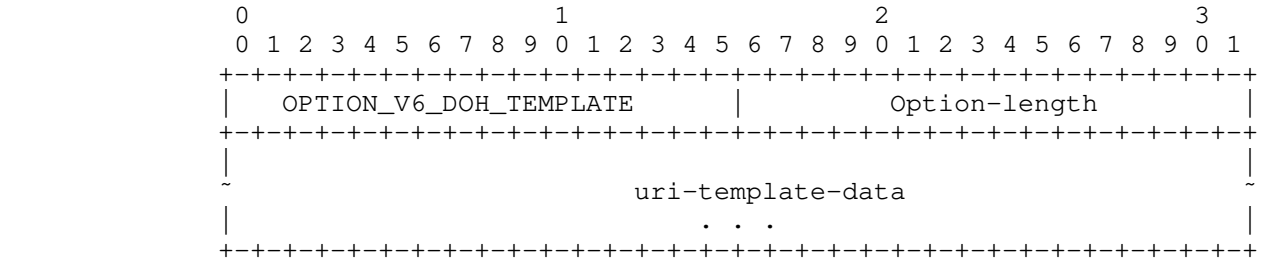

Each instance of the uri-template-data is formatted as follows:

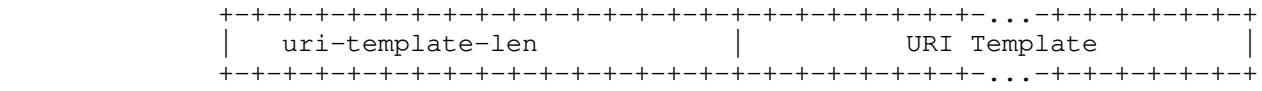

Figure 10: Example of a DHCPv6 URI Template Option

Note: More feedback from the WG is needed to decide which approach(es) to follow.

How a DoH client makes use of the configured DoH services is out of the scope of this document.

8. Hosting Encrypted DNS Forwarder in Local Networks

This section discusses some deployment considerations (not recommendations) to host an encrypted DNS forwarder within a local network.

8.1. Managed CPEs

The section discusses mechanisms that can be used to host an encrypted DNS forwarder in a managed CPE (Appendix A.1).

8.1.1. DNS Forwarders

The managed CPE should support a configuration parameter to instruct the CPE whether it has to relay the encrypted DNS server received from the ISP's network or has to announce itself as a forwarder within the local network. The default behavior of the CPE is to supply the encrypted DNS server received from the ISP's network.

8.1.2. ACME

The ISP can assign a unique FQDN (e.g., "cpe1.example.com") and a domain-validated public certificate to the encrypted DNS forwarder hosted on the CPE. Automatic Certificate Management Environment (ACME) [RFC8555] can be used by the ISP to automate certificate management functions such as domain validation procedure, certificate issuance and certificate revocation.

8.1.3. Auto-Upgrade Based on Domains and their Subdomains

If the ADN conveyed in DHCP/RA (Sections 4, 5, and 6) is preconfigured in popular OSes or browsers as a verified resolver and the auto-upgrade (Appendix B) is allowed for both the preconfigured ADN and its sub-domains, the encrypted DNS client will learn the local encrypted DNS forwarder using DHCP/RA and auto-upgrade because the preconfigured ADN would match the subjectAltName value in the server certificate. For example, if the preconfigured ADN is "*.example.com" and the discovered encrypted DNS forwarder is "cpe1.example.com", auto-upgrade will take place.

In this case, the CPE can communicate the ADN of the local DoH forwarder (Section 8.1.2) to internal hosts using DHCP/RA (Sections 4, 5, and 6).

Let's suppose that "*.example.net" is preconfigured as a verified resolved in the browser or OS. If the encrypted DNS client discovers a local forwarder "cpe1-internal.example.net", the encrypted DNS client will auto-upgrade because the preconfigured ADN would match

subjectAltName value "cpe1-internal.example.net" of type dNSName. As shown in Figure 11, the auto-upgrade to a rogue server advertising "rs.example.org" will fail because it does not match "*.example.net".

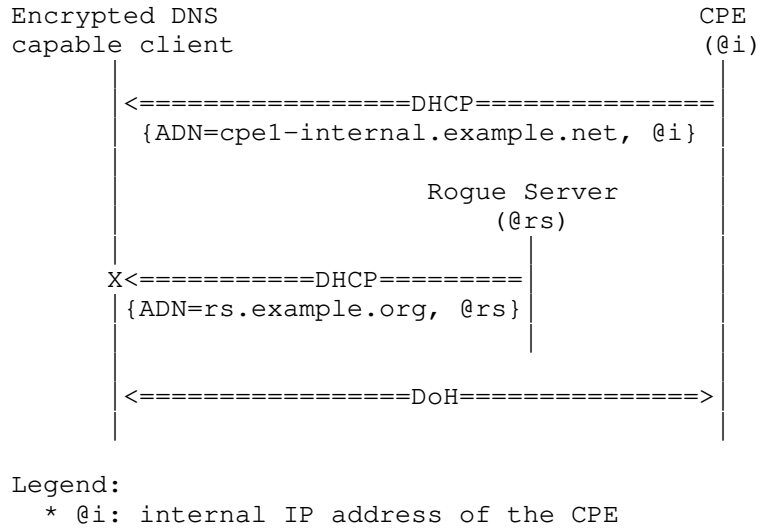

Figure 11: A Simplified Example of Auto-upgrade based on Subdomains

8.2. Unmanaged CPEs

The approach specified in Section 8.1 does not apply for hosting a DNS forwarder in an unmanaged CPE.

The unmanaged CPE administrator (referred to as administrator) can host an encrypted DNS forwarder on the unmanaged CPE. This assumes the following:

- The encrypted DNS server certificate is managed by the entity in-charge of hosting the encrypted DNS forwarder.

  Alternatively, a security service provider can assign a unique FQDN to the CPE. The encrypted DNS forwarder will act like a private encrypted DNS server only be accessible from within the the local network.

- The encrypted DNS forwarder will either be configured to use the ISP's or a 3rd party encrypted DNS server.

- The unmanaged CPE will advertise the encrypted DNS forwarder ADN using DHCP/RA to internal hosts.

Figure 12 illustrates an example of an unmanaged CPE hosting a forwarder which connects to a 3rd party encrypted DNS server. In this example, the DNS information received from the managed CPE (and therefore from the ISP) is ignored by the Internal CPE hosting the forwarder.

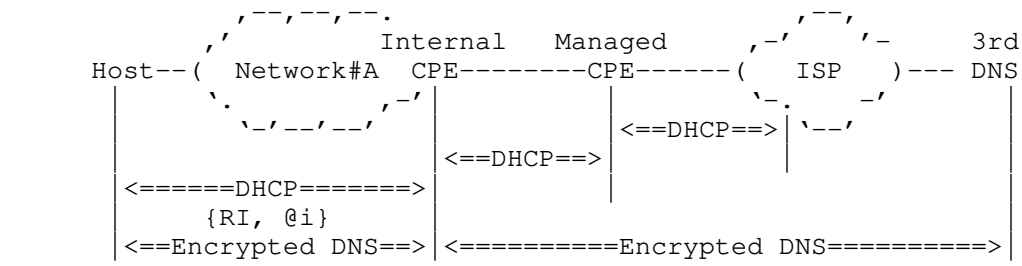

Legend:
* @i: IP address of the DNS forwarder hosted in the Internal CPE.

Figure 12: Example of an Internal CPE Hosting a Forwarder

9. Legacy CPEs

Hosts serviced by legacy CPEs that can't be upgraded to support the options defined in Sections 4, 5, and 6 won't be able to learn the encrypted DNS server hosted by the ISP, in particular. If the ADN is not discovered using DHCP/RA, such hosts will have to fallback to use DEER as defined in [I-D.pauly-add-deer] to discover the encrypted DNS server and to retrieve the list of supported DoH services using the SVCB RRtype [I-D.schwartz-svcb-dns] without verifying the hostname of discovered templates with the ADN. Other guidance in DEER relating to resolver verification must be followed in this case. This will prevent an unencrypted resolver on a local address from referring to an encrypted resolver at a different address without an out-of-band configuration in the client beyond the scope of this document or DEER.

10. Security Considerations

10.1. Spoofing Attacks

DHCP/RA messages are not encrypted or protected against modification within the LAN. Unless mitigated (described below), the content of DHCP and RA messages can be spoofed or modified by active attackers, such as compromised devices within the local network. An active attacker (Section 3.3 of [RFC3552]) can spoof the DHCP/RA response to provide the attacker's Encrypted DNS server. Note that such an

attacker can launch other attacks as discussed in Section 22 of [RFC8415]. The attacker can get a domain name with a domain-validated public certificate from a CA and host an Encrypted DNS server. Also, an attacker can use a public IP address and get an 'IP address'-validated public certificate from a CA to host an Encrypted DNS server.

Attacks of spoofed or modified DHCP responses and RA messages by attackers within the local network may be mitigated by making use of the following mechanisms:

- DHCPv6-Shield described in [RFC7610], the CPEs discards DHCP response messages received from any local endpoint.
- RA-Guard described in [RFC7113], the CPE discards RAs messages received from any local endpoint.
- Source Address Validation Improvement (SAVI) solution for DHCP described in [RFC7513], the CPE filters packets with forged source IP addresses.

Encrypted DNS sessions with rogue servers that spoof the IP address of a DNS server will fail because the DNS client will fail to authenticate that rogue server based upon PKIX authentication [RFC6125], particularly the authentication domain name in the Encrypted DNS Option. DNS clients that ignore authentication failures and accept spoofed certificates will be subject to attacks (e.g., redirect to malicious servers, intercept sensitive data).

Encrypted DNS connections received from outside the local network MUST be discarded by the encrypted DNS forwarder in the CPE. This behavior adheres to REQ#8 in [RFC6092]; it MUST apply for both IPv4 and IPv6.

10.2. Deletion Attacks

If the DHCP responses or RAs are dropped by the attacker, the client can fallback to use a preconfigured encrypted DNS server. However, the use of policies to select servers is out of the scope of this document.

Note that deletion attack is not specific to DHCP/RA.

10.3. Passive Attacks

A passive attacker (Section 3.2 of [RFC3552]) can identify a host is using DHCP/RA to discover an encrypted DNS server and can infer that

host is capable of using DoH/DoT/DoQ to encrypt DNS messages. However, a passive attacker cannot spoof or modify DHCP/RA messages.

10.4. Wireless Security - Authentication Attacks

Wireless LAN (WLAN) as frequently deployed in local networks (e.g., home networks) is vulnerable to various attacks (e.g., [Evil-Twin], [Krack], [Dragonblood]). Because of these attacks, only cryptographically authenticated communications are trusted on WLANs. This means information provided by such networks via DHCP, DHCPv6, or RA (e.g., NTP server, DNS server, default domain) are untrusted because DHCP and RA are not authenticated.

If the pre-shared-key is the same for all clients that connect to the same WLAN, the shared key will be available to all nodes, including attackers, so it is possible to mount an active on-path attack. Man-in-the-middle attacks are possible within local networks because such WLAN authentication lacks peer entity authentication.

This leads to the need for provisioning unique credentials for different clients. Endpoints can be provisioned with unique credentials (username and password, typically) provided by the local network administrator to mutually authenticate to the local WLAN Access Point (e.g., 802.1x Wireless User Authentication on OpenWRT [dot1x], EAP-pwd [RFC8146]). Not all of endpoint devices (e.g., IoT devices) support 802.1x supplicant and need an alternate mechanism to connect to the local network. To address this limitation, unique pre-shared keys can be created for each such device and WPA-PSK is used (e.g., [PSK]).

11. IANA Considerations

11.1. Encrypted DNS Flag Bits

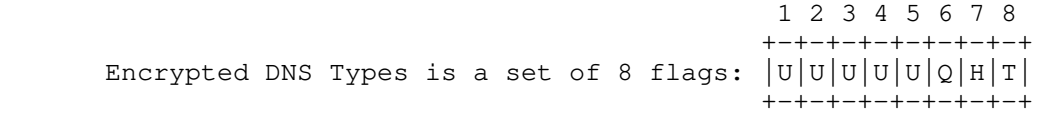

where flag bits in positions 1-5 are for future assignment as additional flag bits.

This document requests IANA to create a new registry called "Encrypted DNS Types". The initial values of the registry are as follows:

Bit Position	Label	Description	Reference
1	U	Unassigned	
2	U	Unassigned	
3	U	Unassigned	
4	U	Unassigned	
5	U	Unassigned	
6	Q	DNS-over-QUIC (DoQ)	[ThisDocument]
7	H	DNS-over-HTTP (DoH)	[ThisDocument]
8	T	DNS-over-TLS (DoT)	[ThisDocument]

New flag bits are assigned via Standards Action [RFC8126].

11.2. DHCPv6 Options

IANA is requested to assign the following new DHCPv6 Option Code in the registry maintained in [DHCPV6].

Value	Description	Client ORO	Singleton Option	Reference
TBA1	OPTION_V6_ENC_ADN	Yes	No	[ThisDocument]
TBA2	OPTION_V6_ENC_ADD	Yes	No	[ThisDocument]

11.3. DHCPv4 Option

IANA is requested to assign the following new DHCP Option Code in the registry maintained in [BOOTP].

Tag	Name	Data Length	Meaning	Reference
TBA3	OPTION_V4_ENC_DNS	N	Encrypted DNS Server	[ThisDocument]

11.4. Neighbor Discovery Options

IANA is requested to assign the following new IPv6 Neighbor Discovery Option type in the "IPv6 Neighbor Discovery Option Formats" sub-registry under the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" registry maintained in [ND].

Type	Description	Reference
TBA4	DNS Encrypted DNS ADN Option	[ThisDocument]
TBA5	DNS Encrypted DNS Address Option	[ThisDocument]

12. Acknowledgements

Many thanks to Christian Jacquenet and Michael Richardson for the review.

Thanks to Stephen Farrell, Martin Thomson, Vittorio Bertola, Stephane Bortzmeyer, Ben Schwartz, and Iain Sharp for the comments.

Thanks to Mark Nottingham for the feedback on HTTP redirection.

The use of DHCP to retrieve an authentication domain name was discussed in Section 7.3.1 of [RFC8310] and [I-D.pusateri-dhc-dns-driu].

13. Contributing Authors

Nicolai Leymann
Deutsche Telekom
Germany

Email: n.leymann@telekom.de

14. References

14.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <https://www.rfc-editor.org/info/rfc2132>.

[RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", RFC 3396, DOI 10.17487/RFC3396, November 2002, <https://www.rfc-editor.org/info/rfc3396>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <https://www.rfc-editor.org/info/rfc8106>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

14.2. Informative References

[Auto-upgrade] The Unicode Consortium, "DoH providers: criteria, process for Chrome", <docs.google.com/document/d/128i2YTV2C7T6Gr3I-81zlQ-_Lprnsp24qzy_20Z1Psw/edit>.

[BOOTP] "BOOTP Vendor Extensions and DHCP Options", <https://www.iana.org/assignments/bootp-dhcp-parameters/bootp-dhcp-parameters.xhtml#options>.

[DHCPV6] "DHCPv6 Option Codes", <https://www.iana.org/assignments/dhcpv6-parameters/dhcpv6-parameters.xhtml#dhcpv6-parameters-2>.

[dot1x] Cisco, "Basic 802.1x Wireless User Authentication", <https://openwrt.org/docs/guide-user/network/wifi/wireless.security.8021x>.

[Dragonblood] The Unicode Consortium, "Dragonblood: Analyzing the Dragonfly Handshake of WPA3 and EAP-pwd", <https://papers.mathyvanhoef.com/dragonblood.pdf>.

[Evil-Twin]
The Unicode Consortium, "Evil twin (wireless networks)", <https://en.wikipedia.org/wiki/Evil_twin_(wireless_networks)>.

[I-D.ietf-dprive-dnsoquic]
Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnsoquic-01 (work in progress), October 2020.

[I-D.ietf-v6ops-rfc7084-bis]
Palet, J., "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-rfc7084-bis-04 (work in progress), June 2017.

[I-D.pauly-add-deer]
Pauly, T., Kinnear, E., Wood, C., McManus, P., and T. Jensen, "Discovery of Equivalent Encrypted Resolvers", draft-pauly-add-deer-00 (work in progress), November 2020.

[I-D.pusateri-dhc-dns-driu]
Pusateri, T. and W. Toorop, "DHCPv6 Options for private DNS Discovery", draft-pusateri-dhc-dns-driu-00 (work in progress), July 2018.

[I-D.schwartz-svcb-dns]
Schwartz, B., "Service Binding Mapping for DNS Servers", draft-schwartz-svcb-dns-01 (work in progress), August 2020.

[Krack]
The Unicode Consortium, "Key Reinstallation Attacks", 2017, <https://www.krackattacks.com/>.

[ND]
"IPv6 Neighbor Discovery Option Formats", <http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xhtml#icmpv6-parameters-5>.

[PSK]
Cisco, "Identity PSK Feature Deployment Guide", <https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_Identity_PSK_Feature_Deployment_Guide.html>.

[RFC3552]
Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/info/rfc3552>.

[RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <https://www.rfc-editor.org/info/rfc3646>.

[RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <https://www.rfc-editor.org/info/rfc6092>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <https://www.rfc-editor.org/info/rfc6125>.

[RFC6731] Savolainen, T., Kato, J., and T. Lemon, "Improved Recursive DNS Server Selection for Multi-Interfaced Nodes", RFC 6731, DOI 10.17487/RFC6731, December 2012, <https://www.rfc-editor.org/info/rfc6731>.

[RFC7113] Gont, F., "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", RFC 7113, DOI 10.17487/RFC7113, February 2014, <https://www.rfc-editor.org/info/rfc7113>.

[RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <https://www.rfc-editor.org/info/rfc7513>.

[RFC7610] Gont, F., Liu, W., and G. Van de Velde, "DHCPv6-Shield: Protecting against Rogue DHCPv6 Servers", BCP 199, RFC 7610, DOI 10.17487/RFC7610, August 2015, <https://www.rfc-editor.org/info/rfc7610>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8146] Harkins, D., "Adding Support for Salted Password Databases to EAP-pwd", RFC 8146, DOI 10.17487/RFC8146, April 2017, <https://www.rfc-editor.org/info/rfc8146>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <https://www.rfc-editor.org/info/rfc8520>.

[RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <https://www.rfc-editor.org/info/rfc8555>.

[TR-069] The Broadband Forum, "CPE WAN Management Protocol", December 2018, <https://www.broadband-forum.org/technical/download/TR-069.pdf>.

[TS.24008] 3GPP, "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3 (Release 16)", December 2019, <http://www.3gpp.org/DynaReport/24008.htm>.

Appendix A. Sample Target Deployment Scenarios

Internet Service Providers (ISPs) traditionally provide DNS resolvers to their customers. To that aim, ISPs deploy the following mechanisms to advertise a list of DNS Recursive DNS server(s) to their customers:

- Protocol Configuration Options in cellular networks [TS.24008].
- DHCPv4 [RFC2132] (Domain Name Server Option) or DHCPv6 [RFC8415][RFC3646] (OPTION_DNS_SERVERS).
- IPv6 Router Advertisement [RFC4861][RFC8106] (Type 25 (Recursive DNS Server Option)).

The communication between a customer's device (possibly via Customer Premises Equipment (CPE)) and an ISP-supplied DNS resolver takes place by using cleartext DNS messages (Do53). Some examples are depicted in Figure 13. In the case of cellular networks, the cellular network will provide connectivity directly to a host (e.g., smartphone, tablet) or via a CPE. Do53 mechanisms used within the Local Area Network (LAN) are similar in both fixed and cellular CPE-based broadband service offerings.

(a) Fixed Networks

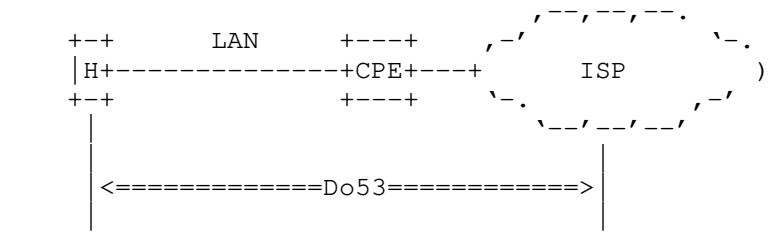

(b) Cellular Networks

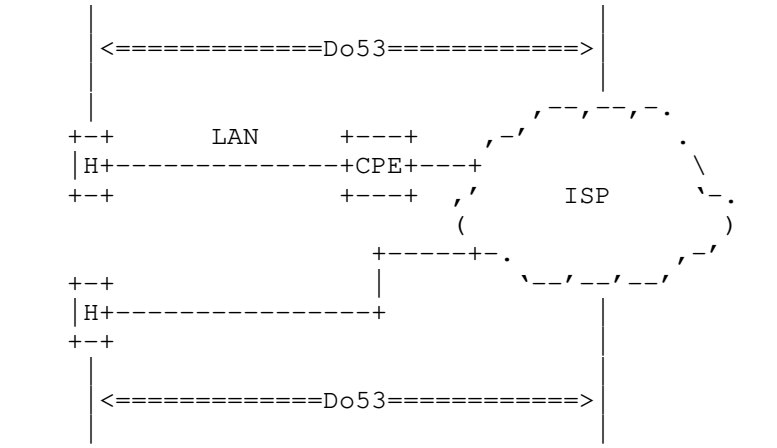

Legend:
* H: refers to a host.

Figure 13: Sample Legacy Deployments

A.1. Managed CPEs

This section focuses on CPEs that are managed by ISPs.

A.1.1. Direct DNS

ISPs have developed an expertise in managing service-specific configuration information (e.g., CPE WAN Management Protocol [TR-069]). For example, these tools may be used to provision the DNS server's ADN to managed CPEs if an encrypted DNS is supported by a local network similar to what is depicted in Figure 14.

For example, DoH-capable (or DoT) clients establish the DoH (or DoT) session with the discovered DoH (or DoT) server.

The DNS client discovers whether the DNS server in the local network supports DoH/DoT/DoQ by using a dedicated field in the discovery message: Encrypted DNS Types (Sections 4, 5, and 6) .

(a) Fixed Networks

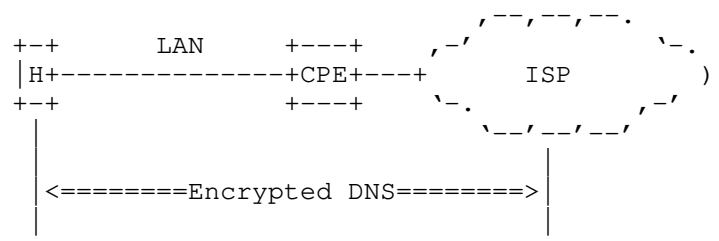

(b) Cellular Networks

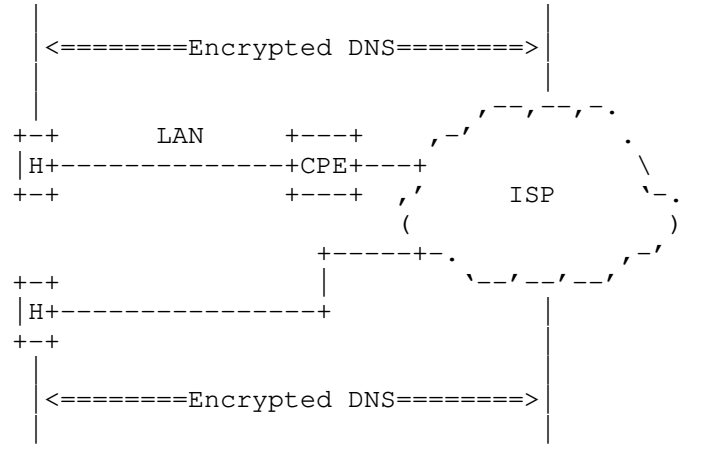

Figure 14: Encrypted DNS in the WAN

Figure 14 shows the scenario where the CPE relays the list of encrypted DNS servers it learns for the network by using mechanisms

like DHCP or a specific Router Advertisement message. In such context, direct encrypted DNS sessions will be established between a host serviced by a CPE and an ISP-supplied encrypted DNS server (see the example depicted in Figure 15 for a DoH/DoT-capable host).

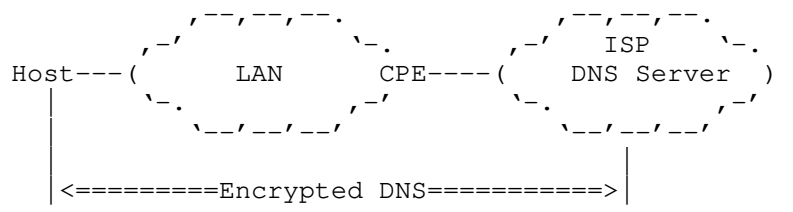

Figure 15: Direct Encrypted DNS Sessions

A.1.2. Proxied DNS

Figure 16 shows a deployment where the CPE embeds a caching DNS forwarder. The CPE advertises itself as the default DNS server to the hosts it serves. The CPE relies upon DHCP or RA to advertise itself to internal hosts as the default DoT/DoH/Do53 server. When receiving a DNS request it cannot handle locally, the CPE forwards the request to an upstream DoH/DoT/Do53 resolver. Such deployment is required for IPv4 service continuity purposes (e.g., Section 5.4.1 of [I-D.ietf-v6ops-rfc7084-bis]) or for supporting advanced services within a local network (e.g., malware filtering, parental control, Manufacturer Usage Description (MUD) [RFC8520] to only allow intended communications to and from an IoT device). When the CPE behaves as a DNS forwarder, DNS communications can be decomposed into two legs:

- The leg between an internal host and the CPE.
- The leg between the CPE and an upstream DNS resolver.

An ISP that offers encrypted DNS to its customers may enable encrypted DNS in one or both legs as shown in Figure 16. Additional considerations related to this deployment are discussed in Section 8.

(a)

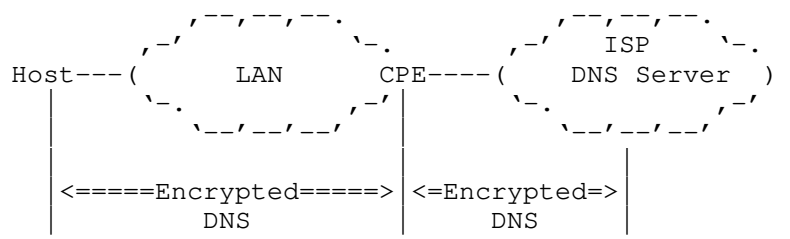

(b)

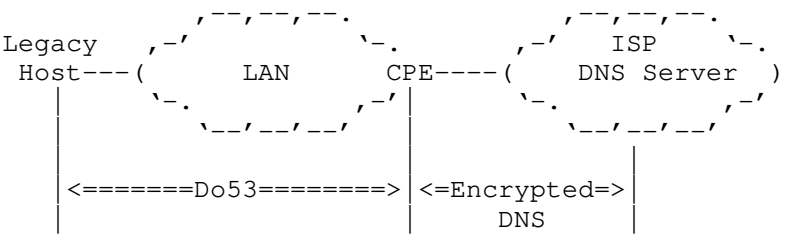

Figure 16: Proxied Encrypted DNS Sessions

A.2. Unmanaged CPEs

A.2.1. ISP-facing Unmanaged CPEs

Customers may decide to deploy unmanaged CPEs (assuming the CPE is compliant with the network access technical specification that is usually published by ISPs). Upon attachment to the network, an unmanaged CPE receives from the network its service configuration (including the DNS information) by means of, e.g., DHCP. That DNS information is shared within the LAN following the same mechanisms as those discussed in Appendix A.1. A host can thus establish DoH/DoT session with a DoH/DoT server similar to what is depicted in Figure 15 or Figure 16.

A.2.2. Internal Unmanaged CPEs

Customers may also decide to deploy internal routers (called hereafter, Internal CPEs) for a variety of reasons that are not detailed here. Absent any explicit configuration on the internal CPE to override the DNS configuration it receives from the ISP-supplied CPE, an Internal CPE relays the DNS information it receives via DHCP/RA from the ISP-supplied CPE to connected hosts. Encrypted DNS sessions can be established by a host with the DNS servers of the ISP (see Figure 17).

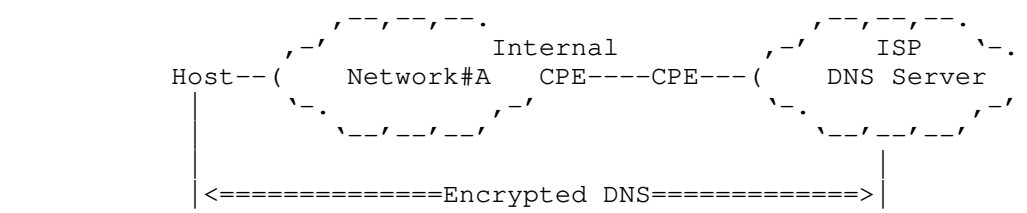

Figure 17: Direct Encrypted DNS Sessions with the ISP DNS Resolver (Internal CPE)

Similar to managed CPEs, a user may modify the default DNS configuration of an unmanaged CPE to use his/her favorite DNS servers instead. Encrypted DNS sessions can be established directly between a host and a 3rd Party DNS server (see Figure 18).

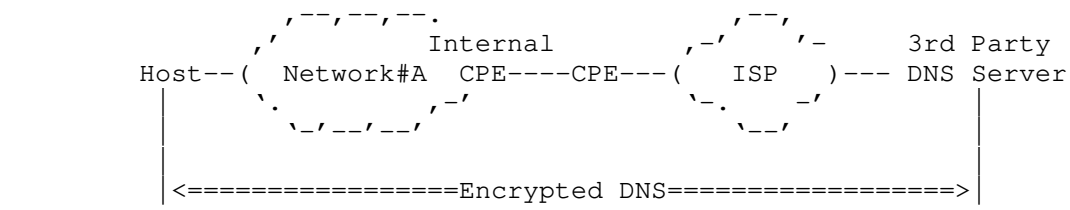

Figure 18: Direct Encrypted DNS Sessions with a Third Party DNS Resolver

Section 8.2 discusses considerations related to hosting a forwarder in the Internal CPE.

Appendix B. Make Use of Discovered Encrypted DNS Servers

Even if the use of a discovered encrypted DNS server is beyond the discovery process and falls under encrypted server selection, the following discusses typical conditions under which discovered encrypted DNS server can be used.

- If the DNS server's IP address discovered by using DHCP/RA is preconfigured in the OS or Browser as a verified resolver (e.g., part of an auto-upgrade program such as [Auto-upgrade]), the DNS client auto-upgrades to use the preconfigured encrypted DNS server tied to the discovered DNS server IP address. In such a case the DNS client will perform additional checks out of band, such as confirming that the Do53 IP address and the encrypted DNS server are owned and operated by the same organisation.

- Similarly, if the ADN conveyed in DHCP/RA (Sections 4, 5, and 6) is preconfigured in the OS or browser as a verified resolver, the

DNS client auto-upgrades to establish an encrypted a DoH/DoT/DoQ session with the ADN.

In such case, the DNS client matches the domain name in the Encrypted DNS DHCP/RA option with the 'DNS-ID' identifier type within subjectAltName entry in the server certificate conveyed in the TLS handshake.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: TirumaleswarReddy_Konda@McAfee.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

Neil Cook
Open-Xchange
UK

Email: neil.cook@noware.co.uk

Tommy Jensen
Microsoft
USA

Email: tojens@microsoft.com

ADD
Internet-Draft
Intended status: Standards Track
Expires: October 9, 2020

M. Boucadair
Orange
T. Reddy
McAfee
D. Wing
Citrix
V. Smyslov
ELVIS-PLUS

April 7, 2020

Internet Key Exchange Protocol Version 2 (IKEv2) Configuration for Encrypted DNS
draft-btw-add-ipsecme-ike-00

Abstract

This document specifies a new Internet Key Exchange Protocol Version 2 (IKEv2) Configuration Payload Attribute Type for encrypted DNS such as DNS-over-HTTPS (DoH) and DNS-over-TLS (DoT).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 9, 2020.

Table of Contents

1. Introduction

This document specifies encrypted DNS configuration for an IKE initiator, particularly the Authentication Domain Name (ADN, defined in [RFC8310]) of DNS-over-HTTPS (DoH) [RFC8484] or DNS-over-TLS (DoT) [RFC7858] server using Internet Key Exchange Protocol Version 2 (IKEv2) [RFC7296].

Particularly, this document introduces a new IKEv2 Configuration Payload Attribute Types (Section 4) for the support of encrypted DNS servers (e.g., DoT, DoH).

Sample use cases are discussed in Section 3. The Configuration Payload Attribute Type defined in Section 4 is not specific to these deployments, but can be used in other deployment contexts.

Note that, for many years, typical designs has often considered that the DNS server was usually located inside the protected domain, but could theoretically be located outside of it. With DoH or DoT, the latter option becomes plausible.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499] and [I-D.ietf-dnsop-terminology-ter].

Also, this document makes use of the terms defined in [RFC7296]. In particular, readers should be familiar with "initiator" and "responder" terms used in that document.

Do53 refers to unencrypted DNS.

'DoH/DoT' refers to DNS-over-HTTPS and/or DNS-over-TLS.

3. Sample Deployment Scenarios

3.1. Roaming Enterprise Users

In this Enterprise scenario (Section 1.1.3 of [RFC7296]), a roaming user connects to the Enterprise network through an IPsec tunnel. The split-tunnel Virtual Private Network (VPN) configuration allows the endpoint to access hosts that resides in the Enterprise network [RFC8598] using that tunnel; other traffic not destined to the Enterprise does not traverse the tunnel. In contrast, a non-split-tunnel VPN configuration causes all traffic to traverse the tunnel into the enterprise.

For both split- and non-split-tunnel configurations, the use of DoT/DoH instead of Do53 provides privacy and integrity protection along the entire path (rather than just to the VPN termination device) and can communicate the DoT/DoH server policies.

For split-tunnel VPN configurations, the endpoint uses the Enterprise-provided DoT/DoH server to resolve internal-only domain names.

For non-split-tunnel VPN configurations, the endpoint uses the Enterprise-provided DoT/DoH server to resolve both internal and external domain names.

Enterprise networks are susceptible to internal and external attacks. To minimize that risk all enterprise traffic is encrypted (Section 2.1 of [I-D.arkko-farrell-arch-model-t]).

3.2. VPN Service Provider

Legacy VPN service providers usually preserve end-users' data confidentiality by sending all communication traffic through an encrypted tunnel. A VPN service provider can also provide guarantees about the security of the VPN network by filtering malware and phishing domains.

Browsers and OSes support DoH/DoT; VPN providers may no longer expect DNS clients to fallback to Do53 just because it is a closed network.

The DoT/DoH server hosted by the VPN service provider can be securely discovered by the endpoint using the IKEv2 Configuration Payload Attribute Type.

3.3. DNS Offload

VPN service providers typically allow split-tunnel VPN configuration in which users can choose applications that can be excluded from the tunnel. For example, users may exclude applications that restrict VPN access.

VPN service providers can also offer publicly accessible DoH/DoT servers. The split-tunnel VPN configuration allows the client to access the DoH/DoT servers hosted by the VPN provider without traversing the tunnel.

The DoT/DoH server hosted by the VPN service provider can be securely discovered by the endpoint using the IKEv2 Configuration Payload Attribute Type.

4. INTERNAL_ENC_DNS Attribute

The INTERNAL_ENC_DNS IKEv2 Configuration Payload Attribute Type is used to configure an encrypted DNS server. The format of this attribute is shown in Figure 1.

