ADD
Internet-Draft
Intended status: Standards Track
Expires: 14 August 2021

T. Pauly
E. Kinnear
Apple Inc.
C.A. Wood
Cloudflare
P. McManus
Fastly
T. Jensen
Microsoft

10 February 2021

Discovery of Designated Resolvers
draft-ietf-add-ddr-00

Abstract

This document defines Discovery of Designated Resolvers (DDR), a mechanism for DNS clients to use DNS records to discover a resolver's encrypted DNS configuration. This mechanism can be used to move from unencrypted DNS to encrypted DNS when only the IP address of an encrypted resolver is known. It can also be used to discover support for encrypted DNS protocols when the name of an encrypted resolver is known. This mechanism is designed to be limited to cases where unencrypted resolvers and their designated resolvers are operated by the same entity.

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 August 2021.

Copyright Notice

Table of Contents

1. Introduction

When DNS clients wish to use encrypted DNS protocols such as DNS-over-TLS (DoT) [RFC7858] or DNS-over-HTTPS (DoH) [RFC8484], they require additional information beyond the IP address of the DNS server, such as the resolver's hostname, non-standard ports, or URL paths. However, common configuration mechanisms only provide the resolver's IP address during configuration. Such mechanisms include network provisioning protocols like DHCP [RFC2132] and IPv6 Router Advertisement (RA) options [RFC8106], as well as manual configuration.

This document defines two mechanisms for clients to discover designated resolvers using DNS server Service Binding (SVCB, [I-D.ietf-dnsop-svcb-https]) records:

1. When only an IP address of an Unencrypted Resolver is known, the client queries a special use domain name to discover DNS SVCB records associated with the Unencrypted Resolver (Section 4).

2. When the hostname of an encrypted DNS server is known, the client requests details by sending a query for a DNS SVCB record. This can be used to discover alternate encrypted DNS protocols supported by a known server, or to provide details if a resolver name is provisioned by a network (Section 5).

Both of these approaches allow clients to confirm that a discovered Encrypted Resolver is designated by the originally provisioned resolver. "Equivalence" in this context means that the resolvers are operated by the same entity; for example, the resolvers are accessible on the same IP address, or there is a certificate that claims ownership over both resolvers.

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:

DDR: Discovery of Designated Resolvers. Refers to the mechanisms defined in this document.

Designated Resolver: A resolver, presumably an Encrypted Resolver, designated by another resolver for use in its own place. This designation can be authenticated with TLS certificates.

Encrypted Resolver: A DNS resolver using any encrypted DNS transport. This includes current mechanisms such as DoH and DoT as well as future mechanisms.

Unencrypted Resolver: A DNS resolver using TCP or UDP port 53.

3. DNS Service Binding Records

DNS resolvers can advertise one or more Designated Resolvers that may offer support over encrypted channels and are controlled by the same entity.

When a client discovers Designated Resolvers, it learns information such as the supported protocols, ports, and server name to use in certificate validation. This information is provided in Service Binding (SVCB) records for DNS Servers, defined by [I-D.schwartz-svcb-dns].

The following is an example of an SVCB record describing a DoH server:

```
_dns.example.net  7200  IN SVCB 1 . (
     alpn=h2 dohpath=/dns-query{?dns} ipv4hint=x.y.z.w )
```

The following is an example of an SVCB record describing a DoT server:

```
_dns.example.net  7200  IN SVCB 1 dot.example.net (
     alpn=dot port=8530 ipv4hint=x.y.z.w )
```

If multiple Designated Resolvers are available, using one or more encrypted DNS protocols, the resolver deployment can indicate a preference using the priority fields in each SVCB record [I-D.ietf-dnsop-svcb-https].

This document focuses on discovering DoH and DoT Designated Resolvers. Other protocols can also use the format defined by [I-D.schwartz-svcb-dns]. However, if any protocol does not involve some form of certificate validation, new validation mechanisms will need to be defined to support validating equivalence as defined in Section 4.1.

4. Discovery Using Resolver IP Addresses

When a DNS client is configured with an Unencrypted Resolver IP address, it SHOULD query the resolver for SVCB records for "dns://resolver.arpa" before making other queries. Specifically, the client issues a query for "_dns.resolver.arpa" with the SVCB resource record type (64) [I-D.ietf-dnsop-svcb-https].

If the recursive resolver that receives this query has one or more Designated Resolvers, it will return the corresponding SVCB records. When responding to these special queries for "dns://resolver.arpa", the SVCB records SHOULD contain at least one "ipv4hint" and/or "ipv6hint" keys. These address hints indicate the address on which the corresponding Encrypted Resolver can be reached and avoid additional DNS lookup for the A and AAAA records of the Encrypted Resolver name.

