Workgroup: ADD
Internet-Draft:
draft-reddy-add-enterprise-split-dns-06
Published: 15 September 2021
Intended Status: Standards Track
Expires: 19 March 2022
Authors: T. Reddy D. Wing K. Smith
Akamai Citrix Vodafone

Split-Horizon DNS Configuration

Abstract

When split-horizon DNS is deployed by a network, certain domains are only resolvable by querying the network-designated DNS server rather than a public DNS server. DNS clients which use DNS servers not provided by the network need to route those DNS domain queries to the network-designated DNS server. This document informs DNS clients of split-horizon DNS, their DNS domains, and is compatible with encrypted DNS.

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 19 March 2022.

Table of Contents

1. Introduction
2. Terminology
3. Split DNS
4. PvD dnsZones
   4.1. Authority over the Domains
5. An Example
6. Split DNS Configuration for IKEv2
7. Security Considerations
8. IANA Considerations
9. Acknowledgements
10. References
   10.1. Normative References
   10.2. Informative References
Authors' Addresses

1. Introduction

Historically, an endpoint would utilize network-designated DNS servers upon joining a network (e.g., DHCP OFFER, IPv6 Router Advertisement). While it has long been possible to configure endpoints to ignore the network's suggestions and use a (public) DNS server on the Internet, this was seldom used because some networks block UDP/53 (in order to enforce their own DNS policies). Also, there has been an increase in the availability of "public resolvers" [RFC8499] which DNS clients may be pre-configured to use instead of the default network resolver for a variety of reasons (e.g., offer a good reachability, support an encrypted transport, provide a claimed privacy policy, (lack of) filtering). With the advent of DoT and DoH, such network blocking is more difficult, but the endpoint is unable to (properly) resolve split-horizon DNS domains which must query the network-designated DNS server.

This document specifies a mechanism to indicate which DNS zones are used for split-horizon DNS. DNS clients can discover and authenticate DNS servers provided by the network, for example using the techniques proposed in [I-D.ietf-add-dnr] and [I-D.ietf-add-ddr].

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. [RFC8801] defines a mechanism for discovering multiple Explicit PvDs on a single network and their Additional Information by means of an HTTP-over-TLS query using a URI derived

from the PvD ID. This set of additional configuration information is referred to as a Web Provisioning Domain (Web PvD). The network lists its claims of authority for DNS domains using the "dnsZones" PvD key (defined in [RFC8801]).

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 terms "Private DNS", "Global DNS" and "Split DNS" are 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 "enterprise network" in this document extends to a wide variety of deployment scenarios. For example, an "enterprise" can be a Small Office, Home Office or Corporation.

3. Split DNS

[RFC2826] "does not preclude private networks from operating their own private name spaces" but notes that if private networks "wish to make use of names uniquely defined for the global Internet, they have to fetch that information from the global DNS naming hierarchy".

There are various DNS deployments outside of the global DNS, including "split horizon" deployments and DNS usages on private (or virtual private) networks. In a split horizon, an authoritative server gives different responses to queries from the Internet than they do to network-designated DNS servers; while some deployments differentiate internal queries from public queries by the source IP address, the concerns in Section 3.1.1 of [RFC6950] relating to trusting source IP addresses apply to such deployments.

When the internal address space range is private [RFC1918], this makes it both easier for the server to discriminate public from private and harder for public entities to impersonate nodes in the private network. The use cases that motivate split-horizon DNS typically involve restricting access to some network services -- intranet resources such as internal web sites, development servers, or directories, for example -- while preserving the ease of use offered by domain names for internal users.

A typical use case is an Enterprise network that requires one or more DNS domains to be resolved via network-designated DNS servers. This can be a special domain, such as "corp.example.com" for an enterprise that is publicly known to use "example.com". In this case, the endpoint needs to be informed what the private domain names are and what the IP addresses of the network-designated DNS servers are. An Enterprise can also run a different version of its global domain on its internal network. In that case, the client is instructed to send DNS queries for the enterprise public domain (e.g., "example.com") to the network-designated DNS servers. A configuration for this deployment scenario is referred to as a Split DNS configuration. Another use case for split-horizon DNS is Cellular and Fixed-access networks (ISPs) typically offer private domains, including account status/controls, and free education initiatives [INS].

The PvD RA option defined in [RFC8801] SHOULD set the H-flag to indicate that Additional Information is available. This Additional Information JSON object SHOULD include the "dnsZones" key to define the DNS domains for which the network claims authority.

4. PvD dnsZones

As discussed in Section 3, internal resources in a network tend to have private DNS names. A network can also run a different version of its global domain on its internal network, and require the use of network-designated DNS servers to get resolved.

The PvD Key dnsZones is defined in [RFC8801]. The PvD Key dnsZones adds support for DNS domains for which the network claims authority. The private domains specified in the dnsZones key are intended to be resolved using network-designated DNS servers. The private domains in dnsZones are only reachable by devices authenticated or attached to the network. The global domains specified in the dnsZones key have a different version in the internal network. DNS resolution for other domains remains unchanged.

The dnsZones PvD Key conveys the specified DNS domains that need to be resolved using a network-designated DNS server. The DNS root zone (".") MUST be ignored if it appears in dnsZones. Other generic or global domains, such as Top-Level Domains (TLDs), similarly MUST be ignored if they appear in dnsZones.