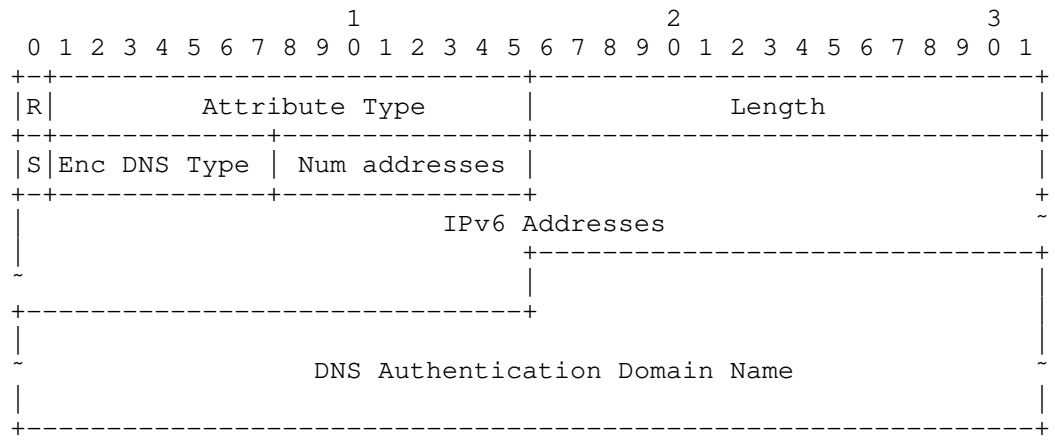

Figure 1: INTERNAL_ENC_DNS Attribute Format

The fields of the attribute shown in Figure 1 are as follows:

- R: Reserved bit; refer to Section 3.15.1 of [RFC7296].

- Attribute Type: MUST be set to TBA (Section 8.1).

- Length: Length of the data in octets. It MUST be set to 1 if the Configuration payload has types CFG_REQUEST or CFG_ACK or to (2 + Length of the ADN + N * 16) if the Configuration payload has types CFG_REPLY or CFG_SET; N being the number of included IP addresses ('Num addresses').

- S: Scope bit. This bit controls whether the DNS queries are sent within the tunnel or outside. If set, this bit instructs the initiator to send encrypted DNS queries outside the tunnel. If the bit is unset, the queries are sent inside the tunnel. The default value of this bit is "0".

- Encrypted DNS Type: Indicates the type of the encrypted DNS server conveyed in this attribute. The following values are defined:

  0: Reserved

  1: DoT

  2: DoH

  See Section 8.2 for future assignment considerations.

- Num addresses: If Length > 1, it indicates the number of enclosed IP addresses.

- IPv6 Address(es): One or more IPv6 addresses to be used to reach the encrypted DNS identified by the name in the DNS Authentication Domain Name.

  IPv4 addresses MUST be encoded using the IPv4-Mapped IPv6 Address format defined in [RFC4291].

- Authentication Domain Name: A fully qualified domain name of the DoT (or DoH) server following the syntax defined in [RFC5890]. The name MUST NOT contain any terminators (e.g., NULL, CR).

  An example of valid ADN for DoH server is "doh1.example.com".

5. URI Template

DoH servers may support more than one URI Template [RFC8484]. The following sub-sections discuss some candidate solutions for a DoH client to retrieve the list of supported templates by a DoH server. Also, if the resolver hosts several DoH services (e.g., no-filtering, blocking adult content, blocking malware), these services can be discovered as templates.

This section will be updated to reflect the outcome of the discussion in [I-D.btw-add-home].

How a DoH client makes use of the configured DoH services is out of the scope of this document.

6. IKEv2 Protocol Exchange

This section describes how an initiator can be configured with an encrypted DNS server (e.g., DoH, DoT) using IKEv2.

Initiators indicate the support of an encrypted DNS in the CFG_REQUEST payloads by including INTERNAL_ENC_DNS attribute, while responders supply the encrypted DNS configuration in the CFG_REPLY payloads. Concretely:

If the initiator supports encrypted DNS, it includes one or more INTERNAL_ENC_DNS attributes in the CFG_REQUEST with the "Encrypted DNS Type" set to the requested encrypted DNS type (Section 4). For each supported encrypted DNS type the initiator MUST include exactly one INTERNAL_ENC_DNS attribute with the Length field set to 1.

If an INTERNAL_ENC_DNS attribute is included in the CFG_REQUEST, the INTERNAL_ENC_DNS attribute MUST NOT include an ADN and list of IP addresses.

For each INTERNAL_ENC_DNS attribute from the CFG_REQUEST, if the responder supports the corresponding encrypted DNS type, then it MAY send back an INTERNAL_ENC_DNS attribute in the CFG_REPLY with this encrypted DNS type and an appropriate list of IP addresses and ADN. The list of IP addresses MUST NOT be empty.

If the CFG_REQUEST includes an INTERNAL_ENC_DNS attribute but the CFG_REPLY does not include an INTERNAL_ENC_DNS, this is an indication that requested encrypted DNS type(s) is not supported by the responder.

The behavior of the responder if it receives both INTERNAL_ENC_DNS and INTERNAL_IP6_DNS (or INTERNAL_IP4_DNS) attributes is policy-based and deployment-specific. However, it is RECOMMENDED that if the responder includes at least one INTERNAL_ENC_DNS attribute in the reply, it should not include any of INTERNAL_IP4_DNS/ INTERNAL_IP6_DNS attributes.

The DNS client establishes a DoH/DoT session with the address(es) conveyed in INTERNAL_ENC_DNS and uses the mechanism discussed in Section 8 of [RFC8310] to authenticate the DNS server certificate using the authentication domain name conveyed in INTERNAL_ENC_DNS.

If the IPsec connection is a split-tunnel configuration and the initiator negotiated INTERNAL_DNS_DOMAIN as per [RFC8598], the DNS client MUST resolve the internal names using INTERNAL_ENC_DNS DNS servers.

Note: [RFC8598] requires INTERNAL_IP6_DNS (or INTERNAL_IP4_DNS) attribute to be mandatory present when INTERNAL_DNS_DOMAIN is included. This specification relaxes that constraint in the presence of INTERNAL_ENC_DNS attribute.

7. Security Considerations

This document adheres to the security considerations defined in [RFC7296]. In particular, this document does not alter the trust on the DNS configuration provided by a responder.

Networks are susceptible to internal attacks as discussed in Section 3.2 of [I-D.arkko-farrell-arch-model-t]. Hosting DoH/DoT server even in case of split-VPN configuration minimizes the attack vector (e.g., a compromised network device cannot monitor/modify DNS

traffic). This specification describes a mechanism to restrict access to the DNS messages to only the parties that need to know.

In most deployment scenarios, the initiator expects that it is using the DoH/DoT server hosted by a specific organization or enterprise. The DNS client can validate the signatory (i.e., cryptographically attested by the organization hosting the DoH/DoT server) using, for example, [I-D.reddy-add-server-policy-selection], and the user can review human-readable privacy policy information of the DNS server and assess whether the DNS server performs DNS-based content filtering. This helps to protect from a compromised IKE server advertising a malicious DoH/DoT server.

The initiator may trust the DoH/DoT servers supplied by means of IKEv2 from a trusted responder more than the locally provided DNS servers, especially in the case of connecting to unknown or untrusted networks (e.g., coffee shops or hotel networks). In addition, the initiator may prefer IKEv2-supplied DoH/DoT servers if they provide additional features (e.g., malware filtering) compared to the pre-configured DNS servers and meets the privacy preserving data policy requirements of the user.

If the DoH/DoT server that was discovered by means of IKEv2 does not meet the privacy preserving data policy and filtering requirements of the user, the user can instruct the DNS client to take appropriate actions. For example, the action can be to use the local DoH/DoT server only to access internal-only DNS names and use another DNS server (that addresses his/her expectations) for public domains. Such actions and their handling is out of scope.

If IKEv2 is being negotiated with an anonymous or unknown endpoint (such as for Opportunistic Security [RFC7435]), the initiator MUST NOT use INTERNAL_ENC_DNS servers unless it is pre-configured in the OS or the browser.

This specification does not extend the scope of accepting DNSSEC trust anchors beyond the usage guidelines defined in Section 6 of [RFC8598].

8. IANA Considerations

8.1. Configuration Payload Attribute Type

This document requests IANA to assign the following new IKEv2 Configuration Payload Attribute Types from the "IKEv2 Configuration Payload Attribute Types" namespace available at https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml#ikev2-parameters-21.

Value	Attribute Type	Multi-Valued	Length	Reference
TBA	INTERNAL_ENC_DNS	YES	1 or more	RFC XXXX

8.2. Encrypted DNS Types

This document requests IANA to create a new registry called "Encrypted DNS Types" under "Internet Key Exchange Version 2 (IKEv2) Parameters" available at https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml#ikev2-parameters-21. The initial values of the registry is as follows:

Value	Description	Reference
0	Reserved	RFC XXXX
1	DNS-over-TLS (DoT)	RFC XXXX
2	DNS-over-HTTPs (DoH)	RFC XXXX

New values are assigned on a First Come, First Served (FCFS) basis (Section 4.4 of [RFC8126]).

9. Acknowledgements

Many thanks to Yoav Nir, Christian Jacquenet, Paul Wouters, and Tommy Pauly for the review and comments.

Yoav and Paul suggested the use of one single attribute carrying both the name and an IP address instead of depending on the existing INTERNAL_IP6_DNS and INTERNAL_IP4_DNS attributes.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <https://www.rfc-editor.org/info/rfc5890>.

[RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <https://www.rfc-editor.org/info/rfc7296>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

10.2. Informative References

[I-D.arkko-farrell-arch-model-t]
Arkko, J. and S. Farrell, "Challenges and Changes in the Internet Threat Model", draft-arkko-farrell-arch-model-t-03 (work in progress), March 2020.

[I-D.btw-add-home]
Boucadair, M., Reddy.K, T., Wing, D., and N. Cook, "DNS-over-HTTPS and DNS-over-TLS Server Discovery and Deployment Considerations for Home and Mobile Networks", draft-btw-add-home-04 (work in progress), March 2020.

[I-D.ietf-dnsop-terminology-ter]
Hoffman, P., "Terminology for DNS Transports and Location", draft-ietf-dnsop-terminology-ter-01 (work in progress), February 2020.

[I-D.reddy-add-server-policy-selection]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "DNS Server Selection: DNS Server Information with Assertion Token", draft-reddy-add-server-policy-selection-00 (work in progress), March 2020.

[RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", RFC 7435, DOI 10.17487/RFC7435, December 2014, <https://www.rfc-editor.org/info/rfc7435>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

[RFC8598] Pauly, T. and P. Wouters, "Split DNS Configuration for the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 8598, DOI 10.17487/RFC8598, May 2019, <https://www.rfc-editor.org/info/rfc8598>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: TirumaleswarReddy_Konda@McAfee.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

Valery Smyslov
ELVIS-PLUS
RU

Email: svan@elvis.ru

ADD
Internet-Draft
Intended status: Standards Track
Expires: March 13, 2021

M. Boucadair
Orange
T. Reddy
McAfee
D. Wing
Citrix
V. Smyslov
ELVIS-PLUS

September 9, 2020

Internet Key Exchange Protocol Version 2 (IKEv2) Configuration for Encrypted DNS
draft-btw-add-ipsecme-ike-01

Abstract

This document specifies a new Internet Key Exchange Protocol Version 2 (IKEv2) Configuration Payload Attribute Types for encrypted DNS with a focus on DNS-over-HTTPS (DoH), DNS-over-TLS (DoT), and DNS-over-QUIC (DoQ).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 13, 2021.

Table of Contents

1. Introduction

This document specifies encrypted DNS configuration for an Internet Key Exchange Protocol Version 2 (IKEv2) [RFC7296] initiator, particularly the Authentication Domain Name (ADN, defined in [RFC8310]) of DNS-over-HTTPS (DoH) [RFC8484], DNS-over-TLS (DoT) [RFC7858], or DNS-over-QUIC (DoQ) [I-D.ietf-dprive-dnsoquic].

This document introduces new IKEv2 Configuration Payload Attribute Types (Section 4) for the support of DoT, DoH, and DoQ DNS servers.

This document targets the deployments discussed in Section 3.3 of [I-D.box-add-requirements]. Sample use cases are discussed in Section 3. The Configuration Payload Attribute Types defined in this document are not specific to these deployments, but can also be used in other deployment contexts.

Note that, for many years, typical designs have often considered that the DNS server was usually located inside the protected domain, but could be located outside of it. With DoH, DoT, or DoQ the latter option becomes plausible.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499] and [I-D.ietf-dnsop-terminology-ter].

Also, this document makes use of the terms defined in [RFC7296]. In particular, readers should be familiar with "initiator" and "responder" terms used in that document.

Do53 refers to unencrypted DNS.

Encrypted DNS refers to as scheme where DNS messages are sent over an encrypted channel. Examples of encrypted DNS are DoT, DoH, and DoQ.

ENCDNS_IP*_* refers to any IKEv2 Configuration Payload Attribute Types defined in Section 4.

ENCDNS_IP4_* refers to an IKEv2 Configuration Payload Attribute Type that carries one or multiple IPv4 addresses of an encrypted DNS server.

ENCDNS_IP6_* refers to an IKEv2 Configuration Payload Attribute Type that carries one or multiple IPv6 addresses of an encrypted DNS server.

3. Sample Deployment Scenarios

3.1. Roaming Enterprise Users

In this Enterprise scenario (Section 1.1.3 of [RFC7296]), a roaming user connects to the Enterprise network through an IPsec tunnel. The split-tunnel Virtual Private Network (VPN) configuration allows the endpoint to access hosts that resides in the Enterprise network [RFC8598] using that tunnel; other traffic not destined to the Enterprise does not traverse the tunnel. In contrast, a non-split-tunnel VPN configuration causes all traffic to traverse the tunnel into the enterprise.

For both split- and non-split-tunnel configurations, the use of encrypted DNS instead of Do53 provides privacy and integrity protection along the entire path (rather than just to the VPN

termination device) and can communicate the encrypted DNS server policies.

For split-tunnel VPN configurations, the endpoint uses the Enterprise-provided encrypted DNS server to resolve internal-only domain names.

For non-split-tunnel VPN configurations, the endpoint uses the Enterprise-provided encrypted DNS server to resolve both internal and external domain names.

Enterprise networks are susceptible to internal and external attacks. To minimize that risk all enterprise traffic is encrypted (Section 2.1 of [I-D.arkko-farrell-arch-model-t]).

3.2. VPN Service Provider

Legacy VPN service providers usually preserve end-users' data confidentiality by sending all communication traffic through an encrypted tunnel. A VPN service provider can also provide guarantees about the security of the VPN network by filtering malware and phishing domains.

Browsers and OSes support DoH/DoT; VPN providers may no longer expect DNS clients to fallback to Do53 just because it is a closed network.

The encrypted DNS server hosted by the VPN service provider can be securely discovered by the endpoint using the IKEv2 Configuration Payload Attribute Type.

3.3. DNS Offload

VPN service providers typically allow split-tunnel VPN configuration in which users can choose applications that can be excluded from the tunnel. For example, users may exclude applications that restrict VPN access.

The encrypted DNS server hosted by the VPN service provider can be securely discovered by the endpoint using the IKEv2 Configuration Payload Attribute Type.

4. IKEv2 Configuration Payload Attribute Types for Encrypted DNS

The ENCDNS_IP*_* IKEv2 Configuration Payload Attribute Types are used to configure a DoT, DoH, or DoQ DNS server. All these attributes share the format shown in Figure 1.

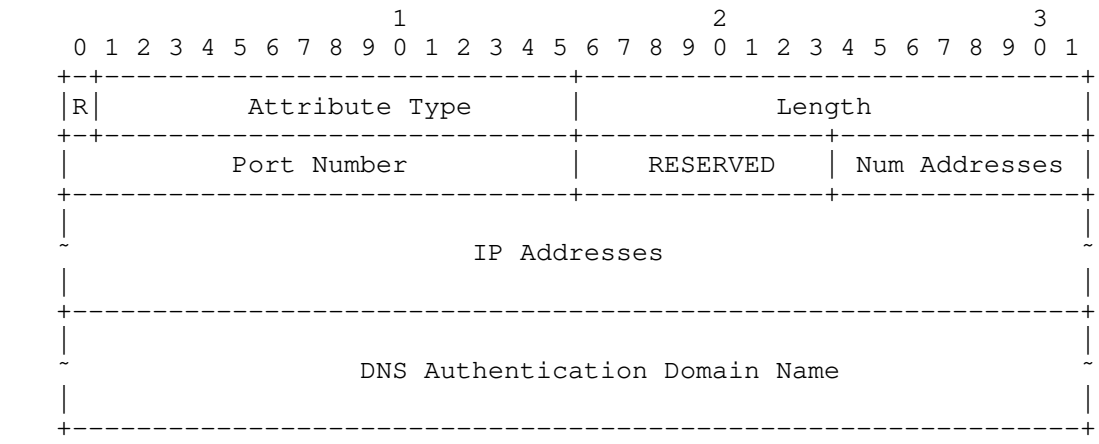

Figure 1: Attributes Format

The fields of the attribute shown in Figure 1 are as follows:

- R (Reserved, 1 bit) - This bit MUST be set to zero and MUST be ignored on receipt (see Section 3.15.1 of [RFC7296] for details).

- Attribute Type (15 bits) - Identifier for Configuration Attribute Type; is set to one of the values listed in Section 8.1.

- Length (2 octets, unsigned integer) - Length of the data in octets. In particular, this field is set to:

  * 0 if the Configuration payload has types CFG_REQUEST or CFG_ACK.

  * (2 + Length of the ADN + N * 4) for ENCDNS_IP4_* attributes if the Configuration payload of a has types CFG_REPLY or CFG_SET; N being the number of included IPv4 addresses ('Num addresses').

  * (2 + Length of the ADN + N * 16) for ENCDNS_IP6_* attributes if the Configuration payload has types CFG_REPLY or CFG_SET; N being the number of included IPv6 addresses ('Num addresses').

- Port Number (2 octets, unsigned integer) - Indicates the port number to be used for the encrypted DNS. As a reminder, the default port number is 853 for DoT and 443 for DoH.

- RESERVED (1 octet) - These bits are reserved for future use. These bits MUST be set to zero by the sender and MUST be ignored by the receiver.

- Num Addresses (1 octet) - Indicates the number of enclosed IPv4 (for ENCDNS_IP4_* attribute types) or IPv6 (for ENCDNS_IP6_* attribute types) addresses.

- IP Address(es) (variable) - One or more IPv4 or IPv6 addresses to be used to reach the encrypted DNS identified by the name in the DNS Authentication Domain Name.

- Authentication Domain Name (variable) - A fully qualified domain name of the DoT, DoH, or DoQ DNS server following the syntax defined in [RFC5890]. The name MUST NOT contain any terminators (e.g., NULL, CR).

  An example of valid ADN for DoH server is "doh1.example.com".

5. IKEv2 Protocol Exchange

This section describes how an initiator can be configured with an encrypted DNS server (e.g., DoH, DoT) using IKEv2.

Initiators indicate the support of an encrypted DNS in the CFG_REQUEST payloads by including one or multiple ENCDNS_IP*_* attributes, while responders supply the encrypted DNS configuration in the CFG_REPLY payloads. Concretely:

If the initiator supports encrypted DNS, it includes one or more ENCDNS_IP*_* attributes in the CFG_REQUEST with the "Attribute Type" set to the requested encrypted DNS type (Section 4). For each supported encrypted DNS type the initiator MUST include exactly one attribute with the Length field set to 0, so that no data is included for these attributes.

For each ENCDNS_IP*_* attribute from the CFG_REQUEST, if the responder supports the corresponding encrypted DNS type, and absent any policy, the responder sends back an ENCDNS_IP*_* attribute in the CFG_REPLY with this encrypted DNS type and an appropriate list of IP addresses, a port number, and an ADN. The list of IP addresses MUST NOT be empty. Multiple instances of the same ENCDNS_IP*_* attribute MAY be returned if distinct ADNs (or port numbers) are to be returned by the responder.

If the CFG_REQUEST includes an ENCDNS_IP*_* attribute but the CFG_REPLY does not include an ENCDNS_IP*_* matching the requested encrypted DNS type, this is an indication that requested encrypted DNS type(s) is not supported by the responder or the responder is not configured to provide corresponding server addresses.

The behavior of the responder if it receives both ENCDNS_IP*_* and INTERNAL_IP6_DNS (or INTERNAL_IP4_DNS) attributes is policy-based and deployment-specific. However, it is RECOMMENDED that if the responder includes at least one ENCDNS_IP*_* attribute in the reply, it should not include any of INTERNAL_IP4_DNS/ INTERNAL_IP6_DNS attributes.

The DNS client establishes an encrypted DNS session (e.g., DoT, DoH, DoQ) with the address(es) conveyed in ENCDNS_IP*_* and uses the mechanism discussed in Section 8 of [RFC8310] to authenticate the DNS server certificate using the authentication domain name conveyed in ENCDNS_IP*_*.

If the IPsec connection is a split-tunnel configuration and the initiator negotiated INTERNAL_DNS_DOMAIN as per [RFC8598], the DNS client MUST resolve the internal names using ENCDNS_IP*_* DNS servers.

Note: [RFC8598] requires INTERNAL_IP6_DNS (or INTERNAL_IP4_DNS) attribute to be mandatory present when INTERNAL_DNS_DOMAIN is included. This specification relaxes that constraint in the presence of ENCDNS_IP*_* attributes.

6. URI Template

DoH servers may support more than one URI Template [RFC8484]. Also, if the resolver hosts several DoH services (e.g., no-filtering, blocking adult content, blocking malware), these services can be discovered as templates.

Upon discovery of a DoH resolver (Section 5), the DoH client contacts that DoH resolver to retrieve the list of supported DoH services using the well-known URI defined in [I-D.btw-add-rfc8484-clarification]. DoH clients re-iterates that request regularly to retrieve an updated list of supported DoH services.

How a DoH client makes use of the configured DoH services is out of the scope of this document.

7. Security Considerations

This document adheres to the security considerations defined in [RFC7296]. In particular, this document does not alter the trust on the DNS configuration provided by a responder.

Networks are susceptible to internal attacks as discussed in Section 3.2 of [I-D.arkko-farrell-arch-model-t]. Hosting encrypted

DNS server even in case of split-VPN configuration minimizes the attack vector (e.g., a compromised network device cannot monitor/ modify DNS traffic). This specification describes a mechanism to restrict access to the DNS messages to only the parties that need to know.

In most deployment scenarios, the initiator expects that it is using the encrypted DNS server hosted by a specific organization or enterprise. The DNS client can validate the signatory (i.e., cryptographically attested by the organization hosting the encrypted DNS server) using, for example, [I-D.reddy-add-server-policy-selection], and the user can review human-readable privacy policy information of the DNS server and assess whether the DNS server performs DNS-based content filtering. This helps to protect from a compromised IKE server advertising a malicious encrypted DNS server.

The initiator may trust the encrypted DNS servers supplied by means of IKEv2 from a trusted responder more than the locally provided DNS servers, especially in the case of connecting to unknown or untrusted networks (e.g., coffee shops or hotel networks).

If the encrypted DNS server that was discovered by means of IKEv2 does not meet the privacy preserving data policy and filtering requirements of the user, the user can instruct the DNS client to take appropriate actions. For example, the action can be to use the local encrypted DNS server only to access internal-only DNS names and use another DNS server (that addresses his/her expectations) for public domains. Such actions and their handling is out of scope.

If IKEv2 responder has used NULL Authentication method [RFC7619] to authenticate itself, the initiator MUST NOT use returned ENCDNS_IP*_* servers configuration unless it is pre-configured in the OS or the browser.

This specification does not extend the scope of accepting DNSSEC trust anchors beyond the usage guidelines defined in Section 6 of [RFC8598].

8. IANA Considerations

8.1. Configuration Payload Attribute Types

This document requests IANA to assign the following new IKEv2 Configuration Payload Attribute Types from the "IKEv2 Configuration Payload Attribute Types" namespace available at https://www.iana.org/assignments/ikev2-parameters/ ikev2-parameters.xhtml#ikev2-parameters-21.

Value	Attribute Type	Multi-Valued	Length	Reference
TBA1	ENCDNS_IP4_DOT	YES	0 or more	RFC XXXX
TBA2	ENCDNS_IP6_DOT	YES	0 or more	RFC XXXX
TBA3	ENCDNS_IP4_DOH	YES	0 or more	RFC XXXX
TBA4	ENCDNS_IP6_DOH	YES	0 or more	RFC XXXX
TBA5	ENCDNS_IP4_DOQ	YES	0 or more	RFC XXXX
TBA6	ENCDNS_IP6_DOQ	YES	0 or more	RFC XXXX

9. Acknowledgements

Many thanks to Yoav Nir, Christian Jacquenet, Paul Wouters, and Tommy Pauly for the review and comments.

Yoav and Paul suggested the use of one single attribute carrying both the name and an IP address instead of depending on the existing INTERNAL_IP6_DNS and INTERNAL_IP4_DNS attributes.

Christian Huitema suiggested to return a port number in the attributes.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <https://www.rfc-editor.org/info/rfc5890>.

[RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <https://www.rfc-editor.org/info/rfc7296>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

10.2. Informative References

[I-D.arkko-farrell-arch-model-t]
Arkko, J. and S. Farrell, "Challenges and Changes in the Internet Threat Model", draft-arkko-farrell-arch-model-t-04 (work in progress), July 2020.

[I-D.box-add-requirements]
Box, C., Pauly, T., Wood, C., and T. Reddy.K, "Requirements for Adaptive DNS Discovery", draft-box-add-requirements-00 (work in progress), September 2020.

[I-D.btw-add-rfc8484-clarification]
Boucadair, M., Cook, N., Reddy.K, T., and D. Wing, "Supporting Redirection for DNS Queries over HTTPS (DoH)", draft-btw-add-rfc8484-clarification-02 (work in progress), July 2020.

[I-D.ietf-dnsop-terminology-ter]
Hoffman, P., "Terminology for DNS Transports and Location", draft-ietf-dnsop-terminology-ter-02 (work in progress), August 2020.

[I-D.ietf-dprive-dnsoquic]
Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnsoquic-00 (work in progress), April 2020.

[I-D.reddy-add-server-policy-selection]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "DNS Server Selection: DNS Server Information with Assertion Token", draft-reddy-add-server-policy-selection-04 (work in progress), July 2020.

[RFC7619] Smyslov, V. and P. Wouters, "The NULL Authentication Method in the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 7619, DOI 10.17487/RFC7619, August 2015, <https://www.rfc-editor.org/info/rfc7619>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

[RFC8598] Pauly, T. and P. Wouters, "Split DNS Configuration for the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 8598, DOI 10.17487/RFC8598, May 2019, <https://www.rfc-editor.org/info/rfc8598>.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: TirumaleswarReddy_Konda@McAfee.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

Valery Smyslov
ELVIS-PLUS
RU

Email: svan@elvis.ru

ADD
Internet-Draft
Updates: 8484 (if approved)
Intended status: Standards Track
Expires: January 7, 2021

M. Boucadair
Orange
N. Cook
Open-Xchange
T. Reddy
McAfee
D. Wing
Citrix

July 6, 2020

Supporting Redirection for DNS Queries over HTTPS (DoH)
draft-btw-add-rfc8484-clarification-02

Abstract

This document clarifies whether DNS-over-HTTPS (DoH) redirection is allowed, describes potential issues with redirection in DoH, and proposes how DoH redirection might be performed.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 7, 2021.

Table of Contents

1. Introduction

This document clarifies the intent of DNS-over-HTTPS (DoH) [RFC8484] whether redirection is allowed (Section 4), potential issues with redirection in DoH (Section 5) and subsequently makes some proposals for how service-level (Section 6) and resource-level (Section 7) redirection might be performed.

This document adheres to Section 4.3 of [I-D.ietf-httpbis-bcp56bis] which discusses the need for protocols using HTTP to specify redirect handling to avoid interoperability problems.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

"A/AAAA" is used to refer to "A and/or AAAA records".

3. Discussion

[RFC8484] indicates that the support of HTTP [RFC7540] redirection is one of DoH design goals (Section 1):

> "The described approach is more than a tunnel over HTTP. It establishes default media formatting types for requests and responses but uses normal HTTP content negotiation mechanisms for selecting alternatives that endpoints may prefer in anticipation of serving new use cases. In addition to this media type negotiation, it aligns itself with HTTP features such as caching, redirection, proxying, authentication, and compression.
>
> The integration with HTTP provides a transport suitable for both existing DNS clients and native web applications seeking access to the DNS."

Nevertheless, Section 3 of [RFC8484] indicates the following:

> "This specification does not extend DNS resolution privileges to URIs that are not recognized by the DoH client as configured URIs."

This looks like an internal inconsistency of [RFC8484] that is worth the clarification: is redirection allowed or not?

Also, Section 3 of [RFC8484] indicates that:

> "A DoH client MUST NOT use a different URI simply because it was discovered outside of the client's configuration (such as through HTTP/2 server push) or because a server offers an unsolicited response that appears to be a valid answer to a DNS query."

Nevertheless, [RFC8484] does not:

- specify under which conditions a discovered different URI can be used.

- describe how a different URI can be discovered using HTTP/2 server push. The only available example in the mailing list archives clarifies that server push is an example of unsolicited responses.

  The text was updated late in the publication process to address this comment: https://mailarchive.ietf.org/arch/msg/doh/f_V-tBgB-KRsLZhttx9tGt75cps/. The example provided in the thread (server push) is related to the second part of the above excerpt.

- clarify that unsolicited messages from a configured DoH server should be excluded.

A clarification is proposed in Section 4. This clarification focuses on a "different URI" that might be discovered while communicating with an HTTP server.

Additionally, assuming that redirection is allowed, this specification recommends how it is achieved. This is required because redirection to a domain-based URI requires DNS resolution of that domain name, which creates a potential bootstrapping problem (e.g., If DoH server is the only configured DNS server, redirecting the client to a new server by presenting a name will fail).

4. RFC8484 Update

OLD:

> A DoH client MUST NOT use a different URI simply because it was discovered outside of the client's configuration (such as through HTTP/2 server push) or because a server offers an unsolicited response that appears to be a valid answer to a DNS query.

NEW

> A DoH client MUST NOT use a different URI that was discovered outside of the client's configuration when communicating with HTTP servers except via HTTP redirection from a configured URI (Section 6.4 of [RFC7231]).
>
> Also, a DoH client MUST ignore an unsolicited response (such as through HTTP/2 server push) that appears to be a valid answer to a DNS query unless that response comes from a configured URI (as described in Section 5.3 of [RFC8484]).

5. Issues with Redirection in DoH

There are several potential issues with redirection in DoH, which are summarized below.

The first issue to be considered is whether a new document considering redirection is needed at all. Redirection in HTTP is done on a per-resource basis; if the only functionality required is to redirect all requests to an entirely different server under the same administrative control, then the alternative service mechanism described in [RFC7838] might be sufficient. However, there are restrictions on the use of alternative services; specifically the certificate presented by the alternative service must be valid for

the origin. This restriction means that alternative services cannot be used for use-cases such as redirecting the client to a locally administered DoH server (e.g., resolver or forwarder) which does not have a certificate valid for the origin. Additionally, alternative services suffer from the bootstrapping issue described below.

The second issue with using HTTP redirection is bootstrapping; any client that is relying solely upon a DoH server for resolution must be able to resolve the domain in the redirect response. Even if a DoH client has a plaintext DNS resolver configured, using that resolver is considered as a minimal privacy leakage [RFC8310]. One possible solution is for the DoH client to use the same server that returned the redirect response to perform the resolution, however that may then lead to a further redirect response. Another solution is for the DoH server to include additional information in the response, similar to the "glue" records as defined in [RFC7719].

The final issue is that HTTP redirection is done on a per-resource basis; this presents several problems for DoH:

1. Every GET request with a new query name will require redirection, which is suboptimal. Indeed, a redirect will only affect a unique request, and the DoH client will thus need to contact the origin server for every new request and get redirected, requiring two roundtrips. Also, permanent redirects [RFC7538] for all these queries would bloat the client's HTTP cache.

2. Using POST requests would solve the issue. Nevertheless POST responses are not widely cached as per Section 4.2.3 of [RFC7231], and mandating the use of POST requests for DoH in order to enable redirection hardly seems reasonable.

The above issues would seem to indicate that despite the intention of [RFC8484] to align itself with HTTP redirection, some additional work is required in order for any other mechanism than alternative services (e.g., [RFC7838]) to be deployed with confidence.

The rest of this document considers the issue of redirection at two levels:

1. Service-level Redirect: Similar to alternative services, this would allow a DoH server to redirect a DoH client to an alternative service for all future queries, rather than on a per-resource basis.

2. Resource-Level Redirect: Solving the bootstrapping problem for regular HTTP redirects. Note that this doesn't solve the caching issues described above, and does raise the question of whether

regular HTTP redirection is desirable or worthwhile (i.e., are there any valid use-cases for resource-level redirection in DoH?).

6. Service-Level Redirect

We considered two possibilities for service-level redirect:

1. Extending [RFC7838] by relaxing the host authentication checks.

2. Using a well-known URI to return information about alternative services.

Extending alternative services was considered, but rejected (see Appendix A for the reasons) in favour of the well-known URI approach.

6.1. Well-Known URI

We propose the use of the well-known URI mechanism [RFC8615], with the name "resinfo" to retrieve resolver information, which could include specifying alternative services, through the use of a JSON object in the response payload. A well-known URI would thus look like "https://doh.example.com/.well-known/resinfo".

The example in Figure 1 shows what a JSON object might look like that specified one or more alternative services. The structure of the response is inspired by Section 4.4.2 of [RFC7975].

Note that the response includes "glue" RR information to allow the alternative service to be accessed without further DNS queries, and includes an authenticated domain name to be used for authenticating the alternative service.

```
{
  "associated-resolvers": {
    "adn": [
      {
        "name": "cpe123.example.net",
        "uri-template": [
          "https://cpe123.example.net/dns-query{?dns}"
        ],
        "a": [
          "192.0.2.1",
          "192.0.2.2"
        ],
        "aaaa": [
          "2001:db8::1",
          "2001:db8::2"
        ],
        "ttl": 3600
      }
    ]
  }
}
```

Figure 1: Response Example with Glue RR Information

7. Resource-Level Redirect

Notwithstanding the issues with resource-level redirects described in Section 5, this section describes a proposal for returning the "glue" RRs required to avoid the bootstrapping issue described in that section (but not the roundtrip or caching issues).

Servers supporting DoH redirect MUST support returning the redirect response body mechanism described hereafter.

> Note: "MUST" is used here because resolving the redirect name using Do53 will fail in some configurations, e.g., https://wiki.mozilla.org/Trusted_Recursive_Resolver (network.trr.mode=3).

Concretely, the DoH server returns in the response body a DNS response with an 'application/dns-message' media type as specified in Section 6 of [RFC8484], containing any A and AAAA records for the domain name in the redirect URI, including any CNAMEs.

For example, if the redirect URI contains the domain name "redirect.example.com", and "redirect.example.com" is a CNAME

pointing to "real.example.com", then an example response body would contain:

- A CNAME record for "redirect.example.com"
- Any A records for "real.example.com"
- Any AAAA records for "real.example.com"

This approach is simple; no client or server support of server push is required, and it is also more efficient in terms of the amount of data transmitted.

8. Security Considerations

DoH-related security considerations are discussed in Section 9 of [RFC8484].

Section 9 of [RFC7838] describes security considerations related to the use of alternate services. Relaxing the host authentication requirements would certainly warrant additional security considerations.

9. IANA Considerations

9.1. resinfo Well-Known URI Suffix

This document requests IANA to assign the following well-known URI from the registry available at https://www.iana.org/assignments/well-known-uris/well-known-uris.xhtml.

URI suffix: resinfo

Change controller: IETF

Specification document(s): This document

Status: permanent

10. Acknowledgements

Many thanks to Christian Jacquenet, Philippe Fouquart, and Ben Schwartz for the comments.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <https://www.rfc-editor.org/info/rfc7231>.

[RFC7538] Reschke, J., "The Hypertext Transfer Protocol Status Code 308 (Permanent Redirect)", RFC 7538, DOI 10.17487/RFC7538, April 2015, <https://www.rfc-editor.org/info/rfc7538>.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <https://www.rfc-editor.org/info/rfc7540>.

[RFC7838] Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", RFC 7838, DOI 10.17487/RFC7838, April 2016, <https://www.rfc-editor.org/info/rfc7838>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8615] Nottingham, M., "Well-Known Uniform Resource Identifiers (URIs)", RFC 8615, DOI 10.17487/RFC8615, May 2019, <https://www.rfc-editor.org/info/rfc8615>.

