

ADD
Internet-Draft
Intended status: Informational
Expires: November 18, 2021

M. Boucadair, Ed.
Orange
T. Reddy, Ed.
McAfee
D. Wing
Citrix
N. Cook
Open-Xchange
T. Jensen
Microsoft
May 17, 2021

Discovery of Encrypted DNS Resolvers: Deployment Considerations
draft-boucadair-add-deployment-considerations-00

Abstract

The document discusses some deployment considerations of the various options to discover encrypted DNS servers (e.g., DNS-over-HTTPS, DNS-over-TLS, DNS-over-QUIC).

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 18, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect 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.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Sample Target Deployment Scenarios	3
3.1. Managed CPEs	5
3.1.1. Direct DNS	5
3.1.2. Proxied DNS	6
3.2. Unmanaged CPEs	7
3.2.1. ISP-facing Unmanaged CPEs	7
3.2.2. Internal Unmanaged CPEs	7
4. Hosting Encrypted DNS Forwarder in Local Networks	8
4.1. Managed CPEs	8
4.1.1. DNS Forwarders	8
4.1.2. ACME	9
4.2. Unmanaged CPEs	9
5. Legacy CPEs	10
6. Security Considerations	10
7. IANA Considerations	10
8. Acknowledgements	10
9. References	10
9.1. Normative References	11
9.2. Informative References	11
Authors' Addresses	12

1. Introduction

[I-D.ietf-add-dnr] 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.

This document discusses deployment considerations for the discovery of encrypted DNS servers such as DNS-over-HTTPS (DoH) [RFC8484], DNS-over-TLS (DoT) [RFC7858], or DNS-over-QUIC (DoQ) [I-D.ietf-dprive-dnsquic] in local networks.

Sample target deployment scenarios are discussed in Section 3; 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 can be used.

Considerations related to hosting a DNS forwarder in a local network are described in Section 4.

2. Terminology

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-dnsquic].

Encrypted DNS options: refers to the options defined in [I-D.ietf-add-dnr].

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. 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:


- o Protocol Configuration Options in cellular networks [TS.24008].
- o DHCPv4 [RFC2132] (Domain Name Server Option) or DHCPv6 [RFC8415][RFC3646] (OPTION_DNS_SERVERS).
- o 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 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.

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 1: Sample Legacy Deployments

3.1. Managed CPEs

This section focuses on CPEs that are managed by ISPs.


3.1.1. 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 2.

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 the service parameters (ALPN).

(a) Fixed Networks



(b) Cellular Networks




Figure 2: Encrypted DNS in the WAN

Figure 2 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 3 for a DoH/DoT-capable host).




Figure 3: Direct Encrypted DNS Sessions


3.1.2. Proxied DNS

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., 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:

- o The leg between an internal host and the CPE.
- o 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 4. Additional considerations related to this deployment are discussed in Section 4.

(a)



(b)




Figure 4: Proxied Encrypted DNS Sessions

3.2. Unmanaged CPEs

3.2.1. 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 Section 3.1. A host can thus establish DoH/DoT session with a DoH/DoT server similar to what is depicted in Figure 3 or Figure 4.

3.2.2. 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 5).




Figure 5: 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 6).




Figure 6: Direct Encrypted DNS Sessions with a Third Party DNS Resolver

Section 4.2 discusses considerations related to hosting a forwarder in the Internal CPE.

4. Hosting Encrypted DNS Forwarder in Local Networks

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

4.1. Managed CPEs

The section discusses mechanisms that can be used to host an encrypted DNS forwarder in a managed CPE (Section 3.1).

4.1.1. 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.

4.1.2. 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.

4.2. Unmanaged CPEs

The approach specified in Section 4.1 does not apply for hosting a DNS forwarder in an unmanaged CPE.

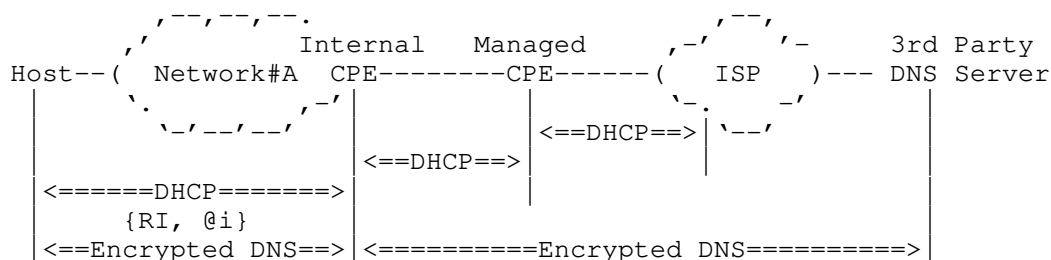
The unmanaged CPE administrator can host an encrypted DNS forwarder on the unmanaged CPE. This assumes the following:

- o 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.

- o The encrypted DNS forwarder will either be configured to use the ISP's or a 3rd party encrypted DNS server.
- o The unmanaged CPE will advertise the encrypted DNS forwarder ADN using DHCP/RA to internal hosts.

Figure 7 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 7: Example of an Internal CPE Hosting a Forwarder

5. Legacy CPEs

Hosts serviced by legacy CPEs that can't be upgraded to support the options defined in Sections 4, 5, and 6 of [I-D.ietf-add-dnr] 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.

6. Security Considerations

DNR-related security considerations are discussed in Section 7 of [I-D.ietf-add-dnr].

7. IANA Considerations

This document does not require any IANA action.

8. Acknowledgements

This text was initially part of [I-D.ietf-add-dnr].

9. References

9.1. Normative References

- [I-D.ietf-add-dnr]
Boucadair, M., Reddy, T., Wing, D., Cook, N., and T. Jensen, "DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)", draft-ietf-add-dnr-00 (work in progress), February 2021.

9.2. Informative References

- [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-dnssoquic]
Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnssoquic-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.
- [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>>.
- [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>>.
- [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>>.

- [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>>.
- [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 (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

ADD
Internet-Draft
Intended status: Standards Track
Expires: 6 October 2022

T. Pauly
E. Kinnear
Apple Inc.
C. A. Wood
Cloudflare
P. McManus
Fastly
T. Jensen
Microsoft
4 April 2022

Discovery of Designated Resolvers
draft-ietf-add-ddr-06

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 a 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 or cooperating entities. It can also be used to discover support for encrypted DNS protocols when the name of an encrypted resolver is known.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Adaptive DNS Discovery Working Group mailing list (add@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/add/>.

Source for this draft and an issue tracker can be found at <https://github.com/ietf-wg-add/draft-ietf-add-ddr>.

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 6 October 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Specification of Requirements	3
2. Terminology	3
3. DNS Service Binding Records	4
4. Discovery Using Resolver IP Addresses	5
4.1. Use of Designated Resolvers	6
4.2. Verified Discovery	7
4.3. Opportunistic Discovery	8
5. Discovery Using Resolver Names	8
6. Deployment Considerations	9
6.1. Caching Forwarders	9
6.2. Certificate Management	10
6.3. Server Name Handling	10
6.4. Handling non-DDR queries for resolver.arpa	10
6.5. Interaction with Network-Designated Resolvers	10
7. Security Considerations	11
8. IANA Considerations	11
8.1. Special Use Domain Name "resolver.arpa"	12
9. References	12
9.1. Normative References	12
9.2. Informative References	13
Appendix A. Rationale for using SVCB records	15
Authors' Addresses	15

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 (SUDN) [RFC6761] to discover DNS SVCB records associated with one or more Encrypted Resolvers the Unencrypted Resolver has designated for use when support for DNS encryption is requested (Section 4).
2. When the hostname of an Encrypted Resolver 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. "Designated" in this context means that the resolvers are operated by the same entity or cooperating entities; for example, the resolvers are accessible on the same IP address, or there is a certificate that claims ownership over the IP address for the original designating resolver.

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 verified 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 and ports. This information is provided in ServiceMode Service Binding (SVCB) records for DNS Servers, although AliasMode SVCB records can be used to direct clients to the needed ServiceMode SVCB record per [I-D.ietf-dnsop-svcb-https]. The formatting of these records, including the DNS-unique parameters such as "dohpath", are defined by [I-D.ietf-add-svcb-dns].

The following is an example of an SVCB record describing a DoH server discovered by querying for `_dns.example.net`:

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

The following is an example of an SVCB record describing a DoT server discovered by querying for `_dns.example.net`:

```
_dns.example.net 7200 IN SVCB 1 dot.example.net (
    alpn=dot port=8530 )
```

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].

If the client encounters a mandatory parameter in an SVCB record it does not understand, it MUST NOT use that record to discover a Designated Resolver. The client can still use others records in the

same response if the client can understand all of their mandatory parameters. This allows future encrypted deployments to simultaneously support protocols even if a given client is not aware of all those protocols. For example, if the Unencrypted Resolver returns three SVCB records, one for DoH, one for DoT, and one for a yet-to-exist protocol, a client which only supports DoH and DoT should be able to use those records while safely ignoring the third record.

To avoid name lookup deadlock, Designated Resolvers SHOULD follow the guidance in Section 10 of [RFC8484] regarding the avoidance of DNS-based references that block the completion of the TLS handshake.

This document focuses on discovering DoH and DoT Designated Resolvers. Other protocols can also use the format defined by [I-D.ietf-add-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 designation as defined in Section 4.2.

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].

Because this query is for an SUDN, which no entity can claim ownership over, the ServiceMode SVCB response MUST NOT use the "." value for the TargetName. Instead, the domain name used for DoT or used to construct the DoH template MUST be provided.

