Workgroup: ADD
Internet-Draft:
draft-reddy-add-enterprise-split-dns-10
Published: 13 April 2022
Intended Status: Standards Track
Expires: 15 October 2022
Authors: T. Reddy D. Wing K. Smith B. Schwartz
Akamai Citrix Vodafone Google

Establishing Local DNS Authority in Split-Horizon Environments

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

Table of Contents

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 RRSet. 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 perfomed 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].

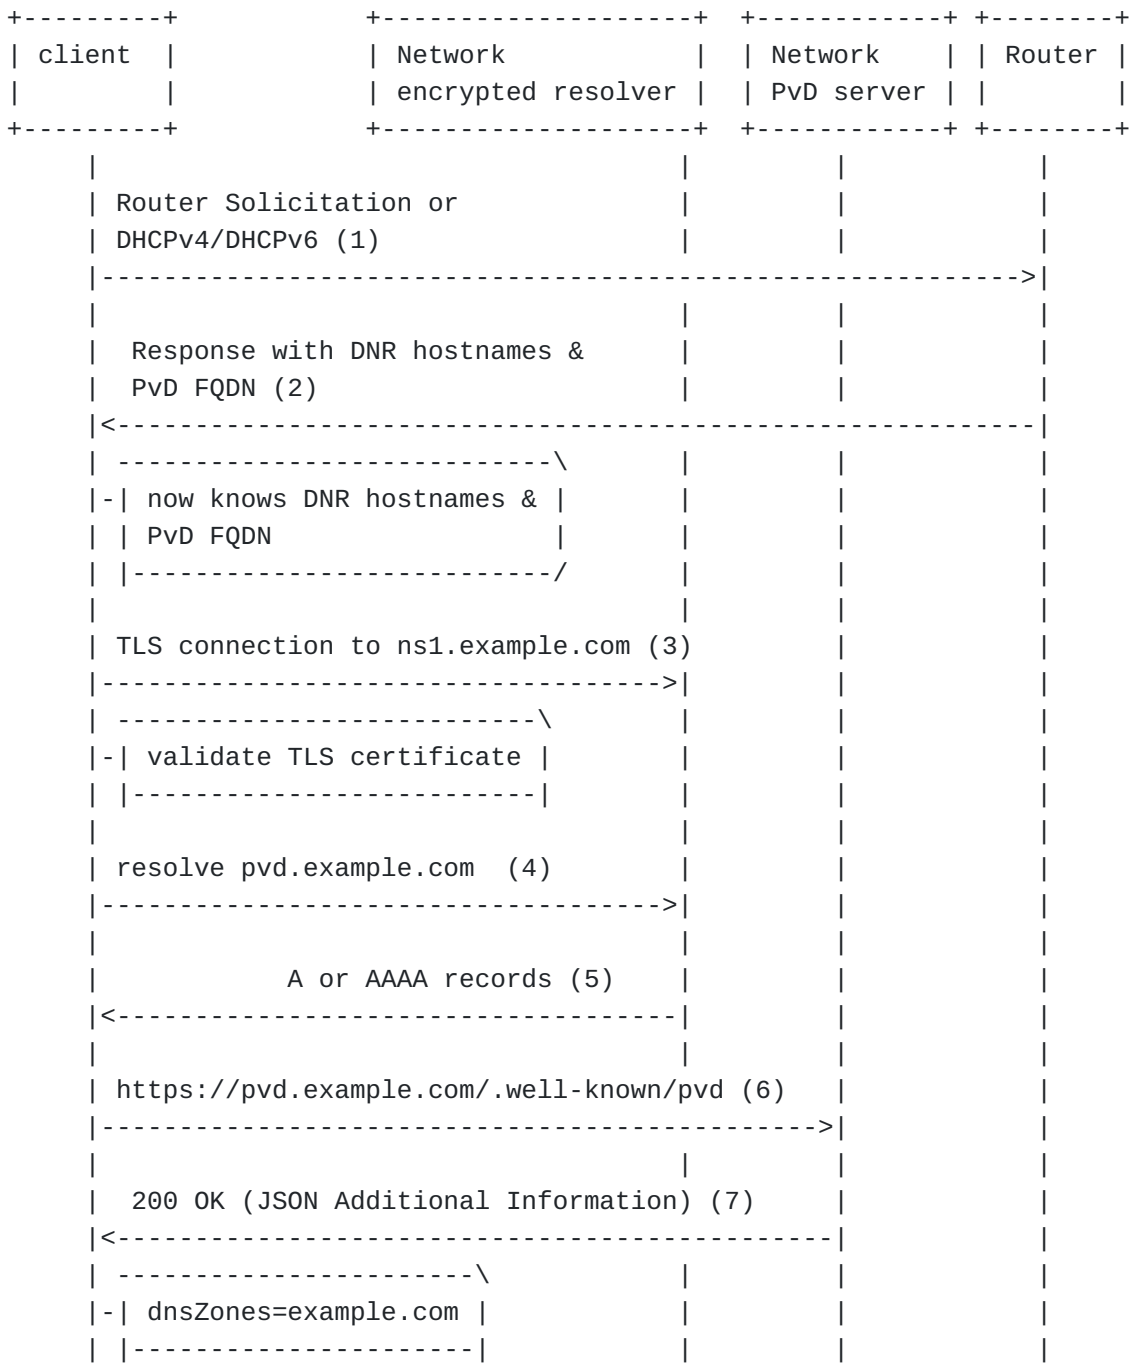

Figure 1: Learning Local Claims of DNS Authority

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.

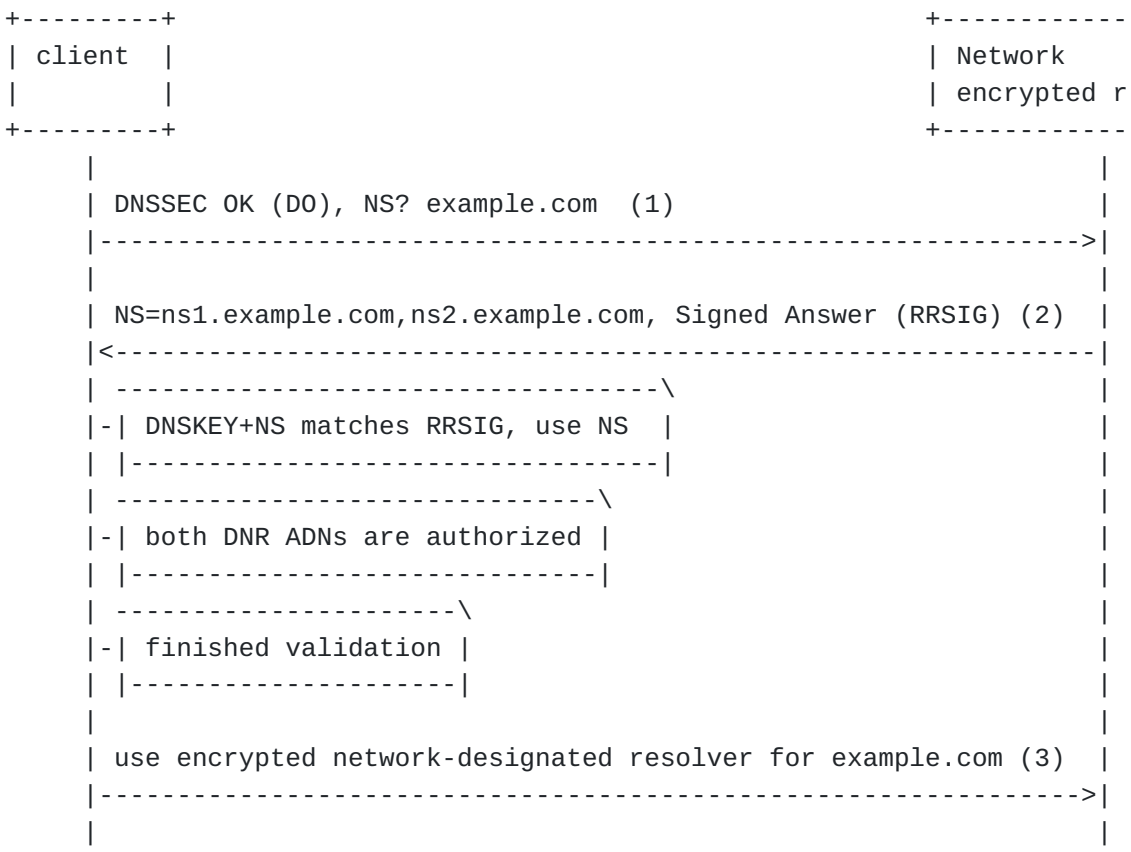

Figure 3: Verifying Claims using DNSSEC

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 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 Transparancy 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