4.1. Authenticated Discovery

In order to be considered an authenticated Designated Resolver, the TLS certificate presented by the Encrypted Resolver MUST contain both the domain name (from the SVCB answer) and the IP address of the designating Unencrypted Resolver within the SubjectAlternativeName certificate field. The client MUST check the SubjectAlternativeName field for both the Unencrypted Resolver's IP address and the advertised name of the Designated Resolver. If the certificate can be validated, the client SHOULD use the discovered Designated Resolver for any cases in which it would have otherwise used the Unencrypted Resolver. If the Designated Resolver has a different IP address than the Unencrypted Resolver and the TLS certificate does not cover the Unencrypted Resolver address, the client MUST NOT use the discovered Encrypted Resolver. Additionally, the client SHOULD suppress any further queries for Designated Resolvers using this Unencrypted Resolver for the length of time indicated by the SVCB record's Time to Live (TTL).

If the Designated Resolver and the Unencrypted Resolver share an IP address, clients MAY choose to opportunistically use the Encrypted Resolver even without this certificate check (Section 4.2).

4.2. Opportunistic Discovery

There are situations where authenticated discovery of encrypted DNS configuration over unencrypted DNS is not possible. This includes Unencrypted Resolvers on non-public IP addresses whose identity cannot be confirmed using TLS certificates.

Opportunistic Privacy is defined for DoT in Section 4.1 of [RFC7858] as a mode in which clients do not validate the name of the resolver presented in the certificate. A client MAY use information from the SVCB record for "dns://resolver.arpa" with this "opportunistic" approach (not validating the names presented in the SubjectAlternativeName field of the certificate) as long as the IP address of the Encrypted Resolver does not differ from the IP address of the Unencrypted Resolver, and that IP address is a private address (such as those defined in [RFC1918]). This approach can be used for DoT or DoH.

If the IP addresses of the Encrypted and Unencrypted Resolvers are not the same, or the shared IP address is not a private IP address, the client MUST NOT use the Encrypted Resolver opportunistically.

5. Discovery Using Resolver Names

A DNS client that already knows the name of an Encrypted Resolver can use DEER to discover details about all supported encrypted DNS protocols. This situation can arise if a client has been configured to use a given Encrypted Resolver, or if a network provisioning protocol (such as DHCP or IPv6 Router Advertisements) provides a name for an Encrypted Resolver alongside the resolver IP address.

For these cases, the client simply sends a DNS SVCB query using the known name of the resolver. This query can be issued to the named Encrypted Resolver itself or to any other resolver. Unlike the case of bootstrapping from an Unencrypted Resolver (Section 4), these records SHOULD be available in the public DNS.

For example, if the client already knows about a DoT server "resolver.example.com", it can issue an SVCB query for "_dns.resolver.example.com" to discover if there are other encrypted DNS protocols available. In the following example, the SVCB answers indicate that "resolver.example.com" supports both DoH and DoT, and that the DoH server indicates a higher priority than the DoT server.

```
_dns.resolver.example.com  7200  IN SVCB 1 . (
     alpn=h2 dohpath=/dns-query{?dns} )
_dns.resolver.example.com  7200  IN SVCB 2 . (
     alpn=dot )
```

Often, the various supported encrypted DNS protocols will be accessible using the same hostname. In the example above, both DoH and DoT use the name "resolver.example.com" for their TLS certficates. If a deployment uses a different hostname for one protocol, but still wants clients to treat both DNS servers as designated, the TLS certificates MUST include both names in the SubjectAlternativeName fields. Note that this name verification is not related to the DNS resolver that provided the SVCB answer.

For example, being able to discover a Designated Resolver for a known Encrypted Resolver is useful when a client has a DoT configuration for "foo.resolver.example.com" but is on a network that blocks DoT traffic. The client can still send a query to any other accessible resolver (either the local network resolver or an accessible DoH server) to discover if there is a designated DoH server for "foo.resolver.example.com".

6. Deployment Considerations

Resolver deployments that support DEER are advised to consider the following points.

6.1. Caching Forwarders

If a caching forwarder consults multiple resolvers, it may be possible for it to cache records for the "resolver.arpa" Special Use Domain Name (SUDN) for multiple resolvers. This may result in clients sending queries intended to discover Designated Resolvers for resolver "foo" and receiving answers for resolvers "foo" and "bar".

A client will successfully reject unintended connections because the authenticated discovery will fail or the resolver addresses do not match. Clients that attempt unauthenticated connections to resolvers discovered through SVCB queries run the risk of connecting to the wrong server in this scenario.

To prevent unnecessary traffic from clients to incorrect resolvers, DNS caching resolvers SHOULD NOT cache results for the "resolver.arpa" SUDN other than for Designated Resolvers under their control.

6.2. Certificate Management

Resolver owners that support authenticated discovery will need to list valid referring IP addresses in their TLS certificates. This may pose challenges for resolvers with a large number of referring IP addresses.

7. Security Considerations

Since client can receive DNS SVCB answers over unencrypted DNS, on-path attackers can prevent successful discovery by dropping SVCB packets. Clients should be aware that it might not be possible to distinguish between resolvers that do not have any Designated Resolver and such an active attack.

