

Network
Internet-Draft
Intended status: Standards Track
Expires: March 25, 2017

T. Pauly
Apple Inc.
P. Wouters
Red Hat
September 21, 2016

Split DNS Configuration for IKEv2 draft-pauly-ipsecme-split-dns-02

Abstract

This document defines two Configuration Payload Attribute Types for the IKEv2 protocol that add support for private DNS domains. These domains should be resolved using DNS servers reachable through an IPsec connection, while leaving all other DNS resolution unchanged. This approach of resolving a subset of domains using non-public DNS servers is referred to as "Split 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 <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on March 25, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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
1.1.	Requirements Language	3
2.	Background	3
3.	Protocol Exchange	3
3.1.	Configuration Request	4
3.2.	Configuration Reply	4
3.3.	Mapping DNS Servers to Domains	5
3.4.	Example Exchanges	5
3.4.1.	Simple Case	5
3.4.2.	Requesting Limited Domains	6
3.4.3.	Requesting Domains and DNSSEC trust anchors	6
4.	Payload Formats	7
4.1.	INTERNAL_DNS_DOMAIN Configuration Attribute Type	7
4.2.	INTERNAL_DNSSEC_TA Configuration Attribute	8
5.	Split DNS Usage Guidelines	8
6.	Security Considerations	9
7.	IANA Considerations	10
8.	References	10
8.1.	Normative References	11
8.2.	Informative References	11
	Authors' Addresses	11

[1.](#) Introduction

Split DNS is a common configuration for secure tunnels, such as Virtual Private Networks in which host machines private to an organization can only be resolved using internal DNS resolvers [[RFC2775](#)]. In such configurations, it is often desirable to only resolve hosts within a set of private domains using the tunnel, while letting resolutions for public hosts be handled by a device's default DNS configuration.

The Internet Key Exchange protocol version 2 [[RFC7296](#)] negotiates configuration parameters using Configuration Payload Attribute Types. This document defines two Configuration Payload Attribute Types that add support for trusted Split DNS domains.

The INTERNAL_DNS_DOMAIN attribute type is used to convey one or more DNS domains that should be resolved only using the provided DNS nameserver IP addresses, causing these requests to use the IPsec connection.

The `INTERNAL_DNSSEC_TA` attribute type is used to convey DNSSEC trust anchors for those domains.

When only a subset of traffic is routed into a private network using an IPsec SA, these Configuration Payload options can be used to define which private domains should be resolved through the IPsec connection without affecting the client's global DNS resolution.

For the purposes of this document, DNS resolution servers accessible through an IPsec connection will be referred to as "internal DNS servers", and other DNS servers will be referred to as "external DNS servers".

A client using these configuration payloads will be able to request and receive Split DNS configurations using the `INTERNAL_DNS_DOMAIN` and `INTERNAL_DNSSEC_TA` configuration attributes. The client device can use the internal DNS server(s) for any DNS queries within the assigned domains, while routing other DNS queries to its regular external DNS server.

1.1. Requirements Language

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

2. Background

Split DNS is a common configuration for enterprise VPN deployments, in which only one or a few private DNS domains are accessible and resolvable via an IPsec based VPN connection.

Other tunnel-establishment protocols already support the assignment of Split DNS domains. For example, there are proprietary extensions to IKEv1 that allow a server to assign Split DNS domains to a client. However, the IKEv2 standard does not include a method to configure this option. This document defines a standard way to negotiate this option for IKEv2.

3. Protocol Exchange

In order to negotiate which domains are considered internal to an IKEv2 tunnel, initiators indicate support for Split DNS in their `CFG_REQUEST` payloads, and responders assign internal domains (and DNSSEC trust anchors) in their `CFG_REPLY` payloads. When Split DNS has been negotiated, the existing DNS server configuration attributes will be interpreted as internal DNS servers that can resolve hostnames within the internal domains.

3.1. Configuration Request

To indicate support for Split DNS, an initiator sending a CFG_REQUEST payload MAY include one or more INTERNAL_DNS_DOMAIN attributes as defined in [Section 4](#). If an INTERNAL_DNS_DOMAIN attribute is included in the CFG_REQUEST, the initiator SHOULD also include one or both of the INTERNAL_IP4_DNS and INTERNAL_IP6_DNS attributes in its CFG_REQUEST.