The following is an example of an SVCB record describing a DoH server discovered by querying for _dns.resolver.arpa:

```
_dns.resolver.arpa 7200 IN SVCB 1 doh.example.net (  
    alpn=h2 dohpath=/dns-query{?dns} )
```

The following is an example of an SVCB record describing a DoT server discovered by querying for _dns.resolver.arpa:

```
_dns.resolver.arpa 7200 IN SVCB 1 dot.example.net (  
    alpn=dot port=8530 )
```

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 recursive resolver SHOULD include the A and AAAA records for the name of the Designated Resolver in the Additional Answers section. This will allow the DNS client to make queries over an encrypted connection without waiting to resolve the Encrypted Resolver name per [I-D.ietf-dnsop-svcb-https]. If no A/AAAA records or SVCB IP address hints are included, clients will be forced to delay use of the Encrypted Resolver until an additional DNS lookup for the A and AAAA records can be made to the Unencrypted Resolver (or some other resolver the DNS client has been configured to use).

If the recursive resolver that receives this query has no Designated Resolvers, it SHOULD return NODATA for queries to the "resolver.arpa" SUDN.

4.1. Use of Designated Resolvers

When a client discovers Designated Resolvers from an Unencrypted Resolver IP address, it can choose to use these Designated Resolvers either automatically, or based on some other policy, heuristic, or user choice.

This document defines two preferred methods to automatically use Designated Resolvers:

- * Verified Discovery Section 4.2, for when a TLS certificate can be used to validate the resolver's identity.
- * Opportunistic Discovery Section 4.3, for when a resolver is accessed using a non-public IP address.

A client MAY additionally use a discovered Designated Resolver without either of these methods, based on implementation-specific policy or user input. Details of such policy are out of scope of this document. Clients SHOULD NOT automatically use a Designated Resolver without some sort of validation, such as the two methods defined in this document or a future mechanism.

A client MUST NOT use a Designated Resolver designated by one Unencrypted Resolver in place of another Unencrypted Resolver. As these are known only by IP address, this means each unique IP address used for unencrypted DNS requires its own designation discovery. This ensures queries are being sent to a party designated by the resolver originally being used.

Generally, clients also SHOULD NOT reuse the Designated Resolver discovered from an Unencrypted Resolver over one network connection in place of the same Unencrypted Resolver on another network connection. Instead, clients SHOULD repeat the discovery process on the other network connection.

However, if a given Unencrypted Resolver designates a Designated Resolver that uses a public IP address and can be verified using the mechanism described in Section 4.2, it MAY be used on different network connections so long as the subsequent connections over other networks can also be successfully verified using the mechanism described in Section 4.2. This is a tradeoff between performance (by having no delay in establishing an encrypted DNS connection on the new network) and functionality (if the Unencrypted Resolver intends to designate different Designated Resolvers based on the network from which clients connect).

4.2. Verified Discovery

Verified Discovery is a mechanism that allows automatic use of a Designated Resolver that supports DNS encryption that performs a TLS handshake.

In order to be considered a verified Designated Resolver, the TLS certificate presented by the Designated Resolver MUST contain the IP address of the designating Unencrypted Resolver in a subjectAltName extension. 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 automatically use the discovered Designated 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 Designated Resolver even without this certificate check (Section 4.3).

If resolving the name of a Designated Resolver from an SVCB record yields an IP address that was not presented in the Additional Answers section or ipv4hint or ipv6hint fields of the original SVCB query, the connection made to that IP address MUST pass the same TLS certificate checks before being allowed to replace a previously known and validated IP address for the same Designated Resolver name.

4.3. Opportunistic Discovery

There are situations where Verified Discovery of encrypted DNS configuration over unencrypted DNS is not possible. This includes Unencrypted Resolvers on non-public IP addresses such as those defined in [RFC1918] 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. Clients SHOULD use this mode only for resolvers using non-public IP addresses. This approach can be used for any encrypted DNS protocol that uses TLS.