While the IP address of the Unencrypted Resolver is often provisioned over insecure mechanisms, it can also be provisioned securely, such as via manual configuration, a VPN, or on a network with protections like RA guard [RFC6105]. An attacker might try to direct Encrypted DNS traffic to itself by causing the client to think that a discovered Designated Resolver uses a different IP address from the Unencrypted Resolver. Such an Encrypted Resolver might have a valid certificate, but be operated by an attacker that is trying to observe or modify user queries without the knowledge of the client or network.

If the IP address of a Designated Resolver differs from that of an Unencrypted Resolver, clients MUST validate that the IP address of the Unencrypted Resolver is covered by the SubjectAlternativeName of the Encrypted Resolver's TLS certificate (Section 4.1).

Opportunistic use of Encrypted Resolvers MUST be limited to cases where the Unencrypted Resolver and Designated Resolver have the same IP address (Section 4.2).

8. IANA Considerations

8.1. 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 designated encrypted 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.

9. References

9.1. 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-02, 2 November 2020, <http://www.ietf.org/internet-drafts/draft-ietf-dnsop-svcb-https-02.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-09, 16 December 2020, <http://www.ietf.org/internet-drafts/draft-ietf-tls-esni-09.txt>.

[I-D.schwartz-svcb-dns]
Schwartz, B., "Service Binding Mapping for DNS Servers", Work in Progress, Internet-Draft, draft-schwartz-svcb-dns-01, 10 August 2020, <http://www.ietf.org/internet-drafts/draft-schwartz-svcb-dns-01.txt>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., 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>.

[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>.

9.2. 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-02, 13 July 2020, <http://www.ietf.org/internet-drafts/draft-schinazi-httpbis-doh-preference-hints-02.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>.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <https://www.rfc-editor.org/info/rfc6105>.

[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 Designated Resolver for clients to use in place of an Unencrypted Resolver (using a SUDN) or another Encrypted Resolver (using its domain name).

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 a DNS resolver operator to define which resolver names and IP addresses are related to one another.

* Using DNS records also does not 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 for discovering Designated Resolvers of different transport types as well as those advertised by Unencrypted Resolvers or another Encrypted Resolver.

* 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 Designated Resolver 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

ADD
Internet-Draft
Intended status: Standards Track
Expires: November 5, 2021

M. Boucadair, Ed.
Orange
T. Reddy, Ed.
McAfee
D. Wing
Citrix
N. Cook
Open-Xchange
T. Jensen
Microsoft
May 4, 2021

DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)
draft-ietf-add-dnr-01

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 set of service parameters to reach such encrypted DNS servers.

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 5, 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 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 a set of service parameters.

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 7.

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].

Encrypted DNS options: refers to the options defined in Sections 4, 5, and 6.

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

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

DHCP: refers to both DHCPv4 and DHCPv6.

3. Overview

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

These options configure an authentication domain name, a list of IPv6 addresses, and a set of service parameters of the encrypted DNS server. More information about the design of these options is provided in the following subsections.

3.1. Configuration Data for Encrypted DNS

In order to allow for PKIX-based authentication between a DNS client and an encrypted DNS server, the Encrypted DNS options are designed to include an authentication domain name. This ADN is presented as a reference identifier for DNS authentication purposes. This design accommodates the current best practices for issuing certificates as per 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.

To avoid adding a dependency on another server to resolve the ADN, the Encrypted DNS options return the IP address(es) 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.

In order to optimize the size of discovery messages when all servers terminate on the same IP address, early versions of this document considered relying 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 a design.

A list of IP addresses to reach an encrypted DNS server may be returned in the Encrypted DNS options to accommodate current deployments relying upon primary and backup servers. Whether one IP address or more are returned in an Encrypted DNS 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 option, these addresses are ordered in the preference for use by the client.

Because distinct Encrypted DNS protocols may be provisioned by a network (e.g., DoT, DoH, and DoQ) and that some of these protocols may make use of customized port numbers instead of default ones, the Encrypted DNS options are designed to return a set of service parameters. These parameters are encoded following the same rules for encoding SvcParams in Section 2.1 of [I-D.ietf-dnsop-svcb-https]. This encoding approach may increase the size of the options but it has the merit to rely upon an existing IANA registry and thus to accommodate new Encrypted DNS protocols and service parameters that may be defined in the future. For example, "dohpath" service parameter (Section 5.1 of [I-D.schwartz-svcb-dns]) supplies a relative DoH URI Template.

A single option is used to convey both the ADN and IP addresses because otherwise means to correlate an IP address with an ADN will be required if, for example, more than one ADN is supported by the network.

3.2. Handling Configuration Data Conflicts

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.

DHCP/RA options to discover encrypted DNS servers (including, DoH URI Templates) takes precedence over DDR [I-D.ietf-add-ddr] since DDR uses unencrypted DNS to an external DNS resolver, which is

susceptible to both internal and external attacks whereas DHCP/RA is typically protected using the mechanisms discussed in Section 8.1.