11.2. Informative References

[I-D.ietf-httpbis-bcp56bis]
Nottingham, M., "Building Protocols with HTTP", draft-ietf-httpbis-bcp56bis-09 (work in progress), November 2019.

[RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <https://www.rfc-editor.org/info/rfc7719>.

[RFC7975] Niven-Jenkins, B., Ed. and R. van Brandenburg, Ed., "Request Routing Redirection Interface for Content Delivery Network (CDN) Interconnection", RFC 7975, DOI 10.17487/RFC7975, October 2016, <https://www.rfc-editor.org/info/rfc7975>.

Appendix A. Extending Alternative Services

Section 9.2 of [RFC7838] discusses the possibilities for attackers to hijack the communication to an origin. This is the justification for the requirement in Section 2.1 of [RFC7838] that "Clients MUST have reasonable assurances that the alternative service is under control of and valid for the whole origin.".

However, when a DoH server presents an alternative DoH service to a DoH client, both the origin and alternative service, as well as the DNS queries and responses, must be, by definition, resistant to MITM attacks. Thus it could be argued that in these circumstances, relaxing the host authentication requirements is justified. The relaxation could be limited, e.g., still requiring some relationship between the origin and the alternative, or unlimited, allowing no such relationship to exist.

However the bootstrapping issues described in Section 5 still apply, and there is no mechanism for the DoH server to specify an authenticated domain name to use to authenticate the alternative service, making this proposal unsuitable for deployment.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Neil Cook
Open-Xchange
UK

Email: neil.cook@noware.co.uk

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: TirumaleswarReddy_Konda@McAfee.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

ADD
Internet-Draft
Intended status: Informational
Expires: January 14, 2021

A. Campling
419 Consulting Limited
N. Kowalewski
Deutsche Telekom
G. Scalone
Vodafone
C. Box
BT Group
A. Winfield
Sky
July 13, 2020

Practical Observations from Encrypted DNS Deployments by Network Operators
draft-campling-operator-observations-00

Abstract

The following document includes observations regarding a variety of implementations of recursive DNS capabilities that are important to network operators in terms of delivering DNS services to their (several tens of millions of) customers. It highlights some of the challenges that need to be addressed to allow the widespread adoption of encrypted DNS by the end-users of network operators.

The information is intended to aid the development of discovery mechanisms for protocols such as DNS-over-HTTPS. It clearly defines problems that need technical solutions to allow the deployment of encrypted DNS by the largest number of operators to the largest number of users in the shortest possible timeframe with little or no disruption to the user experience.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2021.

1. Introduction

The IETF has developed many protocols to improve the security and reliability of DNS over UDP or TCP (Do53) [RFC1035] including DNS over TLS (DoT) [RFC7858], DNS over HTTPS (DoH) [RFC8484] and DNS Security Extensions (DNSSEC) [RFC2535]. To enable the broadest adoption of these technologies, there are issues for consideration of any discovery solutions that are proposed to the Adaptive DNS Discovery [ADD] working group.

Many network operators, including Internet Service Providers (ISPs), whether using fixed or mobile networks, would like to ensure that their encrypted DNS services can be seamlessly discovered and used by applications and operating systems that support encrypted DNS, particularly DoH, in order that encrypted DNS can be deployed to the widest possible community of users. They would particularly like to ensure that any proposed DNS discovery mechanisms take into account ISP use-cases such as DNS forwarders on CPE (Customer Premises Equipment or routers), the use of DNS for CDNs (Content Delivery Networks) with local content caches and the non-public nature of most ISP DNS services.

This document has taken observations and experiences from a number of network operators that have been actively working on adding support for encrypted DNS to their networks. It is intended to make clear the requirements needed by any discovery mechanism developed by the ADD group. It collates and succinctly describes common problems faced by existing stakeholders in adopting encrypted DNS mechanisms.

This document also presents some background information that is relevant to describing the issues and explains concerns around

current proposed solutions. It should also be noted that, in many European countries, some regulations are specific to ISPs. One such requirement is that their customers should be able to connect to the Internet with any home router of their choice even if a router is provided by the ISP as part of its service. Therefore new proocols cannot be accommodated simply by requiring ISP customers to upgrade their routers.

2. Rationale

This document is intended provide information to aid interested parties in developing discovery mechanisms for protocols such as DoH to allow their adoption with minimal disruption to the end user experience, maximising the number of users that can enjoy an easy upgrade path towards DNS encryption.

The information provided will help interested paries develop discovery mechanisms that avoid the unnecessary exclusion of the majority of customers of a significant number of ISPs (including the major ones in Europe that serve several tens of millions of customers) from easy access to this new technology using the secure by design, same-provider auto-upgrade mechanisms.

Such discovery choices will ensure that easy access to encrypted DNS is not dependent on the Internet access network architecture and on the ease of upgrade of any CPE. In addition, it will ensure that users are not forced to change their DNS resolver to a third party, potentially via manual configuration by the user, possibly losing functionality in the process.

3. The 'Same Provider Auto-Upgrade' Model

Both Google Chrome and Microsoft Windows (and perhaps other client software in the future) currently deploy encrypted DNS through a 'same provider auto-upgrade' (SPAU) model. This approach results in the client not needing to prompt the user to change to a different resolver operator to enjoy an encrypted connection. Instead the client will rather try to determine whether an encrypted channel exists for communication with the same resolver operator that the user currently employs for unencrypted DNS resolution. If such a channel can be found, the client will automatically upgrade the connection from the original unencrypted one to the new encrypted one; otherwise, the client will continue sending DNS queries unencrypted.

The current implementation of this model is as follows:

- Out of band, the client software vendor discovers ISPs running DoH services (in the case of Google Chrome, ISPs will more likely apply for inclusion through Google's announced process). The vendor records the existing (Do53) resolver IP addresses, and adds them to a hard-wired table that maps those existing Do53 IP addresses to the DoH service that the vendor discovered to be associated with those resolver IPs.

- When the client starts for the first time, and thereafter whenever it detects a network change, it checks the resolver configuration of the local host. If the configured resolver matches one of the IPs listed in the above table, the client auto-upgrades to use the DoH service instead.

The above method ensures that users are only upgraded to DoH when the vendor is sure that the DoH service is the same service as the Do53 service already used.

4. The Problem with Auto-Upgrade and Forwarders

Automatic upgrades that rely upon the user device being able to know and compare the address of the resolver that is serving the device can fail in some home network environments where the CPE is acting as a DNS proxy. To do this, the CPE will run software like DNSMASQ which acts as a proxy between the client and the DNS resolver, also providing DHCP services and performing DNS caching as well as forwarding. This is often paired with a home network architecture that uses RFC1918 [RFC1918] private IP addresses.

In circumstances where private IP addresses are used, any auto-upgrade on the user device that compares the IP address of the device's DNS resolver against a list of known public DNS resolvers will fail because the client DNS resolver is a RFC1918 private address of the CPE device and not the public address of the actual DNS resolver operated by the network operator.

As can be seen, the existing SPAU model doesn't work with the DNS-forwarder, private IP approach commonly used by network operators. Given that this approach allows for the implementation of the best privacy practices and best latency/engineering reqirements, it shouldn't change, therefore the SPAU model needs to be adapted to work with it.

5. Why DNS Discovery Needs to Support Forwarders

5.1. Loss of Functionality if CPE Doesn't Support DNS Forwarders

If the CPE is upgraded to announce the public resolver to clients, the following functionality will be lost

- Caching/Proxy on the CPE - This leads to more load on the ISP's DNS platform because every client talks directly to the public resolvers (not only the clients which are auto-upgraded to DoH but also all other clients).

- Local DNS routing and resilience - Some deployments segment the user base into regions, with CPE in each region receiving a different IPv4 and IPv6 address for the DNS server, improving latency and load balancing, as well as helping with cyber resilience compared to a single address for a typical anycast implementation.

- Addressing local clients via their names - Often the CPE assigns the name configured to a client to the client's IP address on the CPE (for example, if the hostname is set to 'myhost' on a home network to reach this host on that network under that name). This will not work if the clients communicate directly with a resolver in the carrier network nor would it for auto-upgraded clients because, even if they fallback to Do53, they will still ask a resolver in the carrier network and not the CPE - and in both cases private network details will be leaked.

- The CPE is the only network element that is aware of the local network topology. If the local network information is lost it is not possible to differentiate devices. The Discovery mechanism alone is not enough to solve this use case as additional logic is required on the DoH server to send back the request to the CPE. By using EDNS0 (Extension Mechanisms for DNS) [RFC2671] it is possible for a client running on the CPE to pass EDNS0 information to the ISP's DNS and provide, to the opted-in customers, information on the client that performed the request. This in turn allows the execution, for example, of parental controls on devices belonging to children (there are various ways of doing this that preserves privacy, for example by providing information only about the required filtering profile or by providing only a locally generated ID to distinguish between devices without necessarily identifying them).

- Similar to the above use-case, some CPE can be configured to perform filtering directly, relying on a DNS forwarder's ability to intercept and modify DNS queries to do so. Moving queries to the network DoH server removes this capability, allowing more data

to leave the local network, even if a resolver is available to perform similar filtering.

5.2. Why Not Just Upgrade the CPE to Stop Forwarding?

It may seem easier to simply ignore the loss of functionality detailed above and just upgrade the CPE to stop DNS forwarding. However, such a software upgrade programme, or even the wholesale replacement of CPE, is not without its own challenges.

The following is based on information from various large European ISPs, all of which use a DNS forwarder in their CPE. This configuration applies to operators in multiple countries, each of which has many millions of customers, so is a fair reflection of the envirnment in which any DNS discovery process needs to operate.

- Non-standard CPE - Whilst many ISPs provide their customers with CPE, some customers will elect to use alternative equipment which will not accept software upgrades

- Multiple hardware variants - ISPs typically endeavour to maintain support for legacy CPE. Upgrading the CPE software therefore requires complex and lengthy quality assurance processes to accommodate the many device variants, with some of the larger ISPs having 20-30 variants of devices.

- Large, dispersed customer bases - Cycle times to replace CPE are lengthy due to the costs and numbers involved, and the outcome of any upgrade programme is uncertain due to the aversion of some customers to replace their existing equipment

In summary, the timeframe for non-critical software updates of ISP-supplied CPE can be lengthy. In addition, any such upgrades will only apply to the ISP-supplied CPE so will at best only ever reach between 60-80% of the customer base for many of the largest European ISPs. A replacement programme will also take an extended period without a guaranteed outcome, and that is without considering the cost.

5.3. The Advantage of Supporting Forwarding

The above is intended to illustrate why it is more effective to ensure that DNS discovery methods, including those that support the SPAU model, are developed that work with the hardware and software environments in common use by network operators.

6. Alternative Solutions

Some may be tempted to suggest that the simplest solution to address the issues noted in this document would be for users to connect to global DNS resolvers. Aside from the loss of functionality and significant reduction in user choice that this would imply, it would also result in the further, forced, centralisation of Internet infrastructure, a policy outcome which is out of scope for the ADD working group. It would also, of course, result in the personal data of very large numbers of users to shared with additional parties simply to provide encrypted DNS functionality.

A better approach would be to address the underlying issues so that client software is able to auto-discover and connect to encrypted resolvers on existing network wherever these are available, giving users a seamless upgrade, maintaining full functionality and maximising choice.

7. Extending the Use Case

TO DO

The information in this document is largely based on input from a number of large European network operators, augmented with knowledge of the operations of others, mainly in Europe but with some from North America. It would be beneficial to extend this document with data from additional ISPs to complement the existing content and also to extend the narrative with examples of additional working practices relating to the operation of DNS where possible. This would provide valuable information to inform the development of DNS discovery approaches that will benefit a far broader set of users than would otherwise be possible.

To this end, additional contributions are welcomed as these would ensure that the document is fully representative of the significant use cases globally.

8. Acknowledgements

In addition to the authors, this document is the product of an informal group of experts including the following people:

Andy Fidler, BT plc

Neil Cook, Open-Xchange

Nic Leymann, Deutsche Telekom

Ralf Weber, Akamai

Vittorio Bertola, Open-Xchange

9. Informative References

[ADD] IETF, "Adaptive DNS Discovery (ADD) Working Group", February 2020, <https://datatracker.ietf.org/wg/add/about/>.

[EDDI] EDDI, "Encrypted DNS Deployment Initiative", July 2020, <https://www.encrypted-dns.org/>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

[RFC2535] Eastlake 3rd, D., "Domain Name System Security Extensions", RFC 2535, DOI 10.17487/RFC2535, March 1999, <https://www.rfc-editor.org/info/rfc2535>.

[RFC2671] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, DOI 10.17487/RFC2671, August 1999, <https://www.rfc-editor.org/info/rfc2671>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

Authors' Addresses

Andrew J Campling
419 Consulting Limited

Email: Andrew.Campling@419.Consulting
URI: https://www.419.Consulting/

Normen B Kowalewski
Deutsche Telekom

Email: Normen.Kowalewski@Telecom.DE
URI: https://www.Telecom.DE/

Gianpaolo A Scalone
Vodafone

Email: Gianpaolo-Angelo.Scalone@Vodafone.Com
URI: https://www.Vodafone.Com/

Chris Box
BT Group

Email: Chris.Box@BT.Com
URI: https://www.BT.Com/

Alister Winfield
Sky

Email: Alister.Winfield@Sky.UK
URI: https://www.Sky.Com/

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 14, 2021

N. Cook
V. Bertola
Open-Xchange
A. Fidler
BT plc
N. Leymann
Deutsche Telekom
R. Weber
Akamai

July 13, 2020

A Proposal for a DoH Discovery Trial
draft-cook-doh-discovery-trial-00

Abstract

The following document describes a proposal for a trial of an experimental mechanism for the discovery of DNS-over-HTTPS resolvers provided by Internet Service Providers to their customers.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2021.

Copyright Notice

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

The introduction of encrypted DNS transport protocols like DoH (DNS-over-HTTPS [RFC8484]) can provide additional confidentiality to Internet users that need a DNS resolution service to access online resources. Most end-users currently get their DNS resolution service from the Internet Service Provider that also supplies them with Internet access; thus, to promote a straightforward migration path from unencrypted to encrypted DNS transport and to avoid the issues deriving from a change of DNS provider, it would be useful to establish a mechanism through which stub DNS resolvers on the user's device can discover under appropriate security conditions whether the local network provides a DoH resolver, and if so, start using it automatically. This DoH deployment model will be referred to as "same provider auto-upgrade".

This document describes an experimental mechanism which was developed by a group of Internet Service Providers and DNS implementers for that use case, based on the use of a DNS query for a special use domain name. It is intended as an informational document to support and encourage other parties to join the experiment.

2. Requirements Notation and Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Throughout this document, values are quoted to indicate that they are to be taken literally. When using these values in protocol messages, the quotes MUST NOT be used as part of the value.

3. Rationale

The IETF ADD Working Group was approved by the IESG in February 2020; an extract from the charter [ADD] follows:

"This working group will focus on discovery and selection of DNS resolvers by DNS clients in a variety of networking environments, including public networks, private networks, and VPNs, supporting both encrypted and unencrypted resolvers. It is chartered solely to develop technical mechanisms. Making any recommendations about specific policies for clients or servers is out of scope."

To support the achievement of this technical objective, non-technical considerations also come into play. There is a desire to maximise the number of users that can enjoy an easy upgrade path towards DNS encryption, by making it possible for customers of ISPs that deploy DoH interfaces to their resolvers to get upgraded automatically.

The early discovery mechanisms implemented by some browsers cannot cope with home networks advertising the CPE's private IP address as the endpoint for the local DNS resolver, and thus would exclude the fixed broadband and home mobile router customers of a significant number of ISPs (including the major ones in Europe) from access to this new technology, depending on their Internet access network architecture and on their ease of upgrade of CPEs.

Additionally, in Europe regulation requires all ISPs to allow users to connect to the Internet with any home router of their choice, so it is not even possible for ISPs to prevent the use of home routers that do not support any other DNS resolver mode than dnsmasq over a private IP address.

Regardless of the outcome of IETF and policy discussions, it is likely that any fully fledged, standard discovery protocol will take a relatively long time to reach consensus. Therefore this document proposes an interim solution, which combines in-band discovery with out-of-band protections such as those already used by Google Chrome and Microsoft Windows (i.e. pre-vetting of DoH services, plus some additional protections as discussed below).

This solution would allow participating ISPs using DNS forwarders in their CPEs to provide DoH resolver services to their users in a short timeframe, as long as they used clients (browsers and operating systems) that participate in this trial.

The proposed solution is described in the following section.

4. The Proposed DoH Discovery Trial

4.1. How ISPs Join the Trial

Out-of-band (e.g. through the already established process, or through direct contact at industry venues such as EDDI [EDDI]), ISPs make it known that they would like client vendors to discover their DoH service, but have a significant proportion of users who are using CPEs which act as forwarders.

Each participating vendor, depending on their own security policies, decides if they are fine with an open DNS-based discovery of the local resolver, or if they want to reduce the potential attack

surface by restricting the resolvers that the local network can advertise. Section 5.2. describes the security implications of such a choice.

In the latter case, participating ISPs, regardless of whether they plan to offer public DoH services, guarantee that they (also) offer a DoH service on a URI which is closed, i.e. only accessible from their own network and not from the Internet, and whose hostname is located within a domain name owned by the ISP; this is the DoH service that can be enrolled in the program for vendors that require such restriction. We will refer to these DoH services as "closed resolvers".

This restriction prevents malicious actors from switching a user's DNS resolution to an off-net DNS resolver which is also a trial participant. (However it does not prevent switching from an off-net to an on-net resolver; see section 5). It is up to each DoH client vendor whether they choose to validate (once or continuously) such a guarantee.

The ISPs then provide the vendors with the URI(s) of their (optionally closed) DoH service(s). The URI must contain an FQDN; IP addresses are not acceptable. These URIs are added to a list maintained by the DoH client vendor. For the purposes of this document, we shall call this list the "whitelist" of DoH servers corresponding to ISP resolvers reached via CPEs; it could be a different list from the list of public DoH services used by the current auto-upgrade mechanisms.

4.2. Proposed DoH Resolver Discovery Logic

This process only starts if the configured Do53 resolver is a private ([RFC 1918]) IPv4 or IPv6 link local or unique local address. The auto-upgrade of resolvers with public IP addresses is outside of the scope of this document, though participating vendors, if they want, can use this mechanism for that purpose as well.

The client performs over Do53 (traditional DNS) a TXT record lookup for dohresolver.arpa, a specifically chosen special use domain name (SUDN) [RFC6761]. Eventually, if this mechanism gains adoption, it may be appropriate to register this name with IANA, but we do not anticipate any problems using it on an interim basis, since it is restricted to specific resolvers and does not affect the wider DNS or the arpa TLD. Also, it would be easy to move to any other SUDN that might be standardized by the IETF.

The resolver is configured to respond to that SUDN with a TXT record containing the URI template of the DoH service.

If the Do53 response is anything other than a TXT answer, the discovery is terminated. As a note, some browser makers reported that they sometimes have difficulties with performing lookups on DNS records other than A/AAAA. If CPE routinely filter/drop TXT record lookups, then this approach will not work. To our knowledge, none of the CPE of the ISPs providing data for this document do any modification or filtering of TXT records, and common forwarding software such as dnsmasq does not appear to have issues with arbitrary RRs. Any more facts on this topic would be useful.

In the event of no response being received, the client should decide its own retry policy for the dohresolver.arpa query, but we recommend one or more retries are performed to mitigate packet loss or temporary high load.

In the event of a successful response, the client - if so desires - can check whether the URI in the response matches one of the (optionally closed) DoH URIs that have been added to the "whitelist", and discard it if not.

In the event of a successful response which points to an acceptable DoH resolver, it is up to the DoH client vendor what happens next, for example:

- Auto-upgrade takes place - i.e. connection is attempted to that URI. Assuming that certificate validation and TLS handshake succeeds etc., resolution switches to the DoH service, otherwise the client continues to use the Do53 service.

- The user is presented with a dialog asking them if they'd like to use the newly discovered DoH server. If they accept, then connection proceeds as above.

- The DoH server is added to a list of manually selectable DoH servers.

- Any other suitable logic, e.g. ignoring the response for policy reasons.

The DoH client should respect the TTL of the TXT record returned, and perform a new DNS lookup upon expiry.

4.3. Effect on Possible Use Cases

The basic policy principle for the existing auto-upgrade methods is to avoid changing the resolver that the user has chosen either explicitly (i.e. through manual configuration) or implicitly (e.g. via DHCP).

This logic attempts to preserve that as far as practically possible, through use of the TXT record lookup; the lookup will only return a valid answer if the resolver being used has actively created an authoritative TXT record for the dohresolver.arpa domain. This assumes no malicious actors; see section 5 for security considerations.

We will now discuss the impact of running such an experiment in real world use cases, showing the behaviour that will occur for customers of participating ISPs using participating clients if the DoH client follows the logic described in this proposal. The four possible scenarios for a consumer setup cover all the possible variations of resolver/forwarder configuration and downstream resolver. If either the ISP or the client does not participate in the experiment, no auto-upgrade will ever happen.

This is the effect of the proposed logic in the different scenarios:

1. The user is using an ISP-supplied CPE, which forwards Do53 traffic to the ISP's Do53 resolver. The ISP's Do53 resolver will return a TXT record for dohresolver.arpa, and thus auto-upgrade will take place. The client will then bypass the forwarder, directing queries via DoH directly to the ISP's resolver.

2. The user manually configured a local DNS forwarder themselves (e.g. an off-the-shelf CPE, or they run dnsmasq on a local server) to forward queries to their local ISP resolver. This resolver will return a TXT record for dohresolver.arpa, and thus auto-upgrade will take place, bypassing the forwarder, exactly like in the previous case.

3. The user has manually configured a local DNS forwarder themselves (e.g. an off-the-shelf CPE, or they have modified the ISP-provided CPE) to forward to a resolver that is not the ISP resolver, e.g. 1.1.1.1 or 9.9.9.9. These resolvers will return NXDOMAIN for dohresolver.arpa, and thus no auto-upgrade will take place.

4. The user has manually configured a local DNS resolver themselves (e.g. Raspberry PI or similar). This resolver will return NXDOMAIN for dohresolver.arpa, and thus no auto-upgrade will take place.

From the DoH client's perspective, all four scenarios are the same (i.e. the system resolver has a RFC1918 IPv4 or IPv6 link local or unique local address), and whether that resolver was configured via DHCP or manually would not appear to matter that much. Scenarios 3

and 4 can be discounted because no action takes place, however scenarios 1 and 2 do have the effect of bypassing the forwarder.

There is a semantic difference between scenarios 1 and 2; in scenario 2 the user may have configured a forwarder deliberately, for example to do filtering, caching, logging or innumerable other reasons. For that reason, DoH client vendors need to consider whether the above scenarios, (or any additional scenarios not considered above), justify asking for the user's consent to the upgrade to DoH directly to the ISP resolver (and thus bypassing any intermediate forwarders).

Moreover, in cases 3 and 4 it would be easy for the operator of the alternative resolver (whether it is another DNS operator, as in case 3, or the user themselves, as in case 4) to also allow auto-upgrade to DoH of the connection, simply by configuring their own resolver to reply to the query for dohresolver.arpa. However, if the client vendor chooses to restrict the auto-upgrade mechanism only to whitelisted URIs, then these other operators would need to also join the experiment; in case 3, this could be made impossible by the restriction itself, as by definition the resolver in this case is an "open" one, and the vendor would need to partially lift the restriction and accept known open DoH resolvers from a separate whitelist; in case 4, this could be made impossible by the non-technical requirements of the procedure for joining.

5. Security Considerations

5.1. Existing Auto-Upgrade Mechanisms

The security objective for current same-provider auto-upgrade mechanisms is to ensure that the client is talking to the same resolver operator as before, but now over DoH. On-path attackers have no way to influence this, since the mechanism is based on the client knowing the public IP address of the existing resolver and a pre-configured URI template for the auto-upgrade DoH resolver for that IP address.

5.2. Current Proposal

This proposal does not use the same threat model as the existing auto-upgrade solution. The differences are discussed below.

On-path attackers could perform a downgrade attack. However, given that current auto-upgrade mechanisms do not work for users with forwarders in their CPE, such a downgrade attack would result in the same situation as currently, i.e. the client would continue to use Do53. However, given this threat, the upgrade to DoH can only be considered as opportunistic security.

On-path attackers could change the discovery response from that returned by the actual configured resolver. There are two scenarios that need to be considered, depending on the vendor's policy.

If the DoH client vendor enforces the "closed resolver" restriction, then the following applies:

- Given that the client will only accept auto-upgrade via discovery to a "closed resolver", there is only one resolver that will be accepted by the client via the discovery mechanism - the DoH resolver offered by the user's ISP. (This assumes that the ISP is diligent about ensuring their resolver is actually closed - this could be verified periodically by the vendors).

- There exists a risk that an on-path attacker could redirect a user from a manually selected resolver (configured manually by the user on the CPE/forwarder) to the resolver provided by the local ISP. This would have the effect of moving the user from a resolver that they did select (e.g. 9.9.9.9) to one they did not select. Such a risk is not mitigated by this proposal. Note that this risk exists only when there is an on-path attacker, since the discovery query happens via DNS and thus goes to the resolver originally chosen by the user.

If the DoH client vendor does not enforce the "closed resolver" restriction, then the following applies. An on-path attacker could redirect a user from a manually selected resolver (configured manually by the user on the CPE/forwarder) to any resolver on the "whitelist" of DoH servers. This would have the effect of moving the user from a resolver that they did select to one they did not select. Such a risk is not mitigated by this proposal. Note that this risk exists only when there is an on-path attacker, since the discovery query happens via DNS and thus goes to the resolver originally chosen by the user.

In both of the above use cases, the only possible end result of a successful attack aimed at changing resolvers is that a user has moved from an insecure Do53 service whose results are controlled by a malicious on-path attacker, to a secure DoH service which is on the DoH client's "whitelist". The on-path attacker has only succeeded in moving the user to a vendor-verified resolver over which they have no control, and which they cannot use for further attacks; as long as the vendor's whitelisting process is secure, an attacker wishing to gain control of the user's DNS resolution process for further steps, e.g. to redirect the user's Web requests to a phishing page, would not be able to do so through this attack. This would apply even if the new DoH resolver were open, as long as it is on (the same or another) vendor-verified whitelist. However, the user has still

moved to a resolver operated by a different organization which is almost certainly not what the user "wanted"; hence the possible need for user confirmation.

The security assumptions regarding the "closed resolvers" above are predicated on the participating ISPs performing the appropriate actions to "close" their resolver(s) to the public internet, thus making them only available to customers on their network. DoH client vendors relying on the security assumptions provided by this may wish to make periodic checks (see section 6) to ensure that listed DoH resolvers are indeed not accessible from the public internet, otherwise new attacks would be possible such as an on-path attacker redirecting a user from their currently selected resolver to the resolver of another participating ISP.

In the end, vendors that adopt the approach of vetting and whitelisting DoH resolvers before allowing users to auto-upgrade to them will always enjoy a certain degree of reassurance on the legitimacy of those resolvers (though at the expense of excluding the users of other DoH resolvers, including self-managed resolvers that people may install on their home networks, from automatic upgrade).

Without the additional "closed resolver" restriction, an attacker may succeed in redirecting the user to any of the whitelisted DoH services, while with that additional restriction, an attacker may only succeed in redirecting the user to the ISP’s own closed DoH service, if the user is not already using it. Whether this gain in security is worth the additional organizational complexity is for each vendor to consider; we expect that running this experiment could also allow to evaluate how useful that additional restriction could be in practice.

6. Implications for Vendors

The experiment requires participating vendors to change their current implementation of the auto-upgrade mechanism and add the logic described.

For DoH client vendors enforcing the "closed resolver" restriction, some additional vetting and active checking of "auto-upgrade" DoH providers would be necessary, to ensure that ISP resolvers are indeed "closed" and only accessible to customers on their own networks, as assumed in Section 5 above. This could take the form of periodic attempts to connect to all the DoH URIs on the "whitelist" from a variety of locations known to be outside of any service provider networks. It would be up to the organisation responsible for the DoH client to decide how stringent this check should be. For example, it

may involve automated weekly checks, and alerts to ISPs whose resolvers do not meet the required standards.

DoH client vendors who also support the "auto-upgrade based on public resolver IP" logic need to maintain two "whitelists"; DoH servers could of course be on both lists, or both lists could be merged into one with additional parameters for each featured resolver, as preferred.

7. Acknowledgements

This document is the product of an informal group of experts including the following people:

Alister Winfield, Sky

Andrew Campling, 419 Consulting

Andy Fidler, BT plc

Chris Box, BT plc

Gianpaolo Scalone, Vodafone

Neil Cook, Open-Xchange

Nic Leymann, Deutsche Telekom

Norman Kowalewski, Deutsche Telekom

Ralf Weber, Akamai

Vittorio Bertola, Open-Xchange

The authors would like to thank Kenji Baheux and Eric Orth (Google) and Tommy Jensen (Microsoft) for their feedback and suggestions.

8. References

8.1. Normative References

[ADD] IETF, "Adaptive DNS Discovery (ADD) Working Group", February 2020, <https://datatracker.ietf.org/wg/add/about/>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <https://www.rfc-editor.org/info/rfc6761>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

8.2. Informative References

[EDDI] EDDI, "Encrypted DNS Deployment Initiative", July 2020, <https://www.encrypted-dns.org/>.

Authors' Addresses

Neil Cook
Open-Xchange Ltd
7 Gerard Street
Ashton-in-Makerfield, Wigan, Greater Manchester WN4 9AG
United Kingdom

Email: neil.cook@noware.co.uk
URI: https://www.open-xchange.com/

Vittorio Bertola
Open-Xchange Srl
Via Treviso 12
Torino 10144
Italy

Email: vittorio.bertola@open-xchange.com
URI: https://www.open-xchange.com/

Andy Fidler
BT plc
BT Adastral Park
Martlesham Heath, Ipswich IP5 3RE
United Kingdom

Email: andrew.fidler@bt.com
URI: https://www.bt.com/

Nicolai Leymann
Deutsche Telekom AG
Friedrich-Ebert-Allee 140
Bonn 53113
Germany

Email: N.Leymann@telekom.de
URI: https://www.telekom.com/

Ralf Weber
Akamai Technologies GmbH
Parkring 20-22
Garching 85748
Germany

Email: ralf.weber@akamai.com
URI: https://www.akamai.com/

Adaptive DNS Discovery (ADD)
Internet-Draft
Intended status: Informational
Expires: January 14, 2021

G. Deen
Comcast-NBCUniversal

July 13, 2020

Adaptive DNS Discovery Threats Here
draft-deen-add-threats-00

Abstract

DNS resolver discovery is designed to operate under a variety different levels of trust in the underlying network. This document describes the various trust types that DNS resolver discovery and selection may take place under. Internet Draft.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2021.

Table of Contents

1. Introduction

There are a variety of network environments users may interact with where they will be discovering and selecting a DNS resolver each of which presents a different threat level to the user. This document attempts to establish a common set of threats classifications for reference by Adaptive DNS Discovery (ADD) working group drafts.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Classifications

2.1. Approach

There are many ways to classify and structure threat analysis the approach used here is centered on the perspective of the user and how much subjective trust they can place in different access network situations that they may encounter.

2.2. Green or Trusted Networks

These are networks in which the user has an high sense of trust. These are networks run by a trusted party who is known to the user and is trusted by the user to operate the network with security and operational integrity. While even the best run network can be compromised by attackers or malware, the user has subjective trust that the Green network is very unlikely to be compromised.

The user often has a relationship with the network operator - either personally, as an employee, or by contract they user has entered into such as with an ISP or Mobile Carrier.

Examples of Green Networks

- o User's own home network
- o User's organization, company, or enterprise network
- o Mobile user's mobile network
- o User's ISP network

2.3. Yellow or Unknown Networks

These are networks in which the user does not have any sense of trust and yet has no sense or expectation that the network maybe compromised or hostile. The network's threat level is simply unknown.

These are networks which provided a service to visitors such as public Wifi networks.

Examples of Yellow Networks

- o School network
- o Cafe or coffee shop network
- o Airport network
- o Hotel network
- o Conference or event network

2.4. Red or Hostile Networks

These are networks in which the user has an high sense of potential threats being present, but the use may have no other choice but to use them.

These are networks which the user not only does not trust, but also expects the network maybe doing things that the user does not want.

Red Networks

- o War zone region network

- Hostile regime network

3. IANA Considerations

This memo includes no request to IANA.

All drafts are required to have an IANA considerations section (see Guidelines for Writing an IANA Considerations Section in RFCs [RFC5226] for a guide). If the draft does not require IANA to do anything, the section contains an explicit statement that this is the case (as above). If there are no requirements for IANA, the section will be removed during conversion into an RFC by the RFC Editor.

4. Security Considerations

All drafts are required to have a security considerations section. See RFC 3552 [RFC3552] for a guide.

5. References

5.1. Normative References

[Contributors] Deen, G., "Authors", 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

5.2. Informative References

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/info/rfc3552>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <https://www.rfc-editor.org/info/rfc5226>.

Appendix A. Additional Stuff

This becomes an Appendix.

Author's Address

Glenn Deen
Comcast-NBCUniversal
Universal City, California 91608
USA

Email: glenn_deen@comcast.com

add
Internet-Draft
Intended status: Informational
Expires: November 30, 2020

D. Migault
Ericsson

May 29, 2020

DNS Resolver Discovery Protocol (DRDP)
draft-mglt-add-rdp-02

Abstract

This document describes the DNS Resolver Discovery Protocol (DRDP) that enables a DNS client to discover various available local and global resolving service. The discovery is primarily initiated by a DNS client, but a resolving service may also inform the DNS client other resolving services are available and eventually preferred.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 30, 2020.

Table of Contents

1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Introduction

A DNS client can proceed to DNS resolution using various resolving services. These services can be local or global and can use a wide range of DNS transport protocols such as, for example, standard DNS [RFC1035], DNS over TLS [RFC7858] or DNS over HTTPS [RFC8484]. The local scope of these services may take various forms. For example, it could be associated to a network perspective (restricted to the network the DNS client is connected to) or to an application perspective (restricted to some domain names).

The purpose of the DNS Resolving service Protocol (DRDP) is to discover resolving services available to the DNS client. These available resolving services to a given DNS client may highly depend on its location or browsing activity. The number of resolving services available to the DNS client is expected to remain quite consequent and evolve over time. Similarly, characteristics associated to these resolving services may also evolve over time. As a result, the DNS client is unlikely willing to synchronize such a huge data base of resolving services. DRDP proposes an alternative that consists in adaptively discovering the available resolving services based on the DNS client context.

DRDP adopts a hierarchical approach where the DNS client gets _resolving domains_ from the context. These _resolving domains_ are entry points for resolving services (associated to each of these resolving domains).

The DNS client may obtain the contextual resolving domains via various way, including a configuration or via DHCP Options [I-D.btw-add-home].

This document describes two mechanisms for a DNS client to retrieve resolving domains. Firstly, it is envisioned that these resolving domains will be provided by multiple third party providers which could designate a set of resolving domains. This set is designated by a pointer used by the DNS client to retrieve the resolving domains.

Secondly, resolving domain may be derived from the IP address of the legacy resolving service provided via the Recursive Name Server option [RFC3646]. Such a resolving domain can be seen as a network local scope resolving domain. This resolving domain may then be used by the DNS client to discover he various flavors of resolving services provided by the ISP (DoH, DoT for example), while the legacy IP address provided is reserved to the legacy resolving service.

The discovery process is expected to be followed by a selection process by which the DNS client selects the resolving service it is willing to use for the DNS resolution of the end user or application. How the selection is performed is out of scope of this document.

3. Terminology

DNS client the client that sends DNS queries fro resolution. In this document the DNS client designates also the end entity that is collecting information about the available Resolving Services and then proceed to the selection of a subset them. The selection is processed according to the DNS client’s policy.

Resolving Service designates a service that receives DNS queries from a DNS client and resolves them. A Resolving Service is implemented by one or multiple resolvers.