5. Discovery Using Resolver Names

A DNS client that already knows the name of an Encrypted Resolver can use DDR 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 resolver.example.com. (
  alpn=h2 dohpath=/dns-query{?dns} )
_dns.resolver.example.com. 7200 IN SVCB 1 resolver.example.com. (
  alpn=dot )
```

Clients MUST validate that for any Encrypted Resolver discovered using a known resolver name, the TLS certificate of the resolver contains the known name in a subjectAltName extension. In the example above, this means that both servers need to have certificates that cover the name resolver.example.com. Often, the various supported encrypted DNS protocols will be specified such that the SVCB TargetName matches the known name, as is true in the example above. However, even when the TargetName is different (for example, if the DoH server had a TargetName of doh.example.com), the clients still check for the original known resolver name in the certificate.

Note that this resolver validation is not related to the DNS resolver that provided the SVCB answer.

As another 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 DDR are advised to consider the following points.

6.1. Caching Forwarders

A DNS forwarder SHOULD NOT forward queries for "resolver.arpa" upstream. This prevents a client from receiving an SVCB record that will fail to authenticate because the forwarder's IP address is not in the upstream resolver's Designated Resolver's TLS certificate SAN field. A DNS forwarder which already acts as a completely blind forwarder MAY choose to forward these queries when the operator expects that this does not apply, either because the operator knows the upstream resolver does have the forwarder's IP address in its TLS certificate's SAN field or that the operator expects clients of the unencrypted resolver to use the SVCB information opportunistically.

Operators who choose to forward queries for "resolver.arpa" upstream should note that client behavior is never guaranteed and use of DDR by a resolver does not communicate a requirement for clients to use the SVCB record when it cannot be verified.

6.2. Certificate Management

Resolver owners that support Verified 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.

6.3. Server Name Handling

Clients MUST NOT use "resolver.arpa" as the server name either in the TLS Server Name Indication (SNI) ([RFC8446]) for DoT or DoH connections, or in the URI host for DoH requests.

When performing discovery using resolver IP addresses, clients MUST use the IP address as the URI host for DoH requests.

Note that since IP addresses are not supported by default in the TLS SNI, resolvers that support discovery using IP addresses will need to be configured to present the appropriate TLS certificate when no SNI is present for both DoT and DoH.

6.4. Handling non-DDR queries for resolver.arpa

DNS resolvers that support DDR by responding to queries for `_dns.resolver.arpa` SHOULD treat `resolver.arpa` as a locally served zone per [RFC6303]. In practice, this means that resolvers SHOULD respond to queries of any type other than SVCB for `_dns.resolver.arpa` with NODATA and queries of any type for any domain name under `resolver.arpa` with NODATA.

6.5. Interaction with Network-Designated Resolvers

Discovery of network-designated resolvers (DNR, [I-D.ietf-add-dnr]) allows a network to provide designation of resolvers directly through DHCP [RFC2132] [RFC8415] and IPv6 Router Advertisement (RA) [RFC4861] options. When such indications are present, clients can suppress queries for "resolver.arpa" to the unencrypted DNS server indicated by the network over DHCP or RAs, and the DNR indications SHOULD take precedence over those discovered using "resolver.arpa" for the same resolver if there is a conflict.

The designated resolver information in DNR might not contain a full set of SvcParams needed to connect to an encrypted resolver. In such a case, the client can use an SVCB query using a resolver name, as described in Section 5, to the authentication-domain-name (ADN).

7. Security Considerations

Since clients 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. To limit the impact of discovery queries being dropped either maliciously or unintentionally, clients can re-send their SVCB queries periodically.

DoH resolvers that allow discovery using DNS SVCB answers over unencrypted DNS MUST NOT provide differentiated behavior based on the HTTP path alone, since an attacker could modify the "dohpath" parameter.

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 a Designated 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 applying Verified Discovery (Section 4.2) MUST validate that the IP address of the Unencrypted Resolver is covered by the SubjectAlternativeName of the Designated Resolver's TLS certificate.

Clients using Opportunistic Discovery (Section 4.3) MUST be limited to cases where the Unencrypted Resolver and Designated Resolver have the same IP address.

The constraints on the use of Designated Resolvers specified here apply specifically to the automatic discovery mechanisms defined in this document, which are referred to as Verified Discovery and Opportunistic Discovery. Clients MAY use some other mechanism to verify and use Designated Resolvers discovered using the DNS SVCB record. However, use of such an alternate mechanism needs to take into account the attack scenarios detailed here.

8. IANA Considerations

8.1. Special Use Domain Name "resolver.arpa"

This document calls for the addition of "resolver.arpa" to the Special-Use Domain Names (SUDN) registry established by [RFC6761]. 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.

The "resolver.arpa" SUDN is similar to "ipv4only.arpa" in that the querying client is not interested in an answer from the authoritative "arpa" name servers. The intent of the SUDN is to allow clients to communicate with the Unencrypted Resolver much like "ipv4only.arpa" allows for client-to-middlebox communication. For more context, see the rationale behind "ipv4only.arpa" in [RFC8880].

IANA is requested to add an entry in "Transport-Independent Locally-Served DNS Zones" registry for 'resolver.arpa.' with the description "DNS Resolver Special-Use Domain", listing this document as the reference.

9. References

9.1. Normative References

- [I-D.ietf-add-svcb-dns]
Schwartz, B., "Service Binding Mapping for DNS Servers", Work in Progress, Internet-Draft, draft-ietf-add-svcb-dns-02, 1 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-add-svcb-dns-02>>.
- [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-08, 12 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-svcb-https-08>>.
- [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/rfc/rfc1918>>.

- [RFC6303] Andrews, M., "Locally Served DNS Zones", BCP 163, RFC 6303, DOI 10.17487/RFC6303, July 2011, <<https://www.rfc-editor.org/rfc/rfc6303>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <<https://www.rfc-editor.org/rfc/rfc6761>>.
- [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/rfc/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/rfc/rfc8484>>.

9.2. Informative References

- [I-D.ietf-add-dnr]
Boucadair, M., Reddy, T., Wing, D., Cook, N., and T. Jensen, "DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)", Work in Progress, Internet-Draft, draft-ietf-add-dnr-06, 22 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-add-dnr-06>>.
- [I-D.ietf-tls-esni]
Rescorla, E., Oku, K., Sullivan, N., and C. A. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-14, 13 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-tls-esni-14>>.
- [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, <<https://datatracker.ietf.org/doc/html/draft-schinazi-httpbis-doh-preference-hints-02>>.
- [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/rfc/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/rfc/rfc2132>>.
- [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/rfc/rfc4861>>.
- [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/rfc/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/rfc/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/rfc/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/rfc/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/rfc/rfc8415>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.
- [RFC8880] Cheshire, S. and D. Schinazi, "Special Use Domain Name 'ipv4only.arpa'", RFC 8880, DOI 10.17487/RFC8880, August 2020, <<https://www.rfc-editor.org/rfc/rfc8880>>.

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: 15 October 2022

M. Boucadair, Ed.
Orange
T. Reddy, Ed.
Akamai
D. Wing
Citrix
N. Cook
Open-Xchange
T. Jensen
Microsoft
13 April 2022

DHCP and Router Advertisement Options for the Discovery of Network-
designated Resolvers (DNR)
draft-ietf-add-dnr-07

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 15 October 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Overview	3
3.1. Configuration Data for Encrypted DNS	4
3.2. Handling Configuration Data Conflicts	6
3.3. Connection Establishment	6
3.4. Multihoming Considerations	7
4. DHCPv6 Encrypted DNS Option	7
4.1. Option Format	7
4.2. DHCPv6 Client Behavior	9
5. DHCPv4 Encrypted DNS Option	9
5.1. Option Format	9
5.2. DHCPv4 Client Behavior	11
6. IPv6 RA Encrypted DNS Option	12
6.1. Option Format	12
6.2. IPv6 Host Behavior	14
7. Security Considerations	14
7.1. Spoofing Attacks	14
7.2. Deletion Attacks	15
7.3. Passive Attacks	15
7.4. Wireless Security - Authentication Attacks	15
8. IANA Considerations	16
8.1. DHCPv6 Option	16
8.2. DHCPv4 Option	16
8.3. Neighbor Discovery Option	17
9. Acknowledgements	17
10. Contributing Authors	17
11. References	18
11.1. Normative References	18
11.2. Informative References	18
Authors' Addresses	21

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-dnsquic] in local networks.

In particular, the document specifies how a local encrypted DNS server can be discovered by connected hosts by means of DHCPv4 [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. This procedure is called Discovery of Network-designated Resolvers (DNR).

The options defined in this document can be deployed in a variety of deployments (e.g., local networks with Customer Premises Equipment (CPEs) that may or may not be managed by an Internet Service Provider (ISP), local networks with or without DNS forwarders). It is out of the scope of this document to provide an inventory of such deployments.

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.

DNR: refers to the Discovery of Network-designated Resolvers procedure.

Encrypted DNS: refers to a scheme where DNS exchanges are transported over an encrypted channel. Examples of encrypted DNS are DoT, DoH, or DoQ.

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

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. These encrypted DNS servers may be hosted on the same or distinct IP addresses. Such a decision is deployment specific.

In order to optimize the size of discovery messages when all DNS 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 protocol (e.g., DoH and DoT) to probe that list of IP addresses. To avoid such a probing, the options defined in Sections 4, 5, and 6 associate an IP address with an encrypted DNS protocol. No probing is required in such a design.

A list of IP addresses to reach an encrypted DNS server may be returned in an Encrypted DNS option to accommodate current deployments relying upon primary and backup servers. Whether one or more IP addresses are returned in an Encrypted DNS option is deployment specific. For example, a router embedding a recursive server or a forwarder has to include one single IP address pointing to one of its LAN-facing interfaces. This IP 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. At least the following service parameters are RECOMMENDED to be supported by a DNR implementation:

alpn: Used to indicate the set of supported protocols (Section 7.1 of [I-D.ietf-dnsop-svcb-https]).

port: Used to indicate the target port number for the encrypted DNS connection (Section 7.2 of [I-D.ietf-dnsop-svcb-https]).

ech: Used to enable Encrypted ClientHello (ECH) (Section 7.3 of [I-D.ietf-dnsop-svcb-https]).

dohpath: Used to supply a relative DoH URI Template (Section 5.1 of [I-D.ietf-add-svcb-dns]).

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.

The DHCP options defined in Sections 4 and 5 follow the option ordering guidelines in Section 17 of [RFC7227]. Likewise, the RA option (Section 6) adheres to the recommendations in Section 9 of [RFC4861].

ServiceMode (Section 2.4.3 of [I-D.ietf-dnsop-svcb-https]) SHOULD be used because the Encrypted DNS options are self-contained and do not require any additional DNS queries. The reader may refer to [RFC7969] for an overview of advanced capabilities that are supported by DHCP servers to populate configuration data (e.g., issue DNS queries).

In contexts where putting additional complexity on requesting hosts is acceptable, returning an ADN only can be considered. The supplied ADN will be processed by a host following the procedure in Section 5 of [I-D.ietf-add-ddr]. Note that this mode may be subject to active attacks, which can be mitigated by DNSSEC.

Other mechanisms may be considered in other contexts (e.g., secure discovery) for the provisioning of encrypted DNS servers. It is RECOMMENDED that at least the following DNR information is made available to a requesting host:

- * A service priority whenever the discovery mechanism does not rely on implicit ordering if multiple instances of the encrypted DNS are used.
- * An authentication domain name.
- * A list of IP addresses to locate the encrypted DNS server.
- * A set of service parameters.

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 Discovery of Designated Resolvers (DDR) [I-D.ietf-add-ddr] since DDR uses Do53 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 7.1.

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.

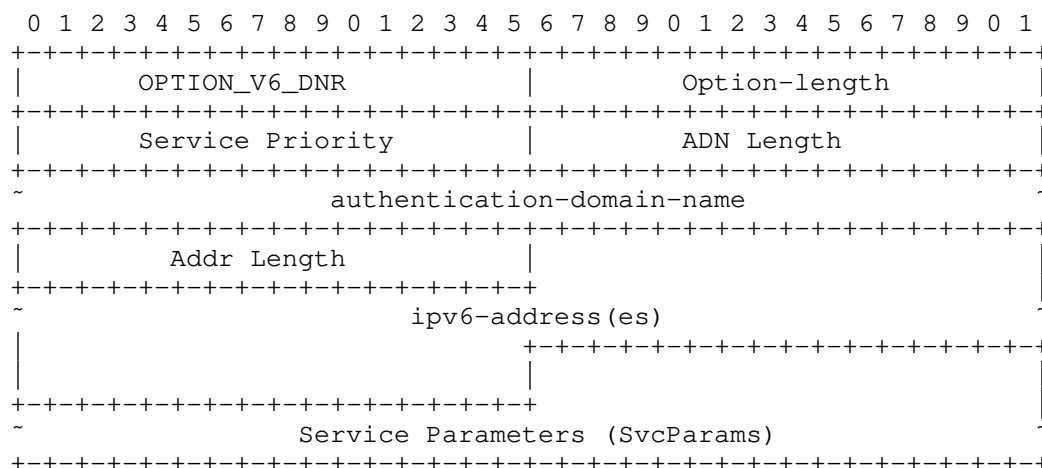


Figure 1: DHCPv6 Encrypted DNS Option

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

Option-code: OPTION_V6_DNR (TBA1, see Section 8.1)

Option-length: Length of the enclosed data in octets. The option length is ('ADN Length' + 4) when only an ADN is included in the option.

Service Priority: The priority of this OPTION_V6_DNR instance compared to other instances. This field is encoded following the rules specified in Section 2.4.1 of [I-D.ietf-dnsop-svcb-https].

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

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, 443 for DoH, and 853 for DoQ.

The length of this field is ('Option-length' - 6 - 'ADN Length' - 'Addr Length').

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 DHCPv6 client MUST be prepared to receive multiple instances of the OPTION_V6_DNR option; each option is to be treated as a separate encrypted DNS server. These instances SHOULD be processed following their service priority (i.e., smaller service priority indicates a higher preference).

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 8.2).

Length: Indicates the length of the enclosed data in octets. The option length is ('ADN Length' + 3) when only an ADN is included in the option.

Service Priority: The priority of this OPTION_V4_DNR instance compared to other instances. This field is encoded following the rules specified in Section 2.4.1 of [I-D.ietf-dnsop-svcb-https].

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 for ServiceMode.

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 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. The service parameters MUST NOT include "ipv4hint" or "ipv6hint" SvcParams as they are superseded by the included IP addresses.

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

The length of this field is ('Option-length' - 4 - '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.

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 instances of the OPTION_V4_DNR option; each option is to be treated as a separate encrypted DNS server. These instances SHOULD be processed following their service priority (i.e., smaller service priority indicates a higher preference).

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.

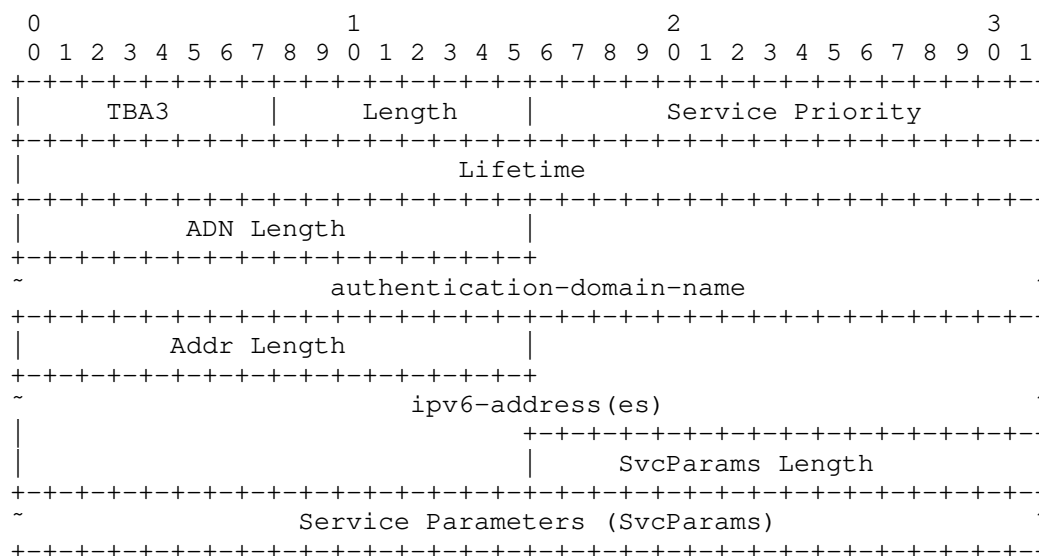


Figure 7: RA Encrypted DNS Option

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 8.3).

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

Service Priority: The priority of this Encrypted DNS option instance compared to other instances. This field is encoded following the rules specified in Section 2.4.1 of [I-D.ietf-dnsop-svcb-https].

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 $3 * \text{MaxRtrAdvInterval}$, 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 for ServiceMode.

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. The service parameters MUST NOT include "ipv4hint" or "ipv6hint" SvcParams as they are superseded by the included IP addresses.

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] with the formatting requirements in Section 6.1 substituted for the length validation.

The host MUST be prepared to receive multiple Encrypted DNS options in RAs. These instances SHOULD be processed following their service priority (i.e., smaller service priority indicates a higher preference).

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

7. Security Considerations

7.1. 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.

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 router (e.g., a border router, a CPE) discards DHCP response messages received from any local endpoint.
- * RA-Guard described in [RFC7113], the router discards RAs messages received from any local endpoint.
- * Source Address Validation Improvement (SAVI) solution for DHCP described in [RFC7513], the router filters packets with forged source IP addresses.

The above mechanisms would ensure that the endpoint receives the correct configuration information of the encrypted DNS servers selected by the DHCP server (or RA sender), but cannot provide any information about the DHCP server or the entity hosting the DHCP server (or RA sender) .

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.

7.2. 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.

7.3. 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.

7.4. 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]).

8. IANA Considerations

8.1. 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]

Table 1

8.2. 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]

8.3. 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]

Table 2

9. Acknowledgements

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

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

Thanks to Mark Nottingham for the feedback on HTTP redirection that was discussed in previous versions of this specification.

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.

10. 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

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-08, 12 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-dnsop-svcb-https-08.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>>.
- [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>>.
- [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>>.

11.2. Informative References

- [BOOTP] "BOOTP Vendor Extensions and DHCP Options",
<[https://www.iana.org/assignments/bootp-dhcp-parameters/
bootp-dhcp-parameters.xhtml#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](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](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_\(wireless_networks\)](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", Work in
Progress, Internet-Draft, draft-ietf-add-ddr-06, 4 April
2022, <[https://www.ietf.org/archive/id/draft-ietf-add-ddr-
06.txt](https://www.ietf.org/archive/id/draft-ietf-add-ddr-06.txt)>.
- [I-D.ietf-add-svcb-dns]
Schwartz, B., "Service Binding Mapping for DNS Servers",
Work in Progress, Internet-Draft, draft-ietf-add-svcb-dns-
02, 1 February 2022, <[https://www.ietf.org/archive/id/
draft-ietf-add-svcb-dns-02.txt](https://www.ietf.org/archive/id/draft-ietf-add-svcb-dns-02.txt)>.
- [I-D.ietf-dprive-dnssoquic]
Huitema, C., Dickinson, S., and A. Mankin, "DNS over
Dedicated QUIC Connections", Work in Progress, Internet-
Draft, draft-ietf-dprive-dnssoquic-11, 21 March 2022,
<[https://www.ietf.org/archive/id/draft-ietf-dprive-
dnssoquic-11.txt](https://www.ietf.org/archive/id/draft-ietf-dprive-dnssoquic-11.txt)>.
- [I-D.pusateri-dhc-dns-driu]
Pusateri, T. and W. Toorop, "DHCPv6 Options for private
DNS Discovery", Work in Progress, Internet-Draft, draft-
pusateri-dhc-dns-driu-00, 2 July 2018,
<[https://www.ietf.org/archive/id/draft-pusateri-dhc-dns-
driu-00.txt](https://www.ietf.org/archive/id/draft-pusateri-dhc-dns-driu-00.txt)>.

- [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>>.

- [RFC7227] Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", BCP 187, RFC 7227, DOI 10.17487/RFC7227, May 2014, <<https://www.rfc-editor.org/info/rfc7227>>.
- [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>>.
- [RFC7969] Lemon, T. and T. Mrugalski, "Customizing DHCP Configuration on the Basis of Network Topology", RFC 7969, DOI 10.17487/RFC7969, October 2016, <<https://www.rfc-editor.org/info/rfc7969>>.
- [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>>.

Authors' Addresses

Mohamed Boucadair (editor)
Orange
35000 Rennes
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy (editor)
Akamai
Embassy Golf Link Business Park
Bangalore 560071
Karnataka
India
Email: kondtir@gmail.com

Dan Wing
Citrix Systems, Inc.
United States of America
Email: dwing-ietf@fuggles.com

Neil Cook
Open-Xchange
United Kingdom
Email: neil.cook@noware.co.uk

Tommy Jensen
Microsoft
United States of America
Email: tojens@microsoft.com

ADD
Internet-Draft
Intended status: Standards Track
Expires: 15 October 2022

T. Reddy
Akamai
D. Wing
Citrix
K. Smith
Vodafone
B. Schwartz
Google
13 April 2022

Establishing Local DNS Authority in Split-Horizon Environments
draft-reddy-add-enterprise-split-dns-10

Abstract

When split-horizon DNS is deployed by a network, certain domains can be resolved authoritatively by the network-provided DNS resolver. DNS clients that don't always use this resolver might wish to do so for these domains. This specification describes how clients can confirm the local resolver's authority over these domains.

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 15 October 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights

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

Table of Contents

1. Introduction	2
2. Terminology	4
2.1. Validated Split-Horizon	4
3. Scope	4
4. Local Domain Hint Mechanisms	4
4.1. DHCP Options	4
4.2. Host Configuration	5
4.3. Provisioning Domains dnsZones	6
4.4. Split DNS Configuration for IKEv2	6
5. Establishing Local DNS Authority	6
6. Validating Authority over Local Domain Hints	6
6.1. Using Pre-configured Public Resolver	7
6.2. Using DNSSEC	7
7. Examples of Split-Horizon DNS Configuration	7
7.1. Split-Horizon Entire Zone	8
7.1.1. Verification using Public Resolver	9
7.1.2. Verification using DNSSEC	10
7.2. Split-Horizon Only Subdomain of Zone	12
8. Validation with IKEv2	12
9. Security Considerations	12
10. IANA Considerations	12
11. Acknowledgements	12
12. References	12
12.1. Normative References	13
12.2. Informative References	13
Authors' Addresses	15

1. Introduction