If the length of the INTERNAL_DNS_DOMAIN attribute is zero, then the initiator is requesting that the attribute be assigned without restricting the subdomains that it will accept.

If the length of the INTERNAL_DNS_DOMAIN is greater than zero, the value is a single DNS domain. The initiator is indicating that it will only allow this domain and any sub-domains within this domain to be resolved using the internal DNS servers. The list of INTERNAL_DNS_DOMAIN attributes in the CFG_REQUEST defines the full set of domains the initiator is willing to resolve using the internal DNS servers.

The absence of INTERNAL_DNS_DOMAIN attributes in the CFG_REQUEST payload indicates that the initiator does not support or is unwilling to accept Split DNS configuration.

To indicate support for DNSSEC, an initiator sending a CFG_REQUEST payload MAY include one or more INTERNAL_DNS_TA attributes as defined in [Section 4](#). These payloads MUST immediately follow a INTERNAL_DNS_DOMAIN attribute, which binds the DNSSEC trust anchor request to the domain.

An initiator MAY convey its current DNSSEC trust anchors for the domain specified in the INTERNAL_DNS_DOMAIN attribute. If it does not wish to convey this information, it MUST use a length of 0.

The absence of INTERNAL_DNS_TA attributes in the CFG_REQUEST payload indicates that the initiator does not support or is unwilling to accept DNSSEC trust anchor configuration.

3.2. Configuration Reply

Responders MAY send one or more INTERNAL_DNS_DOMAIN attributes in their CFG_REPLY payload if the CFG_REQUEST contained at least one INTERNAL_DNS_DOMAIN attribute. If the CFG_REQUEST did not contain an INTERNAL_DNS_DOMAIN attribute, the responder MUST NOT include an INTERNAL_DNS_DOMAIN attribute in the CFG_REPLY. If an INTERNAL_DNS_DOMAIN attribute is included in the CFG_REPLY, the responder SHOULD also include one or both of the INTERNAL_IP4_DNS and

INTERNAL_IP6_DNS attributes in its CFG_REPLY. These DNS server configurations are necessary to define which servers should receive queries for hostnames in internal domains. If the CFG_REQUEST included an INTERNAL_DNS_DOMAIN attribute, but the CFG_REPLY does not include an INTERNAL_DNS_DOMAIN attribute, the initiator should behave as if Split DNS configurations are not supported by the server.

Each INTERNAL_DNS_DOMAIN represents a domain that the DNS servers address listed in INTERNAL_IP4_DNS and INTERNAL_IP6_DNS can resolve.

If the CFG_REQUEST included INTERNAL_DNS_DOMAIN attributes with non-zero lengths, the CFG_REPLY MUST NOT assign any domains in its INTERNAL_DNS_DOMAIN attributes that are not contained within the requested domains. The initiator SHOULD ignore any domains beyond its requested list.

For each DNS domain specified in an INTERNAL_DNS_DOMAIN attribute, one or more INTERNAL_DNSSEC_TA attributes MAY be included by the responder. This attribute lists the corresponding DNSSEC trust anchor in the DNS wire format of a DS record as specified in [\[RFC4034\]](#). The INTERNAL_DNSSEC_TA attribute MUST immediately follow the INTERNAL_DNS_DOMAIN attribute that it applies to.

[3.3.](#) Mapping DNS Servers to Domains

All DNS servers provided in the CFG_REPLY MUST support resolving hostnames within all INTERNAL_DNS_DOMAIN domains. In other words, the INTERNAL_DNS_DOMAIN attributes in a CFG_REPLY payload form a single list of Split DNS domains that applies to the entire list of INTERNAL_IP4_DNS and INTERNAL_IP6_DNS attributes.

[3.4.](#) Example Exchanges

[3.4.1.](#) Simple Case

In this example exchange, the initiator requests INTERNAL_IP4_DNS and INTERNAL_DNS_DOMAIN attributes in its CFG_REQUEST, but does not specify any value for either. This indicates that it supports Split DNS, but has no preference for which DNS requests should be routed through the tunnel.

The responder replies with two DNS server addresses, and one internal domain, "example.com".