Resolver: A resolver designates the software or hardware handling the DNS exchange. See [RFC7719] for more details.

DNS transport designates the necessary parameters a DNS client needs to establish a session with a Resolving Service.

Resolving Domain a DNS domain that hosts one or multiple resolving services.

4. DRDP Requirements

This section lists the DRDP requirements.

REQ 1: DRDP MUST enable a DNS client to discover the available resolving services with their associated characteristics in order to proceeds to a selection process. The selection process takes resolving services identities and associated parameters and proceed to the selection.
Any sort of resolving service selection is outside the scope of DRDP.

REQ 2: While the discovery process is triggered by the DNS client, a third party MUST be able to provide necessary input information so a resolving service discovery process can be triggered within a specific context.
Provisioning protocols to provide this information is not as per say in scope of DRDP. DRDP defines the format of the format for such input as well as a set of such inputs.

REQ 3: Any information used in DRDP MUST be authenticated by its owner. In particular, the characteristics associated to the resolving service MUST be certified by the resolving service operator / owner and MUST be associated a validity period. In addition, a third party providing a set of inputs MUST authenticate that set.

REQ 4: Information associated to the resolving services is intended to enable the selection process that is assumed to match the end user or application policy. The selection process is out of scope of DRDP. Information may carry some characteristics of a resolving service or hints that will help the selection. In particular an operator of multiple resolving service MUST be able to associate a preference to the proposed resolving services. To ease automation of the selection as well as to make multiple applications benefit from DRDP the information MUST be carried over a standardized format.

REQ 5: DRDP MUST be designed to be used indifferently by a DNS client using any DNS transport protocol (Do53, DoT, DoH, ...). However, DRDP SHOULD be able to restrict the information retrieved to a certain type of resolving service, i.e. Do53, DoT, DoH.

REQ 6: DRDP deployment MUST NOT be disruptive for the legacy DNS client or infrastructure and legacy client SHOULD be able to incrementally include DRDP.

5. Resolving Domains

The resolving domain is the input of a discovery process. Section Section 5.1 defines the format resolving domain and exposes why resolving domains seem convenient pointers to resolving service as well as how the relates to resolving service identities. Section Section 5.2 defines the format of a pointer to a set of resolving domains as well as how to retrieve how to handle such set. Such pointer are expected to be used by third party providers to indicate a subset of resolving domains that match a certain context. The use of a pointer is expected to ease the management of the set as opposed to a explicit list. The definition of such a format is expected to favor interoperability with third party providers.

Finally, section Section 5.3 defines a procedure to derive a resolving domain from the IP address provided by Recursive Name Server option [RFC3646]. Such procedure is expected to leverage from the existing and legacy infrastructure.

5.1. Resolving Domain and Resolving Service Identity

A resolving service is identified by a FQDN - such as resolver.example.com - and the domain part (example.com) is designated as the resolving domain. Note that the hostname (resolver) is only considered as a way to distinguish different resolving services but it is not expected to carry any specific information that will be useful for the selection process.

The resolving domain is expected to be representative to the end user of the brand or legal entity of the institution the end user sends its data to. The end user is likely to select a given resolving domain based on the trust he puts into the associated legal entity. The resolving domain follows some DNS encoding rules and as such may not be believed to be so user friendly. On the other hand, the end user may also be familiar with that format and the use or a standardize format helps automation in the selection. Typically, it might be ericsson.com or ericsson which is different from Ericsson (with appropriated police character and color) which would be more meaningful for the end user. Note that a user interface may also use

the resolving domain to derive more user friendly and additional specific information that will be presented to the user. This could include for example additional RDAP queries, favicons of web sites that are shown to the end users. What is presented to the end user is out of scope of this document, but the resolving domain can be used as the key.

The hostname part is only meaningful within the resolving domain. While, it may carry some information that may be interpreted to the end user, the constraint provided by the DNS format may be too restricting. As a result, it is expected that a more user friendly string might be associated with the hostname and that the hostname remain reserved for networking administrators.

5.2. List of Resolving Domains

A resolving domain list is designated by a FQDN example.com and the resolving domains contained in that list are retrieved by sending a query of type PTR for b._dns.example.com.

The zone file below is inspired from DNS-SD where b indicates a browsing domain, _dns indicates the DNS resolving service, example.com designates the list of the resolving domains. resolving domain_0, resolving_domain_n indicates the various resolving domains. The order of the resolving domains is irrelevant, and the zone administrator SHOULD regularly reorder them. The RRsets MUST be signed with DNSSEC.

```
b._dns.example.com  PTR <resolving_domain_0>
[...]
b._dns.example.com  PTR <resolving_domain_n>
```

Using the DNS provides the advantage to retrieve the resolving domain without requiring other libraries than DNS as well as benefit from the DNS caching infrastructure including the use of the TTL.

An EDNS buffer size of 1232 bytes will avoid fragmentation on nearly all current networks. This is based on an MTU of 1280, which is required by the IPv6 specification, minus 48 bytes for the IPv6 and UDP headers. Such size makes lists of a 100 names viable over UDP without fragmentation. Larger lists will require the DNS exchange to be performed over TCP. While there is no hard limits, downloading the full list every TTL may not be appropriated for very large lists where the synchronization mechanisms may be needed.

The current size of such lists [curl][dnsprivacy.org] have less than 50 resolving domains. Other lists such as [public-dns.info] have as

much as 11.000 entries, but such lists seems to contain open resolvers which is out side of the scope of a selection process. Web browser (Google Chrome) also have lists over 10.000 entries, but in case a significant number of entries seems to be IP addresses that have a very limited network scope (e.g. limited to the ISP). The length of the list in scope to the DNS client is in fact significant smaller in term of IP addresses and even smaller if resolving domain are able to represent multiple IP addresses. Overall, the size of such lists are currently due to the absence of discovery protocols.

5.3. Local Network Resolving Domain

Resolving service are currently configured or advertised via IP addresses rather than a FQDN as a DNS resolution would be needed to resolve the IP address. More specifically, networks usually advertise the resolving service via a Recursive Name Server option [RFC3646] that contains an IP address. Similarly application usually configures their resolving services with IP addresses (8.8.8.8, 1.1.1.1, 9.9.9.9,...). As a result, this section indicates a mechanism that would enable a DNS client to derive a resolving domain of a resolver from an IP address of an advertised resolver. The mechanism described here is expected to be used as an hint.

The resolving domain will be derived from the IP address by:

1. performing a reverse resolution

2. assume the resulting FQDN is composed of a hostname appended to the resolving domain. For example, if resolver.example.com is the resulting FQDN from the reverse resolution, then the rdns domain will be example.com.

In most cases local resolving services uses global IP address which does not limit the reverse resolution to an associated local resolver. However the zone associated to the resolving domain might not be available globally and instead be restricted to the local network. As a result, DNS client SHOULD perform DNS resolution associated to the local resolving domain using the local resolver, and resolving service operator SHOULD publish the resolving domain zone to the global Internet.

Legacy DNS client will not be impacted. Upon receiving the IP address they will send their DNS queries to that IP address. DRDP aware DNS client will derive the resolving domain and attempt to perform a discovery within the resolving domain.

If other mechanisms as used to advertise the resolving domains such as those described in [I-D.btw-add-home], and the resolving domain

are different, the DNS client should perform DRDP with both resolving domains.

6. Resolving Service Discovery

6.1. Discovery of all service instances

Given a resolving domain example.com, a DNS client MAY request all possible resolving service instances with a query of type SVCB with the service _dns.

The example below presents the use of an AliasForm followed by a ServiceForm which allows an indirection. The Alias form is not madatory and instead only ServiceForm associated to _dns.example.com could have been used instead.

The SvcFieldPriority indicates the preference of the resolving service instance.

The SvcParamKey alpn MUST be present when TLS is used as its presence and value indicates the DNS transport. The absence of the alpn SvcParamKey indicates Do53, alpn set to dot indicates DoT is served while h* indicates DoH is served. Note that the port value (53, 853, 443) is not used to determine the DNS transport as non standard port MAY be used. The example below uses an non standard port 5353 for illustrative purpose.

Other SvcParam are detailed in Section 6.4 and are optional. A SvcParam not understood by the DNS client MUST be ignored.

The RRsets MUST be protected with DNSSEC and when alpn is provided a TLSA RRset SHOULD be present. These RRsets have been omitted for clarity.

```
## Discovery of all service instances
_dns.example.com. 7200 IN SVCB 0 svc.example.com.
svc.example.com.    7200 IN SVCB 12 ( svc0.example.net.
                                      port="5353" user-display="Legacy Resolver"
)
svc.example.com.    7200 IN SVCB 1 ( svc1.example.net.  alpn="dot"
                                      port="5353" esniconfig="..."
                                      user-display="Preferred Example's Choice" )
svc.example.com.    7200 IN SVCB 3 ( svc2.example.net. alpn="h2"
                                       port="5353" esniconfig="..." user-display=
 )
svc.example.com.    7200 IN SVCB 2 ( svc3.example.net. alpn="h3"
                                       port="5353" esniconfig="..." user-display=
 )
```

6.2. Discovery of specific service instances

To reduce the size of the messages, the DNS client MAY also prefer to query information of resolving services using a specific transport (DNS, DoT, DoH) that are designated as sub sets. A DNS client MAY list the different subsets of that resolving domain with a PTR query. This document defines the following subsets _53._dns for DNS, _853._dns for DoT and _443.__dns for DoH. Other subsets MAY be defined in the future. A DNS client that does not understand a subset SHOULD ignore it and maybe proceed to the discovery as defined in Section 6.1.

All subsets MUST share the same resolving domain and be listed with a PTR RRsets. The DNS client MAY NOT performed a DNS query of type PTR, for example, if it has a previous knowledge of the existence of the subset or if indicated by its policy. In this it MAY directly proceed to the SRVCB resolution.

The same restrictions as defined in section Section 6.1 apply.

Note that while the SvcFieldPriority indicates the priority within a subservice, this field MUST have a coherence across subservices. The priority provided SHOULD be coherent with the case of a _dns SRVCB query of section Section 6.1.

The figure below illustrates an example of zone file. RRSIG and TLSA have been omitted for the purpose of clarity.

```
### Definition of the resolving service subsets
_dns.example.com PTR _53._dns.example.com
_dns.example.com PTR _853._dns.example.com
_dns.example.com PTR _443._dns.example.com

### services instances per service subset
_53._dns.example.com. 7200 IN SVCB 0 svc0.example.com.
svc0.example.com.    7200 IN SVCB 12 ( svc0.example.net.
                                       port="5353" user-display="Legacy Resolver"
 )
_853._dns.example.com.    7200 IN SVCB 0 svc1.example.com.
svc1.example.com.    7200 IN SVCB 1 ( svc1.example.net.  alpn="dot"
                                      port="5353" esniconfig="..."
                                      user-display="Preferred Example's Choice" )

_443_dns.example.com.    7200 IN SVCB 0 svc4.example.net.
svc4.example.com.    7200 IN SVCB 3 ( svc2.example.net. alpn="h2"
                                       port="5353" esniconfig="..." user-display=
 )
svc4.example.com.    7200 IN SVCB 2 ( svc3.example.net. alpn="h3"
                                      port="5353" esniconfig="..."
                                      user-display="Testing QUIC")
```

Some notes:

1. _domain uses SVCB but does not have TLS. While SVCB has been created essentially for TLS based service, this does not appear to be mandatory.

2. Should we have some constraints regarding the SvcDomainName and QNAME ?

3. do we need the service subsets

6.3. TTL

The DNS client SHOULD perform DRDP at regular intervals as indicated by its policy.

The selection process MAY remove resolving services with short TTL lower than a day as it indicates some source of instalbility. Given a subset of selected resolving services, the DNS client may perform DRDP 1 hour before an SVB RRset expires.

6.4. SvcParamKey

This section defines a set of SvcParamKey that MAY be use to carry the necessary informations for the selection process.

alpn :

esniconfig :

port :

user-display indicates a strings in UTF-8 that is expected to be representative to a potential end user. Though there is no restriction in the scope of that string. The string is likely to represent the service within the resolving domain.

uri_template designates the URI template for DoH. This key MUST NOT be present on non DoH services and MUST be ignored by the DNS client on non DoH resolving Services.

auth_domain indicates the list of authoritative domain name the resolving service has strong relation with. It is expected that a resolving service may prefer to perform DNS resolution over these domains to that specific resolving service as to preserve its privacy. This information MUST be verified and validated.

filtering indicates the presence of a filtering service

ip_subnet indicates a subnetwork restriction. This is mostly useful for resolving services that are not globally.

dnssec indicates whether dnssec is enabled or not.

7. Resolver advertising other service sub type

A resolving service receiving a DNS request over a service sub type MAY be willing to advertise the DNS client that other sub service type are available. This is especially useful, when, for example, a resolver wants that the DNS resolver switches to other service sub types that are more secure.

In order to do so the resolver MAY provide in the additional data field the _dns SRVCB of ServiceForm.

8. Migration to service sub types

The principle of the discovery mechanism is that the resolver indicates the available service sub types and let the DNS client chose which sub type it prefers. On the other hand, the resolver MAY also indicate a preference using the priority and weight fields. However, there is no mechanisms that could permit an indirection from one service sub type to another service sub type. This document specifies that weight needs to be considered across sub types. Redirection MAY especially be needed when a DNS client is using the Do53 and the resolver would like to upgrade the DNS client session to a more secure session. This MAY require a specific ERROR code that will request the DNS client to perform service discovery.

It is expected that DRDP MUST always be available via Do53. However, this does not mean that a resolver is expected to implement the Do53 sub type service for a resolving service. If a resolving service provider chooses not to provide a resolving service using Do53, that service MUST NOT be pointed by the _53._dns.example.com search and MUST NOT provide _dns.example.com SRVCB RRsets with no SvcParamKey alpn.

9. Security Considerations

9.1. Use of protected channel is RECOMMENDED

When available, it is recommended to chose a protected version of the rdns service. More specifically, the use of end-to-end protection ensures that the DNS client is connected to the expected platform and that its traffic cannot be intercepted on path. Typically, the selection of resolver on the Internet (and not on your ISP network) and the use of a non protected channel enables an attacker to monitor

your DNS traffic. The similar observation remains true if you are connected to the resolver of your ISP. It is commonly believed that trusting your ISP (that is your first hop) makes encryption unecessary. Trusting your ISP is mandatory in any case, but the associated level of trust with an protected channel is restricted to the operation of the DNS platform. With non protected channel the trust is extended to any segment between the DNS client and the resolver, which is consequently larger. The use of a protected channel is recommended as it will prevent anyone on path to monitor your traffic.

9.2. DNSSEC is RECOMMENDED

The exchanges SHOULD be protected with DNSSEC to ensure integrity of the information between the authoritative servers and the DNS client. Without DNSSEC protection, DNS messages may be tampered typically when they are transmitted over an unprotected channel either between the DNS client and the resolver or between the resolver and the authoritative servers. The messages may be tampered by an online attacker intercepting the messages or by the intermediary devices. It is important to realize that protection provided by TLS is limited to the channel between the DNS client and the resolver. There are a number of cases were the trust in the resolver is not sufficient which justify the generalization of the use of DNSSEC. The following examples are illustrative and are intended to be exhaustive.

First, the discovery exchanges may happen over an unprotected channel, in which case, the messages exchanged may be tampered by anyone on-path between the DNS client and the resolver as well as between the resolver and the authoritative servers - including the resolver. When TLS is used between the DNS client and the resolver, this does not necessarily mean the DNS client trusts the resolver. Typically, the TLS session may be established with a self-signed certificate in which case the session is basically protected by a proof-of-ownership. In other cases, the session may be established based on Certificate Authorities (CA) that have been configured into the TLS client, but that are not necessarily trusted by the DNS client. In such cases, the connected resolver may be used to discover resolvers from another domain. In this case, the resolver is probably interacting with authoritative servers using untrusted and unprotected channels. Integrity protection relies on DNSSEC.

9.3. TLSA is RECOMMENDED

When TLS is used to protect the DNS exchanges, certificates or fingerprint SHOULD be provided to implement trust into the communication between the DNS client and the resolver. The TLS session and the association of the private key to a specific identity

can be based on two different trust model. The Web PKI that will rely on CA provisioned in the TLS library or the TA provided to the DNS client. A DNS client SHOULD be able to validate the trust of a TLS session based on the DNSSEC trust model using DANE.

When the DNS client is protecting its session to the resolver via TLS, the DNS client may initiate an TLS session that is not validated by a CA or a TLSA RRsets. The DNS client MUST proceed to the discovery process and validate the certificate match the TLSA RRset. In case of mismatch the DNS client MUST abort the session.

10. Privacy Considerations

When the discovery protocol is performed using a resolver that belongs to one domain for another domain, or over an unprotected channel, the DNS client must be conscious that its is revealing to the resolver its intention to use another resolver. More specifically, suppose an resolver is complying some legal requirements that DNS traffic must be unencrypted. Using this resolver to perform a resolver discovery reveals the intention of potentially using alternative resolvers. Alternatively, narrowing down the discovery over a specific sub type of resolver (DoT, or DoH) may reveal to that resolver the type of communication. As result, when performing a discovery over a domain that differs to the domain the resolver belongs to, it is RECOMMENDED to request the SRV RRsets associated to all different sub type of proposed services.

The absence of traffic that results from switching completely to a newly discovered resolver right after the discovery process provides an indication to the resolver the DNS client is switching to. It is hard to make that switch unnoticed to the initial resolver and the DNS resolver MUST assume this will be noticed. The information of switching may be limited by sharing the traffic between different resolvers, however, the traffic pattern associated to each resolver may also reveal the switch. In addition, when the initial resolver is provided by the ISP, the ISP is also able to monitor the IP traffic and infer the switch. As a result, the DNS client SHOULD assume the switch will be detected.

With DoT or DoH, the selection of port 443 will make the traffic indistinguishable from HTTPS traffic. This means that an observer will not be able to tell whether the traffic carries web traffic or DNS traffic. Note that it presents an interest if the server offers both a web service as well as a resolution service. Note that many resolvers have a dedicated IP address for the resolution service, in which case, the information will be inferred from the IP address. Note also that traffic analysis may infer this as well. Typically suppose an IP address hosts one or multiple web sites that are not

popular as well as a resolving service. If this IP address is associated frequent short size exchanges, it is likely that these exchanges will be DNS exchanges rather than Web traffic. The size of the packet may also be used as well as many other patterns. As a result, the use port 443 to hide the DNS traffic over web traffic should be considered as providing limited privacy.

11. IANA Considerations

This document requests the IANA the creation of the following underscored node names in the Underscored and Globally Scoped DNS Node Names registry https://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml#dns-parameters-14

RR Type	_NODE NAME	Reference
SRVCB	_dns	RFC-TBD

SvcParamKey	NAME	Meaning	Reference
7	user-display	User friendly string (UTF8) to represent the resolver	RFC-TBD
	uri_template	URI template	
	auth_domain	Domains the resolving service is authoritative	
	filetring	Filetring services provided	
	ip_subnet	ip ranges accepted.	
	dnssec	DNSSEC validation enabled	

12. Acknowledgments

We would like thank Mirja Kuehlewind as well as the GSMA IG for their comments. We also thank Ted Hardie and Paul Hoffman for their feed backs regarding the dns schemes for DoT and DoH.
We thank Ben Schwartz for the comments on the list size. We thank Harald Alvestrand for its recommendation on having a model that enable multiple third party providers to provide their own list of resolving domains.

13. References

13.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3646] Droms, R., Ed., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, DOI 10.17487/RFC3646, December 2003, <https://www.rfc-editor.org/info/rfc3646>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

13.2. Informative References

[curl] "Publicly available servers", n.d., <https://github.com/curl/curl/wiki/DNS-over-HTTPS#publicly-available-servers>.

[dnsprivacy.org] "DNS Privacy Test Servers", n.d., <https://dnsprivacy.org/wiki/display/DP/DNS+Privacy+Test+Servers#DNSPrivacyTestServers-Publicresolvers>.

[I-D.btw-add-home] Boucadair, M., Reddy.K, T., Wing, D., and N. Cook, "Encrypted DNS Discovery and Deployment Considerations for Home Networks", draft-btw-add-home-06 (work in progress), May 2020.

[public-dns.info] "Public DNS Server List", n.d., <https://public-dns.info/>.

[RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <https://www.rfc-editor.org/info/rfc7719>.

Author's Address

Daniel Migault
Ericsson
8275 Trans Canada Route
Saint Laurent, QC 4S 0B6
Canada

EMail: daniel.migault@ericsson.com

add
Internet-Draft
Intended status: Informational
Expires: January 29, 2021

D. Migault
Ericsson

July 28, 2020

DNS Resolving service Discovery Protocol (DRDP)
draft-mglt-add-rdp-03

Abstract

This document describes the DNS Resolver Discovery Protocol (DRDP) that enables a DNS client to discover various available local and global resolving service. The discovery is primarily initiated by a DNS client, but a resolving service may also inform the DNS client other resolving services are available and eventually preferred.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 29, 2021.

Table of Contents

1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Introduction

A DNS client can proceed to DNS resolution using various resolving services. These services can be local or global and can use a wide range of DNS transport protocols such as, for example, standard DNS [RFC1035], DNS over TLS [RFC7858] or DNS over HTTPS [RFC8484]. The local scope of these services may take various forms. For example, it could be associated to a network perspective (restricted to the network the DNS client is connected to) or to an application perspective (restricted to some domain names).

The purpose of the DNS Resolving service Discovery Protocol (DRDP) is to discover resolving services available to the DNS client. These

available resolving services to a given DNS client may highly depend on its location or browsing activity. The number of resolving services available to the DNS client is expected to remain quite consequent and evolve over time. Similarly, characteristics associated to these resolving services may also evolve over time. As a result, the DNS client is unlikely willing to synchronize such a huge data base of resolving services. DRDP proposes an alternative that consists in adaptively discovering the available resolving services based on the DNS client context.

DRDP adopts a hierarchical approach where the DNS client (or DRDP client) discovers the resolving services from resolving domains (RD) or a pointer to a list of resolving domains (Pointer).

The document does not describe how the DNS client is provisioned with RD or RD_list. The DNS client may obtain the contextual resolving domains via various way, including a configuration, via DHCP Options [I-D.btw-add-home] or derived from specific procedures [I-D.mglt-add-drdp-isp]. The DNS client is expected to discover resolving services from all RD or RD_list before proceeding to a selection process. The selection process of the resolving service is out of scope of this document.

3. Terminology

DNS client the client that sends DNS queries fro resolution. In this document the DNS client designates also the end entity that is collecting information about the available Resolving Services and then proceed to the selection of a subset them. The selection is processed according to the DNS client's policy.

Resolving Service designates a service that receives DNS queries from a DNS client and resolves them. A Resolving Service is implemented by one or multiple resolvers.

Resolver: A resolver designates the software or hardware handling the DNS exchange. See [RFC7719] for more details.

DNS transport designates the necessary parameters a DNS client needs to establish a session with a Resolving Service.

Resolving Domain a DNS domain that hosts one or multiple resolving services.

4. Overview

DRDP is a DNS based protocol that addresses the requirements listed in Section 15.1. The figure below represents provides a high level description of DRDP.

1. The DRDP client considers the source of resolving domains (RD). This document defines two ways to collect RDs: RD are directly provisioned - in our case rd.org) , or RD are retrieved from a Pointer - in our case rd_pointer.org. In the later case RD are collected from a DNS request of type PTR. In the case of Pointer being rd_pointer.org, the QNAME of the PTR request would be b._dns.rd_pointer.org.

2. The DRDP client collects the resolving services and associated parameters under the umbrella of each RD. The resolving services offered by the RD are collected via a DNS request of type SVCB and the associated parameters are provided through SvcParameters. In the case of the RD being rd.org, the QNAME of the SVCB request would be _dns.rd.org.

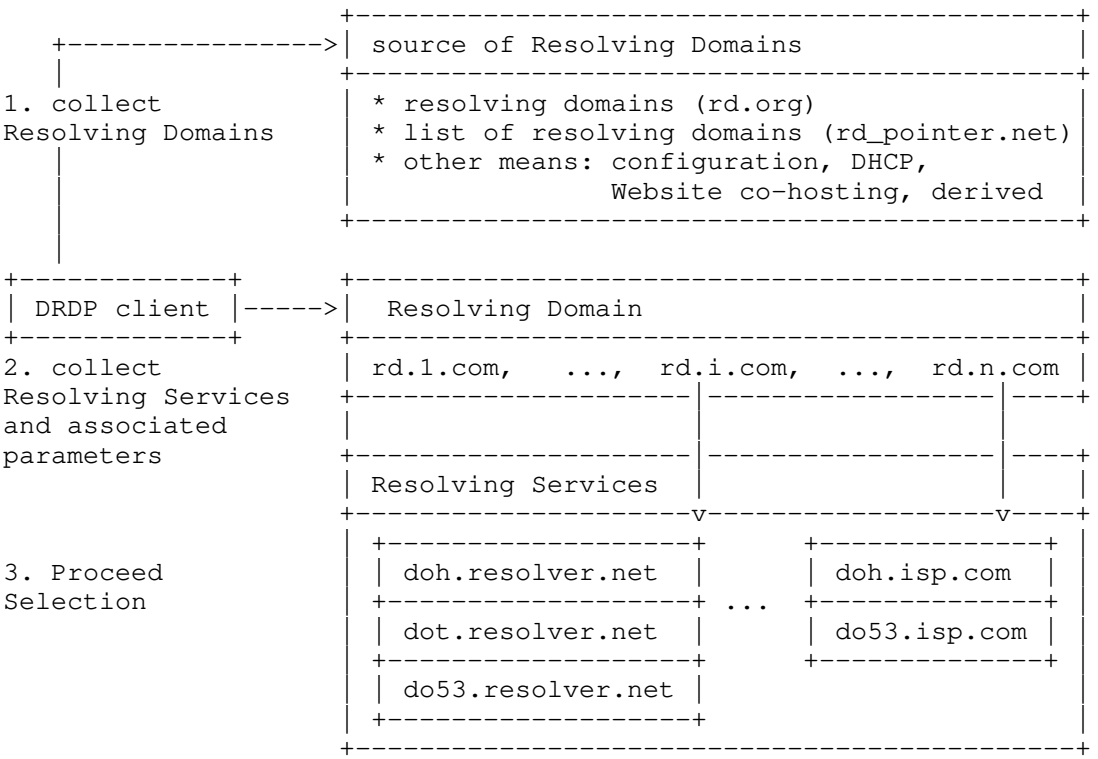

5. Pointer to a list of Resolving Domains

A Pointer is a FQDN that points to a list of FQDN that designates RD. If Pointer is represented by rd_pointer.net, the associated RDs are retrieved by the DNS query of type PTR for b._dns.rd_pointer.org.

The zone file below is inspired from DNS-SD where b indicates a browsing domain, _dns indicates the DNS resolving service, rd_pointer.org the Pointer and rd.1.com, ... rd.i.com the associated RDs. Note that they do not necessarily need to share a TLD. The order of the resolving domains is irrelevant, and the zone administrator SHOULD regularly reorder them. The RRsets MUST be signed with DNSSEC.

```
b._dns.rd_pointer.net  PTR rd.1.com
[...]
b._dns.rd_pointer.net  PTR rd.n.com
```

Using the DNS provides the advantage to retrieve the resolving domain without requiring other libraries than DNS as well as benefit from the DNS caching infrastructure including the use of the TTL.

An EDNS buffer size of 1232 bytes will avoid fragmentation on nearly all current networks. This is based on an MTU of 1280, which is required by the IPv6 specification, minus 48 bytes for the IPv6 and UDP headers. This document RECOMMENDS that the number of RDs associated to a Pointer do not generate fragmentation of the DNS UDP packet. It is believed to address most common needs or expectation from a vast majority of stub DNS client.

When the number of RD exceeds this limit, the DNS client may carry this over TCP which is likely to be supported by DNS client wiling to upgrade to DoH or DoT resolving services. However, the transfer of large number of RDs is considered as an application specificity that would benefit benefit from the compression of the transferred data provided by ftp or http. In such case, these application may define there own specific mechanism to provision the RDs.

As of July 27 2020, 1232 bytes correspond to the 94 first most popular FQDN listed by [moz.com]. The current size of such lists [curl][dnsprivacy.org] have less than 50 resolving domains. Other lists such as [public-dns.info] have as much as 11.000 entries, but such lists seems to contain open resolvers which is out side of the scope of a selection process.
Web browser (Google Chrome) also have lists over 10.000 entries, but in case a significant number of entries seems to be IP addresses that have a very limited network scope (e.g. limited to the ISP). The length of the list in scope to the DNS client is in fact significant

smaller in term of IP addresses and even smaller if resolving domain are able to represent multiple IP addresses. Overall, the size of such lists are currently due to the absence of discovery protocols.

6. Discovery of Resolving Services

The discovery of resolving services is performed by the RDP client with all the available RDs. Given a RD rd.org, a DRDP client sends a DNS request of type SVCB for _dns.rd.org.

The example below presents the use of an AliasForm followed by a ServiceForm which allows an indirection. The Alias form is not madatory and instead only ServiceForm associated to _dns.rd.org could have been used instead.

The SvcFieldPriority indicates the preference of the RD. It typically enables an operator to indicate that an encrypted DNS is preferred.

The SvcParamKey alpn MUST be present when TLS is used as its presence and value indicates the DNS transport. The absence of the alpn SvcParamKey indicates Do53, alpn set to dot indicates DoT is served while h* indicates DoH is served. Note that the port value (53, 853, 443) is not used to determine the DNS transport as non standard port MAY be used. The example below uses an non standard port 5353 for illustrative purpose.

Other SvcParam are detailed in Section 8 and are optional. A SvcParam not understood by the DNS client MUST be ignored.

The RRsets MUST be protected with DNSSEC and when alpn is provided a TLSA RRset SHOULD be present. These RRsets have been omitted for clarity.

```
## Discovery of all service instances
_dns.rd.org. 7200 IN SVCB 0 svc.example.com.
svc.example.com.    7200 IN SVCB 12 ( svc0.example.net.
                                      port="5353" user-display="Legacy Resolver"
)
svc.example.com.    7200 IN SVCB 1 ( svc1.example.net.  alpn="dot"
                                      port="5353" esniconfig="..."
                                      user-display="Preferred Example's Choice" )
svc.example.com.    7200 IN SVCB 3 ( svc2.example.net. alpn="h2"
                                       port="5353" esniconfig="..." user-display=
 )
svc.example.com.    7200 IN SVCB 2 ( svc3.example.net. alpn="h3"
                                       port="5353" esniconfig="..." user-display=
 )
```

Note that Section 15.2 provides another variant to perform RDP. Such variant is left for further discussion and address the need to be able to narrow down the discovery to a subset of resolving services such as DoH-only or DoT-only services.

Some notes:

1. _domain uses SVCB but does not have TLS. While SVCB has been created essentially for TLS based service, this does not appear to be mandatory.

2. Should we have some constraints regarding the SvcDomainName and QNAME ?

3. do we need the service subsets

7. TTL

The DNS client SHOULD perform DRDP at regular intervals as indicated by its policy.

The selection process MAY remove resolving services with short TTL lower than a day as it indicates some source of instalbility. Given a subset of selected resolving services, the DNS client may perform DRDP 1 hour before an SVB RRset expires.

8. SvcParamKey

This section defines a set of SvcParamKey that MAY be use to carry the necessary informations for the selection process.

alpn :

esniconfig :

port :

user-display indicates a strings in UTF-8 that is expected to be representative to a potential end user. Though there is no restriction in the scope of that string. The string is likely to represent the service within the resolving domain.

uri_template designates the URI template for DoH. This key MUST NOT be present on non DoH services and MUST be ignored by the DNS client on non DoH resolving Services.

auth_domain indicates the list of authoritative domain name the resolving service has strong relation with. It is expected that a

resolving service may prefer to perform DNS resolution over these domains to that specific resolving service as to preserve its privacy. This information MUST be verified and validated.

scope_domain indicates the limitation of resolved domains. When present DNS request sent to the resolution service MUST belong to that domain.

filtering indicates the presence of a filtering service

ip_subnet indicates a subnetwork restriction. This is mostly useful for resolving services that are not globally.

dnssec indicates whether dnssec is enabled or not.

9. Resolver advertising other service sub type

A resolving service receiving a DNS request over a service sub type MAY be willing to advertise the DNS client that other sub service type are available. This is especially useful, when, for example, a resolver wants that the DNS resolver switches to other service sub types that are more secure.

In order to do so the resolver MAY provide in the additional data field the _dns SRVCB of ServiceForm.

10. Migration to service sub types

The principle of the discovery mechanism is that the resolver indicates the available service sub types and let the DNS client chose which sub type it prefers. On the other hand, the resolver MAY also indicate a preference using the priority and weight fields. Redirection MAY especially be needed when a DNS client is using the Do53 and the resolver would like to upgrade the DNS client session to a more secure session. This MAY require a specific ERROR code that will request the DNS client to perform service discovery.

It is expected that DRDP MUST always be available via Do53. However, this does not mean that a resolver is expected to implement the Do53 sub type service for a resolving service.

11. Security Considerations

11.1. Use of protected channel is RECOMMENDED

When available, it is recommended to chose a protected version of the rdns service. More specifically, the use of end-to-end protection ensures that the DNS client is connected to the expected platform and

that its traffic cannot be intercepted on path. Typically, the selection of resolver on the Internet (and not on your ISP network) and the use of a non protected channel enables an attacker to monitor your DNS traffic. The similar observation remains true if you are connected to the resolver of your ISP. It is commonly believed that trusting your ISP (that is your first hop) makes encryption unecessary. Trusting your ISP is mandatory in any case, but the associated level of trust with an protected channel is restricted to the operation of the DNS platform. With non protected channel the trust is extended to any segment between the DNS client and the resolver, which is consequently larger. The use of a protected channel is recommended as it will prevent anyone on path to monitor your traffic.

11.2. DNSSEC is RECOMMENDED

The exchanges SHOULD be protected with DNSSEC to ensure integrity of the information between the authoritative servers and the DNS client. Without DNSSEC protection, DNS messages may be tampered typically when they are transmitted over an unprotected channel either between the DNS client and the resolver or between the resolver and the authoritative servers. The messages may be tampered by an online attacker intercepting the messages or by the intermediary devices. It is important to realize that protection provided by TLS is limited to the channel between the DNS client and the resolver. There are a number of cases were the trust in the resolver is not sufficient which justify the generalization of the use of DNSSEC. The following examples are illustrative and are intended to be exhaustive.

First, the discovery exchanges may happen over an unprotected channel, in which case, the messages exchanged may be tampered by anyone on-path between the DNS client and the resolver as well as between the resolver and the authoritative servers - including the resolver. When TLS is used between the DNS client and the resolver, this does not necessarily mean the DNS client trusts the resolver. Typically, the TLS session may be established with a self-signed certificate in which case the session is basically protected by a proof-of-ownership. In other cases, the session may be established based on Certificate Authorities (CA) that have been configured into the TLS client, but that are not necessarily trusted by the DNS client. In such cases, the connected resolver may be used to discover resolvers from another domain. In this case, the resolver is probably interacting with authoritative servers using untrusted and unprotected channels. Integrity protection relies on DNSSEC.

11.3. TLSA is RECOMMENDED

When TLS is used to protect the DNS exchanges, certificates or fingerprint SHOULD be provided to implement trust into the communication between the DNS client and the resolver. The TLS session and the association of the private key to a specific identity can be based on two different trust model. The Web PKI that will rely on CA provisioned in the TLS library or the TA provided to the DNS client. A DNS client SHOULD be able to validate the trust of a TLS session based on the DNSSEC trust model using DANE.

When the DNS client is protecting its session to the resolver via TLS, the DNS client may initiate an TLS session that is not validated by a CA or a TLSA RRsets. The DNS client MUST proceed to the discovery process and validate the certificate match the TLSA RRset. In case of mismatch the DNS client MUST abort the session.