To resolve a DNS query, there are three essential behaviors that an implementation can apply: (1) answer from a local database, (2) query the relevant authorities and their parents, or (3) ask a server to query those authorities and return the final answer. Implementations that use these behaviors are called "authoritative nameservers", "full resolvers", and "forwarders" (or "stub resolvers"). However, an implementation can also implement a mixture of these behaviors, depending on a local policy, for each query. We term such an implementation a "hybrid resolver".

Most DNS resolvers are hybrids of some kind. For example, stub resolvers frequently support a local "hosts file" that preempts query forwarding, and most DNS forwarders and full resolvers can also serve responses from a local zone file. Other standardized hybrid resolution behaviors include Local Root [RFC8806], mDNS [RFC6762], and NXDOMAIN synthesis for .onion [RFC7686].

In many network environments, the network offers clients a DNS server (e.g. DHCP OFFER, IPv6 Router Advertisement). Although this server is formally specified as a recursive resolver (e.g. Section 5.1 of [RFC6106]), some networks provide a hybrid resolver instead. If this resolver acts as an authoritative server for some names, we say that the network has "split-horizon DNS", because those names resolve in this way only from inside the network.

Network clients that use pure stub resolution, sending all queries to the network-provided resolver, will always receive the split-horizon results. Conversely, clients that send all queries to a different resolver or implement pure full resolution locally will never receive them. Clients with either pure resolution behavior are out of scope for this specification. Instead, this specification enables hybrid clients to access split-horizon results from a network-provided hybrid resolver, while using a different resolution method for some or all other names.

There are several existing mechanisms for a network to provide clients with "local domain hints", listing domain names that have special treatment in this network (Section 4). However, none of the local domain hint mechanisms enable clients to determine whether this special treatment is authorized by the domain owner. Instead, these specifications require clients to make their own determinations about whether to trust and rely on these hints.

This specification describes a protocol between domains, networks, and clients that allows the network to establish its authority over a domain to a client (Section 5). Clients can use this protocol to confirm that a local domain hint was authorized by the domain (Section 6), which might influence its processing of that hint.

This specification relies on securely identified local DNS servers and globally valid NS records. Use of this specification is therefore limited to servers that support authenticated encryption and split-horizon DNS names that are properly rooted in the global DNS.

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 term "Global DNS" is defined in [RFC8499].

'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 term 'Validated Split-Horizon' is also defined.

2.1. Validated Split-Horizon

A split horizon configuration for some name is considered "validated" if the network client has confirmed that a parent of that name has authorized the local network to serve its own responses for that name. Such authorization generally extends to the entire subtree of names below the authorization point.

3. Scope

The protocol in this document allows the domain owner to create a split-horizon DNS. Other entities which do not own the domain are detected by the client. Thus, DNS filtering is not enabled by this protocol.

4. Local Domain Hint Mechanisms

There are various mechanisms by which a network client might learn "local domain hints", which indicate a special treatment for particular domain names upon joining a network. This section provides a review of some common and standardized mechanisms for receiving domain hints.

4.1. DHCP Options

There are several DHCP options that convey local domain hints of different kinds. The most directly relevant is "RDNSS Selection" [RFC6731], which provides "a list of domains ... about which the RDNSS has special knowledge", along with a "High", "Medium", or "Low" preference for each name. The specification notes the difficulty of relying on these hints without validation:

Trustworthiness of an interface and configuration information received over the interface is implementation and/or node deployment dependent, and the details of determining that trust are beyond the scope of this specification.

Other local domain hints in DHCP include the "Domain Name" [RFC2132], "Access Network Domain Name" [RFC5986], "Client FQDN" [RFC4702][RFC4704], and "Name Service Search" [RFC2937] options. This specification may help clients to interpret these hints. For example, a rogue DHCP server could use the "Client FQDN" option to assign a client the name "www.example.com" in order to prevent the client from reaching the true "www.example.com". A client could use this specification to check the network's authority over this name, and adjust its behavior to avoid this attack if authority is not established.

The Domain Search option [RFC3397] [RFC3646], which offers clients a way to expand short names into Fully Qualified Domain Names, is not a "local domain hint" by this definition, because it does not modify the processing of any specific domain. (The specification notes that this option can be a "fruitful avenue of attack for a rogue DHCP server", and provides a number of cautions against accepting it unconditionally.)

4.2. Host Configuration

A host can be configured with DNS information when it joins a network, including when it brings up VPN (which is also considered joining a(n additional) network, detailed in Section 8). Existing implementations determine the host has joined a certain network via SSID, IP subnet assigned, DNS server IP address or name, and other similar mechanisms. For example, one existing implementation determines the host has joined an internal network because the DHCP-assigned IP address belongs to the company's IP address (as assigned by the regional IP addressing authority) and the DHCP-advertised DNS IP address is one used by IT at that network. Other mechanisms exist in other products but are not interesting to this specification; rather what is interesting is this step to determine "we have joined the internal corporate network" occurred and the DNS server is configured as authoritative for certain DNS zones (e.g., *.example.com).

Because a rogue network can simulate all or most of the above characteristics this specification details how to validate these claims in Section 6.

4.3. Provisioning Domains dnsZones

Provisioning Domains (PvDs) are defined in [RFC7556] as sets of network configuration information that clients can use to access networks, including rules for DNS resolution and proxy configuration. The PvD Key "dnsZones" is defined in [RFC8801] as a list of "DNS zones searchable and accessible" in this provisioning domain. Attempting to resolve these names via another resolver might fail or return results that are not correct for this network.

4.4. Split DNS Configuration for IKEv2

In IKEv2 VPNs, the INTERNAL_DNS_DOMAIN configuration attribute can be used to indicate that a domain is "internal" to the VPN [RFC8598]. To prevent abuse, the specification notes various possible restrictions on the use of this attribute:

"If a client is configured by local policy to only accept a limited set of INTERNAL_DNS_DOMAIN values, the client MUST ignore any other INTERNAL_DNS_DOMAIN values."

"IKE clients MAY want to require whitelisted domains for Top-Level Domains (TLDs) and Second-Level Domains (SLDs) to further prevent malicious DNS redirections for well-known domains."

Within these guidelines, a client could adopt a local policy of accepting INTERNAL_DNS_DOMAIN values only when it can validate the local DNS server's authority over those names as described in this specification.

5. Establishing Local DNS Authority

To establish its authority over some DNS zone, a participating network MUST offer one or more encrypted resolvers via DNR [I-D.ietf-add-dnr] or an equivalent mechanism (see Section 8). At least one of these resolvers' Authentication Domain Names (ADNs) MUST appear in an NS record for that zone. This arrangement establishes this resolver's authority over the zone.

6. Validating Authority over Local Domain Hints

To validate the network's authority over a domain name, participating clients MUST resolve the NS record for that name. If the resolution result is NODATA, the client MUST remove the last label and repeat the query until a response other than NODATA is received.

Once the NS record has been resolved, the client MUST check if each local encrypted resolver's Authentication Domain Name appears in the NS record. The client SHALL regard each such resolver as authoritative for the zone of this NS record.

Each validation of authority applies only to the specific resolvers whose names appear in the NS RSet. If a network offers multiple encrypted resolvers, each DNS entry may be authorized for a distinct subset of the network-provided resolvers.

A zone is termed a "Validated Split-Horizon zone" after successful validation using a "tamperproof" NS resolution method, i.e. a method that is not subject to interference by the local network operator. Two possible tamperproof resolution methods are presented below.

6.1. Using Pre-configured Public Resolver

The client sends the NS query to a pre-configured resolver that is external to the network, over a secure transport. Clients SHOULD apply whatever acceptance rules they would otherwise apply when using this resolver (e.g. checking the AD bit, validating RRSIGs).

6.2. Using DNSSEC

The client resolves the NS record using any resolution method of its choice (e.g. querying one of the network-provided resolvers, performing iterative resolution locally), and performs full DNSSEC validation locally [RFC6698]. The result is processed based on its DNSSEC validation state (Section 4.3 of [RFC4035]):

Secure: the response is used for validation.

Bogus or Indeterminate: the response is rejected and validation is considered to have failed.

Insecure: the client SHOULD retry the validation process using a different method, such as the one in Section 6.1, to ensure compatibility with unsigned names.

7. Examples of Split-Horizon DNS Configuration

Two examples are shown below. The first example showing an company with an internal-only DNS server resolving the entire zone for that company (e.g., *.example.com) the second example resolving only a subdomain of the company's zone (e.g., *.internal.example.com).

7.1. Split-Horizon Entire Zone

Consider an organization that operates "example.com", and runs a different version of its global domain on its internal network. Today, on the Internet it publishes two NS records, "ns1.example.com" and "ns2.example.com".

The host and network first need mutual support one of the mechanisms described in learning (Section 4). Shown in Figure 1 is learning using DNR and PvD.

Validation is then performed using either Public DNS (Section 7.1.1) or DNSSEC (Section 7.1.2).

steps 1-2: The client determines the network's DNS server (ns1.example.com) and Provisioning Domain (pvd.example.com) using DNR [I-D.ietf-add-dnr] and PvD [RFC8801], using one of DNR Router Solicitation, DHCPv4, or DHCPv6.

step 3-5: The client connects to the DNR-learned DNS server (ns1.example.com), validates its certificate, and queries for pvd.example.com.

steps 6-7: The client connects to the PvD server, validates its certificate, and retrieves the provisioning domain JSON information indicated by the associated PvD. The PvD contains:

```
{
  "identifier": "pvd.example.com",
  "expires": "2020-05-23T06:00:00Z",
  "prefixes": ["2001:db8:1::/48", "2001:db8:4::/48"],
  "dnsZones": ["example.com"]
}
```

The JSON keys "identifier", "expires", and "prefixes" are defined in [RFC8801].

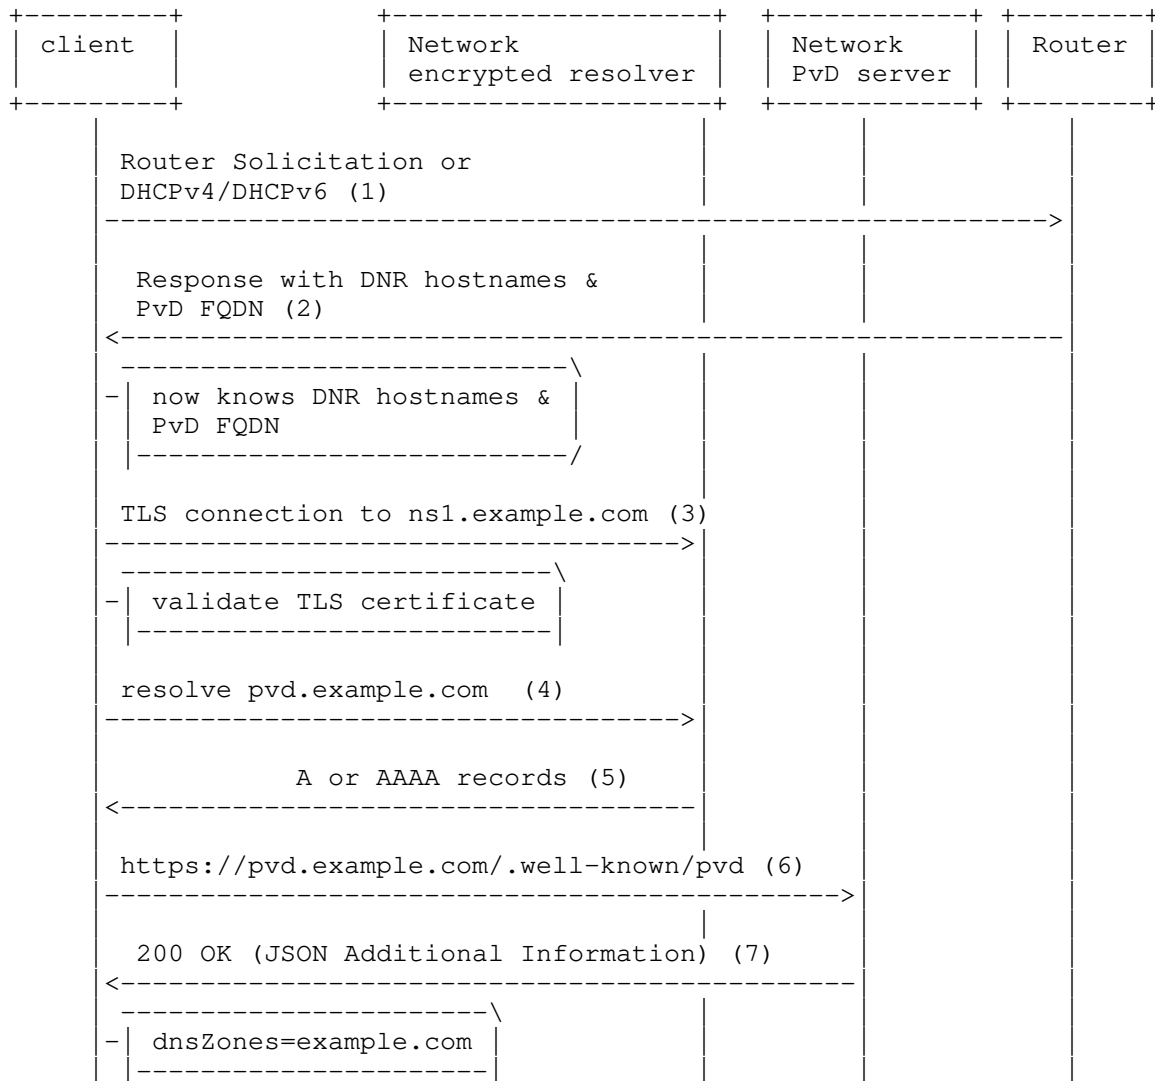


Figure 1: Learning Local Claims of DNS Authority

7.1.1. Verification using Public Resolver

The figure below shows the steps performed to verify the local claims of DNS authority using a public resolver.

Steps 1-2: The client uses an encrypted DNS connection to a public

resolver (e.g., 1.1.1.1) to issue NS queries for the domains in dnsZones. The NS lookup for "example.com" will return "ns1.example.com" and "ns2.example.com".

Step 3: As the network-provided nameservers are the same as the names retrieved from the public resolver and the network-designated resolver's certificate includes at least one of the names retrieved from the public resolver, the client has finished validation that the nameservers signaled in [I-D.ietf-add-dnr] and [RFC8801] are owned and managed by the same entity that published the NS records on the Internet. The endpoint will then use that information from [I-D.ietf-add-dnr] and [RFC8801] to resolve names within dnsZones.

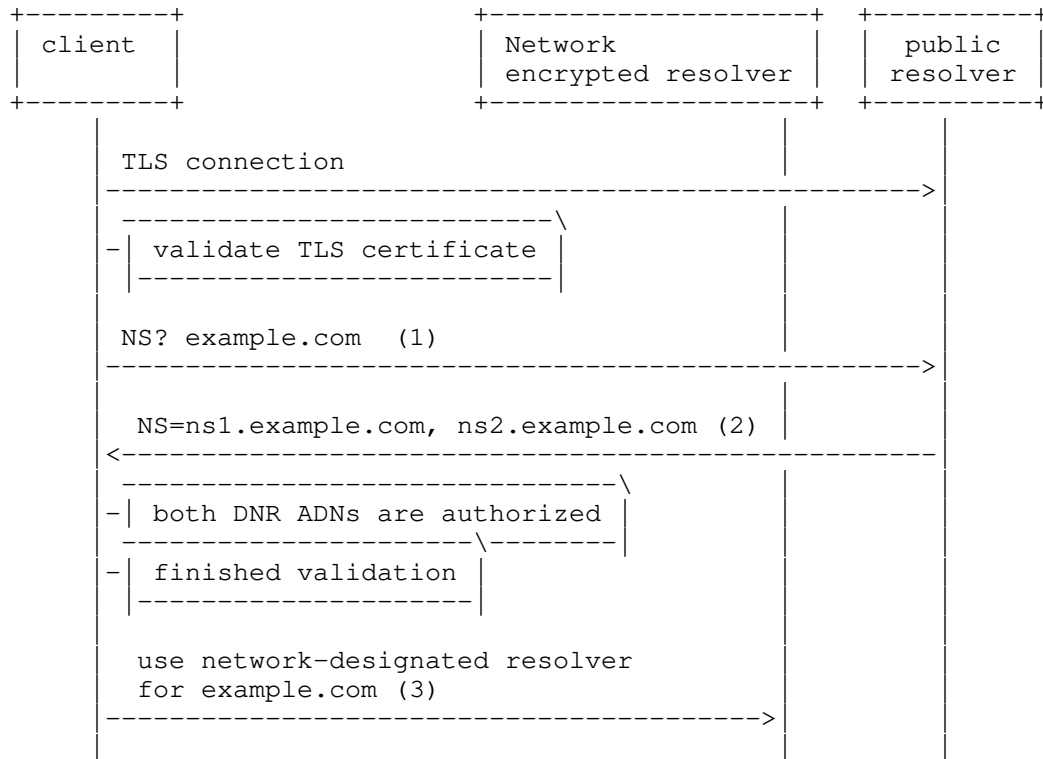


Figure 2: Verifying Claims using Public Resolver

7.1.2. Verification using DNSSEC

The figure below shows the steps performed to verify the local claims of DNS authority using DNSSEC.

Steps 1-2: The DNSSEC-validating client queries the network encrypted resolver to issue NS queries for the domains in dnsZones. The NS lookup for "example.com" will return a signed response containing "ns1.example.com" and "ns2.example.com". The client then performs full DNSSEC validation locally.

Step 3: As the DNSSEC validation is successful and the network-provided nameservers are the same as the names in the DNSSEC response, and the network-designated resolver's certificate includes at least one of the names returned in the DNSSEC response, the client has finished validation that the nameservers signaled in [I-D.ietf-add-dnr] and [RFC8801] are owned and managed by the same entity that published the NS records on the Internet. The endpoint will then use that information from [I-D.ietf-add-dnr] and [RFC8801] to resolve names within dnsZones.

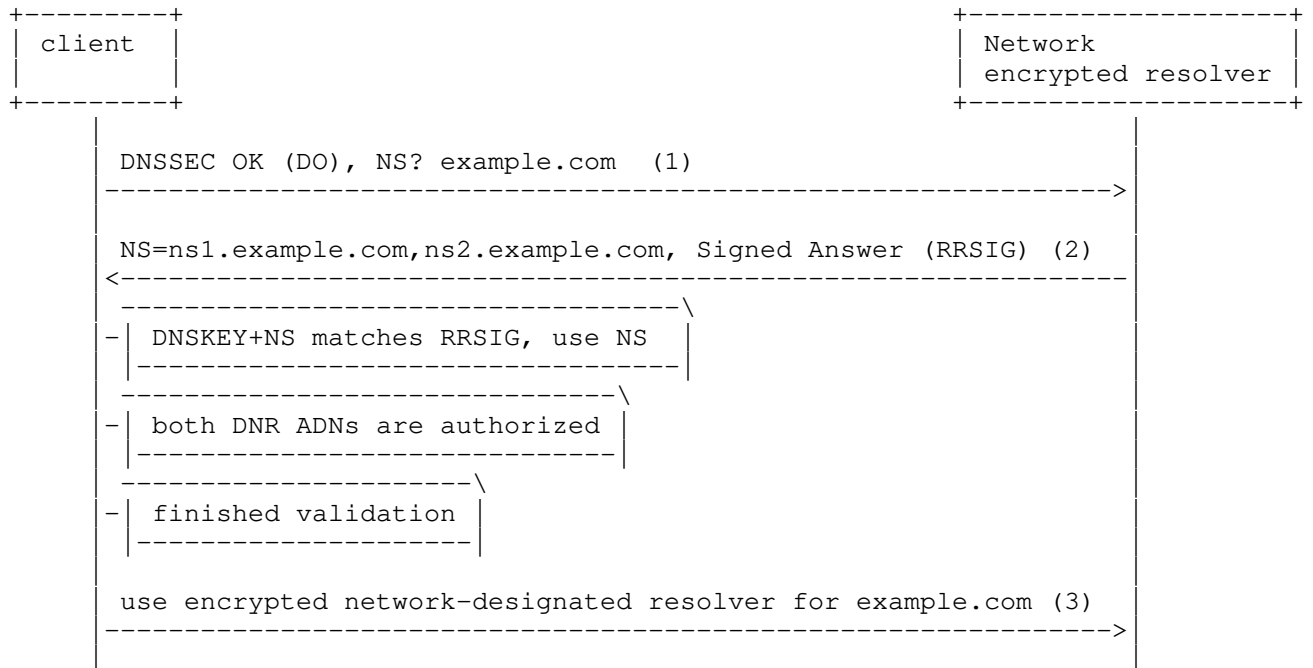


Figure 3: Verifying Claims using DNSSEC

7.2. Split-Horizon Only Subdomain of Zone

A subdomain can also be used for all internal DNS names (e.g., the zone `internal.example.com` exists only on the internal DNS server). For successful validation described in this document the internal DNS server will need a certificate signed by a CA trusted by the client.

For such a name `internal.example.com` the message flow is similar to Section 7.1 the difference is that queries for hosts not within the subdomain (`www.example.com`) are sent to the public resolver rather than resolver for `internal.example.com`.

8. Validation with IKEv2

When the VPN tunnel is IPsec, the encrypted DNS resolver hosted by the VPN service provider can be securely discovered by the endpoint using the `ENCDNS_IP*_*` IKEv2 Configuration Payload Attribute Types defined in [I-D.ietf-ipsecme-add-ike].

Other VPN tunnel types have similar configuration capabilities, not detailed here.

9. Security Considerations

This specification does not alter DNSSEC validation behaviour. To ensure compatibility with validating clients, network operators **MUST** ensure that names under the split-horizon are correctly signed or place them in an unsigned zone.

If an internal zone name (e.g., `internal.example.com`) is used with in conjunction with this specification and a public certificate is obtained for validation, that internal zone name will exist in Certificate Transparency [RFC9162] logs. It should be noted, however, that this specification does not leak individual host names (e.g., `www.internal.example.com`) into the Certificate Transparency logs or to public DNS resolvers.

10. IANA Considerations

This document has no IANA actions.

11. Acknowledgements

Thanks to Mohamed Boucadair, Jim Reid, Tommy Pauly, Paul Vixie, Paul Wouters and Vinny Parla for the discussion and comments.

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>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/info/rfc4035>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <<https://www.rfc-editor.org/info/rfc6698>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [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>>.
- [RFC8801] Pfister, P., Vyncke, É., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", RFC 8801, DOI 10.17487/RFC8801, July 2020, <<https://www.rfc-editor.org/info/rfc8801>>.

12.2. Informative References

- [I-D.ietf-add-dnr]
Boucadair, M., Reddy, T., Wing, D., Cook, N., and T. Jensen, "DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)", Work in Progress, Internet-Draft, draft-ietf-add-dnr-06, 22 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-add-dnr-06.txt>>.
- [I-D.ietf-ipsecme-add-ike]
Boucadair, M., Reddy, T., Wing, D., and V. Smyslov, "Internet Key Exchange Protocol Version 2 (IKEv2) Configuration for Encrypted DNS", Work in Progress, Internet-Draft, draft-ietf-ipsecme-add-ike-01, 22 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-ipsecme-add-ike-01.txt>>.

- [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>>.
- [RFC2937] Smith, C., "The Name Service Search Option for DHCP", RFC 2937, DOI 10.17487/RFC2937, September 2000, <<https://www.rfc-editor.org/info/rfc2937>>.
- [RFC3397] Aboba, B. and S. Cheshire, "Dynamic Host Configuration Protocol (DHCP) Domain Search Option", RFC 3397, DOI 10.17487/RFC3397, November 2002, <<https://www.rfc-editor.org/info/rfc3397>>.
- [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>>.
- [RFC4702] Stapp, M., Volz, B., and Y. Rekhter, "The Dynamic Host Configuration Protocol (DHCP) Client Fully Qualified Domain Name (FQDN) Option", RFC 4702, DOI 10.17487/RFC4702, October 2006, <<https://www.rfc-editor.org/info/rfc4702>>.
- [RFC4704] Volz, B., "The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN) Option", RFC 4704, DOI 10.17487/RFC4704, October 2006, <<https://www.rfc-editor.org/info/rfc4704>>.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, DOI 10.17487/RFC5986, September 2010, <<https://www.rfc-editor.org/info/rfc5986>>.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, DOI 10.17487/RFC6106, November 2010, <<https://www.rfc-editor.org/info/rfc6106>>.
- [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>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <<https://www.rfc-editor.org/info/rfc7556>>.

- [RFC7686] Appelbaum, J. and A. Muffett, "The ".onion" Special-Use Domain Name", RFC 7686, DOI 10.17487/RFC7686, October 2015, <<https://www.rfc-editor.org/info/rfc7686>>.
- [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>>.
- [RFC8806] Kumari, W. and P. Hoffman, "Running a Root Server Local to a Resolver", RFC 8806, DOI 10.17487/RFC8806, June 2020, <<https://www.rfc-editor.org/info/rfc8806>>.
- [RFC9162] Laurie, B., Messeri, E., and R. Stradling, "Certificate Transparency Version 2.0", RFC 9162, DOI 10.17487/RFC9162, December 2021, <<https://www.rfc-editor.org/info/rfc9162>>.

Authors' Addresses

Tirumaleswar Reddy
Akamai
Embassy Golf Link Business Park
Bangalore 560071
Karnataka
India
Email: kondtir@gmail.com

Dan Wing
Citrix Systems, Inc.
4988 Great America Pkwy
Santa Clara, CA 95054
United States of America
Email: danwing@gmail.com

Kevin Smith
Vodafone Group
One Kingdom Street
London
United Kingdom
Email: kevin.smith@vodafone.com

Benjamin Schwartz
Google LLC
Email: bemasc@google.com

dprive
Internet-Draft
Intended status: Informational
Expires: 25 April 2022

B. Schwartz
Google LLC
C. Box
BT
22 October 2021

Discovery of Designated Resolvers in the Presence of Legacy Forwarders
draft-schwartz-add-ddr-forwarders-01

Abstract