Any subsequent DNS queries from the initiator for domains such as "www.example.com" should use 198.51.100.2 or 198.51.100.4 to resolve.


```
CP(CFG_REQUEST) =  
  INTERNAL_IP4_ADDRESS()  
  INTERNAL_IP4_DNS()  
  INTERNAL_DNS_DOMAIN()  
  
CP(CFG_REPLY) =  
  INTERNAL_IP4_ADDRESS(198.51.100.234)  
  INTERNAL_IP4_DNS(198.51.100.2)  
  INTERNAL_IP4_DNS(198.51.100.4)  
  INTERNAL_DNS_DOMAIN(example.com)
```

3.4.2. Requesting Limited Domains

In this example exchange, the initiator requests `INTERNAL_IP4_DNS` and `INTERNAL_DNS_DOMAIN` attributes in its `CFG_REQUEST`, specifically requesting only "example.com" and "other.com". The responder replies with two DNS server addresses, 198.51.100.2 and 198.51.100.4, and two domains, "example.com" and "city.other.com". Note that one of the domains in the `CFG_REPLY`, "city.other.com", is a subset of the requested domain, "other.com". This indicates that hosts within "other.com" that are not within "city.other.com" should be resolved using an external DNS server. The `CFG_REPLY` would not be allowed to respond with "com" or "example.net", however, since these were contained within the limited set of requested domains.

Any subsequent DNS queries from the initiator for domains such as "www.example.com" or "city.other.com" should use 198.51.100.2 or 198.51.100.4 to resolve.

```
CP(CFG_REQUEST) =  
  INTERNAL_IP4_ADDRESS()  
  INTERNAL_IP4_DNS()  
  INTERNAL_DNS_DOMAIN(example.com)  
  INTERNAL_DNS_DOMAIN(other.com)  
  
CP(CFG_REPLY) =  
  INTERNAL_IP4_ADDRESS(198.51.100.234)  
  INTERNAL_IP4_DNS(198.51.100.2)  
  INTERNAL_IP4_DNS(198.51.100.4)  
  INTERNAL_DNS_DOMAIN(example.com)  
  INTERNAL_DNS_DOMAIN(city.other.com)
```

3.4.3. Requesting Domains and DNSSEC trust anchors

In this example exchange, the initiator requests `INTERNAL_IP4_DNS`, `INTERNAL_DNS_DOMAIN` and `INTERNAL_DNS_TA` attributes in its `CFG_REQUEST`

- o Reserved (1 bit) - Defined in IKEv2 RFC [[RFC7296](#)].
- o Attribute Type (15 bits) [TBD IANA] - INTERNAL_DNS_DOMAIN.
- o Length (2 octets, unsigned integer) - Length of domain name.
- o Domain Name (0 or more octets) - A domain or subdomain used for Split DNS rules, such as example.com. This is a string of ASCII

For each INTERNAL_DNS_DOMAIN entry in a CFG_REPLY payload, the client SHOULD use the provided INTERNAL_IP4_DNS or INTERNAL_IP6_DNS DNS servers as the only resolvers for the listed domains and its sub-domains and it SHOULD NOT attempt to resolve the provided DNS domains using its external DNS servers.

If the initiator host is configured to block DNS answers containing IP addresses from special IP address ranges such as those of [\[RFC1918\]](#), the initiator SHOULD allow the DNS domains listed in the INTERNAL_DNS_DOMAIN attributes to contain these IP addresses.

If a CFG_REPLY contains one or more INTERNAL_DNS_DOMAIN attributes, the client SHOULD configure its DNS resolver to resolve those domains and all their subdomains using only the DNS resolver(s) listed in that CFG_REPLY message. If those resolvers fail, those names SHOULD NOT be resolved using any other DNS resolvers. All other domain names MUST be resolved using some other external DNS resolver(s), configured independently, and SHOULD NOT be sent to the internal DNS resolver(s) listed in that CFG_REPLY message. For example, if the INTERNAL_DNS_DOMAIN attribute specifies "example.com", then "example.com", "www.example.com" and "mail.eng.example.com" MUST be resolved using the internal DNS resolver(s), but "anotherexample.com" and "ample.com" MUST be resolved using the system's external DNS resolver(s).

An initiator SHOULD ignore INTERNAL_DNS_DOMAIN attributes containing domains that are designated Special Use Domain Names in [\[RFC6761\]](#), such as "local", "localhost", "invalid", etc. Although it may explicitly wish to support some Special Use Domain Names.