3.3. Connection Establishment

If the local DNS client supports one of the discovered Encrypted DNS protocols identified by Application Layer Protocol Negotiation (ALPN) protocol identifiers, the DNS client establishes an encrypted DNS session following the order of the discovered servers. The client follows 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. ALPN-related considerations can be found in Section 6.1 of [I-D.ietf-dnsop-svcb-https].

3.4. Multihoming Considerations

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.

4. DHCPv6 Encrypted DNS Option

4.1. Option Format

The format of the DHCPv6 Encrypted DNS option is shown in Figure 1.

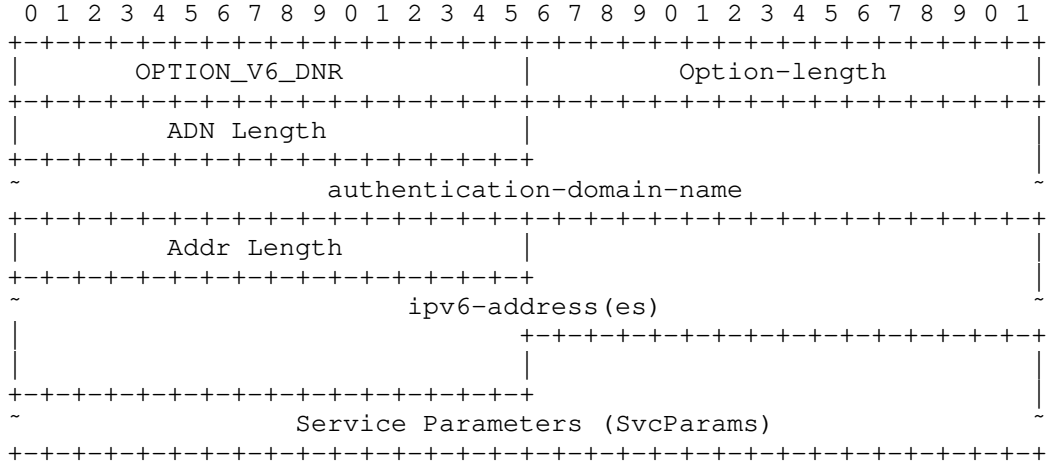

Figure 1: DHCPv6 Encrypted DNS Option

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

Option-code: OPTION_V6_DNR (TBA1, see Section 9.1)

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

ADN Length: Length of the authentication-domain-name field in octets.

authentication-domain-name (variable length): 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 2. This example conveys the FQDN "doh1.example.com.", and the resulting Option-length field is 18.

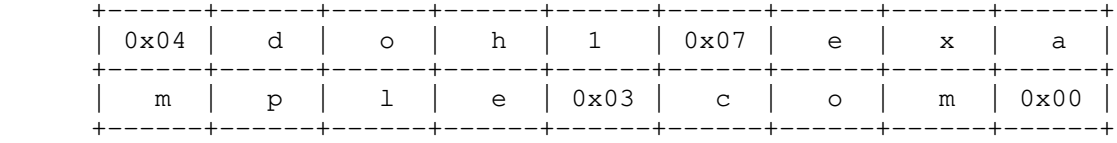

Figure 2: An Example of the DNS authentication-domain-name Encoding

Addr Length: Length of enclosed IPv6 addresses in octets. It MUST be a multiple of 16.

ipv6-address(es) (variable length): Indicates one or more IPv6 addresses to reach the encrypted DNS server. An address can be link-local, ULA, or GUA. The format of this field is shown in Figure 3.

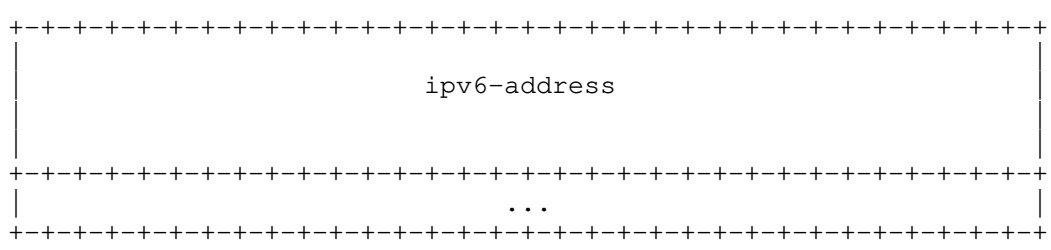

Figure 3: Format of the IPv6 Addresses Field

Service Parameters (SvcParams) (variable length): Specifies a set of service parameters that are encoded following the rules in Section 2.1 of [I-D.ietf-dnsop-svcb-https]. Service parameters

may include, for example, a list of ALPN protocol identifiers or alternate port numbers.

If no port service parameter is included, this indicates that default port numbers should be used. As a reminder, the default port number is 853 for DoT and 443 for DoH.

The length of this field is ('Option-length' - 4 - 'ADN Length' - 'Addr Length').

Multiple instances of OPTION_V6_DNR may be returned to a DHCPv6 client; each pointing to a distinct encrypted DNS server. These instances are ordered in the preference for use by the client.

4.2. DHCPv6 Client Behavior

To discover an encrypted DNS server, the DHCPv6 client MUST include OPTION_V6_DNR 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 DHCP client MUST be prepared to receive multiple OPTION_V6_DNR 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_DNR.