12. Privacy Considerations

When the discovery protocol is performed using a resolver that belongs to one domain for another domain, or over an unprotected channel, the DNS client must be conscious that its is revealing to the resolver its intention to use another resolver. More specifically, suppose an resolver is complying some legal requirements that DNS traffic must be unencrypted. Using this resolver to perform a resolver discovery reveals the intention of potentially using alternative resolvers. Alternatively, narrowing down the discovery over a specific sub type of resolver (DoT, or DoH) may reveal to that resolver the type of communication. As result, when performing a discovery over a domain that differs to the domain the resolver belongs to, it is RECOMMENDED to request the SRV RRsets associated to all different sub type of proposed services.

The absence of traffic that results from switching completely to a newly discovered resolver right after the discovery process provides an indication to the resolver the DNS client is switching to. It is hard to make that switch unnoticed to the initial resolver and the DNS resolver MUST assume this will be noticed. The information of switching may be limited by sharing the traffic between different resolvers, however, the traffic pattern associated to each resolver may also reveal the switch. In addition, when the initial resolver is provided by the ISP, the ISP is also able to monitor the IP traffic and infer the switch. As a result, the DNS client SHOULD assume the switch will be detected.

With DoT or DoH, the selection of port 443 will make the traffic indistinguishable from HTTPS traffic. This means that an observer will not be able to tell whether the traffic carries web traffic or

DNS traffic. Note that it presents an interest if the server offers both a web service as well as a resolution service. Note that many resolvers have a dedicated IP address for the resolution service, in which case, the information will be inferred from the IP address. Note also that traffic analysis may infer this as well. Typically suppose an IP address hosts one or multiple web sites that are not popular as well as a resolving service. If this IP address is associated frequent short size exchanges, it is likely that these exchanges will be DNS exchanges rather than Web traffic. The size of the packet may also be used as well as many other patterns. As a result, the use port 443 to hide the DNS traffic over web traffic should be considered as providing limited privacy.

13. IANA Considerations

This document requests the IANA the creation of the following underscored node names in the Underscored and Globally Scoped DNS Node Names registry https://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml#dns-parameters-14

RR Type	_NODE NAME	Reference
SRVCB	_dns	RFC-TBD

SvcParamKey	NAME	Meaning	Reference
7	user-display	User friendly string (UTF8) to represent the resolver	RFC-TBD
	uri_template	URI template	
	auth_domain	Domains the resolving service is authoritative	
	filetring	Filetring services provided	
	ip_subnet	ip ranges accepted.	
	dnssec	DNSSEC validation enabled	

14. Acknowledgments

We would like thank Mirja Kuehlewind as well as the GSMA IG for their comments. We also thank Ted Hardie and Paul Hoffman for their feed backs regarding the dns schemes for DoT and DoH.
We thank Ben Schwartz for the comments on the list size. We thank Harald Alvestrand for its recommendation on having a model that enable multiple third party providers to provide their own list of resolving domains. We thank Stephan Bortzmeyer, Ralf Weber, Chris Box for its clarifications.

15. Appendices

15.1. DRDP Requirements

This section lists the DRDP requirements.

REQ 1: DRDP MUST enable a DNS client to discover the available resolving services with their associated characteristics in order to proceeds to a selection process. The selection process takes resolving services identities and associated parameters and proceed to the selection.
Any sort of resolving service selection is outside the scope of DRDP.

REQ 2: While the discovery process is triggered by the DNS client, a third party MUST be able to provide necessary input information so a resolving service discovery process can be triggered within a specific context.
Provisioning protocols to provide this information is not as per say in scope of DRDP. DRDP defines the format of the format for such input as well as a set of such inputs.

REQ 3: Any information used in DRDP MUST be authenticated by its owner. In particular, the characteristics associated to the resolving service MUST be certified by the resolving service operator / owner and MUST be associated a validity period. In addition, a third party providing a set of inputs MUST authenticate that set.

REQ 4: Information associated to the resolving services is intended to enable the selection process that is assumed to match the end user or application policy. The selection process is out of scope of DRDP. Information may carry some characteristics of a resolving service or hints that will help the selection. In particular an operator of multiple resolving service MUST be able to associate a preference to the proposed resolving services. To ease automation of the selection as well as to make multiple applications benefit from DRDP the information MUST be carried over a standardized format.

REQ 5: DRDP MUST be designed to be used indifferently by a DNS client using any DNS transport protocol (Do53, DoT, DoH, ...). However, DRDP SHOULD be able to restrict the information retrieved to a certain type of resolving service, i.e. Do53, DoT, DoH.

REQ 6: DRDP deployment MUST NOT be disruptive for the legacy DNS client or infrastructure and legacy client SHOULD be able to incrementally include DRDP.

15.2. Discovery of specific service instance

To reduce the size of the messages, the DNS client MAY also prefer to query information of resolving services using a specific transport (DNS, DoT, DoH) that are designated as sub sets. A DNS client MAY list the different subsets of that resolving domain with a PTR query. This document defines the following subsets _53._dns for DNS, _853._dns for DoT and _443.__dns for DoH. Other subsets MAY be defined in the future. A DNS client that does not understand a subset SHOULD ignore it and maybe proceed to the discovery as defined in Section 6.

All subsets MUST share the same resolving domain and be listed with a PTR RRsets. The DNS client MAY NOT performed a DNS query of type PTR, for example, if it has a previous knowledge of the existence of the subset or if indicated by its policy. In this it MAY directly proceed to the SRVCB resolution.

The same restrictions as defined in section Section 6 apply.

Note that while the SvcFieldPriority indicates the priority within a subservice, this field MUST have a coherence across subservices. The priority provided SHOULD be coherent with the case of a _dns SRVCB query of section Section 6.

The figure below illustrates an example of zone file. RRSIG and TLSA have been omitted for the purpose of clarity.

```
### Definition of the resolving service subsets
_dns.example.com PTR _53._dns.example.com
_dns.example.com PTR _853._dns.example.com
_dns.example.com PTR _443._dns.example.com

### services instances per service subset
_53._dns.example.com. 7200 IN SVCB 0 svc0.example.com.
svc0.example.com.    7200 IN SVCB 12 ( svc0.example.net.
                                       port="5353" user-display="Legacy Resolver"
 )
_853._dns.example.com.    7200 IN SVCB 0 svc1.example.com.
svc1.example.com.    7200 IN SVCB 1 ( svc1.example.net.  alpn="dot"
                                      port="5353" esniconfig="..."
                                      user-display="Preferred Example's Choice" )

_443_dns.example.com.    7200 IN SVCB 0 svc4.example.net.
svc4.example.com.    7200 IN SVCB 3 ( svc2.example.net. alpn="h2"
                                       port="5353" esniconfig="..." user-display=
 )
svc4.example.com.    7200 IN SVCB 2 ( svc3.example.net. alpn="h3"
                                      port="5353" esniconfig="..."
                                      user-display="Testing QUIC")
```

16. References

16.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

16.2. Informative References

[curl] "Publicly available servers", n.d., <https://github.com/curl/curl/wiki/DNS-over-HTTPS#publicly-available-servers>.

[dnsprivacy.org] "DNS Privacy Test Servers", n.d., <https://dnsprivacy.org/wiki/display/DP/DNS+Privacy+Test+Servers#DNSPrivacyTestServers-Publicresolvers>.

[I-D.btw-add-home] Boucadair, M., Reddy.K, T., Wing, D., and N. Cook, "Encrypted DNS Discovery and Deployment Considerations for Home Networks", draft-btw-add-home-07 (work in progress), July 2020.

[moz.com] "The Moz Top 500 Websites", n.d., <https://moz.com/top500>.

[public-dns.info]
"Public DNS Server List", n.d.,
<https://public-dns.info/>.

[RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <https://www.rfc-editor.org/info/rfc7719>.

Author's Address

Daniel Migault
Ericsson
8275 Trans Canada Route
Saint Laurent, QC 4S 0B6
Canada

EMail: daniel.migault@ericsson.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 14 January 2021

T. Pauly
E. Kinnear
Apple Inc.
C.A. Wood
Cloudflare
P. McManus
Fastly
T. Jensen
Microsoft

13 July 2020

Adaptive DNS Resolver Discovery
draft-pauly-add-resolver-discovery-01

Abstract

This document defines a method for dynamically discovering resolvers that support encrypted transports, and introduces the concept of designating a resolver to be used for a subset of client queries based on domain. This method is intended to work both for locally-hosted resolvers and resolvers accessible over the broader Internet.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 14 January 2021.

Table of Contents

1. Introduction

When clients need to resolve names into addresses in order to establish networking connections, they traditionally use by default the DNS resolver that is provisioned by the local network along with their IP address [RFC2132] [RFC8106]. Alternatively, they can use a resolver indicated by a tunneling service such as a VPN.

However, privacy-sensitive clients might prefer to use an encrypted DNS service other than the one locally provisioned in order to prevent interception, profiling, or modification by entities other than the operator of the name service for the name being resolved.

Protocols that can improve the transport security of a client when using DNS or creating TLS connections include DNS-over-TLS (DoT) [RFC7858], DNS-over-HTTPS (DoH) [RFC8484], and Encrypted TLS Client Hellos [I-D.ietf-tls-esni].

This document defines a method for dynamically discovering resolvers that support encrypted transports, and introduces the concept of designating a resolver to be used for a subset of client queries based on domain. This method is intended to work both for locally-hosted resolvers and resolvers accessible over the broader Internet.

1.1. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

This document defines the following terms:

Direct Resolver: A DNS resolver using any transport, encrypted or unencrypted, that is provisioned directly by a local router or a VPN.

Designated Resolver: A DNS resolver that is designated as a responsible resolver for a given domain or zone. Designated resolvers use encrypted transports.

Companion DoH Server: A DNS resolver that provides connectivity over HTTPS (DoH) that is designated as equivalent to querying a particular Direct Resolver.

3. Designated Resolvers

An encrypted DNS resolver, such as a DoH or DoT server, can be designated for use in resolving names within one or more zones. This means that clients can learn about an explicit mapping from a given domain or zone to one or more Designated Resolvers, and use that mapping to select the best resolver for a given query.

Designating a resolver MUST rely on agreement between the entity managing a zone (the Domain Owner) and the entity operating the resolver, such that clients can securely validate this designation. These entities can be one and the same, or a Domain Owner can choose to designate a third-party resolver to handle its traffic. Proof of

this agreement asserts to clients that sending any query to the designated resolver exposes no more information than sending that query to the entity managing the corresponding zone.

As an example with only one entity, a company that runs many sites within "enterprise.example.com" can provide its own DoH resolver, "doh.enterprise.example.com", and designate only that resolver for all names that fall within "enterprise.example.com". This means that no other resolver would be designated for those names, and clients would only resolve names with the same entity that would service TLS connections.

As an example with several entities, the organization that operates sites within "example.org" may work with two different Content Delivery Networks (CDNs) to serve its sites. It might designate names under "example.com" to two different entities, "doh.cdn-a.net" and "doh.cdn-b.net". These are CDNs that have an existing relationship with the organization that runs "example.org", and have agreements with that organization about how data with information on names and users is handled.

There are several methods that can be used to discover and validate a resolver designation:

* Discovery using SVCB DNS records (Section 3.1), and validation using DNSSEC

* Discovery using information in a provisioning domain (PvD) file from the Designated DoH Resolver (Section 3.2)

* Validation using a file hosted on a well-known HTTPS URI based on a zone apex (Section 3.3)

* Validation using TLS certificates to confirm of domain name ownership (Section 3.4)

Note that clients MUST NOT accept designations for effective top-level domains (eTLDs), such as ".com".

3.1. Designating with Service Binding DNS Records

The primary source for discovering Designated DoH Server configurations is from properties stored in a SVCB DNS resource record, or a SVCB-conformant resource record type, like HTTPS [I-D.ietf-dnsop-svcb-https]. This record provides the URI Template of a DoH server that is designated for a specific domain. A specific domain may have more than one such record.

The rationale for using SVCB records for recolver discovery is discussed in Appendix A.

In order to designate a DoH server for a domain, a SVCB record can contain the "dohuri" (Section 9). The value stored in the parameter is a URI, which is the DoH URI template [RFC8484].

The following example shows a record containing a DoH URI, as returned by a query for the HTTPS variant of the SVCB record type on "foo.example.com".

```
foo.example.com.  7200  IN HTTPS 1 . (
                           dohuri=https://doh.example.net/dns-query )
```

If this record is DNSSEC-signed [RFC4033], clients can immediately create a mapping that indicates the server (doh.example.net) as a Designated Resolver for the name in the SVCB record (foo.example.com).

Once a record that designated a DoH server has expired, the client SHOULD issue another SVCB/HTTPS query whenever issuing queries within the designated domain. This query SHOULD still be performed using the designated DoH server. If the response designates a different DoH server, the client should verify and use the new designation.

If this record is not DNSSEC-signed, clients MUST perform other validation to determine that the zone designation is permitted, as described in Section 3.3.

3.2. Additional Designation with PvD JSON

A provisioning domain (PvD) defines a coherent set of information that can be used to access a network and resolve names. Section 4.3 of [I-D.ietf-intarea-provisioning-domains] defines a JSON dictionary format that can be fetched over HTTPS at the well-known URI "/.well-known/pvd".

Designated Resolvers that support DoH SHOULD provide a PvD JSON dictionary available at the well-known PvD URI with the path of the DoH server's URI template appended.

For example, the PvD JSON for the DoH server "https://doh.example.net/dns-query" would be available at "https://doh.example.net/.well-known/pvd/dns-query".

The key "dohTemplate" is defined within the JSON dictionary (Section 9) to point back to the DoH URI Template itself. This is used for confirming the DoH server when the PvD is discovered locally or during zone apex confirmation (Section 3.3).

Names that are listed in the "dnsZones" key in the JSON dictionary indicate a set of zones that designate the resolver. These are the zones that are available to resolve through the associated DoH server. Note that this list does not need to be exhaustive, but is the set of common zones managed by the resolver that all clients should be aware of. Before using DNS results for these names, clients MUST validate the designation either with a DNSSEC-signed SVCB record (Section 3.1), or the confirmation methods described in Section 3.3 and Section 3.4. DNS queries for validating records SHOULD be sent to the DoH resolver. In order to optimize the validation of these domains, servers MAY use HTTP Server Push to deliver the signed SVCB answers prior to requests being made.

The "expires" key indicates a time after which the content of the PvD file is no longer valid. Clients SHOULD re-fetch PvD information if the expiration time has passed before using any designations that were based on the PvD content.

```
{
      "identifier": "doh.example.net.",
      "dohTemplate": "https://doh.example.net/dns-query",
      "dnsZones": ["example.com"],
      "expires": "2020-08-23T06:00:00Z"
}
```

3.3. Confirmation of Designation with Zone Apex PvD

Designated DoH Resolvers that provide the PvD JSON described in Section 3.2 can also provide information to validate of zone's designation without DNSSEC. In order to confirm the designation, the client requests a well-known HTTPS URI based on a zone apex name, and checks a PvD file to ensure that it matches the DoH resolver. This ensures that a DoH resolver cannot claim a designation for a given zone without cooperation from the entity that owns the certificate for the apex of that zone.

In order to enumerate the zone apex names that confirm designation in this manner, the DoH resolver's PvD JSON dictionary can contain an array of strings, with the key "trustedNames". Clients can validate the resolver designation by checking a resource hosted by a name indicated in "trustedNames". The client first issues an HTTP GET request by appending "/.well-known/pvd" to the trusted name, using the "https" scheme. In order to validate the designation, the PvD

JSON MUST contain a "dohTemplate" key pointing to the correct DoH resolver. The client's query for the IP addresses of the trusted name MAY use the DoH resolver prior to fully validating the designation, since the validation uses HTTPS to authenticate the designation.

Note that the names listed in "trustedNames" are only useful for confirming a designation that was indicated either by a non-DNSSEC-signed SVCB designation (Section 3.1), or an additional designation provided by the DoH resolver's PvD (Section 3.2). A trusted name MUST be an exact match of a designating name, or else a parent of a designating name.

If a name has more specific sub-domains that should not be allowed to designate a given DoH resolver, this method of confirmation MUST NOT be used.

As an example of this process, the JSON dictionary for the DoH server "https://doh.example.net/dns-query", which is retrieved from "https://doh.example.net/.well-known/pvd/dns-query", could contain the following contents:

```
{
  "identifier": "doh.example.net.",
  "dohTemplate": "https://doh.example.net/dns-query",
  "dnsZones": ["example.com"],
  "trustedNames": ["example.com"],
  "expires": "2020-08-23T06:00:00Z"
}
```

This indicates that "example.com" should be treated as a designated domain, and that it can be validated by checking with the "example.com" server rather than using DNSSEC.

In this example, the well-known URI used for validation is "https://example.com/.well-known/pvd". In order to trust the designation, this request must return valid JSON with the "dohTemplate" key matching the original DoH resolver. For example, this dictionary could contain the following contents:

```
{
  "identifier": "example.com.",
  "dohTemplate": "https://doh.example.net/dns-query",
  "expires": "2020-08-23T06:00:00Z"
}
```

A client MUST NOT trust a designation if the JSON content is not present, does not contain a "dohTemplate" key, or the value in the "dohTemplate" key does not match. The following result would not be acceptable for the example above:

```
{
  "identifier": "example.com.",
  "dohTemplate": "https://not-the-doh-youre-looking-for.example.net/dns-query"
  "expires": "2020-08-23T06:00:00Z"
}
```

Note that the domains listed in "trustedNames" may be broader than the zones that designate the resolver. In the following example, names under "foo.example.com" and "bar.example.com" designate the DoH server "https://doh.example.net/dns-query", and use the PvD JSON from "example.com" to validate the designation. However, the client would not designate the DoH server for all names under "example.com".

```
{
  "identifier": "doh.example.net.",
  "dohTemplate": "https://doh.example.net/dns-query",
  "dnsZones": ["foo.example.com", "bar.example.com"],
  "trustedNames": ["example.com"],
  "expires": "2020-08-23T06:00:00Z"
}
```

3.4. Confirmation of Designation with TLS Certificates

A DoH server designation can also be validated by checking the SubjectAlternativeName field in the DoH server's own TLS certificate. When a client wants to confirm the validity of the designation in this situation, it can check the TLS certificate of the DoH server for the name of the domain which triggered the original designation query.

The following example shows an HTTPS variant of the SVCB record type for "foo.example.com". If this record was received without DNSSEC, the client can confirm its validity by establishing a connection to "doh.example.net" and verifying the TLS certificate contains an exact match for the "foo.example.com" name. If the queried domain is not present in the TLS certificate of the designated DoH server, the client may confirm the validity by an alternate method such as zone apex confirmation (Section 3.3) but MUST NOT use the record until otherwise validated.

```
foo.example.com.  7200  IN HTTPS 1 . (
                        dohuri=https://doh.example.net/dns-query )
```

4. Explicit Discovery of Local Resolvers

If the local network provides configuration with an Explicit Provisioning Domain (PvD), as defined by [I-D.ietf-intarea-provisioning-domains], clients can learn about domains for which the local network's resolver is authoritative. The keys for DoH resolvers described in Section 3.2 also allow this local PvD to be used for resolver discovery.

If an RA provided by the router on the network defines an Explicit PvD that has additional information, and this additional information JSON dictionary contains the key "dohTemplate", then the client SHOULD add this DoH server to its list of known DoH configurations. The domains that the DoH server claims authority for are listed in the "dnsZones" key.

Local deployments that want to designate a resolver for a private name that is not easily signed with DNSSEC MUST provide an alternate method of validating a designation, such as described in Section 3.3 or Section 3.4.

5. Discovery of DoH Capabilities for Direct Resolvers

Direct Resolvers can advertise a Companion DoH server that offers equivalent services and is controlled by the same entity. To do this, a DNS server returns an SVCB record for "dns://resolver.arpa" with "ipv4hint" and/or "ipv6hint" set to a valid IP address and the "dohuri" key set to a valid DoH URI template as with the Designated DoH Server SVCB record. The TLS certificate used with the DoH URI MUST have the IP addresses for each of its DNS endpoints, classic or DoH, within the SubjectAlternativeName field to allow the client to verify ownership.

Once a client is configured to query a Direct Resolver, it SHOULD query the resolver for SVCB records for "dns://resolver.arpa" before making other queries. This will help the client avoid leaking queries that could go over DoH once the Companion DoH Server is discovered. If an SVCB record is returned, its "dohip" field designates an IP address the client can send DoH queries to in lieu of sending classic DNS queries to the Direct Resolver. The "dohuri" field contains the DoH URI similarly to the SVCB record for a Designated DoH Server.

To validate the Companion DoH Server and the resolver that advertised it are related, the client MUST check the SubjectAlternativeName field of the Companion DoH Server's TLS certificate for the original resolver's IP address and the advertised IP address for the Companion DoH server. If both are present, the discovered Companion DoH Server

MUST be used whenever the original Direct Resolver would be used. Otherwise, the client SHOULD suppress queries for Companion DoH Servers against this resolver for the TTL of the negative or invalid response and continue to use the original Direct Resolver.

The following example shows a record containing a Companion DoH URI, as returned by a query for an SVCB record for "dns://resolver.arpa":

```
_dns.resolver.arpa  7200  IN SVCB 1 doh.example.net (
                      ipv4hint=x.y.z.w
                      dohuri=https://doh.example.net/dns-query )
```

A DNS resolver MAY return more than one SVCB record of this form to advertise multiple Companion DoH Servers that are valid as a replacement for itself. Any or all of these servers may have the same IP address as the DNS resolver itself. In this case, clients will only have one IP address to check for when verifying ownership of the Companion DoH server.

6. Server Deployment Considerations

When servers designate DoH servers for their names, the specific deployment model can impact the effective privacy and performance characteristics.

6.1. Single Content Provider

If a name always resolves to server IP addresses that are hosted by a single content provider, the name ought to designate a single DoH server. This DoH server will be most optimal when it is designated by many or all names that are hosted by the same content provider. This ensures that clients can increase connection reuse to reduce latency in connection setup.

A DoH server that corresponds to the content provider that hosts content has an opportunity to tune the responses provided to a client based on the location inferred by the client IP address.

6.2. Multiple Content Providers

Some hostnames may resolve to server IP addresses that are hosted by multiple content providers. In such scenarios, the deployment may want to be able to control the percentage of traffic that flows to each content provider.

In these scenarios, there can either be:

* multiple designated DoH servers that are advertised via SVCB DNS Records; or,

* a single designated DoH server that can be referenced by one or more SVCB DNS Records, operated by a party that is aware of both content providers and can manage splitting the traffic.

If a server deployment wants to easily control the split of traffic between different content providers, it ought to use the latter model of using a single designated DoH server that can better control which IP addresses are provided to clients. Otherwise, if a client is aware of multiple DoH servers, it might use a single resolver exclusively, which may lead to inconsistent behavior between clients that choose different resolvers.

6.3. Avoid Narrow Deployments

Using designated DoH servers can improve the privacy of name resolution whenever a DoH server is designated by many different names within one or more domains. This limits the amount of information leaked to an attacker observing traffic between a client and a DoH server: the attacker only learns that the client might be resolving one of the many names for which the server is designated.

However, if a deployment designates a given DoH server for only one name, or a very small set of names, then it becomes easier for an attacker to infer that a specific name is being accessed by a client. For this reason, deployments are encouraged to avoid deploying a DoH server that is only designated by a small number of names. Clients can also choose to only allow DoH servers that are associated with many names.

Beyond the benefits to privacy, having a larger number of names designate a given DoH server improves the opportunity for DoH connection reuse, which can improve the performance of name resolutions.

7. Security Considerations

In order to avoid interception and modification of the information sent between clients and Designated Resolvers, all exchanges between clients and servers are performed over encrypted connections, e.g., TLS.

Malicious adversaries may block client connections to a Designated Resolver as a Denial-of-Service (DoS) measure. Clients which cannot connect these resolvers may be forced to, if local policy allows, fall back to unencrypted DNS if this occurs.

8. Privacy Considerations

Clients must be careful in determining to which DoH servers they send queries directly. A malicious resolver that can direct queries to itself can track or profile client activity. In order to avoid the possibility of a spoofed SVCB record designating a malicious DoH server for a name, clients MUST ensure that such records validate using DNSSEC (Section 3.1), using zone apex confirmation (Section 3.3), or using domain names in TLS certificates (Section 3.4).

Even servers that are validly designated can risk leaking or logging information about client lookups. Such risk can be mitigated by further restricting the list of resolvers that are allowed for direct use based on client policy.

An adversary able to see traffic on each path segment of a DoH query (e.g., from client to a Designated Resolver, and the Designated Resolver to an authoritative DNS server) can link queries to specific clients with high probability. Failure to observe traffic on any one of these path segments makes this linkability increasingly difficult. For example, if an adversary can only observe traffic between a client and proxy and egress traffic from a target, then it may be difficult identify a specific client's query among the recursive queries generated by the target.

9. IANA Considerations

9.1. DoH Template PvD Key

This document adds a key to the "Additional Information PvD Keys" registry [I-D.ietf-intarea-provisioning-domains].

JSON key	Description	Type	Example
dohTemplate	DoH URI Template [RFC8484]	String	"https://dnsserver.example.net/dns-query{?dns}"

Table 1

9.2. Trusted Names PvD Key

This document adds a key to the "Additional Information PvD Keys" registry [I-D.ietf-intarea-provisioning-domains].

JSON key	Description	Type	Example
trustedNames	Names of servers that can validate resolver designation.	Array of Strings	["example.com"]

Table 2

9.3. DoH URI Template DNS Service Parameter

This document adds a parameter to the "Service Binding (SVCB) Parameter" registry. The allocation request is 32768, taken from the to the First Come First Served range.

If present, this parameters indicates the URI template of a DoH server that is designated for use with the name being resolved. This is a string encoded as UTF-8 characters.

Name: dohuri

SvcParamKey: 32768

Meaning: URI template for a designated DoH server

Reference: This document.

9.4. Special Use Domain Name "resolver.arpa"

This document calls for the creation of the "resolver.arpa" SUDN. This will allow resolvers to respond to queries directed at themselves rather than a specific domain name. While this document uses "resolver.arpa" to return SVCB records indicating DoH capability, the name is generic enough to allow future reuse for other purposes where the resolver wishes to provide information about itself to the client.

10. Acknowledgments

Thanks to Erik Nygren, Lorenzo Colitti, Mikael Abrahamsson, Ben Schwartz, Ask Hansen, Leif Hedstrom, Tim McCoy, Stuart Cheshire, Miguel Vega, Joey Deng, Ted Lemon, and Elliot Briggs for their feedback and input on this document.

11. References

11.1. Normative References

[I-D.ietf-dnsop-svcb-https]
Schwartz, B., Bishop, M., and E. Nygren, "Service binding and parameter specification via the DNS (DNS SVCB and HTTPS RRs)", Work in Progress, Internet-Draft, draft-ietf-dnsop-svcb-https-00, 12 June 2020, <http://www.ietf.org/internet-drafts/draft-ietf-dnsop-svcb-https-00.txt>.

[I-D.ietf-intarea-provisioning-domains]
Pfister, P., Vyncke, E., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", Work in Progress, Internet-Draft, draft-ietf-intarea-provisioning-domains-11, 31 January 2020, <http://www.ietf.org/internet-drafts/draft-ietf-intarea-provisioning-domains-11.txt>.

[I-D.ietf-tls-esni]
Rescorla, E., Oku, K., Sullivan, N., and C. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-07, 1 June 2020, <http://www.ietf.org/internet-drafts/draft-ietf-tls-esni-07.txt>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/info/rfc4033>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

11.2. Informative References

[I-D.schinazi-httpbis-doh-preference-hints]
Schinazi, D., Sullivan, N., and J. Kipp, "DoH Preference Hints for HTTP", Work in Progress, Internet-Draft, draft-schinazi-httpbis-doh-preference-hints-01, 8 January 2020, <http://www.ietf.org/internet-drafts/draft-schinazi-httpbis-doh-preference-hints-01.txt>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <https://www.rfc-editor.org/info/rfc2132>.

[RFC5507] IAB, Faltstrom, P., Ed., Austein, R., Ed., and P. Koch, Ed., "Design Choices When Expanding the DNS", RFC 5507, DOI 10.17487/RFC5507, April 2009, <https://www.rfc-editor.org/info/rfc5507>.

[RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <https://www.rfc-editor.org/info/rfc8106>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Appendix A. Rationale for using SVCB records

This mechanism uses SVCB/HTTPS resource records [I-D.ietf-dnsop-svcb-https] to communicate that a given domain designates a particular DoH resolver for clients to use for subsequent queries to within the domain.

There are various other proposals for how to provide similar functionality. There are several reasons that this mechanism has chosen SVCB records:

* Discovering encrypted resolver using DNS records keeps client logic for DNS self-contained, and allows an operator of a DNS zone to define exactly which names should use a given DoH server.

* Using DNS records also doesn't rely on bootstrapping with higher-level application operations (such as [I-D.schinazi-httpbis-doh-preference-hints]).

* SVCB records are extensible and allow definition of parameter keys. This makes them a superior mechanism for extensibility, as compared to approaches such as overloading TXT records. The same keys can be used both for upgrading direct resolvers to DoH through an explicit query (Section 5) and for discovering designated resolvers when issuing standard HTTPS queries (Section 3.1).

* Clients and servers that are interested in privacy of names will already need to support SVCB records in order to use Encrypted TLS Client Hello [I-D.ietf-tls-esni]. Without encrypting names in TLS, the value of encrypting DNS is reduced, so pairing the solutions provides the largest benefit.

* Clients that support SVCB will generally send out three queries when accessing web content on a dual-stack network: A, AAAA, and HTTPS queries. Discovering a resolver designation for a zone as part of one of these queries, without having to add yet another query, minimizes the total number of queries clients send. While [RFC5507] recommends adding new RRTypes for new functionality, SVCB provides an extension mechanism that simplifies client behavior.

Authors' Addresses

Tommy Pauly
Apple Inc.
One Apple Park Way
Cupertino, California 95014,
United States of America

Email: tpauly@apple.com

Eric Kinnear
Apple Inc.
One Apple Park Way
Cupertino, California 95014,
United States of America

Email: ekinnear@apple.com

Christopher A. Wood
Cloudflare
101 Townsend St
San Francisco,
United States of America

Email: caw@heapingbits.net

Patrick McManus
Fastly

Email: mcmanus@ducksong.com

Tommy Jensen
Microsoft

Email: tojens@microsoft.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 29, 2020

P. Sood
Google
P. Hoffman
ICANN

June 27, 2020

DNS Resolver Information Self-publication
draft-pp-add-resinfo-02

Abstract

This document describes methods for DNS resolvers to self-publish information about themselves. The information is returned as a JSON object. The names in this object are defined in an IANA registry that allows for light-weight registration. Applications and operating systems can use the methods defined here to get the information from resolvers in order to make choices about how to send future queries to those resolvers.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 29, 2020.

Table of Contents

1. Introduction

Historically, DNS stub resolvers typically communicated with the recursive resolvers in their configuration without needing to know anything about the features of the recursive resolvers. More recently, recursive resolvers have different features that may cause stub resolvers to make choices about which configured resolver from its configuration to use, and also how to communicate with the recursive resolver (such as over different transports). Thus stub resolvers need a way to get information from recursive resolvers about features that might affect the communication.

This document specifies a method for stub resolvers to ask recursive resolvers for such information. In short, a new RRtype is defined for stub resolvers to query using the DNS to a special-use domain name.

The response from this method is a JSON object. The JSON object MUST use the I-JSON message format defined in [RFC7493]. Note that [RFC7493] was based on RFC 7159, but RFC 7159 was replaced by [RFC8259]. Requiring the use of I-JSON instead of more general JSON format greatly increases the likelihood of interoperability.

The information that a resolver might want to give to a recursive resolver is not defined in this document; instead other documents

will follow that will specify that information and the format that it comes in.

In nearly every common scenario today, a DNS stub resolver gets the IP addresses of the recursive resolvers that it will use in an insecure fashion, such as from DHCP. Because these addresses were obtained insecurely, the protocol specified here does not try to use authenticated communication. If, in the future, more stub resolvers get the addresses of their recursive resolvers in a secure fashion, this protocol can be enhanced to include authenticated ways of getting information from the resolver.

1.1. Definitions

In the rest of this document, the term "resolver" without qualification means "recursive resolver" as defined in [RFC8499]. Also, the term "stub" is used to mean "stub resolver".

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Retrieving Resolver Information by DNS

A stub that wants to use the DNS to get information about a resolver can use the DNS query defined here. The query a stub resolver uses is resolver-info.arpa/IN/RESINFO. The RRtype "RESINFO" is defined in this document, and the IANA assignment is given in Section 4.1. The contents of the Rdata in the response to this query is defined in Section 3. If the resolver understands the RESINFO RRtype, the RRset in the Answer section MUST have exactly one record.

The name resolver-info.arpa is defined in this document, and the IANA assignment is given in Section 4.3. As described in Section 4.3, the zone resolver-info.arpa is not actually delegated and never will be. The resolver that receives this query acts as if it is delegated, and responds with its own RESINFO data in the Answer section.

A resolver that receives a query with the RRtype of RESINFO with a QNAME of resolver-info.arpa acts as if it is delegated, and responds with its own RESINFO data in the Answer section.

A resolver MAY be configured to respond to queries for the RESINFO RRtype on names other than resolver-info.arpa. For example, a resolver might be known to some of its clients by both an IP address and a few domain names, and be configured to be authoritative for

those names. For all names other than resolver-info.arpa or names that the resolver is configured to be authoritative for, a query for the RESINFO RRtype is meaningless and MUST result in a NODATA or NXDOMAIN response.

3. Contents of the Returned I-JSON Object

The JSON object returned by a DNS query or an HTTPS query MAY contain any name/value pairs.

All names in the returned object MUST either be defined in the IANA registry or, if for local use only, begin with the substring "temp-". The IANA registry (Section 4.2) will never register names that begin with "temp-".

All names MUST consist only of lower-case ASCII characters, digits, and hyphens (that is, Unicode characters U+0061 through 007A, U+0030 through U+0039, and U+002D), and MUST be 63 characters or shorter. As defined in Section 4.2, the IANA registry will not register names that begin with "temp-", so these names can be used freely by any implementer.

Note that the message returned by the resolver MUST be in I-JSON format. I-JSON requires that the message MUST be encoded in UTF8.

3.1. Example

The I-JSON object that a resolver returns might look like the following:

```
{
   "temp-field2": 42
}
```

As specified in [RFC7493], the I-JSON object is encoded as UTF8. [RFC7493] explicitly allows the returned objects to be in any order.

4. IANA Considerations

4.1. RESINFO RRtype

This document defines a new DNS RR type, RESINFO, whose value TBD will be allocated by IANA from the "Resource Record (RR) TYPEs" sub-registry of the "Domain Name System (DNS) Parameters" registry:

Type: RESINFO

Value: TBD

Meaning: Information self-published by a resolver as an I-JSON (RFC 7493) object

Reference: This document

4.2. Registry for DNS Resolver Information

IANA will create a new registry titled "DNS Resolver Information" that will contain definitions of the names that can be used with the protocols defined in this document. The registration procedure is by Expert Review and Specification Required, as defined in [RFC8126].

The specification that is required for registration can be either an Internet-Draft or an RFC. The reviewer for this registry is instructed to generally be liberal in what they accept into the registry: as long as the specification that comes with the registration request is reasonably understandable, the registration should be accepted.

The registry has the following fields for each element:

Name: The name to be used in the JSON object. This name MUST NOT begin with "temp-". This name MUST conform to the definition of "string" in I-JSON [RFC7493] message format.

Value type: The type of data to be used in the JSON object.

Specification: The name of the specification for the registered element.

4.3. resolver-info.arpa Special-Use Domain Name

IANA will record the domain name "resolver-info.arpa" in the "Special-Use Domain Names" registry [SUDN]. IANA MUST NOT delegate resolver-info.arpa in the .arpa zone.