When an IPsec connection is terminated, the DNS forwarding must be unconfigured. The DNS forwarding itself MUST be deleted. All cached data of the INTERNAL_DNS_DOMAIN provided DNS domains MUST be flushed. This includes negative cache entries. Obtained DNSSEC trust anchors MUST be removed from the list of trust anchors. The outstanding DNS request queue MAY be cleared.

A domain that is served via INTERNAL_DNS_DOMAIN MUST NOT have indirect references to DNS records that point to other Split DNS domains that are not served via INTERNAL_DNS_DOMAIN attributes. Indirect reference RRtypes include CNAME, DNAME, MX and SRV RR's.

INTERNAL_DNS_DOMAIN and INTERNAL_DNSSEC_TA attributes SHOULD only be used on split tunnel configurations where only a subset of traffic is routed into a private remote network using the IPsec connection. If all traffic is routed over the IPsec connection, the existing global INTERNAL_IP4_DNS and INTERNAL_IP6_DNS can be used without creating specific DNS exemptions.

6. Security Considerations

The use of Split DNS configurations assigned by an IKEv2 responder is predicated on the trust established during IKE SA authentication. However, if IKEv2 is being negotiated with an anonymous or unknown

endpoint (such as for Opportunistic Security [[RFC7435](#)]), the initiator MUST ignore Split DNS configurations assigned by the responder.

If a host connected to an authenticated IKE peer is connecting to another IKE peer that attempts to claim the same domain via the INTERNAL_DNS_DOMAIN attribute, the IKE connection should be terminated.

If the IP address value of the received INTERNAL_IP4_DNS or INTERNAL_IP6_DNS attribute is not covered by the proposed IPsec connection, then the local DNS should not be reconfigured until a CREATE_CHILD Exchange is received that covers these IP addresses.

INTERNAL_DNSSEC_TA directives MUST immediately follow an INTERNAL_DNS_DOMAIN directive. As the INTERNAL_DNSSEC_TA format itself does not contain the domain name, it relies on the preceding INTERNAL_DNS_DOMAIN to provide the domain for which it specifies the trust anchor.

If the initiator is using DNSSEC validation for a domain in its public DNS view, and it requests and receives an INTERNAL_DNS_DOMAIN attribute without an INTERNAL_DNSSEC_TA, it will need to reconfigure its DNS resolver to allow for an insecure delegation. It SHOULD NOT accept insecure delegations for domains that are DNSSEC signed in the public DNS view, for which it has not explicitly requested such delegation by specifying the domain specifically using a INTERNAL_DNS_DOMAIN(domain) request.

7. IANA Considerations

This document defines two new IKEv2 Configuration Payload Attribute Types, which are allocated from the "IKEv2 Configuration Payload Attribute Types" namespace.

Value	Attribute Type	Multi-Valued	Length	Reference
-----	-----	-----	-----	-----
[TBD]	INTERNAL_DNS_DOMAIN	YES	0 or more	[this document]
[TBD]	INTERNAL_DNSSEC_TA	YES	0 or more	[this document]

Figure 1

8. References

8.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), DOI 10.17487/RFC1918, February 1996, <<http://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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", [RFC 4034](#), DOI 10.17487/RFC4034, March 2005, <<http://www.rfc-editor.org/info/rfc4034>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", [RFC 5890](#), DOI 10.17487/RFC5890, August 2010, <<http://www.rfc-editor.org/info/rfc5890>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, [RFC 7296](#), DOI 10.17487/RFC7296, October 2014, <<http://www.rfc-editor.org/info/rfc7296>>.

8.2. Informative References

- [RFC2775] Carpenter, B., "Internet Transparency", [RFC 2775](#), DOI 10.17487/RFC2775, February 2000, <<http://www.rfc-editor.org/info/rfc2775>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", [RFC 6761](#), DOI 10.17487/RFC6761, February 2013, <<http://www.rfc-editor.org/info/rfc6761>>.
- [RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", [RFC 7435](#), DOI 10.17487/RFC7435, December 2014, <<http://www.rfc-editor.org/info/rfc7435>>.

Authors' Addresses

Tommy Pauly
Apple Inc.
1 Infinite Loop
Cupertino, California 95014
US

Email: tpauly@apple.com

Paul Wouters
Red Hat

Email: pwouters@redhat.com