

ADD
Internet-Draft
Intended status: Standards Track
Expires: May 12, 2022

M. Boucadair
Orange
T. Reddy
McAfee
D. Wing
Citrix
V. Smyslov
ELVIS-PLUS
November 8, 2021

**Internet Key Exchange Protocol Version 2 (IKEv2) Configuration for
Encrypted DNS
draft-btw-add-ipsecme-ike-04**

Abstract

This document specifies a new Internet Key Exchange Protocol Version 2 (IKEv2) Configuration Payload Attribute Types for encrypted DNS protocols such as DNS-over-HTTPS (DoH), DNS-over-TLS (DoT), and DNS-over-QUIC (DoQ).

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 May 12, 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.	Introduction	2
2.	Terminology	3
3.	IKEv2 Configuration Payload Attribute Types for Encrypted DNS	3
3.1.	ENCDNS_IP* Configuration Payload Attributes	3
3.2.	ENCDNS_DIGEST_INFO Configuration Payload Attribute	5
4.	IKEv2 Protocol Exchange	7
5.	Security Considerations	8
6.	IANA Considerations	9
6.1.	Configuration Payload Attribute Types	9
7.	Acknowledgements	9
8.	References	9
8.1.	Normative References	9
8.2.	Informative References	10
Appendix A.	Sample Deployment Scenarios	11
A.1.	Roaming Enterprise Users	11
A.2.	VPN Service Provider	12
A.3.	DNS Offload	12
	Authors' Addresses	12

[1.](#) Introduction

This document specifies encrypted DNS configuration for an Internet Key Exchange Protocol Version 2 (IKEv2) [[RFC7296](#)] initiator, particularly the Authentication Domain Name (ADN) of encrypted DNS protocols such as DNS-over-HTTPS (DoH) [[RFC8484](#)], DNS-over-TLS (DoT) [[RFC7858](#)], or DNS-over-QUIC (DoQ) [[I-D.ietf-dprive-dnsquic](#)].

This document introduces new IKEv2 Configuration Payload Attribute Types ([Section 3](#)) for the support of encrypted DNS servers. These attributes can be used to provision authentication domain names, a list of IP addresses, and a set of service parameters.

Sample use cases are discussed in [Appendix A](#). The Configuration Payload Attribute Types defined in this document are not specific to these deployments, but can also be used in other deployment contexts. It is out of the scope of this document to provide a comprehensive list of deployment contexts.

The encrypted DNS server hosted by the VPN provider can get a domain-validate certificate from a public CA. The VPN client does not need

to be provisioned with the root certificate of a private CA to authenticate the certificate of the encrypted DNS server. The encrypted DNS server can run on private IP addresses and its access can be restricted to clients connected to the VPN.

Note that, for many years, typical designs have often considered that the DNS server was usually located inside the protected domain, but could be located outside of it. With encrypted DNS, the latter option becomes plausible.

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 uses of the terms defined in [\[RFC8499\]](#).

Also, this document uses of the terms defined in [\[RFC7296\]](#). In particular, readers should be familiar with "initiator" and "responder" terms used in that document.

This document makes use of the following terms:

'Do53': refers to unencrypted DNS.

'Encrypted DNS': refers to a scheme where DNS messages are sent over an encrypted channel. Examples of encrypted DNS are DoT, DoH, and DoQ.

'ENCDNS_IP*': refers to any IKEv2 Configuration Payload Attribute Types defined in [Section 3.1](#).

3. IKEv2 Configuration Payload Attribute Types for Encrypted DNS

3.1. ENCDNS_IP* Configuration Payload Attributes

The ENCDNS_IP* IKEv2 Configuration Payload Attribute Types are used to configure encrypted DNS servers. All these attributes share the format shown in Figure 1.