5. Security Considerations

Unless a DNS request for resolver-info.arpa/IN/RESINFO as described in Section 2 is sent over DNS-over-TLS (DoT) [RFC7858] or DNS-over-HTTPS (DoH) [RFC8484], the response is susceptible to forgery. Given that one of the first expected uses for the protocol in this document is to find out whether DoT or DoH is available for the resolver, it is thus expected that most if not all such DNS requests will be sent without any chance of authentication. Stubs and resolvers SHOULD use normal DNS methods for avoiding forgery such as query ID randomization and source port randomization.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7493] Bray, T., Ed., "The I-JSON Message Format", RFC 7493, DOI 10.17487/RFC7493, March 2015, <https://www.rfc-editor.org/info/rfc7493>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <https://www.rfc-editor.org/info/rfc8259>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

[SUDN] "Special-Use Domain Names", n.d., <https://www.iana.org/assignments/special-use-domain-names/>.

6.2. Informative References

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

Appendix A. Ideas From Earlier Work that was Abandoned

This document is based on work done earlier in the DNSOP working group, and personal drafts before that.

In that earlier work, "<reverse-ip>.{in-addr,ip6}.arpa" was proposed as the domain name to allow for the possibility of DNSSEC-signed responses. However, it was pointed out that people often do not control their reverse IP names and thus their ISP (or their ISP's ISP) could spoof responses and make them look legitimate by signing with DNSSEC.

In an earlier version of this draft, a second way to get the resolver information was specified: using a query to a well-known URI over HTTPS, possibly with authentication. Many participants in the ADD Working Group in early 2020 disagreed with specifying this transport because the IP address being used was most likely obtained by the stub resolver in an insecure fashion, so using an authenticated method could lead to inappropriate assumptions about the security of the answer.

Acknowledgments

The idea of various types of servers publishing information about themselves has been around for decades. However this idea has not been used in the DNS. This document aims to fix this omission.

Roy Arends contributed many ideas to an earlier version of this draft before it was moved to the ADD working group.

Authors' Addresses

Puneet Sood
Google

Email: puneets@google.com

Paul Hoffman
ICANN

Email: paul.hoffman@icann.org

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 1, 2021

P. Sood
Google
P. Hoffman
ICANN

June 30, 2020

Upgrading Communication from Stub Resolvers to DoT or DoH
draft-pp-add-stub-upgrade-02

Abstract

This document describes methods for a DNS stub resolver to upgrade its communications with a known recursive resolver to include encrytion using DoT or DoH. This protocol is designed for the scenario where the stub resolver already has the IP address of the recursive resolver.

Other protocols under develpment address scenarios where the stub resolver wants to discover recursive resolvers that use DoT or DoH. This document does not cover such discovery.

Table of Contents

1. Introduction

A stub resolver (hereafter called "a stub") using traditional DNS over port 53 may wish to use encrypted communication with the recursive resolver (hereafter called "a resolver"). In such a scenario, the stub needs to know how to probe the resolver to find out if it can use encrypted communication. This document describes a mechanism for a stub that knows the IP address of the resolver to do so. It is assumed that the IP address was received insecurely, such as through DHCP.

The method in this document assumes that a stub wants to attempt to upgrade its communication with the resolver to either DNS-over-TLS (DoT, [RFC7858]) or DNS-over-HTTPS (DoH, [RFC8484]). The method is basically to use a DNS request as defined in [I-D.pp-add-resinfo] to get information about whether the resolver supports DoT or DoH. The method can later be extended to other secure transports for stub-to-resolver communication transports.

1.1. Definitions

In the rest of this document, the term "resolver" without qualification means "recursive resolver" as defined in [RFC8499]. Also, the term "stub" is used to mean "stub resolver".

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Using RESINFO Responses for Upgrade

This document defines two entries for the IANA DNS Resolver Information Registry that is defined in [I-D.pp-add-resinfo].

2.1. Contacting This Resolver Using DoH

The "doh-templates" name is used to specify the URI template or templates that can be used by the stub resolver for DoH queries. The value MUST be an array of URI templates. Each element of the array in the value is a JSON string. The host part of the URI template MUST be an IP address.

[[For future: maybe drop the "MUST be an IP address" restriction and say that it can be either an IP address or host name.]]

The array in the value can be empty, which indicates that the resolver does not offer DoH service. An empty array and the absence of a name/value pair for "doh-templates" have identical meanings.

The value of "doh-templates" is an array of strings instead of just one string because a resolver might have more than one IP address or URL paths. The order of the elements in the array has no meaning; that is, the array could instead be considered a set.

[[This section needs to be updated to handle DoH over HTTP/3. These updates then need to be reflected in Section 3.]]

2.2. Contacting This Resolver Using DoT

The "dot-ports" name is used to specify the port(s) that can be used by the stub resolver for DoT queries. The value MUST be an array of port numbers. Each element of the array in the value is a JSON number.

The value of "dot-ports" is an array of numbers instead of just one number because a resolver might support DoT on more than one port. The order of the elements in the array has no meaning; that is, the array could instead be considered a set.

The array in the value can be empty, which indicates that the resolver does not offer DoT service. An empty array and the absence of a name/value pair for "dot-ports" have identical meanings.

[[For future: maybe add "dot-hostnames" to enable authentication.]]

2.3. Examples

A resolver has two IP addresses, 192.0.2.222 and 203.0.113.77. It offers DoH service, and offers DoT service on the default port. It's response to the RESINFO query might be either one of:

```
{ "dot-ports": [ 853 ], "doh-templates":
   [ "https://203.0.113.77//dns-query{?dns}",
   "https://192.0.2.222//dns-query{?dns}" ] }
```

A resolver does not offer DoH service, but does offer DoT service on the default port. It's response to the RESINFO query might be either one of:

```
{ "dot-ports": [ 853 ], "doh-templates": [] }
```

or

```
{ "dot-ports": [ 853 ] }
```

3. Method Overview

The pseudocode for the method is:

```
# Things the stub resolver knows
# dohCapable    Does the stub know how to do DoH
# dotCapable    Does the stub know how to do DoT
# resIP         IP address of resolver
# upgradeNoAuth Does the stub want to upgrade even if it can't
                authenticate the TLS session
# insecureOK    Does the stub want to use unauthenticated classic
                DNS if DoH/DoT upgrades fail
```

[[Need to fix dohCapable to deal with DoH templates that point to resolvers other than the one queried.]]

```
   if dohCapable:
       send a DNS query of resolver-info.arpa/IN/RESINFO
       if there is a non-empty "doh-templates" name in the response:
           for each template in the name/value pair:
               start TLS session on resIP, port from DoH template
               if it succeeds
                   if it authenticates correctly
                       resolve the URI template
                       if 200-level response
                           use result to do DoH; finished
                       else if 300-level response
                               follow redirect, act appropriately
                       else if 400-level response
                               continue
                   else if upgradeNoAuth:
                       resolve the URI template
                       if 200-level response
                           use result to do DoH; finished
                       else if 300-level response
                           follow redirect, act appropriately
                       else if 400-level response
                           continue
                   else
                       continue
               else
                   continue
           # no DoH template worked

   if dotCapable:
       send a DNS query of resolver-info.arpa/IN/RESINFO
       if there is a non-empty "dot-ports" name in the response:
           for each port in the name/value pair:
               start TLS session on resIP and the port number
               if it succeeds
                   if it authenticates correctly
                       start doing DoT; finished
                   else if upgradeNoAuth:
                       start doing DoT; finished
                   else
                       continue
               else
                   continue
           # no DoT port worked

   if insecureOK:
       Use unencrypted DNS on port 53
   else
       DNS transport setup failed
```

3.1. Order of Desired Protocols

The pseudocode in the previous section attempts to use DoH, DoT, and unencrypted DNS, in that order. This is done to keep the pseudocode simple while demonstrating one possible order of transport selection. A stub implementation could attempt some or all of the available DNS transports in an implementation-specific or user-defined order. For example, possible lists of transports to attempt might be:

- DoH, DoT, classic DNS
- DoT, DoH
- DoT, classic DNS
- Classic DNS

4. Method Details

4.1. Inputs to the Process

The method described here requires the following information. It is listed with variable names from the pseudocode in Section 3.

resIP The IP address of resolver. This can be either an IPv4 or IPv6 address.

dohCapable Set to true if the stub knows how to be a DoH client

dotCapable Set to true if the stub knows how to be a DoT client

upgradeNoAuth Set to true the stub wants to use unauthenticated DoT or DoH if it is available. Note that using unauthenticated DoT or DoH is inherently insecure because an on-path attacker can impersonate the resolver.

insecureOK Set to true if the stub wants to keep using classic (unencrypted) DNS on port 53 if the attempt to upgrade fails. Note that setting this to false will cause further DNS queries to fail if upgrade fails.

[[Add some possible implementation examples. Here's one.]]

For example, if an OS implementation's design is "just try TLS on port 853 of the current resolver", resIP is the resolver address, dohCapable is false, dotCapable is true, and upgradeNoAuth is set to true.

4.2. TLS Authentication

In this mechanism, the stub has an IP address of the resolver. It does not necessarily have a domain name associated with that IP address.

In order to authenticate TLS sessions, the stub resolver must have a set of TLS trust anchors, such as those maintained by some operating systems.

If the stub has a domain name associated with the resolver's IP address, and if the resolver uses that domain name in one of the subject identifiers in its certificate during the TLS exchange, the stub can use the domain name for authentication of the TLS session.

The stub always has an IP address for the resolver. If the resolver uses the same IP address used by the stub in one of the subject identifiers in its certificate during the TLS exchange, the stub can use the IP address for authentication of the TLS session.

A resolver that uses this method to publish its information SHOULD, if possible, have a TLS certificate whose subject identifiers contain any of the IP addresses that stubs might be using for the resolver. At the time that this document is published, getting IP addresses in TLS certificates is possible, but there are only a few widely-trusted CAs that issue such certificates. [RFC8738] describes a protocol that may cause IP address certificates to become more common.

5. IANA Considerations

This document defines two entries for the IANA DNS Resolver Information Registry that is defined in [I-D.pp-add-resinfo].

5.1. Registration for doh-templates in the IANA DNS Resolver Information Registry

Name: doh-templates

Value type: Array of strings

Specification: This document, Section 2.1

5.2. Registration for dot-ports in the IANA DNS Resolver Information Registry

Name: dot-ports

Value type: Array of numbers

Specification: This document, Section 2.2

6. Security Considerations

The method described in this document explicitly allows a stub to perform DNS communications over traditional unencrypted, unauthenticated DNS on port 53.

The method described in this document explicitly allows a stub to choose to allow unauthenticated TLS. In this case, the resulting communication will be susceptible to obvious and well-understood attacks from an attacker in the path of the communications.

7. References

7.1. Normative References

[I-D.pp-add-resinfo]
Sood, P. and P. Hoffman, "DNS Resolver Information Self-publication", draft-pp-add-resinfo-01 (work in progress), May 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

7.2. Informative References

[RFC8738] Shoemaker, R., "Automated Certificate Management Environment (ACME) IP Identifier Validation Extension", RFC 8738, DOI 10.17487/RFC8738, February 2020, <https://www.rfc-editor.org/info/rfc8738>.

Authors' Addresses

Puneet Sood
Google

Email: puneets@google.com

Paul Hoffman
ICANN

Email: paul.hoffman@icann.org

ADD
Internet-Draft
Intended status: Informational
Expires: December 25, 2020

T. Reddy
McAfee
D. Wing
Citrix

June 23, 2020

DNS-over-HTTPS and DNS-over-TLS Server Deployment Considerations for Enterprise Networks
draft-reddy-add-enterprise-00

Abstract

This document discusses DoH/DoT deployment considerations for Enterprise networks. It particularly sketches the required steps to use DNS-over-TLS (DoT) and/or DNS-over-HTTPS (DoH) server provided by the Enterprise network.

One of the goals of the document is to assess to what extent existing tools can be used to provide such service.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 25, 2020.

Table of Contents

1. Introduction

[RFC7626] discusses DNS privacy considerations in both "on the wire" (Section 2.4 of [RFC7626]) and "in the server" (Section 2.5 of [RFC7626]) contexts. In recent years there has also been an increase in the availability of "public resolvers" [RFC8499] which DNS clients may be pre-configured to use instead of the default network resolver for a variety of reasons (e.g., offer a good reachability, support an encrypted transport, provide a strong privacy policy, (lack of) filtering).

If public (DoT) [RFC7858] or DNS-over-HTTPS (DoH) [RFC8484] servers are used instead of using local DNS servers, it can adversely impact Enterprise network-based security. Various network security services are provided by Enterprise networks to protect endpoints (e.g., laptops, printers, IoT devices), and to enforce enterprise policies. These policies may be necessary to protect employees, customers, or citizens. They are not the subject of this memo.

Enterprise DNS servers in place for these purpose act on DNS requests originating from endpoints. However, if an endpoint uses public DoT or DoH servers, the desired enterprise protection and enforcement can be bypassed.

In order to act on DNS requests from endpoints, network security services can block DoT traffic by dropping outgoing packets to destination port 853. Identifying DoH traffic is far more challenging than DoT traffic. Network security services may try to identify the well-known DoH resolvers by their domain name, and DNS-over-HTTPS traffic can be blocked by dropping outgoing packets to these domains. However, DoH traffic can not be fully identified without acting as a TLS proxy.

If a network security service blocks access to the public DoH/DoT server, there are incompatibilities with the privacy profiles discussed in [RFC8310]:

- If an endpoint has enabled strict privacy profile (Section 5 of [RFC8310]), the endpoint cannot resolve DNS names.

- If an endpoint has enabled opportunistic privacy profile (Section 5 of [RFC8310]), the endpoint will either fallback to an encrypted connection without authenticating the DNS server provided by the local network or fallback to clear text DNS, and cannot exchange encrypted DNS messages. The fallback adversely impacts security and privacy as internal attacks are possible in Enterprise networks. For example, an internal attacker can modify the DNS responses to re-direct the client to malicious servers or pervasively monitor the DNS traffic. The reader may refer to Section 3.2.1 of [I-D.arkko-farrell-arch-model-t] for a discussion on the need for more awareness about attacks from within closed domains.

To overcome the above threats, this document specifies mechanisms to configure endpoints to use Enterprise provided DoT and DoH servers, and bootstrap IoT devices and unmanaged endpoints to discover and authenticate the DoT and DoH servers provided by the Enterprise network.

A common usage pattern for an IoT device is for it to "call home" to a service that resides on the public Internet, where that service is referenced through a domain name (A or AAAA record). As discussed in Manufacturer Usage Description Specification [RFC8520], because these devices tend to require access to very few sites, all other access should be considered suspect. However, if the query is not accessible for inspection, it becomes quite difficult for the infrastructure to suspect anything.

This document focuses on DoH/DoT deployment considerations for Enterprise networks, DoH/DoT sever discovery and deployment considerations for home networks are discussed in [I-D.btw-add-home].

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499] and [I-D.ietf-dnsop-terminology-ter].

'DoH/DoT' refers to DNS-over-HTTPS and/or DNS-over-TLS.

3. IT-owned devices

If a device is managed by an enterprise's IT department, the device can be configured to use Enterprise-provided DoH/DoT servers. This configuration might be manual or rely upon whatever deployed device management tool in an Enterprise. For example, customizing Firefox using Group Policy to use the Enterprise DoH server is discussed in [Firefox-Policy] for Windows and MacOS, and setting Chrome policies is discussed in [Chrome-Policy] and [Chrome-DoH].

4. IoT Devices

The solution described in this document is aimed in general at non-constrained IoT devices (i.e., class 2+ [RFC7228]) operating on a Enterprise network without a device management tool and require agentless or standardized approaches. The basis for trust, therefore, is quite different from that of a laptop, tablet, or smart phone. The following bootstrapping mechanisms can be used to securely provision IoT devices to use Enterprise provided DoT and DoH servers:

- IoT devices supporting Bootstrapping Remote Secure Key Infrastructures (BRSKI) discussed in [I-D.ietf-anima-bootstrapping-keyinfra] can be bootstrapped with the Enterprise-provided DoH/DoT servers using the mechanism discussed in Section 5 of [I-D.reddy-add-iot-byod-bootstrap].

- [RFC8572] defines a bootstrapping strategy for enabling devices to securely obtain the required configuration information with no installer input. DHCP/RA [I-D.btw-add-home] can be used to discover the DoH/DoT information. If the insecurely discovered DoH/DoT information is not pre-configured in the IoT device, the client can validate the Policy Assertion Token signature (Section 7 [I-D.reddy-add-server-policy-selection]) using the owner certificate (Section 3.2 of [RFC8572]).

- When IoT devices connect to a network via EAP methods such as Tunnel Extensible Authentication Protocol (TEAP) [RFC7170], it would be possible to extend these methods to return additional configuration elements as part of completion of the authentication transaction. One simple approach would be after successful completion of the EAP method in Phase 2 for a TEAP server to return a new TLV that indicates the local DoH/DoT information.

- Not all of IoT devices support 802.1x supplicant and need an alternate mechanism to connect to the Enterprise network. To address this limitation, unique pre-shared keys are created for each IoT device and WPA-PSK is used [PSK]. In other words, WPA-PSK is used with unique pre-shared keys for different IoT devices to deal security issues.

  * The IoT device needs to be provisioned with a Pre-Shared Key (PSK) for mutual authentication. The PSK is only known to the IoT device and the WPA server. In this case, the bootstrapping mechanism discussed in Section 4 of [I-D.reddy-add-iot-byod-bootstrap] may be used to securely bootstrap IoT device with the authentication domain name (ADN) and DNS server certificate of the local network's DoH/ DoT server. It uses password-based authenticated key exchange (PAKE) scheme to authenticate the EST server and fetch the DoH/ DoT server certificate. Note that provisioning massive number of IoT devices with PSK is not a scalable onboarding mechanism but will work in Small Office/Home Office (SOHO) and Small/ Medium Enterprise (SME).

- If Device Provisioning Protocol (DPP) [dpp] is used, the configurator can securely configure IoT devices with the local DoH/DoT server by extending the content of the configuration elements provided by the configurator. Because DPP can provide a private shared key for use with WPA-PSK, it can be combined with the above methods.

- The OMA LWM2M specification [oma] defines an architecture where a new device (LWM2M client) contacts a Bootstrap-server which is responsible for "provisioning" essential bootstrap information. The current standard defines the following four bootstrapping modes (1) Factory Bootstrap (2) Bootstrap from Smartcard (3) Client Initiated Bootstrap (4) Server Initiated Bootstrap. The bootstrap information can be extended to include the local DoH/DoT server details.

- The Open Connectivitiy Foundation [ocf] defines the onboarding process before a device is operational. Once the onboarding tool and the new device have authenticated and established secure

communication, the onboarding tool can provision the IoT device with the local DoH/DoT server.

This document does not discuss opportunistic or leap-of-faith bootstrapping methods, they are susceptible to security issues (e.g., IoT device can be configured with the attacker's DoH/DoT server or disable the use of DoH/DoT).

5. BYOD

The following mechanisms can be used to bootstrap BYOD (bring your own device) with the DoH/DoT server used by the Enterprise network:

- If mobile device management (MDM) [MDM-Apple] is used to secure BYOD, MDM can be used to configure OS/browser with the Enterprise provided DoH/DoT server.

- If an endpoint is on-boarded, for example, using Over-The-Air (OTA) enrollment [OTA] to provision the device with a certificate and configuration profile, the configuration profile can include the authentication domain name (ADN) of the DoH/DoT server. The OS/Browser can use the configuration profile to use the Enterprise provided DoH/DoT server. In this case, MDM is not installed on the device.

- If an endpoint uses the credentials (username and password) provided by the IT admin to mutually authenticate to the Enterprise WiFi Access Point (e.g., PEAP-MSCHAPv2 [PEAP], EAP-pwd [RFC8146], EAP-PSK [RFC4764]), the boostrapping mechanism discussed in Section 4 of [I-D.reddy-add-iot-byod-bootstrap] can be used to securely bootstrap the endpoint with the ADN and DNS server certificate of the local network's DoH/DoT server.

  The DNS client uses PAKE scheme to authenticate the EST server using the credentials to authenticate to the network. In this case, the endpoint is neither provisioned with a configuration profile or MDM is installed on the device. Many users have privacy and personal data sovereignty concerns with employers installing MDM on their personal devices; they are concerned that admin can glean personal information and could control how they use their devices. Yet when users do not install MDM on their devices, IT admins do not get visibility into the security posture of those devices.

  To overcome this problem, a host agent can cryptographically attest the security status associated with device, such as minimum passcode length, biometric login enabled, OS version etc. This approach is fast gaining traction especially with the advent of

closed OS like Windows 10 in S mode [win10s] or Chromebook [Chromebook], where applications are sandboxed (e.g., ransomware attack is not possbile) and applications can only be installed via the OS store.

When attached to the enterprise network yet needing to use the enterprise's DoH server only to access the internal-only DNS names, the client device can learn about domains for which the local network's resolver is authoritative via dnsZones key defined in Section 4.3 of [I-D.ietf-intarea-provisioning-domains] (as other DoH/DoT servers will be unaware of the internal-only DNS names).

6. Roaming Enterprise Users

6.1. VPN tunnel

In this Enterprise scenario (Section 1.1.3 of [RFC7296]), a roaming user connects to the Enterprise network through an VPN tunnel (e.g., IPsec, SSL, Wireguard). The split-tunnel Virtual Private Network (VPN) configuration allows the endpoint to access hosts that reside in the Enterprise network [RFC8598] using that tunnel; other traffic not destined to the Enterprise does not traverse the tunnel. In contrast, a non-split- tunnel VPN configuration causes all traffic to traverse the tunnel into the enterprise.

When the VPN tunnel is IPsec, The DoH/DoT server hosted by the Enterprise network can be securely discovered by the endpoint using the INTERNAL_ENC_DNS IKEv2 Configuration Payload Attribute Type defined in [I-D.btw-add-ipsecme-ike]. For split-tunnel VPN configurations, the endpoint uses the Enterprise-provided DoT/DoH server to resolve internal-only domain names. For non-split-tunnel VPN configurations, the endpoint uses the Enterprise-provided DoT/DoH server to resolve both internal and external domain names.

Other VPN tunnel types have similar configuration capabilities, not detailed here.

6.2. Client Authentication

When not on the local enterprise network (e.g., at home or coffee shop) yet needing to access the enterprise DoH/DoT server but not through a tunnel, roaming users can use client authentication to access the Enterprise provided DoH/DoT server. For example, Firefox DoH setting accepts user credentials [Firefox-TRR] to authenticate the client to access the DoH server. The exact client authentication mechanism to authenticate to the DoH/DoT server is outside the scope of this specification.

7. Upstream Encryption

If the Enterprise network is using the local DoH/DoT server configured as a Forwarding DNS server [RFC8499] relying on the upstream resolver (e.g., at an ISP) to perform recursive DNS lookups, DNS messages exchanged between the local DoH/DoT server and recursive resolver MUST be encrypted. If the Enterprise network is using the local DoH/DoT server configured as a recursive DNS server, DNS messages exchanges between the recursive resolver and authoritative servers SHOULD be encrypted to conform to the requirements discussed in [I-D.ietf-dprive-phase2-requirements].

8. Security Considerations

Security and privacy considerations in [I-D.reddy-add-iot-byod-bootstrap] need to be taken into consideration.

The mechanism defined in [I-D.reddy-add-server-policy-selection] can be used by the DNS server to communicate its privacy statement URL and filtering policy to a DNS client. This communication is cryptographically signed to attest to its authenticity.

The DNS client can validate the signatory (i.e., cryptographically attested by the Organization hosting the DoH/DoT server) and the user can review human-readable privacy policy information of the DNS server and assess whether the DNS server performs DNS-based content filtering.

If the discovered DoH/DoT server does not meet the privacy preserving data policy and filtering requirements of the user, the user can instruct the DNS client to take appropriate actions. For example, the action can be to use the local DNS server only to access internal-only DNS names and use another DNS server (adhering with his/her expectations) for public domains.

9. IANA Considerations

This document has no actions for IANA.

10. Acknowledgements

Thanks to Mohamed Boucadair, Sandeep Rao, Vinny Parla, Nancy Cam-Winget and Eliot Lear for the discussion and comments.

11. References

11.1. Normative References

[I-D.reddy-add-iot-byod-bootstrap]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "A Bootstrapping Procedure to Discover and Authenticate DNS-over-TLS and DNS-over-HTTPS Servers for IoT and BYOD Devices", draft-reddy-add-iot-byod-bootstrap-00 (work in progress), May 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

11.2. Informative References

[Chrome-DoH]
The Unicode Consortium, "Chrome DNS over HTTPS (aka DoH)", <https://www.chromium.org/developers/dns-over-https>.

[Chrome-Policy]
The Unicode Consortium, "Chrome policies for users or browsers", <https://support.google.com/chrome/a/answer/2657289?hl=en>.

[Chromebook]
Microsoft, "Chromebook security", <https://support.google.com/chromebook/answer/3438631?hl=en>.

[dpp] Wi-Fi Alliance, "Wi-Fi Device Provisioning Protocol (DPP)", Wi-Fi Alliance , 2018, <https://www.wi-fi.org/download.php?file=/sites/default/files/private/Device_Provisioning_Protocol_Specification_v1.1_1.pdf>.

[Firefox-Policy]
"Policy templates for Firefox", <https://github.com/mozilla/policy-templates/blob/master/README.md#dnsoverhttps>.

[Firefox-TRR]
"Trusted Recursive Resolver", <https://wiki.mozilla.org/Trusted_Recursive_Resolver>.

[I-D.arkko-farrell-arch-model-t]
Arkko, J. and S. Farrell, "Challenges and Changes in the Internet Threat Model", draft-arkko-farrell-arch-model-t-03 (work in progress), March 2020.

[I-D.btw-add-home]
Boucadair, M., Reddy.K, T., Wing, D., and N. Cook, "Encrypted DNS Discovery and Deployment Considerations for Home Networks", draft-btw-add-home-06 (work in progress), May 2020.

[I-D.btw-add-ipsecme-ike]
Boucadair, M., Reddy.K, T., Wing, D., and V. Smyslov, "Internet Key Exchange Protocol Version 2 (IKEv2) Configuration for Encrypted DNS", draft-btw-add-ipsecme-ike-00 (work in progress), April 2020.

[I-D.ietf-anima-bootstrapping-keyinfra]
Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-41 (work in progress), April 2020.

[I-D.ietf-dnsop-terminology-ter]
Hoffman, P., "Terminology for DNS Transports and Location", draft-ietf-dnsop-terminology-ter-01 (work in progress), February 2020.

[I-D.ietf-dprive-phase2-requirements]
Livingood, J., Mayrhofer, A., and B. Overeinder, "DNS Privacy Requirements for Exchanges between Recursive Resolvers and Authoritative Servers", draft-ietf-dprive-phase2-requirements-01 (work in progress), June 2020.

[I-D.ietf-intarea-provisioning-domains]
Pfister, P., Vyncke, E., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", draft-ietf-intarea-provisioning-domains-11 (work in progress), January 2020.

[I-D.reddy-add-server-policy-selection]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "DNS Server Selection: DNS Server Information with Assertion Token", draft-reddy-add-server-policy-selection-03 (work in progress), June 2020.

[MDM-Apple]
Apple, "Mobile Device Management", <https://developer.apple.com/documentation/devicemanagement>.

[ocf] Open Connectivity Foundation, "OCF Security Specification", Open Connectivitiy Foundation , June 2017, <https://openconnectivity.org/specs/OCF_Security_Specification_v1.0.0.pdf>.

[oma] Open Mobile Alliance, "Lightweight Machine to Machine Technical Specification: Core", Open Mobile Alliance , June 2019, <http://www.openmobilealliance.org/release/LightweightM2M/V1_1_1-20190617-A/OMA-TS-LightweightM2M_Core-V1_1_1-20190617-A.pdf>.

[OTA] Apple, "Over-the-Air Profile Delivery Concepts", <https://developer.apple.com/library/archive/documentation/NetworkingInternet/Conceptual/iPhoneOTAConfiguration/OTASecurity/OTASecurity.html>.

[PEAP] Microsoft, "[MS-PEAP]: Protected Extensible Authentication Protocol (PEAP)", <https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-peap/5308642b-90c9-4cc4-beec-fb367325c0f9>.

[PSK] Cisco, "Identity PSK Feature Deployment Guide", <https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_Identity_PSK_Feature_Deployment_Guide.html>.

[RFC4764] Bersani, F. and H. Tschofenig, "The EAP-PSK Protocol: A Pre-Shared Key Extensible Authentication Protocol (EAP) Method", RFC 4764, DOI 10.17487/RFC4764, January 2007, <https://www.rfc-editor.org/info/rfc4764>.

[RFC7170] Zhou, H., Cam-Winget, N., Salowey, J., and S. Hanna, "Tunnel Extensible Authentication Protocol (TEAP) Version 1", RFC 7170, DOI 10.17487/RFC7170, May 2014, <https://www.rfc-editor.org/info/rfc7170>.

[RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <https://www.rfc-editor.org/info/rfc7228>.

[RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <https://www.rfc-editor.org/info/rfc7296>.

[RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <https://www.rfc-editor.org/info/rfc7626>.

[RFC8146] Harkins, D., "Adding Support for Salted Password Databases to EAP-pwd", RFC 8146, DOI 10.17487/RFC8146, April 2017, <https://www.rfc-editor.org/info/rfc8146>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <https://www.rfc-editor.org/info/rfc8520>.

[RFC8572] Watsen, K., Farrer, I., and M. Abrahamsson, "Secure Zero Touch Provisioning (SZTP)", RFC 8572, DOI 10.17487/RFC8572, April 2019, <https://www.rfc-editor.org/info/rfc8572>.

[RFC8598] Pauly, T. and P. Wouters, "Split DNS Configuration for the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 8598, DOI 10.17487/RFC8598, May 2019, <https://www.rfc-editor.org/info/rfc8598>.

[win10s] Microsoft, "Windows 10 in S mode", <https://www.microsoft.com/en-us/windows/s-mode>.

Authors' Addresses

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: kondtir@gmail.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

ADD WG
Internet-Draft
Intended status: Standards Track
Expires: January 27, 2021

T. Reddy
McAfee
D. Wing
Citrix
M. Richardson
Sandelman Software Works
M. Boucadair
Orange

July 26, 2020

A Bootstrapping Procedure to Discover and Authenticate DNS-over-TLS and DNS-over-HTTPS Servers for IoT and BYOD Devices
draft-reddy-add-iot-byod-bootstrap-01

Abstract

This document specifies mechanisms to bootstrap endpoints (e.g., hosts, IoT devices) to discover and authenticate DNS-over-TLS and DNS-over-HTTPS servers provided by a local network for IoT/BYOD devices in Enterprise networks.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 27, 2021.

Copyright Notice

Table of Contents

1. Introduction

Traditionally a caching DNS server has been provided by local networks. This provides benefits such as low latency to reach that DNS server (owing to its network proximity to the endpoint). However, if an endpoint is configured to use Internet-hosted or public DNS-over-TLS (DoT) [RFC7858] or DNS-over-HTTPS (DoH) [RFC8484] servers, any available local DNS server cannot serve DNS requests from local endpoints. If public DNS servers are used instead of using local DNS servers, some operational problems can occur such as those listed below:

- "Split DNS" [RFC2775] to use the special internal-only domain names (e.g., "internal.example.com") in enterprise networks will not work, and ".local" and "home.arpa" names cannot be locally resolved in home networks.

- Content Delivery Networks (CDNs) that map traffic based on DNS may lose the ability to direct end-user traffic to a nearby service-specific cluster in cases where a DNS service is being used that

is not affiliated with the local network and which does not send "EDNS Client Subnet" (ECS) information [RFC7871] to the CDN's DNS authorities [CDN].

If public DNS servers are used instead of local DNS servers, the following discusses the impacts on network-based security:

- Various network security services are provided by Enterprise networks to protect endpoints (e.g,. Hosts, IoT devices). Network-based security solutions such as firewalls (FW) and Intrusion Prevention Systems (IPS) rely on network traffic inspection to implement perimeter-based security policies. The network security services may for example prevent malware download, block known malicious URLs, enforce use of strong ciphers, stop data exfiltration, etc. These network security services act on DNS requests originating from endpoints. However, if an endpoint is configured to use public DoH/DoT servers, network security services cannot act on DNS requests from these endpoints.

- In order to act on DNS requests from endpoints, network security services can block DoT traffic by dropping outgoing packets to destination port 853. Identifying DoH traffic is far more challenging than DoT traffic. Network security services may try to identify the domains offering DoH servers, and DoH traffic can be blocked by dropping outgoing packets to these domains. If an endpoint has enabled strict privacy profile (Section 5 of [RFC8310]), and the network security service blocks the traffic to the public DNS server, the DNS service won't be available to the endpoint and ultimately the endpoint cannot access Internet-reachable services.

- If an endpoint has enabled opportunistic privacy profile (Section 5 of [RFC8310]), and the network security service blocks traffic to the public DNS server, the endpoint will either fallback to an encrypted connection without authenticating the DNS server provided by the local network or fallback to clear text DNS, and cannot exchange encrypted DNS messages.

If the network security service fails to block DoH/DoT traffic, this can compromise the endpoint security; some of the potential security threats are listed below:

- The network security service cannot prevent an endpoint from accessing malicious domains.

- If the endpoint is an IoT device which is configured to use public DoH/DoT servers, and if a policy enforcement point in the local

network is programmed using, for example, a Manufacturer Usage Description (MUD) file [RFC8520] by a MUD manager to only allow intended communications to and from the IoT device, the policy enforcement point cannot enforce the network Access Control List (ACL) rules based on domain names (Section 8 of [RFC8520]).

If the network security service successfully blocks DoT and DoH traffic, this can still compromise the endpoint security and privacy; some of the potential security threats are listed below:

- Networks are susceptible to internal attacks as discussed in Section 3.2 of [I-D.arkko-farrell-arch-model-t]. An internal attacker can modify the DNS responses to re-direct the client to malicious servers.

- Pervasive monitoring of DNS traffic.

In addition, the local network's DNS server is advertised using DHCP/RA which is insecure and also provides no mechanism to securely authenticate the DNS server. To overcome the above threats, this document specifies a mechanism to bootstrap endpoints to discover and authenticate the DoT and DoH servers provided by their local network. The overall procedure can be structured into the following steps:

- Bootstrapping (Section 4) is necessary only when connecting to a new network or when the network's DNS certificate has changed. Bootstrapping procedure authenticates the Enrollment over Secure Transport (EST) [RFC7030] server to the endpoint. After authenticating the EST server, DNS server certificate used by the local network is downloaded to the endpoint. This DNS server certificate enables subsequent authenticated encrypted communication with the local DNS server (e.g., DoH) during in the connection phase.

- Connection handshake and service invocation (Section 6): The DNS client initiates a TLS handshake with the DNS server learned in the discovery phase, and validates the DNS server's identity using the credentials obtained in the bootstrapping phase.

Note: The strict and opportunistic privacy profiles as defined in [RFC8310] only applies to DoT protocol, there has been no such distinction made for DoH protocol.

2. Scope

The problems discussed in Section 1 will be encountered in Enterprise networks. Typically Enterprise networks do not assume that all devices in their network are managed by the IT team or Mobile Device

Management (MDM) devices, especially in the quite common BYOD ("Bring Your Own Device") scenario. The mechanisms specified in this document can be used by BYOD devices to discover and authenticate DoT and DoH servers provided by the Enterprise network. This mechanism can also be used by IoT devices (managed by IT team) after onboarding to discover and authenticate DoT and DoH servers provided by the Enterprise network.

Wireless LAN as frequently deployed is vulnerable to various attacks ([Evil-Twin],[Krack] and [Dragonblood]). Because of these attacks, only cryptographically authenticated communications are trusted on Wireless LAN networks. This means information provided by such networks via DHCP, DHCPv6, or RA (e.g., NTP server, DNS server, default domain) are untrusted because DHCP and RA are not authenticated. [I-D.btw-add-home] discusses DoH/DoT server discovery using DHCP/RA but requires the DoH/DoT server to be pre-configured in the endpoint (OS or Browser) or the DNS client must be able cryptographically identify it is connecting to a DoT/DoH server hosted by a specific organization (e.g., ISP or Enterprise) (see [I-D.reddy-add-server-policy-selection]) to prevent the client from connecting to a attackers server.