5. DHCPv4 Encrypted DNS Option

5.1. Option Format

The format of the DHCPv4 Encrypted DNS option is illustrated in Figure 4.

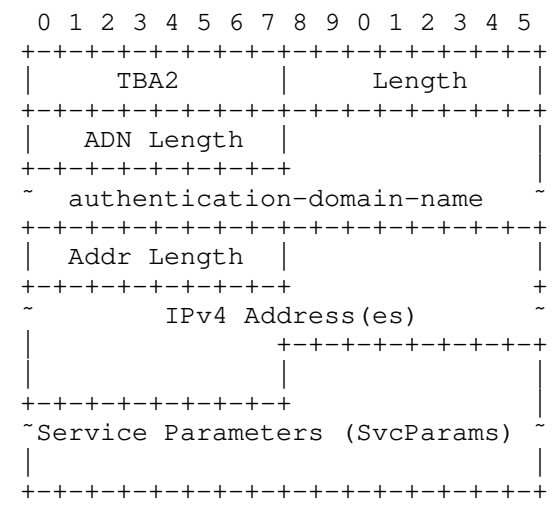

Figure 4: DHCPv4 Encrypted DNS Option

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

Code: OPTION_V4_DNR (TBA2, see Section 9.2).

Length: Indicates the length of the enclosed data in octets.

ADN Length: Indicates the length of the authentication-domain-name in octets.

authentication-domain-name (variable length): Includes the authentication 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 5. The values s1, s2, s3, etc. represent the domain name labels in the domain name encoding.

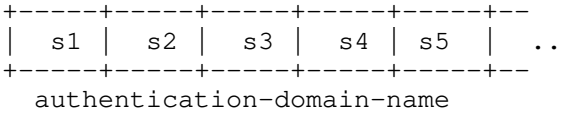

Figure 5: Format of the Authentication Domain Name Field

Addr Length: Indicates the length of included IPv4 addresses in octets. It MUST be a multiple of 4.

IPv4 Address(es) (variable length): 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.

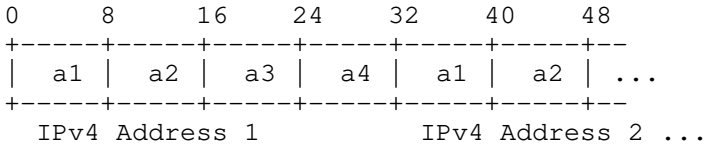

Figure 6: Format of the IPv4 Addresses Field

Service Paramters (SvcParams) (variable length): Specifies a set of service parameters that are encoded following the rules in Section 2.1 of [I-D.ietf-dnsop-svcb-https]. Service parameters may include, for example, a list of ALPN protocol identifiers or alternate port numbers.

If no port service parameter is included, this indicates that default port numbers should be used.

The length of this field is ('Option-length' - 2 - 'ADN Length' - 'Addr Length').

OPTION_V4_DNR is a concatenation-requiring option. As such, the mechanism specified in [RFC3396] MUST be used if OPTION_V4_DNR exceeds the maximum DHCPv4 option size of 255 octets.

Multiple instances of OPTION_V4_DNR may be returned to a DHCPv4 client; each pointing to a distinct encrypted DNS server. These instances are ordered in the preference for use by the client.

5.2. DHCPv4 Client Behavior

To discover an encrypted DNS server, the DHCPv4 client requests the Encrypted DNS server by including OPTION_V4_DNR in a Parameter Request List option [RFC2132].

The DHCPv4 client MUST be prepared to receive multiple DHCP OPTION_V4_DNR 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_DNR.

6. IPv6 RA Encrypted DNS Option

6.1. Option Format

This section defines a new Neighbor Discovery option [RFC4861]: IPv6 RA Encrypted DNS option. This option is useful in contexts similar to those discussed in Section 1.1 of [RFC8106].

The format of the IPv6 RA Encrypted DNS option is illustrated in Figure 7.