This draft describes how the Discovery of Designated Resolvers (DDR) standard interacts with legacy DNS forwarders, including potential incompatibilities and relevant mitigations.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the mailing list (add@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/add/>.

Source for this draft and an issue tracker can be found at <https://github.com/bemasc/ddr-forwarders>.

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 25 April 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect 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.

Table of Contents

1. Conventions and Definitions	2
2. Introduction	3
2.1. Background	3
2.2. Scope	3
3. Relaxed Validation client policy	4
4. Naturally compatible behaviors	4
4.1. Compatible behaviors in the local network	4
4.1.1. Malware and threat domain filtering	4
4.1.2. Service category restrictions	4
4.1.3. Time of use restrictions	5
4.2. Upstream resolver services	5
5. Privacy Considerations	5
6. Security Considerations	5
6.1. Transient attackers	6
6.1.1. Solution: DNR	6
6.1.2. Mitigation: Frequent refresh	6
6.1.3. Mitigation: Resolver reputation	6
6.2. Forensic logging	6
6.2.1. Network-layer logging	6
6.2.2. DNS-layer logging	7
7. Compatibility Considerations	7
7.1. Split-horizon namespaces	7
7.1.1. Mitigation: NXDOMAIN Fallback	7
7.2. Interposable domains	8
7.2.1. Mitigation: Exemption list	8
7.3. Caching	8
7.3.1. Mitigation: Stub caches	8
7.4. General mitigation: User controls	9
8. Informative References	9
Acknowledgments	11
Authors' Addresses	11

1. Conventions and Definitions

Private IP Address - Any IP address reserved for loopback [RFC1122], link-local [RFC3927], private [RFC1918], local [RFC4193], or Carrier-Grade NAT [RFC6598] use.

Legacy DNS Forwarder - An apparent DNS resolver, known to the client only by a private IP address, that forwards the client's queries to an upstream resolver, and has not been updated with any knowledge of DDR.

Cross-Forwarder Upgrade - Establishment of a direct, encrypted connection between the client and the upstream resolver.

2. Introduction

2.1. Background

The Discovery of Designated Resolvers specification [DDR] describes a mechanism for clients to learn about the encrypted protocols supported by a DNS server. It also describes a conservative client validation policy that has strong security properties and is unlikely to create compatibility problems.

On the topic of client validation of encrypted DNS transports, the DDR specification says:

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

As TLS certificates cannot cover private IP addresses, this prevents clients that are behind a legacy DNS forwarder from connecting directly to the upstream resolver ("cross-forwarder upgrade").

Recent estimates suggest that a large fraction, perhaps a majority, of residential internet users in the United States and Europe rely on local DNS forwarders that are not compatible with DDR.

2.2. Scope

This informational document describes the interaction between DDR and legacy DNS forwarders. It discusses possible client policies, problems that might arise, and relevant mitigations.

DNS forwarders and resolvers that are implemented with awareness of DDR are out of scope, as they are not affected by this discussion (although see Security Considerations, Section 6).

IPv6-only networks whose default DNS server has a Global Unicast Address are out of scope, even if this server is actually a simple forwarder. If the DNS server does not use a private IP address, it is not a "legacy DNS forwarder" under this draft's definition.

3. Relaxed Validation client policy

We define a "relaxed validation" client policy as a client behavior that removes the certificate validation requirement when the Unencrypted Resolver is identified by a private IP address, regardless of the Designated Resolver's IP address. Instead, under this condition, the client connects using the Opportunistic Privacy Profile of encrypted DNS ([RFC7858], Section 4.1).

The Opportunistic Privacy Profile is a broad category, including clients that "might or might not validate" the TLS certificate chain even though there is no authentication identity for the server. This kind of validation can be valuable when combined with a reputation system or a user approval step (see Section 6.1.3 and Section 7.4).

This client policy is otherwise identical to the one described in Section 4 of [DDR].

4. Naturally compatible behaviors

The following system behaviors are naturally compatible with relaxed validation.

4.1. Compatible behaviors in the local network

4.1.1. Malware and threat domain filtering

Certain DNS forwarders block access to domains associated with malware and other threats. Such threats rely on frequently changing domains, so these forwarders necessarily maintain an actively curated list of domains to block. To ensure that this service is not lost due to a cross-forwarder upgrade, the maintainers can simply add "resolver.arpa" to the list.

This pattern has been deployed by Mozilla, with the domain "use-application-dns.net" [MOZILLA-CANARY].

4.1.2. Service category restrictions

Certain DNS forwarders may block access to domains based on the category of service provided by those domains, e.g. domains hosting services that are not appropriate for a work or school environment. As in the previous section, this requires an actively curated list of domains, because the set of domains that offer a given type of service is constantly changing. An actively managed blocking list can easily be revised to include "resolver.arpa".

4.1.3. Time of use restrictions

Certain networks may impose restrictions on the time or duration of use by certain users. This behavior is necessarily implemented below the DNS layer, because DNS-based blocking would be ineffective due to stub resolver caching, so it is not affected by changes in the DNS resolver.

4.2. Upstream resolver services

The forwarder's upstream resolver might provide additional services, such as filtering. These services are generally independent of cross-forwarder upgrade, and hence naturally compatible.

In special cases where the upstream resolver requires cooperation from a legacy forwarder (e.g. for marking certain queries), one solution is for the upstream resolver to choose not to deploy DDR until all cooperating forwarders have been upgraded. Alternatively, each legacy forwarder can block "resolver.arpa" as described above.

5. Privacy Considerations

The conservative validation policy results in no encryption when a legacy DNS forwarder is present. This leaves the user's query activity vulnerable to passive monitoring [RFC7258], either on the local network or between the user and the upstream resolver.

The relaxed validation policy allows the use of encrypted transport in these configurations, reducing exposure to a passive surveillance adversary.

6. Security Considerations

When the client uses the conservative validation policy described in [DDR], and a DDR-enabled resolver is identified by a private IP address, the client can establish a secure DDR connection only in the absence of an active attacker. An on-path attacker can impersonate the resolver and intercept all queries, by preventing the DDR upgrade or advertising their own DDR endpoint.

These basic security properties also apply if the client uses the relaxed validation policy described in Section 3. Nonetheless, there are some subtle but important differences in the security properties of these two policies.

6.1. Transient attackers

With the conservative validation policy, a transient on-path attacker can only intercept queries for the duration of their active presence on the network, because the client will only send queries to the original (private) server IP address.

With the relaxed validation behavior, a transient on-path attacker could implant a long-lived DDR response in the client's cache, directing its queries to an attacker-controlled server on the public internet. This would allow the attack to continue long after the attacker has left the network.