Users have to indicate to their system in some way that they desire bootstrapping to be performed only when connecting to a specific network (e.g., organization for which a user works or a user works temporarily within another corporation), similar to the way users disable VPN connection in specific network (e.g., Enterprise network) and enable VPN connection by default in other networks. If the discovered DNS server meets the privacy preserving data policy requirements of the user, the user can select to use the discovered DoT and DoH servers.

3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499] and [I-D.ietf-dnsop-terminology-ter].

'DoH/DoT' refers to DNS-over-HTTPS and/or DNS-over-TLS.

4. Bootstrapping Endpoint Devices

If a device is managed by an enterprise's IT department, the device can be configured to use Enterprise-provided DoH/DoT servers. This configuration might be manual or rely upon whatever deployed device management tool in an Enterprise. For example, customizing Firefox using Group Policy to use the Enterprise DoH server is discussed in [Firefox-Policy] for Windows and MacOS, and setting Chrome policies is discussed in [Chrome-Policy] and [Chrome-DoH].

If mobile device management (MDM) (e.g, [MDM-Apple]) is used to secure endpoint, MDM can be used to configure OS/browser with the Enterprise provided DoH/DoT server. If an endpoint is on-boarded, for example, using Over-The-Air (OTA) enrollment [OTA] to provision the device with a certificate and configuration profile, the configuration profile can include the authentication domain name (ADN) of the DoH/DoT server. The OS/Browser can use the configuration profile to use the Enterprise provided DoH/DoT server. In this case, MDM is not installed on the device.

4.1. Bootstrapping BYOD

This section focuses on bootstrapping Bring your own device (BYOD) to discover and authenticate DoH/DoT server provided by the enterprise network but without MDM or configuration profile on the endpoint. If an endpoint uses the credentials (username and password) provided by the IT admin to mutually authenticate to the Enterprise WLAN Access Point (e.g., PEAP-MSCHAPv2 [PEAP], EAP-pwd [RFC8146], EAP-PSK [RFC4764]), the following steps can be used to securely bootstrap the endpoint with the authentication domain name (ADN, defined in [RFC8310]) and DNS server certificate of the local network's DoH/DoT server:

1. The endpoint authenticates to the local network and discovers the Enrollment over Secure Transport (EST) [RFC7030] server using the procedure discussed in Section 7.

2. The endpoint establishes provisional TLS connection with that EST server, i.e., the endpoint provisionally accepts the unverified TLS server certificate. However, the endpoint MUST authenticate the EST server before it accepts the DNS server certificate. The endpoint either uses password-based authenticated key exchange (PAKE) with TLS 1.3 [I-D.barnes-tls-pake] as an authentication method or uses the mutual authentication protocol for HTTP [RFC8120] to authenticate the discovered EST server.

   As a reminder, PAKE is an authentication method that allows the use of usernames and passwords over unencrypted channels without

revealing the passwords to an eavesdropper. Similarly, the mutual authentication for HTTP is based on PAKE and provides mutual authentication between an HTTP client and an HTTP server using username and password as credentials. The cryptographic algorithms to use with the mutual authentication protocol for HTTP are defined in [RFC8121].

Note that the Crypto Forum Research Group (cfrg) has selected draft-haase-cpace and draft-krawczyk-cfrg-opaque drafts to recommend for balanced and augmented password-based authenticated key establishment in IETF protocols. This step will be further updated.

3. The endpoint needs to use PAKE scheme to perform authentication the first time it connects to an EST server. If the EST server authentication is successful, the server's identity can be used to authenticate subsequent TLS connections to that EST server. The endpoint configures the reference identifier for the EST server using the DNS-ID identifier type in the EST server certificate. On subsequent connections to the EST server, the endpoint MUST validate the EST server certificate using the Implict Trust Anchor database (i.e, the EST server certificate must pass PKIX certification path validation [RFC6125]) and match the reference identifier against the EST server's identity according to the rules specified in Section 6.4 of [RFC6125].

4. The endpoint learns the End-Entity certificates [RFC8295] from the EST server. The certificate provisioned to the DNS server in the local network will be treated as a End-Entity certificate. As a reminder, the End-Entity certificates must be validated by the endpoint using an authorized trust anchor (Section 3.2 of [RFC8295]). The endpoint needs to identify the certificate provisioned to the DNS server. The SRV-ID identifier type [RFC6125] within subjectAltName entry MUST be used to identify the DNS server certificate.

   For example, DNS server certificate will include SRV-ID "_domain-s.example.net" along with DNS-ID "example.net". The SRV service label "domain-s" is defined in Section 6 of [RFC7858] for DoT protocol. The SRV service label "doh" is defined in Section 11 for DoH protocol.

5. The endpoint configures the authentication domain name (ADN) (defined in [RFC8310]) for the DNS server from the DNS-ID identifier type within subjectAltName entry in the DNS server certificate. The DNS server certificate is associated with the ADN to be matched with the certificate given by the DNS server in

TLS. To some extent, this approach is similar to certificate usage PKIX-EE(1) defined in [RFC7671].

Figure 1 illustrates a sequence diagram for bootstrapping an endpoint with the local network's ADN and DNS server certificate.

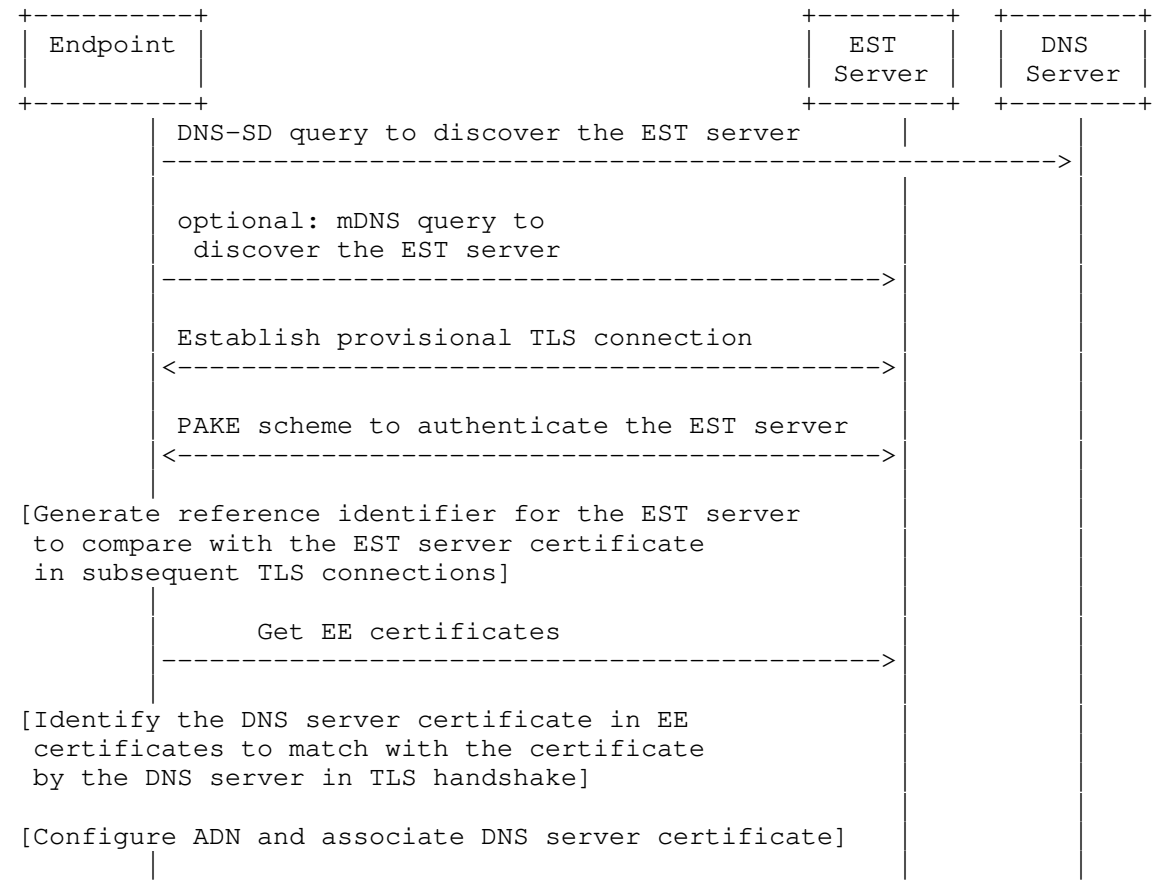

Figure 1: Bootstrapping Endpoint Devices

5. Bootstrapping IoT Devices

The following steps explain the mechanism to bootstrap IoT devices supporting Bootstrapping Remote Secure Key Infrastructures (BRSKI) discussed in [I-D.ietf-anima-bootstrapping-keyinfra] with local network's CA certificates, ADN and DNS server certificate:

- Bootstrapping Remote Secure Key Infrastructures (BRSKI) discussed in [I-D.ietf-anima-bootstrapping-keyinfra] provides a solution for

secure automated bootstrap of devices. BRSKI specifies means to provision credentials on devices to be used to operationally access networks. In addition, BRSKI provides an automated mechanism for the bootstrap distribution of CA certificates from the EST server. The IoT device can use BRSKI to bootstrap the IoT device using the IoT manufacturer provisioned X.509 certificate, in combination with a registrar provided by the local network and IoT device manufacturer's authorizing service (MASA):

1. The IoT device authenticates to the local network using the IoT manufacturer provisioned X.509 certificate. The IoT device can request and get a voucher from the MASA service via the registrar. The voucher is signed by the MASA service and includes the local network's CA public key.

2. The IoT device validates the signed voucher using the manufacturer installed trust anchor associated with the MASA, stores the CA's public key and validates the provisional TLS connection to the registrar.

3. The IoT device requests the full EST distribution of current CA certificates (Section 5.9.1 in [I-D.ietf-anima-bootstrapping-keyinfra]) from the registrar operating as a BRSKI-EST server. The IoT devices stores the CA certificates as Explicit Trust Anchor database entries. The IoT device uses the Explicit Trust Anchor database to validate the DNS server certificate.

4. The IoT device learns the End-Entity certificates from the BRSKI-EST server. The certificate provisioned to the DNS server in the local network will be treated as an End-Entity certificate. The IoT device needs to identify the certificate provisioned to the DNS server. The SRV-ID identifier type within subjectAltName entry MUST be used to identify the DNS server certificate (see Step 4 in Section 4.1).

5. The endpoint configures the ADN for the DNS server from the DNS-ID identifier type within subjectAltName entry in the DNS server certificate. The DNS server certificate is associated with the ADN to be matched with the certificate given by the DNS server in TLS.

6. Connection Handshake and Service Invocation

The DNS client resolves the ADN using the mechanism discussed in Section 7.2 of [RFC8310]. The DNS client initiates TLS handshake with the DNS server, the DNS server presents its certificate in ServerHello message, and the DNS client MUST match the DNS server

certificate downloaded in Step 4 in Section 4.1 or Section 5 with the certificate provided by the DNS server in TLS handshake. If the match is successful, the DNS client MUST validate the server certificate using an authorized trust anchor.

If the match is successful and server certificate is successfully validated, the client continues with the connection as normal. Otherwise, the client MUST treat the server certificate validation failure as a non-recoverable error. If the DNS client cannot reach or establish an authenticated and encrypted connection with the privacy-enabling DNS server provided by the local network, the DNS client can fallback to a privacy-enabling public DNS server.

The DoH client contacts the DoH resolver to retrieve the list of supported DoH services using the well-known URI defined in [I-D.btw-add-rfc8484-clarification].

7. EST Service Discovery Procedure

An EST client discovers the EST server in the local network by using DNS-based Service Discovery (DNS-SD) [RFC6763] or Multicast DNS (mDNS) [RFC6762]. The <Domain> portion specifies the DNS sub-domain where the service instance is registered. It may be "local.", indicating the mDNS local domain, or it may be a conventional domain name such as "example.com.". The <Service> portion of the EST service instance name MUST be "_est._tcp".

A EST client application can proactively discover an EST server being advertised in the site by multicasting a PTR query to the following:

"_est._tcp.local"

An EST server can send out gratuitous multicast DNS answer packets whenever it starts up, wakes from sleep, or detects a change in EST server configuration. EST client application can receive these gratuitous packets and cache information contained in them.

8. Network Reattachment

On subsequent attachments to the network, the endpoint initiates TLS handshake with the DoH/DoT server (configured in Step 5 of Section 4.1 or Section 5) and follows the mechanism discussed in Section 6 to validate the DNS server certificate.

If the DNS server certificate is invalid (e.g., revoked or expired), the endpoint discovers and initiates TLS handshake with the EST server, and uses the validation techniques described in [RFC6125] to compare the reference identifier (created in Step 2 of Section 4.1 in

this document) to the EST server certificate and verifies the entire certification path as per [RFC5280]. The endpoint then gets the DNS server certificate from the EST server. If the DNS-ID identifier type within subjectAltName entry in the DNS server certificate does not match the configured ADN, the ADN is replaced with the DNS-ID identifier type. The DNS server certificate associated with the ADN is replaced with the one provided by the EST server. The endpoint initiates TLS handshake with the newly discovered ADN and follows the mechanism discussed in Section 6 to validate the DNS server certificate.

Figure 2 illustrates a sequence diagram for re-configuring an endpoint with ADN and local network's DNS server certificate on subsequent attachments to the network.

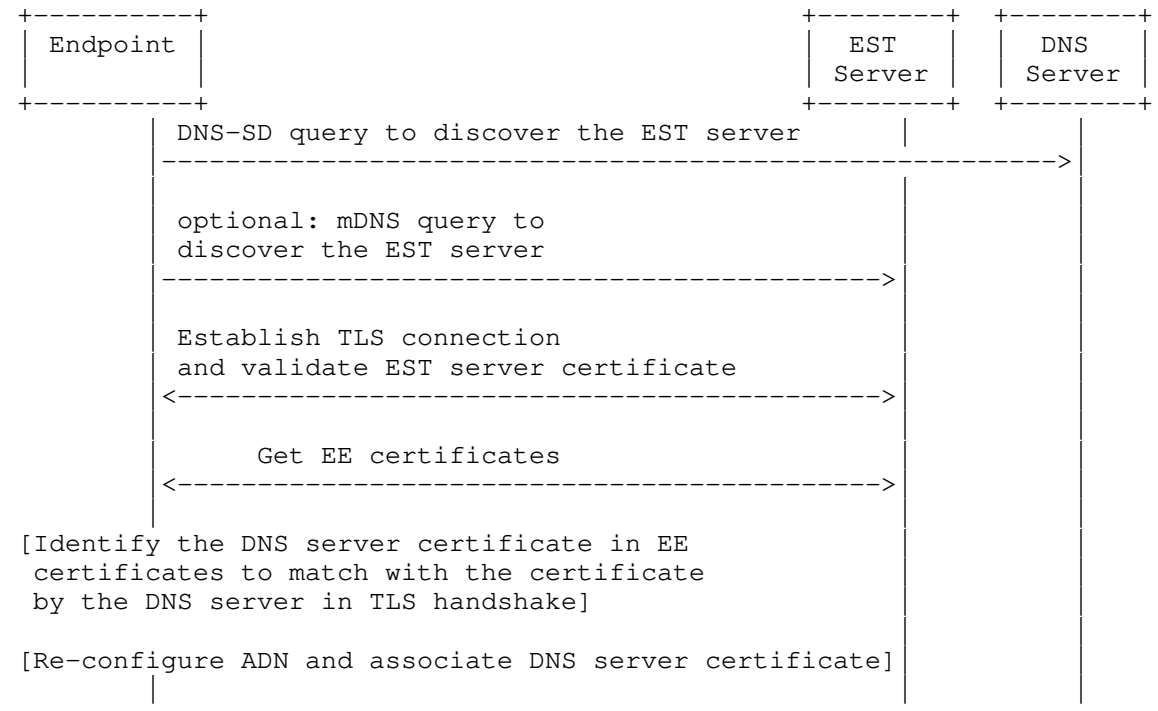

Figure 2: Bootstrapping Endpoint Devices on subsequent attachments to the network

9. Privacy Considerations

[RFC7626] discusses DNS privacy considerations in both "on the wire" (Section 2.4 of [RFC7626]) and "in the server" (Section 2.5 of [RFC7626] contexts. The mechanism defined in [I-D.reddy-add-server-policy-selection] can be used by the DNS server to communicate its privacy statement URL and filtering policy to a DNS client. This communication is cryptographically signed to attest to its authenticity. By evaluating the DNS privacy statement, filtering policy and the signatory, the client can use the discovered DNS server if it meets privacy preserving data policy and filtering requirements of the user.

10. Security Considerations

The bootstrapping procedure to obtain the certificate of the local network's DNS server uses a client identity and password to authenticate the EST server using PAKE schemes. Security considerations such as those discussed in [I-D.barnes-tls-pake] or [RFC8120] and [RFC8121] need to be taken into consideration.

Users cannot be expected to enable or disable the bootstrapping or the discovery procedure as they switch networks. Thus, it is RECOMMENDED that users indicate to their system in some way that they desire bootstrapping to be performed when connecting to a specific network, similar to the way users disable VPN connection in specific network (e.g., Enterprise network) and enable VPN connection by default in other networks.

If an endpoint has enabled strict privacy profile, and the network security service blocks the traffic to the privacy-enabling public DNS server, a hard failure occurs and the user is notified. The user has a choice to switch to another network or if the user trusts the network, the user can enable strict privacy profile with the DoH/DoT server discovered in the network instead of downgrading to opportunistic privacy profile.

The primary attacks against the methods described in Section 7 are the ones that would lead to impersonation of a EST server and spoofing the DNS response to indicate that the network does not support any EST server. To protect against DNS-vectored attacks, secured DNS (DNSSEC) can be used to ensure the validity of the DNS records received. Impersonation of the EST server is prevented by authenticating the EST server using the PAKE scheme. The PAKE scheme is only used once to configure the reference identifier of the EST server and the server certificate is validated for subsequent TLS connections to the EST server.

Security considerations in [I-D.ietf-anima-bootstrapping-keyinfra] need to be taken into consideration for IoT devices.

11. IANA Considerations

11.1. Service Name for EST

IANA is requested to allocate the following service name from the registry available at: https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml.

Service Name:	est
Port Number:	N/A
Transport Protocol(s):	TCP
Description:	Enrollment over Secure Transport (EST)
Assignee:	IESG <iesg@ietf.org>
Contact:	IETF Chair <chair@ietf.org>
Reference:	[ThisDocument]

11.2. Service Name for DoH

IANA is requested to allocate the following service name from the registry available at: https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml.

Service Name:	doh
Port Number:	N/A
Transport Protocol(s):	TCP
Description:	DNS-over-HTTPS
Assignee:	IESG <iesg@ietf.org>
Contact:	IETF Chair <chair@ietf.org>
Reference:	[ThisDocument]

12. Acknowledgments

Thanks to Joe Hildebrand, Harsha Joshi, Shashank Jain, Patrick McManus, Bob Harold, Livingood Jason, Winfield Alister, Eliot Lear, Stephane Bortzmeyer, Ted Lemon and Sara Dickinson for the discussion and comments.

13. References

13.1. Normative References

[I-D.ietf-anima-bootstrapping-keyinfra]
Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-41 (work in progress), April 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <https://www.rfc-editor.org/info/rfc5280>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <https://www.rfc-editor.org/info/rfc6125>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <https://www.rfc-editor.org/info/rfc6762>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <https://www.rfc-editor.org/info/rfc6763>.

[RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <https://www.rfc-editor.org/info/rfc7030>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8121] Oiwa, Y., Watanabe, H., Takagi, H., Maeda, K., Hayashi, T., and Y. Ioku, "Mutual Authentication Protocol for HTTP: Cryptographic Algorithms Based on the Key Agreement Mechanism 3 (KAM3)", RFC 8121, DOI 10.17487/RFC8121, April 2017, <https://www.rfc-editor.org/info/rfc8121>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8295] Turner, S., "EST (Enrollment over Secure Transport) Extensions", RFC 8295, DOI 10.17487/RFC8295, January 2018, <https://www.rfc-editor.org/info/rfc8295>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

13.2. Informative References

[CDN] "End-User Mapping: Next Generation Request Routing for Content Delivery", 2015, <https://conferences.sigcomm.org/sigcomm/2015/pdf/papers/p167.pdf>.

[Chrome-DoH]
The Unicode Consortium, "Chrome DNS over HTTPS (aka DoH)", <https://www.chromium.org/developers/dns-over-https>.

[Chrome-Policy]
The Unicode Consortium, "Chrome policies for users or browsers", <https://support.google.com/chrome/a/answer/2657289?hl=en>.

[Dragonblood]
The Unicode Consortium, "Dragonblood: Analyzing the Dragonfly Handshake of WPA3 and EAP-pwd", <https://papers.mathyvanhoef.com/dragonblood.pdf>.

[Evil-Twin]
The Unicode Consortium, "Evil twin (wireless networks)", <https://en.wikipedia.org/wiki/Evil_twin_(wireless_networks)>.

[Firefox-Policy]
"Policy templates for Firefox", <https://github.com/mozilla/policy-templates/blob/master/README.md#dnsoverhttps>.

[I-D.arkko-farrell-arch-model-t]
Arkko, J. and S. Farrell, "Challenges and Changes in the Internet Threat Model", draft-arkko-farrell-arch-model-t-04 (work in progress), July 2020.

[I-D.barnes-tls-pake]
Barnes, R. and O. Friel, "Usage of PAKE with TLS 1.3", draft-barnes-tls-pake-04 (work in progress), July 2018.

[I-D.btw-add-home]
Boucadair, M., Reddy.K, T., Wing, D., and N. Cook, "Encrypted DNS Discovery and Deployment Considerations for Home Networks", draft-btw-add-home-07 (work in progress), July 2020.

[I-D.btw-add-rfc8484-clarification]
Boucadair, M., Cook, N., Reddy.K, T., and D. Wing, "Supporting Redirection for DNS Queries over HTTPS (DoH)", draft-btw-add-rfc8484-clarification-02 (work in progress), July 2020.

[I-D.ietf-dnsop-terminology-ter]
Hoffman, P., "Terminology for DNS Transports and Location", draft-ietf-dnsop-terminology-ter-01 (work in progress), February 2020.

[I-D.reddy-add-server-policy-selection]
Reddy.K, T., Wing, D., Richardson, M., and M. Boucadair, "DNS Server Selection: DNS Server Information with Assertion Token", draft-reddy-add-server-policy-selection-03 (work in progress), June 2020.

[Krack]
The Unicode Consortium, "Key Reinstallation Attacks", 2017, <https://www.krackattacks.com/>.

[MDM-Apple]
Apple, "Mobile Device Management", <https://developer.apple.com/documentation/devicemanagement>.

[OTA]
Apple, "Over-the-Air Profile Delivery Concepts", <https://developer.apple.com/library/archive/documentation/NetworkingInternet/Conceptual/iPhoneOTAConfiguration/OTASecurity/OTASecurity.html>.

[PEAP] Microsoft, "[MS-PEAP]: Protected Extensible Authentication Protocol (PEAP)", <https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-peap/5308642b-90c9-4cc4-beec-fb367325c0f9>.

[RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, DOI 10.17487/RFC2775, February 2000, <https://www.rfc-editor.org/info/rfc2775>.

[RFC4764] Bersani, F. and H. Tschofenig, "The EAP-PSK Protocol: A Pre-Shared Key Extensible Authentication Protocol (EAP) Method", RFC 4764, DOI 10.17487/RFC4764, January 2007, <https://www.rfc-editor.org/info/rfc4764>.

[RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <https://www.rfc-editor.org/info/rfc7626>.

[RFC7671] Dukhovni, V. and W. Hardaker, "The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance", RFC 7671, DOI 10.17487/RFC7671, October 2015, <https://www.rfc-editor.org/info/rfc7671>.

[RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", RFC 7871, DOI 10.17487/RFC7871, May 2016, <https://www.rfc-editor.org/info/rfc7871>.

[RFC8120] Oiwa, Y., Watanabe, H., Takagi, H., Maeda, K., Hayashi, T., and Y. Ioku, "Mutual Authentication Protocol for HTTP", RFC 8120, DOI 10.17487/RFC8120, April 2017, <https://www.rfc-editor.org/info/rfc8120>.

[RFC8146] Harkins, D., "Adding Support for Salted Password Databases to EAP-pwd", RFC 8146, DOI 10.17487/RFC8146, April 2017, <https://www.rfc-editor.org/info/rfc8146>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <https://www.rfc-editor.org/info/rfc8520>.

Authors' Addresses

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: kondtir@gmail.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

Michael C. Richardson
Sandelman Software Works
USA

Email: mcr+ietf@sandelman.ca

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

ADD WG
Internet-Draft
Intended status: Standards Track
Expires: March 4, 2021

T. Reddy
McAfee
D. Wing
Citrix
M. Richardson
Sandelman Software Works
M. Boucadair
Orange

August 31, 2020

DNS Server Selection: DNS Server Information with Assertion Token draft-reddy-add-server-policy-selection-05

Abstract

The document defines a mechanism that allows communication of DNS resolver information to DNS clients for use in server selection decisions. In particular, the document defines a mechanism for a DNS server to communicate its filtering policy and privacy statement URL to DNS clients. This information is cryptographically signed to attest its authenticity. Such information is used for the selection of DNS resolvers. Typically, evaluating the DNS privacy statement, filtering policy, and the signatory, DNS clients with minimum human intervention can select the DNS server that best supports the user's desired privacy and filtering policy.

This assertion is useful for encrypted DNS (e.g., DNS-over-TLS, DNS-over-HTTPS, DNS-over-QUIC) servers that are either public resolvers or discovered in a local network.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 4, 2021.

Table of Contents

1. Introduction

[RFC7626] discusses DNS privacy considerations in both "on the wire" (Section 2.4 of [RFC7626]) and "in the server" (Section 2.5 of [RFC7626]) contexts. Examples of protocols that provide encrypted channels between DNS clients and servers are DNS-over-HTTPS (DoH) [RFC8484], DNS-over-TLS (DoT) [RFC7858], and DNS-over-QUIC (DoQ) [I-D.ietf-dprive-dnsoquic].

DNS clients can discover and authenticate encrypted DNS servers provided by a local network, for example using the techniques proposed in [I-D.btw-add-home]. If the mechanism used to discover the encrypted DNS server is insecure, the DNS client needs evidence about the encrypted server to assess its trustworthiness and a way to appraise such evidence. The mechanism specified in this document can be used by the DNS client to cryptographically identify if it is connecting to an encrypted DNS server hosted by a specific organization (e.g., ISP or Enterprise).

The DNS Recursive Operator Privacy (DROP) statement explained in [I-D.ietf-dprive-bcp-op] outlines the recommended contents a DNS operator should publish, thereby providing a means for users to evaluate the privacy properties of a given DNS service. While a human can review the privacy statement of a DNS server operator, the challenge is the user has to search to find the URL that points to the human-readable privacy policy information of the DNS server. Also, a user does not know if a DNS server (public or local) performs DNS-based content filtering.

This document simplifies the user experience by supporting a mechanism to retrieve the DNS server policy permitting the user to review human-readable privacy policy information of the DNS server and to assess whether that DNS server performs DNS-based content filtering.

This document also defines a mechanism for DNS clients to gather a set of information related to discovered (or pre-configured) servers and use that information to feed a DNS server selection procedure. The following parameters are supported in this version:

Malware blocking: Indicates whether the DNS server offers malware blocking service.

Phishing blocking: Indicates whether the DNS server offers phishing blocking service.

Policy blocking: Indicates whether the DNS server maintains a block-list due to a policy by the operator of the DNS server.

Censored blocking: Indicates whether the DNS server maintains a block-list based on a requirement from an external entity.

QNAME minimization: Indicates whether the DNS server implements QNAME minimisation [RFC7816].

The cryptographically signed policy allows a DNS client to, e.g., connect to multiple DNS servers and prompt the user to review the DNS privacy statements to select the DNS server that adheres to the privacy preserving data policy and DNS filtering expectations of the user. How a user instructs a DNS client about his/her preferences and how/whether the DNS client prompts a user are out of scope.

2. Sample Use Cases

The mechanism for a DNS server to communicate its cryptographically signed policies to DNS clients contributes to solve the following problems encountered in various deployments:

- The encrypted DNS server discovered using DHCP/RA in Home and Mobile networks is insecure. The mechanism specified in this document can be used by the DNS client to validate the signatory (e.g., cryptographically attested by the ISP).

- Typically, Enterprise networks do not assume that all devices in their network are managed by the IT team or Mobile Device Management (MDM), especially in the quite common BYOD (Bring Your Own Device) scenario. The mechanism specified in this document can be used by users of BYOD devices to determine if the DNS server on the local network complies with their user's privacy policy and DNS filtering expectations.

- The user selects specific well-known networks (e.g., organization for which a user works or a user works temporarily within another

corporation) to learn the privacy policy statement and filtering policy of the local DNS server. If the discovered encrypted DNS server does not meet the privacy preserving data policy and filtering requirements of the user, the DNS client can take appropriate actions. For example, the action can be: use the discovered DNS server only to access internal-only DNS names and use another DNS server for public domains. Such a policy would adhere to the user's expectations.

- The policy information signals the presence of DNS-based content filtering in the attached network. If the network is well-known to the DNS client and the local DNS server meets the privacy requirements of the user, the DNS client can continue to use an encrypted connection with the local encrypted DNS server. If the error code returned by the DNS server indicates access to the domain is blocked because of internal security policy [I-D.ietf-dnsop-extended-error], the DNS client can securely identify that access to the domain is censored by the network.

- The signed policy contains an URL that points to a human-readable privacy policy information of the DNS server for the user to review. The user can then make an informed decision whether the DNS server is trustworthy to honor the privacy requirements of the user. The DNS Push Notifications mechanism defined in [RFC8765] can be used by the DNS client to be asynchronously notified when a policy change occurs. The client automatically learns updates to the policy of the DNS server. If the privacy statement of the DNS server changes, the client can notify the user to re-evaluate the updated privacy statement. As a reminder, DNS Push Notification is only defined for TLS over TCP. DNS client implementations that do not support DNS Push Notifications can use the mechanism discussed in Section 6.1.2 to identify policy updates.

3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC8499] and [I-D.ietf-dnsop-terminology-ter].

'Encrypted DNS' refers to a DNS protocol that provides an encrypted channel between a DNS client and server (e.g., DoT, DoH, or DoQ).

The terms Evidence, Verifier, Background Check, Relying Party, Attestation Results, and Appraisal Policy are defined in [I-D.ietf-rats-architecture].

4. Policy Assertion Token (PAT): Overview

The mechanism used in this specification resembles the Background-Check Model discussed in Sections 5.2 and 5.3 of Remote attestation procedure (RATS) Architecture [I-D.ietf-rats-architecture]. RATS enables a relying party to establish a level of confidence in the trustworthiness of a remote peer through the creation of Evidence to assess the peer's trustworthiness, and an Apprasal Policy for such Evidence. In this document, the Relying Party is the DNS client and the Attester is the encrypted DNS server. The Encrypted DNS servers may use "Domain Validation" (DV) certificates.

In a simpler situation, the Verifier is also the operator of the DNS server. It creates Attestation Results based upon its own claims, signing them using an OV (or EV) certificate provided by a public CA.

A more trustworthy situation, the Evidence is reviewed by an external Verifier (e.g., an Auditor who performed security and privacy audit of the Encrypted DNS server), and this Verifier produces higher confidence Attestation Results.

The background check of the organization hosting the Encrypted DNS server is done by a public CA. An OV/EV certificate is issued only after verification of the requesting organization's legal identity.

JSON Web Token (JWT) [RFC7519] and JSON Web Signature (JWS) [RFC7515] and related specifications define a standard token format that can be used as a way of encapsulating claimed or asserted information with an associated digital signature using X.509 based certificates. JWT provides a set of claims in JSON format that can accommodate asserted policy information of the Encrypted DNS server. Additionally, JWS provides a path for updating methods and cryptographic algorithms used for the associated digital signatures.

JWS defines the use of JSON data structures in a specified canonical format for signing data corresponding to JOSE header, JWS Payload, and JWS Signature. The next sections define the header and claims that MUST be minimally used with JWT and JWS for privacy assertion token.

The Policy Assertion Token (PAT) specifically uses this token format and defines claims that convey the policy information of Encrypted DNS server.

If the DoT session is established, the client can retrieve the PAT object using the RESINFO RRtype defined in [I-D.pp-add-resinfo] and QNAME of the domain name that is used to authenticate the privacy-enabling DNS server (referred to as ADN in [RFC8310]). If a DoH session is established, the DoH client can retrieve the PAT object using the well-known URI defined in [I-D.btw-add-rfc8484-clarification].

If the special-use domain name "resolver-info.arpa/IN" defined in [I-D.pp-add-resinfo] is used to discover the Encrypted DNS server, the client can retrieve the PAT object using the RESINFO RRtype and QNAME of the special-use domain name.

The signature of PAT object can be validated by the DNS client. If the signer and the contents of the PAT object comply with the user's requirements, the user's client can use that DNS server.

The PAT object is signed by the DNS server's domain that is authoritative to assert the DNS server policy information. This authority is represented by the certificate credentials and the signature.

For example, the PAT object could be created by the organization hosting the Encrypted DNS server and optionally by a third party who performed privacy and security audit of the Encrypted DNS server. The DNS client needs to have the capability to verify the digital signature and to parse the PAT object.

5. PAT Header

The JWS token header is a JOSE header (Section 4 of [RFC7515]) that defines the type and encryption algorithm used in the token.

The PAT header MUST include, at a minimum, the header parameters defined in Sections 5.1, 5.2, and 5.3.

5.1. 'typ' (Type) Header Parameter

The 'typ' (Type) Header Parameter is defined in Section 4.1.9 of [RFC7515] to declare the media type of the complete JWS.

For PAT Token the 'typ' header MUST be the string 'pat'. This represents that the encoded token is a JWT of type pat.

5.2. 'alg' (Algorithm) Header Parameter

The 'alg' (Algorithm) Header Parameter is defined in Section 4.1.1 of [RFC7515]. It specifies the JWS signature cryptographic algorithm. It also refers to a list of defined 'alg' values as part of a registry established by JSON Web Algorithms (JWA) [RFC7518] Section 3.1.

For the creation and verification of PAT tokens and their digital signatures, implementations MUST support ES256 as defined in Section 3.4 of [RFC7518]. Implementations MAY support other algorithms registered in the JSON Web Signature and Encryption Algorithms registry created by [RFC7518]. The content of that registry may be updated in the future depending on cryptographic strength requirements guided by current security best practice. The mandatory-to-support algorithm for PAT tokens may likewise be updated in the future.

Implementations of PAT digital signatures using ES256 as defined above SHOULD use deterministic ECDSA when supported for the reasons stated in [RFC6979].

5.3. 'x5u' (X.509 URL) Header Parameter

As defined in Section 4.1.5 of [RFC7515], the 'x5u' header parameter defines a URI [RFC3986] referring to the resource for the X.509 public key certificate or certificate chain [RFC5280] corresponding to the key used to digitally sign the JWS. Generally, as defined in Section 4.1.5 of [RFC7515] this corresponds to an HTTPS or DNSSEC resource using integrity protection.

5.4. An Example of PAT Header

An example of the PAT header is shown in Figure 1. It includes the specified PAT type, ES256 algorithm, and an URI referencing the network location of the certificate needed to validate the PAT signature.

```
{
  "typ":"pat",
  "alg":"ES256",
  "x5u":"https://cert.example.com/pat.cer"
}
```

Figure 1: A PAT Header Example

6. PAT Payload

The token claims consist of the policy information of the DNS server that needs to be verified at the DNS client. These claims follow the definition of a JWT claim (Secion 4 of [RFC7519]) and are encoded as defined by the JWS Payload (Section 3 of [RFC7515]).

PAT defines the use of a standard JWT-defined claim as well as custom claims corresponding to the DoT or DoH servers.

Claim names MUST use the US-ASCII character set. Claim values MAY contain characters that are outside the ASCII range, however they MUST follow the default JSON serialization defined in Section 7 of [RFC7519].