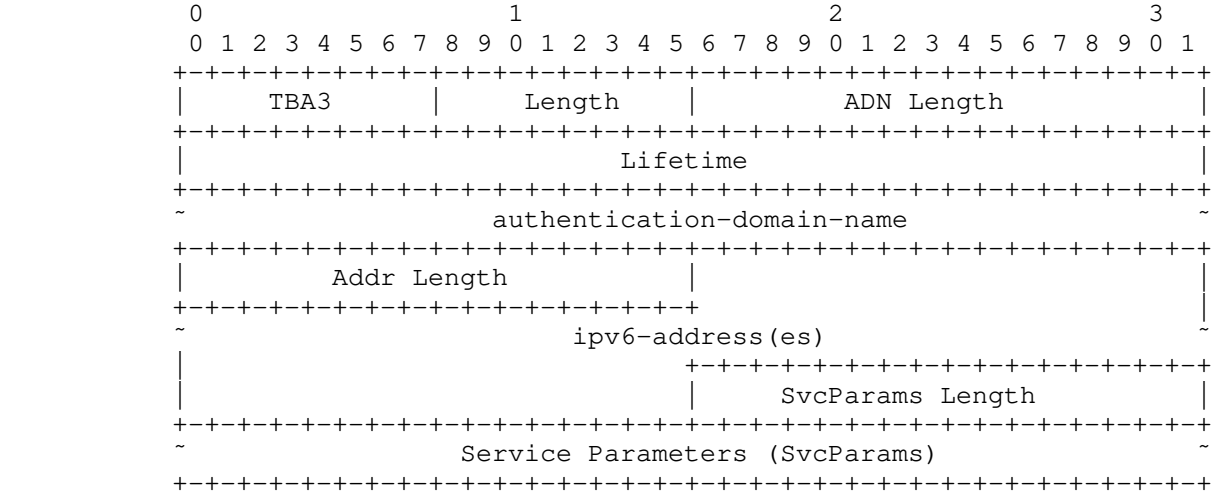

Figure 7: RA Encrypted DNS Option

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

Type: 8-bit identifier of the Encrypted DNS Option as assigned by IANA (TBA3, see Section 9.3).

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

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.

ADN Length: 16-bit unsigned integer. This field indicates the length of the authentication-domain-name field in octets.

authentication-domain-name (variable length): The domain name of the encrypted DNS server. This field is formatted as specified in Section 10 of [RFC8415].

Addr Length: 16-bit unsigned integer. This field indicates the length of enclosed IPv6 addresses in octets. It MUST be a multiple of 16.

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

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 options may be used.

The format of this field is shown in Figure 3.

SvcParams Length: 16-bit unsigned integer. This field indicates the length of the Service Parameters field in octets.

Service Paramters (SvcParams) (variable length): Specifies a set of service parameters that are encoded following the rules in Section 2.1 of [I-D.ietf-dnsop-svcb-https]. Service parameters may include, for example, a list of ALPN protocol identifiers or alternate port numbers.

If no port service parameter is included, this indicates that default port numbers should be used.

The option MUST be padded with zeros so that the full enclosed data is a multiple of 8 octets (Section 4.6 of [RFC4861]).

6.2. IPv6 Host Behavior

The procedure for DNS configuration is the same as it is with any other Neighbor Discovery option [RFC4861]. In addition, the host follows the procedure described in Section 5.3.1 of [RFC8106].

The host MUST silently discard multicast and host loopback addresses conveyed in the Encrypted DNS options.

7. Hosting Encrypted DNS Forwarder in Local Networks

This section discusses some deployment considerations to host an encrypted DNS forwarder within a local network.

7.1. Managed CPEs

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

7.1.1. 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.

7.1.2. 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.

7.1.3. 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 7.1.2) 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 8, the auto-upgrade to a rogue server advertising "rs.example.org" will fail because it does not match "*.example.net".

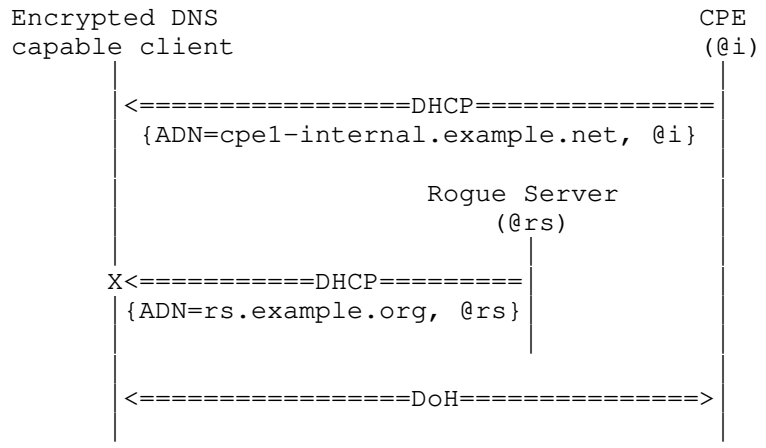

Figure 8: A Simplified Example of Auto-upgrade based on Subdomains

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 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 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 9 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.

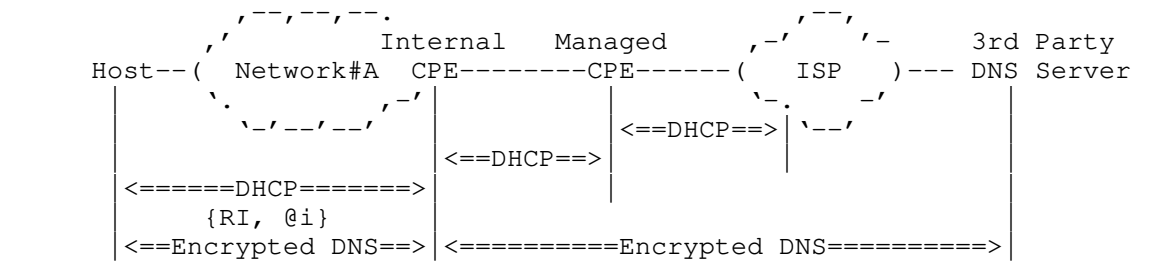

Legend:
* @i: IP address of the DNS forwarder hosted in the Internal CPE.