Solving or mitigating this attack is of great importance for the user's security.

6.1.1. Solution: DNR

This attack does not apply if the client and network implement support for Discovery of Network-designated Resolvers, as that mechanism takes precedence over DDR (see Section 3.2 of [DNR]).

6.1.2. Mitigation: Frequent refresh

The client can choose to refresh the DDR record arbitrarily frequently, e.g. by limiting the TTL. For example, by limiting the TTL to 5 minutes, a client could ensure that any attacker can continue to monitor queries for at most 5 minutes after they have left the local network.

6.1.3. Mitigation: Resolver reputation

A relaxed-validation client might choose to accept a potential cross-forwarder upgrade only if the designated encrypted resolver has sufficient reputation, according to some proprietary reputation scheme (e.g. a locally stored list of respectable resolvers). This limits the ability of a DDR forgery attack to cause harm.

Major DoH client implementations already include lists of known resolvers [CHROME-DOH] [MICROSOFT-DOH] [MOZILLA-TRR].

6.2. Forensic logging

6.2.1. Network-layer logging

With the conservative validation policy, a random sample of IP packets is likely sufficient for manual retrospective detection of an active attack.

With the relaxed validation policy, forensic logs must capture a specific packet (the attacker's DDR designation response) to enable retrospective detection.

6.2.1.1. Mitigation: Log all DDR responses

Network-layer forensic logs that are not integrated with the resolver can enable detection of these attacks by logging all DDR responses, or more generally all DNS responses. This makes retrospective attack detection straightforward, as the attacker's DDR response will indicate an unexpected server.

6.2.2. DNS-layer logging

DNS-layer forensic logging conducted by a legacy DNS forwarder would be lost in a cross-forwarder upgrade.

6.2.2.1. Solution: Respond for resolver.arpa

Forwarders that want to observe all queries from relaxed validation clients will have to synthesize their own response for resolver.arpa, either implementing DDR or disabling it.

7. Compatibility Considerations

Using DDR with legacy DNS forwarders also raises several potential concerns related to loss of existing network services.

7.1. Split-horizon namespaces

Some network resolvers contain additional names that are not resolvable in the global DNS. If these local resolvers are also legacy DNS forwarders, a client that performs a cross-forwarder upgrade might lose access to these local names.

7.1.1. Mitigation: NXDOMAIN Fallback

In "NXDOMAIN Fallback", the client repeats a query to the unencrypted resolver if the encrypted resolver returns NXDOMAIN. This allows the resolution of local names, provided they do not collide with globally resolvable names (as required by [RFC2826]).

This is similar to the fallback behavior currently deployed in Mozilla Firefox [FIREFOX-FALLBACK].

NXDOMAIN Fallback results in slight changes to the security and privacy properties of encrypted DNS. Queries for nonexistent names no longer have protection against a local passive adversary, and local names are revealed to the upstream resolver.

NXDOMAIN Fallback is only applicable when a legacy DNS forwarder might be present, i.e. the unencrypted resolver has a private IP address, and the encrypted resolver has a different IP address. In the other DDR configurations, any local names are expected to resolve similarly on both resolvers.

7.2. Interposable domains

An "interposable domain" is a domain whose owner deliberately allows resolvers to forge certain responses. This arrangement is most common for search engines, which often support a configuration where resolvers forge a CNAME record to direct all clients to a child-appropriate instance of the search engine [DUCK-CNAME] [BING-CNAME] [GOOGLE-CNAME].

Future deployments of interposable domains can instruct administrators to enable or disable DDR when adding the forged record, but forged records in legacy DNS forwarders could be lost due to a cross-forwarder upgrade.

7.2.1. Mitigation: Exemption list

There are a small number of pre-existing interposable domains, largely of interest only to web browsers. Clients can maintain a list of relevant interposable domains and resolve them only via the network's resolver.

7.3. Caching

Some legacy DNS forwarders also provide a shared cache for all network users. Cross-forwarder upgrades will bypass this cache, resulting in slower DNS resolution.

7.3.1. Mitigation: Stub caches

Clients can compensate partially for any loss of shared caching by implementing local DNS caches. This mitigation is already widely deployed in browsers and operating systems.

7.4. General mitigation: User controls

For these and other compatibility concerns, a possible mitigation is to provide users or administrators with the ability to control whether DDR is used with legacy forwarders. For example, this control could be provided via a general preference, or via a notification upon discovering a new upstream resolver.

8. Informative References

[BING-CNAME]

"Block adult content with SafeSearch - Map at a network level", n.d., <<https://help.bing.microsoft.com/#apex/bing/en-us/10003/0>>.

[CHROME-DOH]

"DoH providers: criteria, process for Chrome", n.d., <https://docs.google.com/document/d/128i2YTV2C7T6Gr3I-81zlQ-_Lprnsp24qzy_20ZlPsw/edit>.

[DDR]

Pauly, T., Kinnear, E., Wood, C. A., McManus, P., and T. Jensen, "Discovery of Designated Resolvers", Work in Progress, Internet-Draft, draft-ietf-add-ddr-03, 1 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-add-ddr-03>>.

[DNR]

Boucadair, M., Reddy, T., Wing, D., Cook, N., and T. Jensen, "DHCP and Router Advertisement Options for the Discovery of Network-designated Resolvers (DNR)", Work in Progress, Internet-Draft, draft-ietf-add-dnr-02, 17 May 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-add-dnr-02>>.

[DUCK-CNAME]

"Force Safe Search at a Network Level", n.d., <<https://help.duckduckgo.com/duckduckgo-help-pages/features/safe-search/>>.

[FIREFOX-FALLBACK]

"About our rollout of DNS over HTTPS", n.d., <https://support.mozilla.org/en-US/kb/firefox-dns-over-https#w_about-our-rollout-of-dns-over-https>.

[GOOGLE-CNAME]

"Keep SafeSearch turned on for your school, workplace, or home network", n.d., <<https://support.google.com/websearch/answer/186669?hl=en>>.

[MICROSOFT-DOH]

"Determine which DoH servers are on the known server list", n.d., <<https://docs.microsoft.com/en-us/windows-server/networking/dns/doh-client-support#determine-which-doh-servers-are-on-the-known-server-list>>.

[MOZILLA-CANARY]

"Canary domain - use-application-dns.net", n.d., <<https://support.mozilla.org/en-US/kb/canary-domain-use-application-dnsnet>>.

[MOZILLA-TRR]

"Mozilla Policy Requirements for DNS over HTTPS Partners", n.d., <https://wiki.mozilla.org/Security/DOH-resolver-policy#Mozilla_Policy_Requirements_for_DNS_over_HTTPs_Partners>.

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/rfc/rfc1122>>.

[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/rfc/rfc1918>>.

[RFC2826] Internet Architecture Board, "IAB Technical Comment on the Unique DNS Root", RFC 2826, DOI 10.17487/RFC2826, May 2000, <<https://www.rfc-editor.org/rfc/rfc2826>>.

[RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, DOI 10.17487/RFC3927, May 2005, <<https://www.rfc-editor.org/rfc/rfc3927>>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/rfc/rfc4193>>.

[RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, DOI 10.17487/RFC6598, April 2012, <<https://www.rfc-editor.org/rfc/rfc6598>>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/rfc/rfc7258>>.

[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/rfc/rfc7858>>.

Acknowledgments

Thanks to Anthony Lieuallen and Eric Orth for early reviews.

Authors' Addresses

Benjamin Schwartz
Google LLC

Email: bemasc@google.com

Chris Box
BT

Email: chris.box@bt.com