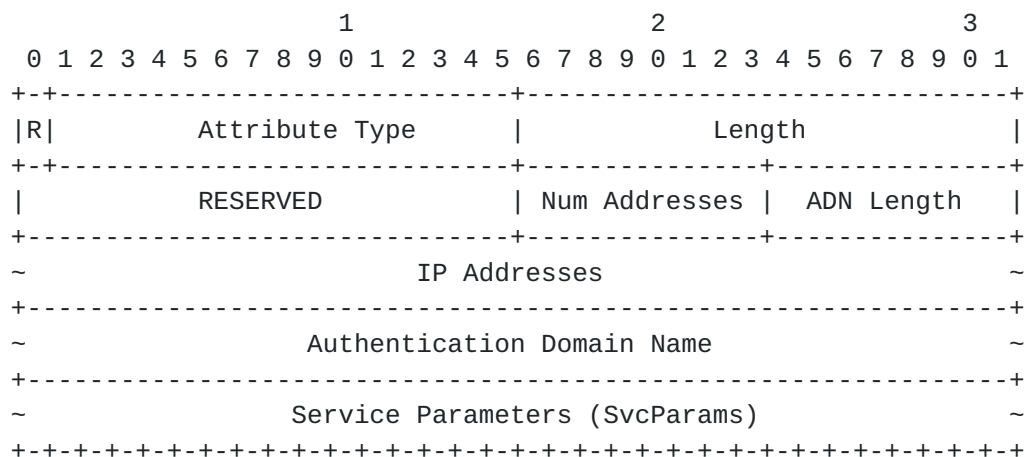


Figure 1: Attributes Format

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

- o R (Reserved, 1 bit) - This bit MUST be set to zero and MUST be ignored on receipt (see [Section 3.15.1 of \[RFC7296\]](#) for details).
- o Attribute Type (15 bits) - Identifier for Configuration Attribute Type; is set to TBA1 or TBA2 values listed in [Section 6.1](#).
- o Length (2 octets, unsigned integer) - Length of the data in octets. In particular, this field is set to:
 - * 0 if the Configuration payload has types CFG_REQUEST or CFG_ACK.
 - * $(4 + \text{Length of the ADN} + N * 4 + \text{Length of SvcParams})$ for ENCDNS_IP4 attributes if the Configuration payload has types CFG_REPLY or CFG_SET; N being the number of included IPv4 addresses ('Num addresses').
 - * $(4 + \text{Length of the ADN} + N * 16 + \text{Length of SvcParams})$ for ENCDNS_IP6 attributes if the Configuration payload has types CFG_REPLY or CFG_SET; N being the number of included IPv6 addresses ('Num addresses').
- o RESERVED (2 octets) - These bits are reserved for future use. These bits MUST be set to zero by the sender and MUST be ignored by the receiver.
- o Num Addresses (1 octet) - Indicates the number of enclosed IPv4 (for ENCDNS_IP4 attribute type) or IPv6 (for ENCDNS_IP6 attribute type) addresses. It MUST NOT be set to 0.

- o ADN Length (1 octet) - Indicates the length of the authentication-domain-name field in octets.
- o IP Address(es) (variable) - One or more IPv4 or IPv6 addresses to be used to reach the encrypted DNS server that is identified by the name in the Authentication Domain Name.
- o Authentication Domain Name (variable) - A fully qualified domain name of the encrypted DNS server following the syntax defined in [\[RFC5890\]](#). The name MUST NOT contain any terminators (e.g., NULL, CR).

An example of a valid ADN for DoH server is "doh1.example.com".

- o Service Parameters (SvcParams) (variable) - 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. As a reminder, the default port number is 853 for DoT and 443 for DoH.

The service parameters apply to all IP addresses in the ENCDNS_IP* Configuration Payload Attribute.

[3.2.](#) ENCDNS_DIGEST_INFO Configuration Payload Attribute

The format of ENCDNS_DIGEST_INFO configuration payload attribute is shown in Figure 2.