Figure 9: Example of an Internal CPE Hosting a Forwarder

8. Security Considerations

8.1. 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 CPE 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.

8.2. 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.

8.3. 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.

8.4. 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 that an information (e.g., NTP server, DNS server, default domain) provided by such networks via DHCP, DHCPv6, or RA are untrusted because DHCP and RA messages 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. As such, 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 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]).

9. IANA Considerations

9.1. DHCPv6 Option

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_DNR	Yes	No	[ThisDocument]

9.2. 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
TBA2	OPTION_V4_DNR	N	Encrypted DNS Server	[ThisDocument]

9.3. Neighbor Discovery 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 [ND].

Type	Description	Reference
TBA3	DNS Encrypted DNS Option	[ThisDocument]

10. 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].

Thanks to Bernie Volz for the review of the DHCP part.

11. Contributing Authors

Nicolai Leymann
Deutsche Telekom
Germany

Email: n.leymann@telekom.de

Zhiwei Yan
CNNIC
No.4 South 4th Street, Zhongguancun
Beijing 100190
China

EMail: yan@cnnic.cn

12. References

12.1. 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)", draft-ietf-dnsop-svcb-https-05 (work in progress), April 2021.

[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>.

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>.

[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-add-ddr] Pauly, T., Kinnear, E., Wood, C. A., McManus, P., and T. Jensen, "Discovery of Designated Resolvers", draft-ietf-add-ddr-00 (work in progress), February 2021.

[I-D.ietf-dprive-dnsoquic] Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnsoquic-02 (work in progress), February 2021.

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

[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-03 (work in progress), April 2021.

[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

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 10. 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.

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.

(a) Fixed Networks

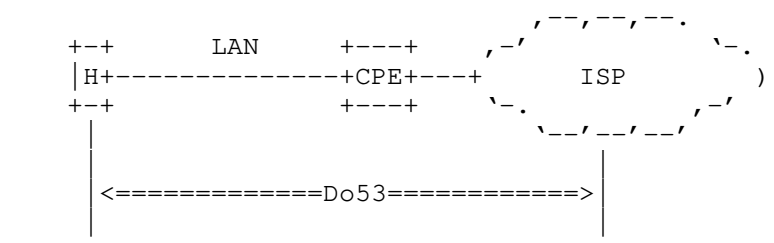

(b) Cellular Networks

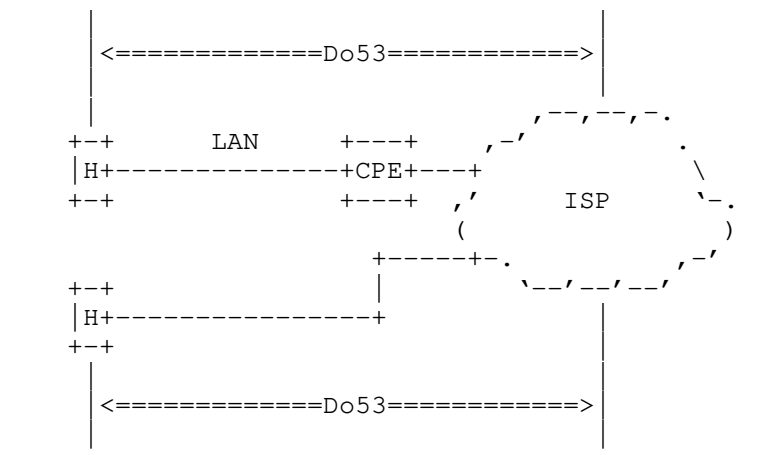

Legend:
 * H: refers to a host.

Figure 10: 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 11.

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

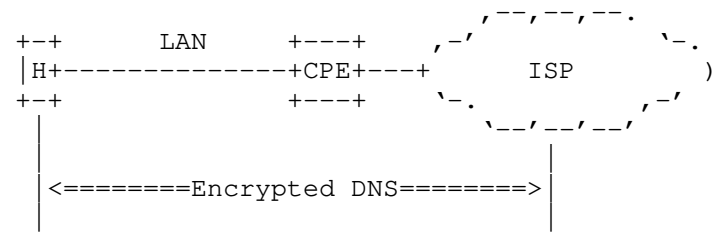

(b) Cellular Networks

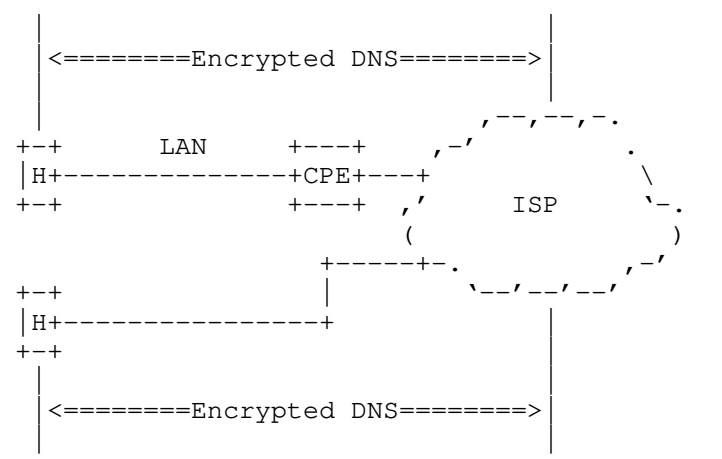

Figure 11: Encrypted DNS in the WAN

Figure 11 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 12 for a DoH/DoT-capable host).