6.1. JWT Defined Claims

6.1.1. 'iat' - Issued At Claim

The JSON claim MUST include the 'iat' (Section 4.1.6 of [RFC7519]) defined claim "Issued At". The 'iat' should be set to the date and time of issuance of the JWT. The time value should be of the format (NumericDate) defined in Section 2 of [RFC7519].

6.1.2. 'exp' - Expiration Time Claim

The JSON claim MUST include the 'exp' (Section 4.1.4 of [RFC7519]) defined "claim Expiration Time". The 'exp' should be set to specify the expiration time on or after which the JWT is not accepted for processing. The PAT object should expire after a reasonable duration. A short expiration time for the PAT object periodically reaffirms the policy information of the DNS server to the DNS client and ensures the DNS client does not use outdated policy information. If the DNS client knows the PAT object has expired, it should make another request to get the new PAT object from the DNS server. For example, the client can compute a hash of the resolver information, retreive the information after the expiration time, computes the hash of the newly retrieved resolver information, and compares with the old hash to detect policy updates. A quality implementation can perform automatic analysis and avoid presenting this information to the user if the DNS server's policies have not changed.

6.2. PAT Specific Claims

6.2.1. DNS Server Identity Claims

The DNS server identity is represented by a claim that is required for PAT: the 'server' claim. The 'server' MUST contain claim values that are identity claim JSON objects where the child claim name represents an identity type and the claim value is the identity string, both defined in subsequent subsections.

These identities can be represented as either authentication domain name (ADN) (defined in [RFC8310]) or Uniform Resource Indicators (URI).

The DNS client constructs a reference identifier for the DNS server based on the ADN or the domain portion in the URI of the DNS server identity. The domain name in the DNS-ID identifier type within subjectAltName entry in the DNS server certificate conveyed in the TLS handshake is matched with the reference identifier. If the match is not successful, the client MUST not accept the PAT for further processing.

6.2.1.1. 'adn' - Authentication Domain Name Identity

If the DNS server identity is an ADN, the claim name representing the identity MUST be 'adn'. The claim value for the 'adn' claim is the ADN.

6.2.1.2. 'uri' - URI Identity

If the DNS server identity is of the form URI Tempate, as defined in [RFC6570], the claim name representing the identity MUST be 'uri' and the claim value is the URI Template form of the DNS server identity.

As a reminder, if DoH is supported by the DNS server, the DNS client uses the URI Template (Section 3 of [RFC8484]).

6.2.2. 'policyinfo' (Policy Information) Claim

The 'policyinfo' claim MUST be formatted as a JSON object. The 'policyinfo' claim contains the policy information of the DNS server, it includes the following attributes:

filtering: If the DNS server changes some of the answers that it returns or failure codes are returned based on policy criteria, such as to prevent access to malware sites or objectionable content (e.g., legal obligation). This optional attribute has the following structure:

malwareblocking: The DNS server offers malware blocking service. If access to domains is blocked on threat data, the parameter value is set to 'true'. Note that some of the commonly known types of malware are viruses, worms, trojans, bots, ransomware, backdoors, spyware, and adware.

phishingblocking: The DNS server offers phishing blocking service. If access to phishing domains is blocked, the parameter value is set to 'true'.

policyblocking: If access to domains is blocked due to an internal policy imposed by the operator of the DNS server, the parameter value is set to 'true'. Note that the extended error code "Blocking" defined in Section 4.16 of [I-D.ietf-dnsop-extended-error] identifies access to domains is blocked due to an policy by the operator of the DNS server.

censoredblocking: If access to domains is blocked due to an external requirement imposed by an external entity, the parameter value is set to 'true'. Note that the extended error code "Censored" defined in Section 4.17 of [I-D.ietf-dnsop-extended-error] identifies access to domains is blocked based on a requirement from an external entity. Similar to the definition of "Censored" blocking in [I-D.ietf-dnsop-extended-error], this version of the specification does not distinguish blocking from regulatory bodies (e.g., Law Enforcement Agency) vs. arbitrary blocking. Such differentiation may be defined if required.

qnameminimization: If the DNS server supports QNAME minimisation [RFC7816] to improve DNS privacy, the parameter value is set to true. This is a mandatory attribute.

clientauth: If the DNS server policy requires client authentication, the parameter value is set to true. For example, when not on the enterprise network (e.g., at home or coffeeshop) yet needing to access the enterprise Encrypted DNS server, roaming users can use client authentication to access the Enterprise provided Encrypted DNS server. This is an optional attribute.

privacyurl: A URL that points to the privacy policy information of the DNS server. This is a mandatory attribute.

auditurl: A URL that points to the security (including privacy) assessment report of the DNS server by a third party auditor. This is an optional attribute.

6.2.3. Example

Figure 2 shows an example of policy information.

```
{
  "server":{
      "adn":["example.com"]
  },
  "iat":1443208345,
  "exp":1443640345,
  "policyinfo": {
     "filtering": {
         "malwareblocking": true,
         "policyblocking": false
     },
     "qnameminimization":false,
     "privacyurl": "https://example.com/commitment-to-privacy/"
  }
}
```

Figure 2: An Example of Policy Information

7. PAT Signature

The signature of the PAT is created as specified in Section 5.1 of [RFC7515] (Steps 1 through 6). PAT MUST use the JWS Protected Header.

For the JWS Payload and the JWS Protected Header, the lexicographic ordering and white space rules described in Section 5 and Section 6, and JSON serialization rules in Section 9 MUST be followed.

The PAT is cryptographically signed by the domain hosting the DNS server and optionally by a third party who performed privacy and security audit of the DNS server.

The policy information is attested using "Organization Validation" (OV) or "Extended Validation" (EV) certificates to avoid bad actors taking advantage of this mechanism to advertise encrypted DNS servers for illegitimate and fraudulent purposes meant to trick DNS clients into believing that they are using a legitimate encrypted DNS server hosted to provide privacy for DNS transactions.

Alternatively, a DNS client has to be configured to trust the leaf of the signer of the PAT object. That is, trust of the signer MUST NOT be determined by validating the signer via the OS or the browser trust chain because that would allow any arbitrary entity to operate a DNS server and assert any sort of policy.

Appendix A provides an example of how to follow the steps to create the JWS Signature.

JWS JSON serialization (Step 7 in Section 5.1 of [RFC7515]) is supported for PAT to enable multiple signatures to be applied to the PAT object. For example, the PAT object can be cryptographically signed by the domain hosting the DNS server and by a third party who performed privacy and security audit of the DNS server.

Appendix B includes an example of the full JWS JSON serialization representation with multiple signatures.

Section 5.1 of [RFC7515] (Step 8) describes the method to create the final JWS Compact Serialization form of the PAT Token.

8. Extending PAT

PAT includes the minimum set of claims needed to securely assert the policy information of the DNS server. JWT supports a mechanism to add additional asserted or signed information by simply adding new claims. PAT can be extended beyond the defined base set of claims to represent other DNS server information requiring assertion or validation. Specifying new claims follows the baseline JWT procedures (Section 10.1 of [RFC7519]). Understanding new claims on the DNS client is optional. The creator of a PAT object cannot assume that the DNS client will understand the new claims.

9. Deterministic JSON Serialization

JSON objects can include spaces and line breaks, and key value pairs can occur in any order. It is therefore a non-deterministic string format. In order to make the digital signature verification work deterministically, the JSON representation of the JWS Protected Header object and JWS Payload object MUST be computed as follows.

The JSON object MUST follow the following rules. These rules are based on the thumbprint of a JSON Web Key (JWK) as defined in Section 3 of [RFC7638] (Step 1).

1. The JSON object MUST contain no whitespace or line breaks before or after any syntactic elements.

2. JSON objects MUST have the keys ordered lexicographically by the Unicode [UNICODE] code points of the member names.

3. JSON value literals MUST be lowercase.

4. JSON numbers are to be encoded as integers unless the field is defined to be encoded otherwise.

5. Encoding rules MUST be applied recursively to member values and array values.

9.1. Example PAT Deterministic JSON Form

This section demonstrates the deterministic JSON serialization for the example PAT Payload shown in Section 6.2.3.

The initial JSON object is shown in Figure 3.

```
{
  "server":{
      "adn":["example.com"]
  },
  "iat":1443208345,
  "exp":1443640345,
  "policyinfo": {
     "qnameminimization":false,
     "privacyurl": "https://example.com/commitment-to-privacy/"
  }
}
```

Figure 3: Initial JSON Object

The parent members of the JSON object are as follows, in lexicographic order: "exp", "iat", "policyinfo", "server".

The final constructed deterministic JSON serialization representation, with whitespace and line breaks removed, (with line breaks used for display purposes only) is:

```
{"exp":1443640345,"iat":1443208345,
"policyinfo":{"privacyurl":"https://example.com/commitment-to-privacy/",
"qnameminimization":false},"server":{"adn":["example.com"]}}
```

Figure 4: Deterministic JSON Form

10. Using RESINFO responses

This document defines the following entries for the IANA DNS Resolver Information Registry that is defined in [I-D.pp-add-resinfo].

1. The "server" name containing the DNS server identity discussed in Section 4.

2. The sub-attribute "adn" discussed in Section 4 contained in the "server" attribute is used to specify the DNS server identity in the form of ADN.

3. The sub-attribute "uri" discussed in Section 4 contained in the "server" attribute is used to specify the DNS server identity in the form of URI template.

4. The "filtering", "qnameminimization", "privacyurl" and "auditurl" names containing the resolver information of the DNS server discussed in Section 4.

5. The sub-attributes "malwareblocking", "phishingblocking", "policyblocking" and "censoredblocking" discussed in Section 4 contained in the "filtering" attribute are used to specify the reasons for performing DNS-based content filtering.

6. The "attested-resinfo" name contains a base64 encoding of a PAT Section 4. If the "attested-resinfo" name is conveyed to the client, the server need not convey the above attributes (1 to 5) seperately as that resolver information will be extracted by the client from the PAT payload.

11. Privacy Considerations

Users are expected to indicate to their system in some way that they trust certain PAT signers (e.g., if working for Example, Inc., the user's system is configured to trust "example.com" signing the PAT). By doing so, the DNS client can automatically discover encrypted DNS server in specific networks, validate the PAT signature and the user can check if the human readable privacy policy information of the DNS server complies with user's privacy needs, prior to using that encrypted DNS server for DNS queries.

The DNS client MUST retrieve the human-readable privacy statement from the 'privacyurl' attribute to assist with that decision (e.g., display the privacy statement when it changes, show differences in previously-retrieved version, etc.). With the steps above, user can review the human-readable privacy policy information of the Encrypted DNS server.

Another scenario is bootstrapping a networking device to use the encrypted DNS server in the local network. Secure Zero Touch Provisioning [RFC8572] defines a bootstrapping strategy for enabling devices to securely obtain the required configuration information with no user input. If the encrypted DNS server is insecurely discovered and not pre-configured in the networking device, the

client can validate the Policy Assertion Token signature using the owner certificate as per Section 3.2 of [RFC8572].

12. Security Considerations

The use of PAT object based on the validation of the digital signature and the associated certificate requires consideration of the authentication and authority or reputation of the signer to attest the policy information of the DNS server being asserted. Bad actors can host encrypted DNS servers, and claim the servers offer privacy but exactly do the opposite to invade the privacy of the user. Bad actor can get a domain name, host encrypted DNS servers, and get the DNS server certificate signed by a CA. The policy information will have to be attested using OV/EV certificates or a PAT object signer trusted by the DNS client to prevent the attack.

The CA that issued the OV/EV certificate does not attest the resolver information. The organization hosting the DNS server attests the resolver information using the OV/EV certificate and the client uses the OV/EV certificate to identify the organization (e.g., ISP or Enterprise) hosting the DNS server.

If the PAT object is asserted by a third party, it can do a "time of check" but the DNS server is susceptible of "time of use" attack. For example, changes to the policy of the DNS server can cause a disagreement between the auditor and the DNS server operation, hence the PAT object needs to be also asserted by the domain hosting the DNS server. In addition, the PAT object needs to have a short expiration time (e.g., 7 days) to ensure the DNS server's domain re-asserts the policy information and limits the damage from change in policy and mis-issuance.

13. IANA Considerations

13.1. Media Type Registration

13.1.1. Media Type Registry Contents Additions Requested

This section registers the 'application/pat' media type [RFC2046] in the 'Media Types' registry in the manner described in [RFC6838], which can be used to indicate that the content is a PAT defined JWT.

- Type name: application
- Subtype name: pat
- Required parameters: n/a

- Optional parameters: n/a
- Encoding considerations: 8bit; application/pat values are encoded as a series of base64url-encoded values (some of which may be the empty string) separated by period ('.') characters..
- Security considerations: See the Security Considerations Section of [RFC7515].
- Interoperability considerations: n/a
- Published specification: [TODO this document]
- Applications that use this media type: DNS
- Fragment identifier considerations: n/a
- Additional information:

  Magic number(s): n/a File extension(s): n/a Macintosh file type code(s): n/a
- Person & email address to contact for further information: Tirumaleswar Reddy, kondtir@gmail.com
- Intended usage: COMMON
- Restrictions on usage: none
- Author: Tirumaleswar Reddy, kondtir@gmail.com
- Change Controller: IESG
- Provisional registration? No

13.2. JSON Web Token Claims Registration

13.2.1. Registry Contents Additions Requested

- Claim Name: 'server'
- Claim Description: DNS server identity
- Change Controller: IESG
- Specification Document(s): Section 6.2.1 of [TODO this document]
- Claim Name: 'policyinfo'

- Claim Description: Policy information of DNS server.
- Change Controller: IESG
- Specification Document(s): Section 6.2.2 of [TODO this document]

13.3. DNS Resolver Information Registration

IANA will add the names attested-resinfo, server, filtering, qnameminimization, privacyurl and auditurl to the DNS Resolver Information registry defined in Section 4.2 of [I-D.pp-add-resinfo]. IANA will add the the sub-attributes "malwareblocking", "phishingblocking", "policyblocking" and "censoredblocking" contained in the "filtering" attribute to the DNS Resolver Information registry. IANA will add the sub-attributes "adn" and "uri" contained in the "server" attribute to the DNS Resolver Information registry.

14. Acknowledgments

This specification leverages some of the work that has been done in [RFC8225]. Thanks to Ted Lemon, Paul Wouters, Neil Cook, Vittorio Bertola, Vinny Parla, Chris Box and Shashank Jain for the discussion and comments.

15. References

15.1. Normative References

[RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, DOI 10.17487/RFC2046, November 1996, <https://www.rfc-editor.org/info/rfc2046>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <https://www.rfc-editor.org/info/rfc3986>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <https://www.rfc-editor.org/info/rfc5280>.

[RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template", RFC 6570, DOI 10.17487/RFC6570, March 2012, <https://www.rfc-editor.org/info/rfc6570>.

[RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <https://www.rfc-editor.org/info/rfc6838>.

[RFC6979] Pornin, T., "Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)", RFC 6979, DOI 10.17487/RFC6979, August 2013, <https://www.rfc-editor.org/info/rfc6979>.

[RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <https://www.rfc-editor.org/info/rfc7515>.

[RFC7518] Jones, M., "JSON Web Algorithms (JWA)", RFC 7518, DOI 10.17487/RFC7518, May 2015, <https://www.rfc-editor.org/info/rfc7518>.

[RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <https://www.rfc-editor.org/info/rfc7519>.

[RFC7638] Jones, M. and N. Sakimura, "JSON Web Key (JWK) Thumbprint", RFC 7638, DOI 10.17487/RFC7638, September 2015, <https://www.rfc-editor.org/info/rfc7638>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <https://www.rfc-editor.org/info/rfc8499>.

15.2. Informative References

[I-D.btw-add-home]
Boucadair, M., Reddy.K, T., Wing, D., and N. Cook, "Encrypted DNS Discovery and Deployment Considerations for Home Networks", draft-btw-add-home-08 (work in progress), August 2020.

[I-D.btw-add-rfc8484-clarification]
Boucadair, M., Cook, N., Reddy.K, T., and D. Wing, "Supporting Redirection for DNS Queries over HTTPS (DoH)", draft-btw-add-rfc8484-clarification-02 (work in progress), July 2020.

[I-D.ietf-dnsop-extended-error]
Kumari, W., Hunt, E., Arends, R., Hardaker, W., and D. Lawrence, "Extended DNS Errors", draft-ietf-dnsop-extended-error-16 (work in progress), May 2020.

[I-D.ietf-dnsop-terminology-ter]
Hoffman, P., "Terminology for DNS Transports and Location", draft-ietf-dnsop-terminology-ter-02 (work in progress), August 2020.

[I-D.ietf-dprive-bcp-op]
Dickinson, S., Overeinder, B., Rijswijk-Deij, R., and A. Mankin, "Recommendations for DNS Privacy Service Operators", draft-ietf-dprive-bcp-op-14 (work in progress), July 2020.

[I-D.ietf-dprive-dnsoquic]
Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnsoquic-00 (work in progress), April 2020.

[I-D.ietf-rats-architecture]
Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", draft-ietf-rats-architecture-05 (work in progress), July 2020.

[I-D.pp-add-resinfo]
Sood, P. and P. Hoffman, "DNS Resolver Information Self-publication", draft-pp-add-resinfo-02 (work in progress), June 2020.

[RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <https://www.rfc-editor.org/info/rfc7626>.

[RFC7816] Bortzmeyer, S., "DNS Query Name Minimisation to Improve Privacy", RFC 7816, DOI 10.17487/RFC7816, March 2016, <https://www.rfc-editor.org/info/rfc7816>.

[RFC8225] Wendt, C. and J. Peterson, "PASSporT: Personal Assertion Token", RFC 8225, DOI 10.17487/RFC8225, February 2018, <https://www.rfc-editor.org/info/rfc8225>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <https://www.rfc-editor.org/info/rfc8310>.

[RFC8572] Watsen, K., Farrer, I., and M. Abrahamsson, "Secure Zero Touch Provisioning (SZTP)", RFC 8572, DOI 10.17487/RFC8572, April 2019, <https://www.rfc-editor.org/info/rfc8572>.

[RFC8765] Pusateri, T. and S. Cheshire, "DNS Push Notifications", RFC 8765, DOI 10.17487/RFC8765, June 2020, <https://www.rfc-editor.org/info/rfc8765>.

[UNICODE] The Unicode Consortium, "The Unicode Standard", June 2016, <http://www.unicode.org/versions/latest/>.

Appendix A. Example ES256 based PAT JWS Serialization and Signature

For PAT, there will always be a JWS with the following members:

- 'protected', with the value BASE64URL(UTF8(JWS Protected Header))
- 'payload', with the value BASE64URL (JWS Payload)
- 'signature', with the value BASE64URL(JWS Signature)

This example will follow the steps in JWS [RFC7515] Section 5.1, steps 1-6 and 8 and incorporates the additional serialization steps required for PAT.

Step 1 for JWS references the JWS Payload, an example PAT Payload is as follows:

```
{
  "server":{
      "adn":["example.com"]
  },
  "iat":1443208345,
  "exp":1443640345,
  "policyinfo": {
     "filtering": {
         "malwareblocking": true,
         "policyblocking": false
     },
     "qnameminimization":false,
     "privacyurl": "https://example.com/commitment-to-privacy/"
  }
}
```

This would be serialized to the form (with line break used for display purposes only):

```
{"exp":1443640345,"iat":1443208345,"policyinfo":{
"filtering":{"malwareblocking": true,"policyblocking": false},
"privacyurl":"https://example.com/commitment-to-privacy/",
"qnameminimization":false},"server":{"adn":["example.com"]}}
```

Step 2 Computes the BASE64URL(JWS Payload) producing this value (with line break used for display purposes only):

```
eyJleHAiOjE0NDM2NDAzNDUsImlhdCI6MTQ0MzIwODM0NSwicG9saWN5aW5mbyI6e
yJmaWx0ZXJpbmciOnsibWFsd2FyZWJsb2NraW5nIjp0cnVlLCJwb2xpY3libG9ja2
luZyI6ZmFsc2V9LCJwcml2YWN5dXJsIjoiaHR0cHM6Ly9leGFtcGxlLmNvbS9jb21
taXRtZW50LXRvLXByaXZhY3kvIiwicW5hbWVtaW5pbWl6YXRpb24iOmZhbHNlfSwi
c2VydmVyIjp7ImFkbiI6WyJleGFtcGxlLmNvbSJdfX0
```

For Step 3, an example PAT Protected Header comprising the JOSE Header is as follows:

```
{
  "alg":"ES256",
  "typ":"pat",
  "x5u":"https://cert.example.com/pat.cer"
}
```

This would be serialized to the form (with line break used for display purposes only):

```
{"alg":"ES256","typ":"pat","x5u":"https://cert.example.com
/pat.cer"}
```

Step 4 Performs the BASE64URL(UTF8(JWS Protected Header)) operation and encoding produces this value (with line break used for display purposes only):

```
eyJhbGciOiJFUzI1NiIsInR5cCI6InBhdCIsIng1dSI6Imh0dHBzOi8vY2VydC5l
eGFtcGxlLmNvbS9wYXQuY2VyIn0
```

Step 5 and Step 6 performs the computation of the digital signature of the PAT Signing Input ASCII(BASE64URL(UTF8(JWS Protected Header)) || '.' || BASE64URL(JWS Payload)) using ES256 as the algorithm and the BASE64URL(JWS Signature).

```
4vQEAF_Vlp1Tr6sJmS4pnIKDRmIjH8EZzY5BMT2qJCHD8PmjBktWVnlmbmyHs05G
KauRBdIFnfp3oDPbE0Jq4w
```

Step 8 describes how to create the final PAT token, concatenating the values in the order Header.Payload.Signature with period ('.') characters. For the above example values this would produce the following (with line breaks between period used for readability purposes only):

```
eyJhbGciOiJFUzI1NiIsInR5cCI6InBhdCIsIng1dSI6Imh0dHBzOi8vY2VydC5l
eGFtcGxlLmNvbS9wYXQuY2VyIn0
.
eyJleHAiOjE0NDM2NDAzNDUsImlhdCI6MTQ0MzIwODM0NSwicG9saWN5aW5mbyI6e
yJmaWx0ZXJpbmciOnsibWFsd2FyZWJsb2NraW5nIjp0cnVlLCJwb2xpY3libG9ja2
luZyI6ZmFsc2V9LCJwcml2YWN5dXJsIjoiaHR0cHM6Ly9leGFtcGxlLmNvbS9jb21
taXRtZW50LXRvLXByaXZhY3kvIiwicW5hbWVtaW5pbWl6YXRpb24iOmZhbHNlfSwi
c2VydmVyIjp7ImFkbiI6WyJleGFtcGxlLmNvbSJdfX0
.
4vQEAF_Vlp1Tr6sJmS4pnIKDRmIjH8EZzY5BMT2qJCHD8PmjBktWVnlmbmyHs05G
KauRBdIFnfp3oDPbE0Jq4w
```

A.1. X.509 Private Key in PKCS#8 Format for ES256 Example**

```
-----BEGIN PRIVATE KEY-----
MIGHAgEAMBMGByqGSM49AgEGCCqGSM49AwEHBG0wawIBAQQgevZzL1gdAFr88hb2
OF/2NxApJCzGCEDdfSp6VQO30hyhRANCAAQRWz+jn65BtOMvdyHKcvjBeBSDZH2r
1RTwjmYSi9R/zpBnuQ4EiMnCqfMPWiZqB4QdbAd0E7oH50VpuZ1P087G
-----END PRIVATE KEY-----
```

A.2. X.509 Public Key for ES256 Example**

```
-----BEGIN PUBLIC KEY-----
MFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEEVs/o5+uQbTjL3chynL4wXgUg2R9
q9UU8I5mEovUf86QZ7kOBIjJwqnzD1omageEHWwHdBO6B+dFabmdT9POxg==
-----END PUBLIC KEY-----
```

Appendix B. Complete JWS JSON Serialization Representation with multiple Signatures

The JWS payload used in this example as follows.

```
{
  "server":{
      "adn":["example.com"]
  },
  "iat":1443208345,
  "exp":1443640345,
  "policyinfo": {
     "filtering": {
         "malwareblocking": true,
         "policyblocking": false
     },
     "qnameminimization":false,
     "privacyurl": "https://example.com/commitment-to-privacy/"
  }
}
```

This would be serialized to the form (with line break used for display purposes only):

```
{"exp":1443640345,"iat":1443208345,"policyinfo":{
"filtering":{"malwareblocking": true,"policyblocking": false},
"privacyurl":"https://example.com/commitment-to-privacy/",
"qnameminimization":false},"server":{"adn":["example.com"]}}
```

The JWS protected Header value used for the first signature is same as that used in the example in Appendix A. The X.509 private key used for generating the first signature is same as that used in the example in Appendix A.1.

The JWS Protected Header value used for the second signature is:

```
{
  "alg":"ES384",
  "typ":"pat",
  "x5u":"https://cert.audit-example.com/pat.cer"
}
```

The complete JWS JSON Serialization for these values is as follows (with line breaks within values for display purposes only):


```
{
  "payload":
       "eyJleHAiOjE0NDM2NDAzNDUsImlhdCI6MTQ0MzIwODM0NSwicG9saWN5aW5mbyI6
        eyJmaWx0ZXJpbmciOnsibWFsd2FyZWJsb2NraW5nIjp0cnVlLCJwb2xpY3libG9j
        a2luZyI6ZmFsc2V9LCJwcml2YWN5dXJsIjoiaHR0cHM6Ly9leGFtcGxlLmNvbS9j
        b21taXRtZW50LXRvLXByaXZhY3kvIiwicW5hbWVtaW5pbWl6YXRpb24iOmZhbHNl
        fSwic2VydmVyIjp7ImFkbiI6WyJleGFtcGxlLmNvbSJdfX0",
  "signatures":[
       {"protected":"eyJhbGciOiJFUzI1NiIsInR5cCI6InBhdCIsIng1dSI6Imh0dHB
        zOi8vY2VydC5leGFtcGxlLmNvbS9wYXQuY2VyIn0",
        "signature": "4vQEAF_Vlp1Tr6sJmS4pnIKDRmIjH8EZzY5BMT2qJCHD8PmjBk
        tWVnlmbmyHs05GKauRBdIFnfp3oDPbE0Jq4w"},
       {"protected":"eyJhbGciOiJFUzM4NCIsInR5cCI6InBhdCIsIng1dSI6Imh0dHB
        zOi8vY2VydC5hdWRpdC1leGFtcGxlLmNvbS9wYXQuY2VyIn0",
        "signature":666ag_mAqDa3Oyxo1DGXUocr0MmRjpXwq8kWp1S21mvs2-kPCIq3
        0xsBJt4apy-sq3VyJgIqzjijoFYURhHvupF0obo-IFUGSZ1YHBCX_MiyBwJQJjtp
        S91ujDatRTtZ"}]
}
```


B.1. X.509 Private Key in PKCS#8 format for E384 Example**


```
-----BEGIN PRIVATE KEY-----
MIGHAgEAMBMGByqGSM49AgEGCCqGSM49AwEHBG0wawIBAQQgevZzL1gdAFr88hb2
OF/2NxApJCzGCEDdfSp6VQO30hyhRANCAAQRWz+jn65BtOMvdyHKcvjBeBSDZH2r
1RTwjmYSi9R/zpBnuQ4EiMnCqfMPWiZqB4QdbAd0E7oH50VpuZ1P087G
-----END PRIVATE KEY-----
```


B.2. X.509 Public Key for ES384 Example**


```
-----BEGIN PUBLIC KEY-----
MFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEEVs/o5+uQbTjL3chynL4wXgUg2R9
q9UU8I5mEovUf86QZ7kOBIjJwqnzD1omageEHWwHdBO6B+dFabmdT9POxg==
-----END PUBLIC KEY-----
```


Authors' Addresses

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: kondtir@gmail.com

Dan Wing
Citrix Systems, Inc.
USA

Email: dwing-ietf@fuggles.com

Michael C. Richardson
Sandelman Software Works
USA

Email: mcr+ietf@sandelman.ca

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 14 January 2021

D. Schinazi
Google LLC
N. Sullivan
J. Kipp
Cloudflare

13 July 2020

DoH Preference Hints for HTTP
draft-schinazi-httpbis-doh-preference-hints-02

Abstract

When using a publicly available DNS-over-HTTPS (DoH) server, some clients may suffer poor performance when the authoritative DNS server is located far from the DoH server. For example, a publicly available DoH server provided by a Content Delivery Network (CDN) should be able to resolve names hosted by that CDN with good performance but might take longer to resolve names provided by other CDNs, or might provide suboptimal results if that CDN is using DNS-based load balancing and returns different address records depending or where the DNS query originated from. This document attempts to lessen these issues by allowing the web server to indicate to the client which DoH server can best resolve its addresses. This document defines an HTTP header field that enables web host operators to inform user agents of the preferred DoH servers to use for subsequent DNS lookups for the host's domain.

Discussion of this work is encouraged to happen on the ADD IETF mailing list add@ietf.org or on the GitHub repository which contains the draft: https://github.com/DavidSchinazi/draft-httpbis-doh-preference-hints.

This Internet-Draft will expire on 14 January 2021.

Table of Contents

1. Introduction

When using a publicly available DNS-over-HTTPS (DoH) server, some clients may suffer poor performance when the authoritative DNS server is located far from the DoH server. For example, a publicly available DoH server provided by a Content Delivery Network (CDN) should be able to resolve names hosted by that CDN with good performance but might take longer to resolve names provided by other CDNs, or might provide suboptimal results if that CDN is using DNS-based load balancing and returns different address records depending or where the DNS query originated from. This document attempts to lessen these issues by allowing the web server to indicate to the client which DoH server can best resolve its addresses. This document defines an HTTP header field that enables web host operators

to inform user agents of the preferred DoH servers to use for subsequent DNS lookups for the host's domain.

When a web server wishes its client to use a specific DoH server to resolve its addresses, it can send the DoH-Preference header to indicate that preference to the user agent. The header is not prescriptive, it only indicates the server's preference to the user. It also only applies to the web server's current hostname.

The header defined in this document is not intended to be used as a discovery mechanism for clients learning about the existence of new DoH servers. Instead, it is intended to be used as an optimization technique for clients with support for multiple DoH servers who wish to choose the most performant DNS server for a given query.

Discussion of this work is encouraged to happen on the ADD IETF mailing list add@ietf.org or on the GitHub repository which contains the draft: https://github.com/DavidSchinazi/draft-httpbis-doh-preference-hints.

1.1. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document uses the Augmented BNF defined in [RFC5234] and updated by [RFC7405] along with the "#rule" extension defined in Section 7 of [RFC7230]. The rules below are defined in [RFC5234], [RFC7230], and [RFC7234]:

```
OWS           = <OWS, see {{RFC7230}}, Section 3.2.3>
delta-seconds = <delta-seconds; see {{RFC7234}}, Section 1.2.1>
quoted-string = <quoted-string, see {{RFC7230}}, Section 3.2.6>
token         = <token, see {{RFC7230}}, Section 3.2.6>
```

2. The DoH-Preference Header Field

An HTTPS origin can indicate its preference regarding DoH servers to the client by adding an DoH-Preference header field to responses.

```
DoH-Preference = doh-uri *( OWS ";" OWS parameter )
doh-uri        = quoted-string
parameter      = token "=" ( token / quoted-string )
```

The "doh-uri" component consists of the DoH URI Template as defined in [RFC8484].

Sending multiple DoH-Preference header fields indicates that the server prefers multiple DoH servers. They are sent in decreasing order of preference.

2.1. The max-age Directive

The REQUIRED "max-age" directive specifies the number of seconds, after the reception of the DoH-Preference header field, during which the user agent caches the server's DoH preferences.

The syntax of the max-age directive's REQUIRED value (after quoted-string unescaping, if necessary) is defined as:

```
max-age-value = delta-seconds
```

A max-age value of zero (i.e., "max-age=0") signals the user agent to remove the DoH URI template from its cache.

2.2. Examples

The header below indicates that the user agent should consider querying DNS results for the web server's hostname using "dnsserver.example.net" for approximately six months. (Lines are folded to fit.)

```
DoH-Preference: "https://dnsserver.example.net/dns-query{?dns}";
    max-age=15768000
```

3. Server Behavior

Web servers MAY send a DoH-Preference header to indicate to clients that it would prefer they use that DoH server when resolving addresses for the hostname of the web server. Web servers MAY send multiple DoH-Preference headers. Web servers MUST NOT send the DoH-Preference header in HTTP responses conveyed over a non-secure transport.

3.1. Considerations For Choosing a Preferred DoH Server

The choice of DoH server can affect overall performance and responsiveness as perceived by the client. Some example considerations in choosing a preferred DoH server are:

* A DoH host specified as a host name rather than an IP address will require one or more additional DNS resolutions when the cached DNS entries for the resolver or resolvers expire.

* Support for extension mechanisms (e.g. EDNS(0)) may be desired.

* Clients, particularly mobile device clients, may frequently move between networks that have different network paths to the DoH server.

4. Client Behavior

If a client chooses to act on received DoH-Preference headers, it SHOULD cache the server's hostname and the corresponding DoH URI template and lifetime. It SHOULD then send subsequent DNS requests for A and AAAA records for that host name to the provided DoH server, until the cache entry expires after the time specified in the "max-age" directive. Any received DoH-Preference header replaces and overrides any and all information received in a previous DoH-Preference header for the same host name and DoH URI template.

Clients MAY decide to only respect the DoH-Preference header for a subset of vetted DoH servers.

Clients MUST NOT use the contents of the DoH-Preference header to impact how it resolves other domain names. Clients MUST ignore the DoH-Preference header in HTTP responses conveyed over a non-secure transport.

If the DoH-Preference URI contains a host expressed as a host name rather than as an IP address and that host name is resolved via DoH, the DoH server might also specify a DoH-Preference header. This means that respecting the DoH server recommendation could result in an excessively long chain of DoH queries or a loop of DoH servers. Clients SHOULD be capable of detecting a loop or an excessively long chain of DoH servers and treat these conditions as a query failure.

4.1. Fallback

If resolution using the recommended DoH server fails, clients MUST fall back and retry their query using another DNS resolution mechanism.

5. Internationalization Considerations

An internationalized domain name that appears in the header field MUST be expressed using A-labels; see Section 2.3.2.1 of [RFC5890].

6. Security Considerations

The DoH-Preference header allows a web server to impact how a user agent resolves DNS A and AAAA records for its own host name. Since the web server has proven ownership of the domain name via its TLS certificate and the DNS result that led to the initial connection, impacting future DNS resolutions to the same host name has limited security impact.

The potential exists for the DoH-Preference header to be used as a form of web tracking. Because a DoH URI is chosen by the server, cached by the client, and then subsequently contacted by the client, a uniquely chosen DoH URI could identify a client even after other client-side state has expired or been removed. Clients SHOULD expire cached DoH URIs when other client state expires or is cleared by the user unless the URIs refer to vetted DoH servers or match common DoH URI patterns that preclude client-unique URIs.

7. IANA Considerations

This document, if approved, requests IANA to register the DoH-Preference header in the "Permanent Message Header Field Names" registry maintained at https://www.iana.org/assignments/message-headers/ (https://www.iana.org/assignments/message-headers/).

Header Field Name	Protocol	Status	Reference
DoH-Preference	http	standard	Section 2

The change controller is: "IETF (iesg@ietf.org) - Internet Engineering Task Force".

8. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <https://www.rfc-editor.org/info/rfc5234>.

[RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <https://www.rfc-editor.org/info/rfc5890>.

[RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <https://www.rfc-editor.org/info/rfc7230>.

[RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, DOI 10.17487/RFC7234, June 2014, <https://www.rfc-editor.org/info/rfc7234>.

[RFC7405] Kyzivat, P., "Case-Sensitive String Support in ABNF", RFC 7405, DOI 10.17487/RFC7405, December 2014, <https://www.rfc-editor.org/info/rfc7405>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

Acknowledgments

The authors would like to thank many members of the IETF community, as this document is the fruit of many hallway conversations.

Authors' Addresses

David Schinazi
Google LLC
1600 Amphitheatre Parkway
Mountain View, California 94043,
United States of America

Email: dschinazi.ietf@gmail.com

Nick Sullivan
Cloudflare

Email: nick@cloudflare.com

Jesse Kipp
Cloudflare

Email: jkipp@cloudflare.com