For each dnsZones entry, the client can use the network-designated DNS servers to resolve the listed domains and its subdomains. Other domain names may be resolved using some other DNS servers that are configured independently. For example, if the dnsZones key specifies "example.test", then "example.test", "www.example.test", and "mail.eng.example.test" can be resolved using the network-designated

DNS resolver(s), but "otherexample.test" and "ple.test" can be resolved using the system's public resolver(s).

4.1. Authority over the Domains

To comply with [RFC2826] the split-horizon DNS zone must either not exist in the global DNS hierarchy or must be authoritatively delegated to the split-horizon DNS server to answer. The client can use the mechanism described in [I-D.ietf-add-dnr] to discover the network-designated resolvers. To determine if the network-designated encrypted resolvers are authoritative over the domains in DnsZones, the client performs the following steps for each domain in DnsZones:

1. The client sends an NS query for the domain in DnsZones. This query MUST only be sent over encrypted DNS session to a public resolver that is configured independently or to a network-designated resolver whose response will be validated using DNSSEC as described in [RFC6698].

2. The client checks that the NS RRset matches any one of the ADN of the discovered network-designated encrypted DNS resolvers.

    a. If the match fails, the client determines the network is not authoritative for the indicated domain. It might log an error, reject the network entirely (because the network lied about its authority over a domain) or other action.

    b. If the match succeeds, the client can then establish a secure connection to that network-designated resolver and validates its certificate.

        *If the server certificate does not validate and a secure connection cannot be established to the network designated resolver, the client can proceed as discussed in step 3 (A).

        *If the server certificate validation is successful and a secure connection is established, the client can subsequently resolve the domains in that subtree using the network-designated resolver.

3. As an exception to this rule, the client need not perform the above validation for domains reserved for special use [RFC6761] or [RFC6762] such as ".home.arpa" or ".local".

4. If the client uses a public resolver, authenticated denial of existence using NSEC3 or NSEC records can be used by a client to identify that the domain name does not exist in the global DNS.

For example, if in a network the private domain names are defined under "internal.corp1.example.com". The DnsZones PvD Key would indicate that "*.internal.corp1.example.com" are private domain names. The client can trigger a NS query of "internal.corp1.example.com" and the NS RRset returns that the nameserver is "ns1.corp2.example.com". The client would then connect to the network-designated encrypted resolver whose name is "ns1.corp2.example.com", authenticate it using server certificate validation in TLS handshake, and use it for resolving the domains in the subtree of "*.internal.corp1.example.com".

5. An Example

The following example shows how the JSON keys defined in this document can be used:

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

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

6. Split DNS Configuration for IKEv2

The split-tunnel Virtual Private Network (VPN) configuration allows the endpoint to access resources that reside in the VPN [RFC8598] via the tunnel; other traffic not destined to the VPN does not traverse the tunnel. In contrast, a non-split- tunnel VPN configuration causes all traffic to traverse the tunnel into the VPN.

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.btw-add-ipsecme-ike]. For split-tunnel VPN configurations, the endpoint uses the discovered encrypted DNS server to resolve domain names for which the VPN provider claims authority. For non-split-tunnel VPN configurations, the endpoint uses the discovered encrypted DNS server to resolve both global and private domain names. For split-tunnel VPN configurations, the IKE client can use the steps discussed in Section 4.1 to determine if the VPN service provider is authoritative over the INTERNAL_DNS_DOMAIN domains.

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

7. Security Considerations

The content of dnsZones may be passed to another (DNS) program for processing. As with any network input, the content SHOULD be considered untrusted and handled accordingly. The client must perform the steps discussed in Section 4.1 to determine if the network-designated encrypted resolvers are authoritative over the domains in DnsZones. If the network is lying, the client can take appropriate action like disconnecting from the network.

As an additional precaution, clients may want to preconfigure global domains for TLDs and Second-Level Domains (SLDs) to prevent malicious DNS redirections for well-known domains. This prevents users from unknowingly giving DNS queries to third parties. This is even more important if those well-known domains are not deploying DNSSEC, as the attached network could then even modify the DNS answers without detection. It is similar to the mechanism discussed in Section 8 of [RFC8598].

8. IANA Considerations

This document has no IANA actions..

9. Acknowledgements

Thanks to Mohamed Boucadair, Jim Reid, Ben Schwartz, Tommy Pauly, Paul Vixie and Vinny Parla for the discussion and comments. The authors would like to give special thanks to Ben Schwartz for his help.

10. References

10.1. Normative References

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

[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>.

[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/info/rfc2826>.

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

[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>.

10.2. Informative References

[I-D.btw-add-ipsecme-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-btw-add-ipsecme-ike-03, 17 May 2021, <https://www.ietf.org/archive/id/draft-btw-add-ipsecme-ike-03.txt>.

[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-02, 8 July 2021, <https://www.ietf.org/archive/id/draft-ietf-add-ddr-02.txt>.

[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-02, 17 May 2021, <https://www.ietf.org/archive/id/draft-ietf-add-dnr-02.txt>.

[INS] The Unicode Consortium, "Vodafone Foundation Instant Schools for Sub-Saharan Africa", <https://www.vodafone.com/about/vodafone-foundation/focus-areas/instant-schools>.

[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>.

[RFC6950] Peterson, J., Kolkman, O., Tschofenig, H., and B. Aboba, "Architectural Considerations on Application Features in the DNS", RFC 6950, DOI 10.17487/RFC6950, October 2013, <https://www.rfc-editor.org/info/rfc6950>.

[RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <https://www.rfc-editor.org/info/rfc7556>.

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

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

Authors' Addresses

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