Network
Internet-Draft
Intended status: Standards Track
Expires: March 27, 2016

T. Pauly
Apple Inc.
P. Wouters
Red Hat

September 24, 2015

Split-DNS Configuration for IKEv2
draft-pauly-ipsecme-split-dns-00

Abstract

This document defines two new Configuration Payload Attribute Types for the IKEv2 protocol that together define a set of private DNS domains which should be resolved by DNS servers reachable through an IPsec connection, while leaving all other DNS resolution unchanged. This allows for split-DNS views for multiple domains and includes support for private DNSSEC trust anchors. The information obtained via the new attribute types can be used to reconfigure a locally running DNS server with DNS forwarding for specific private 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 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 27, 2016.

Copyright Notice

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

The Internet Key Exchange protocol version 2 [RFC7296] negotiates configuration parameters using Configuration Payload Attribute Types. This document adds two new Configuration Payload Attribute Types to support trusted split-DNS domains. The INTERNAL_DNS_DOMAIN attribute type is used to convey one or more local DNS domains. 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, this Configuration Payload option 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 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 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

3.1. Configuration Request

To indicate support for split-DNS, initiators sending a CFG_REQUEST payload MAY include one or more of the INTERNAL_DNS_DOMAIN configuration attribute in their configuration payloads. See Section 4 for details on the payload format. 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 attribute length is zero, then the initiator is only

requesting that the attribute be assigned, without restricting the subdomains that it will accept.

If the attribute length is non-zero, it contains a single DNS domain . The initiator is indicating that it will allow this domain and its sub-domains to be resolved over the IPsec connection. The list of INTERNAL_DNS_DOMAIN attributes in the CFG_REQUEST defines the full set of domains the intiator is willing to resolve over the tunnel.

3.2. Configuration Reply

Responders MAY send one or more INTERNAL_DNS_DOMAIN configuration attributes in their CFG_REPLY payload if the CFG_REQUEST contained at least one INTERNAL_DNS_DOMAIN configuration attribute. If the CFG_REQUEST did not contain an INTERNAL_DNS_DOMAIN configuration attribute, the responder MUST NOT include an INTERNAL_DNS_DOMAIN configuration attribute in the CFG_REPLY. If an INTERNAL_DNS_DOMAIN configuration attribute is included in the CFG_REPLY, the responder SHOULD also include one or both of the INTERNAL_IP4_DNS and INTERNAL_IP6_DNS configuration attributes in its CFG_REPLY. If the CFG_REQUEST included an INTERNAL_DNS_DOMAIN configuration attribute, but the CFG_REPLY does 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 configuration attribute, an INTERNAL_DNSSEC_TA configuration attribute may be included by the responder. This attribute lists the corresponding DSSNEC trust anchor in the presentation format of a DS record as

specified in [RFC4034].