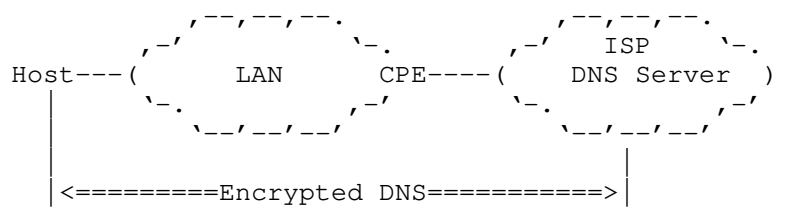

Figure 12: Direct Encrypted DNS Sessions

A.1.2. Proxied DNS

Figure 13 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 13. Additional considerations related to this deployment are discussed in Section 7.

(a)

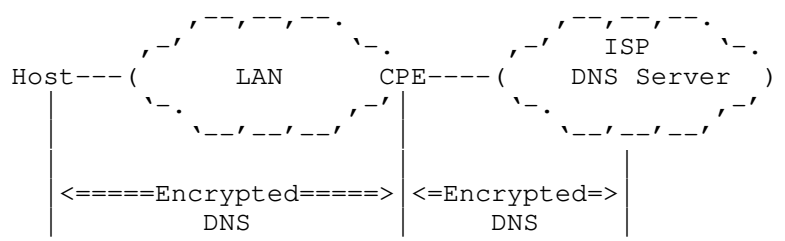

(b)

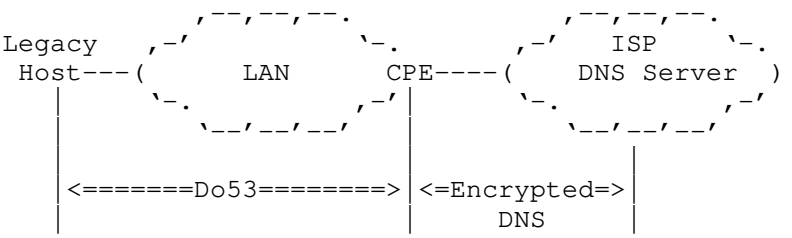

Figure 13: Proxied Encrypted DNS Sessions

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 12 or Figure 13.

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 14).

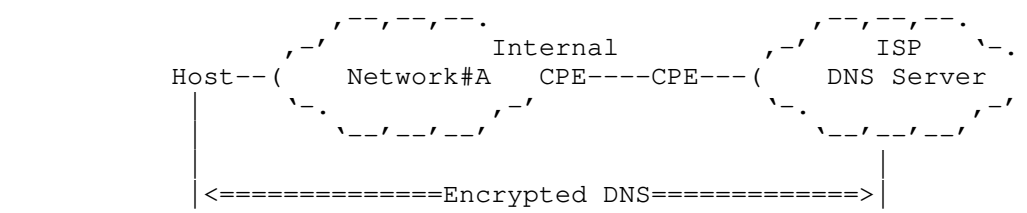

Figure 14: Direct Encrypted DNS Sessions with the ISP DNS Resolver (Internal CPE)

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 15).

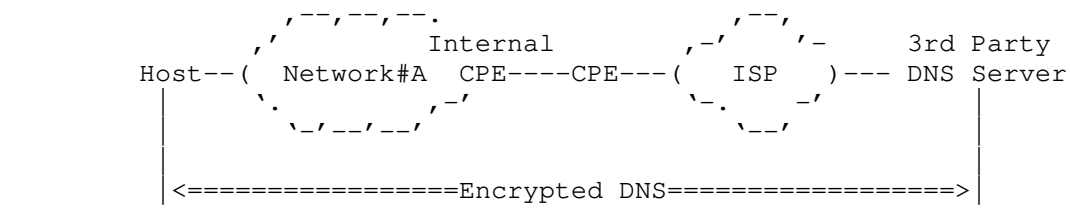

Figure 15: Direct Encrypted DNS Sessions with a Third Party DNS Resolver

Section 7.2 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.

Appendix C. 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 discovery using the resolver IP address as defined in Section 4 of [I-D.ietf-add-ddr] to discover the designated resolvers. The guidance in Sections 4.1 and 4.2 of [I-D.ietf-add-ddr] related to the designated resolver verification has to be followed in such case.

Authors' Addresses

Mohamed Boucadair (editor)
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy (editor)
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