3.3. Mapping DNS Servers to Domains

All DNS servers provided in the CFG_REPLY MUST support all domains. The INTERNAL_DNS_DOMAIN attributes in a CFG_REPLY payload form a single list of split-DNS domains, that apply 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)
```


4. Payload Formats

4.1. INTERNAL_DNS_DOMAIN Configuration Attribute Type

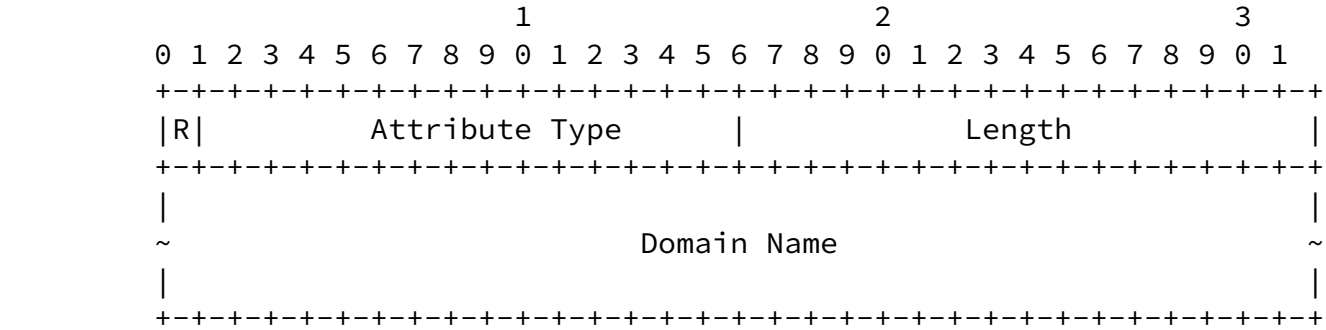

- Reserved (1 bit) - Defined in IKEv2 RFC [RFC7296].

- Attribute Type (15 bits) [TBD IANA] - INTERNAL_DNS_DOMAIN.

- Length (2 octets, unsigned integer) - Length of domain name.

- 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 characters with labels separated by dots, with no trailing dot, using IDNA [RFC5890] for non-ASCII DNS domains. The value is NOT null-terminated.

4.2. INTERNAL_DNSSEC_TA Configuration Attribute

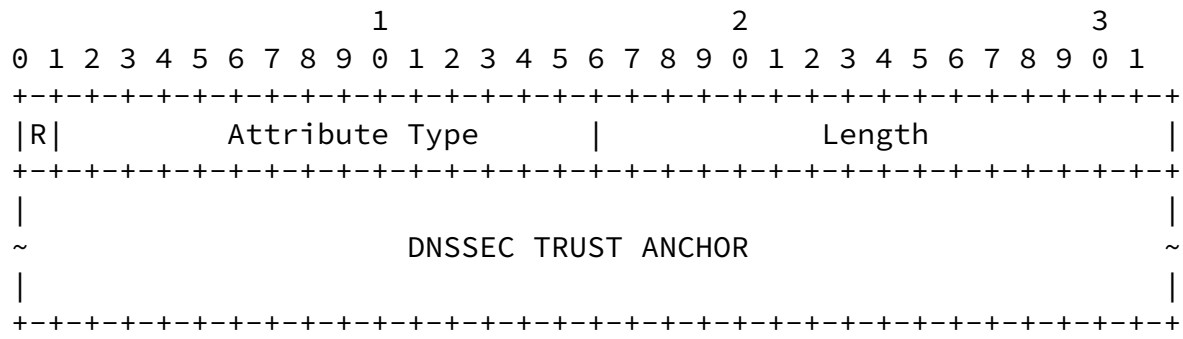

- Reserved (1 bit) - Defined in IKEv2 RFC [RFC7296].

- Attribute Type (15 bits) [TBD IANA] - INTERNAL_DNSSEC_TA.

- Length (2 octets, unsigned integer) - Length of DNSSEC Trust Anchor data.

- DNSSEC Trust anchor (multiple octets) - The presentation format of one DS record as specified in [RFC4034]. The TTL value MAY be omited and when present MUST be ignored. The domain name is

specified as a Fully Qualified Domain Name (FQDN) - irrespective of the presence of a trailing dot, and consits of a string of ASCII characters with labels separated by dots and uses IDNA [RFC5890] for non-ASCII DNS domains. The value is NOT null-terminated.

5. Split-DNS Usage Guidelines

If a CFG_REPLY payload contains no INTERNAL_DNS_DOMAIN configuration attributes, the client MAY use the provided INTERNAL_IP4_DNS or INTERNAL_IP6_DNS servers as the default DNS server(s) for all queries.

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 configuration attributes to contain these IP addresses.

If a CFG_REPLY contains one or more INTERNAL_DNS_DOMAIN configuration 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 SHOULD 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 configuration attribute specifies "example.com", then "example.com", "www.example.com" and "mail.eng.example.com" SHOULD be resolved using the internal DNS resolver(s), but "anotherexample.com" and "ample.com" SHOULD be resolved using the system's external DNS resolver(s).

An initiator MAY ignore INTERNAL_DNS_DOMAIN configuration 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, for example "onion" [I-D.ietf-dnsop-onion-tld].

When an IPsec connection is terminated, the DNS forwarding must be unconfigured. The DNS forwarding itself MUST be be deleted. All cached data of the INTERNAL_DNS_DOMAIN provided DNS domainis 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 configuration attributes. Indirect reference resource record types include CNAME, DNAME, MX and SRV resource records.

INTERNAL_DNS_DOMAIN and INTERNAL_DNSSEC_TA configuration 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 configuration attribute, the IKE connection should be terminated.

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

INTERNAL_DNSSEC_TA directives MUST have an accompanying INTERNAL_DNS_DOMAIN directive. This prevents the insertion of rogue DNSSEC trust anchors for domains that have not been yielded to the IPsec connection.

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

[I-D.ietf-dnsop-onion-tld] Appelbaum, J. and A. Muffett, "The .onion Special-Use Domain Name", draft-ietf-dnsop-onion-tld-01 (work in progress), September 2015.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013,

[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