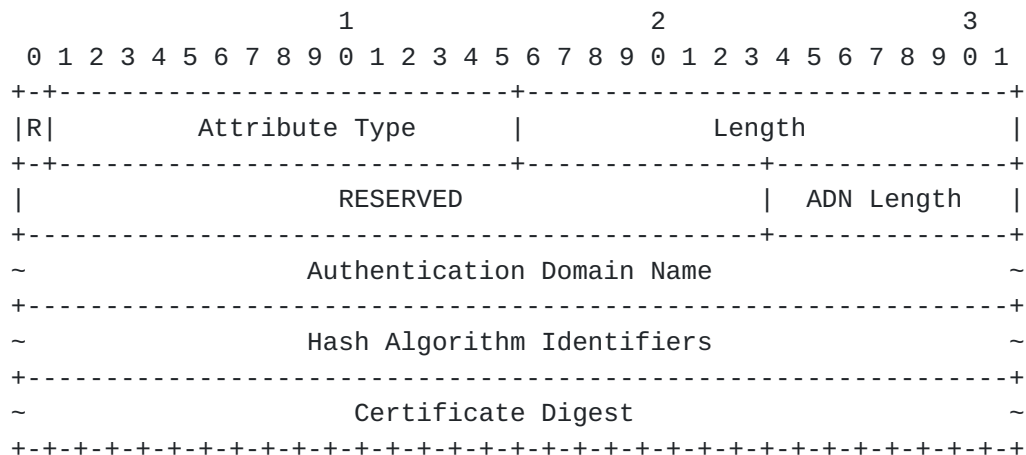


Figure 2: ENCDNS_DIGEST_INFO Attribute Format

- o R (Reserved, 1 bit) - This bit MUST be set to zero and MUST be ignored on receipt (see [Section 3.15.1 of \[RFC7296\]](#) for details).
- o Attribute Type (15 bits) - Identifier for Configuration Attribute Type; is set to TBA3 value listed in [Section 6.1](#).
- o Length (2 octets, unsigned integer) - Length of the data in octets.
- o RESERVED (2 octets) - These bits are reserved for future use. These bits MUST be set to zero by the sender and MUST be ignored by the receiver.
- o ADN Length (1 octet) - Indicates the length of the authentication-domain-name field in octets. When set to '0', this means that the digest applies on the ADN conveyed in the ENCDNS_IP* Configuration Payload Attribute(s).
- o Authentication Domain Name (variable) - A fully qualified domain name of the encrypted DNS server following the syntax defined in [\[RFC5890\]](#). The name MUST NOT contain any terminators (e.g., NULL, CR). A name is included only when multiple ADNs are included in the ENCDNS_IP* Configuration Payload Attributes.
- o Hash Algorithm Identifiers (variable) - In a request, this field specifies a list of 16-bit hash algorithm identifiers that are supported by the Encrypted DNS client. In a response, this field specified the 16-bit hash algorithm identifier selected by the server to generate the digest of its certificate. The values of this field are taken from the Hash Algorithm Identifiers of IANA's "Internet Key Exchange Version 2 (IKEv2) Parameters" registry [\[Hash\]](#).

There is no padding between the hash algorithm identifiers.

Note that SHA2-256 is mandatory to implement.

- o Certificate Digest (variable) - MUST only be present in a response. This field includes the digest of the Encrypted DNS server certificate using the algorithm identified in the 'Hash Algorithm Identifiers' field.

4. IKEv2 Protocol Exchange

This section describes how an initiator can be configured with an encrypted DNS server (e.g., DoH, DoT) using IKEv2.

Initiators indicate the support of an encrypted DNS in the CFG_REQUEST payloads by including one or two ENCDNS_IP* attributes, while responders supply the encrypted DNS configuration in the CFG_REPLY payloads. Concretely:

If the initiator supports encrypted DNS, it includes one or two ENCDNS_IP* attributes in the CFG_REQUEST. For each IP address family the initiator MUST include exactly one attribute with the Length field set to 0, so that no data is included for these attributes. The initiator MAY include the ENCDNS_DIGEST_INFO attribute with a list of hash algorithms that are supported by the Encrypted DNS client.

For each ENCDNS_IP* attribute from the CFG_REQUEST, if the responder supports the corresponding address family, and absent any policy, the responder sends back ENCDNS_IP* attribute(s) in the CFG_REPLY with an appropriate list of IP addresses, service parameters, and an ADN. The list of IP addresses MUST include at least one IP address. Multiple instances of the same ENCDNS_IP* attribute MAY be returned if distinct ADNs or service parameters are to be returned by the responder. The same or distinct IP addresses can be returned in such instances. In addition, the responder MAY return the ENCDNS_DIGEST_INFO attribute to convey a digest of the certificate of the Encrypted DNS and the identifier of the hash algorithm that is used to generate the digest.

The behavior of the responder if it receives both ENCDNS_IP* and INTERNAL_IP6_DNS (or INTERNAL_IP4_DNS) attributes is policy-based and deployment-specific. However, it is RECOMMENDED that if the responder includes at least one ENCDNS_IP* attribute in the reply, it should not include any of INTERNAL_IP4_DNS/INTERNAL_IP6_DNS attributes.

If the CFG_REQUEST includes an ENCDNS_IP* attribute but the CFG_REPLY does not include an ENCDNS_IP* matching the requested address family, this is an indication that requested address family is not supported by the responder or the responder is not configured to provide corresponding server addresses.

The DNS client establishes an encrypted DNS session (e.g., DoT, DoH, DoQ) with the address(es) conveyed in ENCDNS_IP* and uses the mechanism discussed in [Section 8 of \[RFC8310\]](#) to authenticate the DNS server certificate using the authentication domain name conveyed in ENCDNS_IP*.

If the CFG_REPLY includes an ENCDNS_DIGEST_INFO attribute, the DNS client has to create a digest of the DNS server certificate received in the TLS handshake using the negotiated hash algorithm in the ENCDNS_DIGEST_INFO attribute. If the computed digest for an ADN matches the one sent in the ENCDNS_DIGEST_INFO attribute, the encrypted DNS server certificate is successfully validated. If so, the client continues with the TLS connection as normal. Otherwise, the client MUST treat the server certificate validation failure as a non-recoverable error. This approach is similar to certificate usage PKIX-EE(1) defined in [\[RFC7671\]](#).

If the IPsec connection is a split-tunnel configuration and the initiator negotiated INTERNAL_DNS_DOMAIN as per [\[RFC8598\]](#), the DNS client MUST resolve the internal names using ENCDNS_IP* DNS servers.

Note: [\[RFC8598\]](#) requires INTERNAL_IP6_DNS (or INTERNAL_IP4_DNS) attribute to be mandatory present when INTERNAL_DNS_DOMAIN is included. This specification relaxes that constraint in the presence of ENCDNS_IP* attributes.

5. Security Considerations

This document adheres to the security considerations defined in [\[RFC7296\]](#). In particular, this document does not alter the trust on the DNS configuration provided by a responder.

Networks are susceptible to internal attacks as discussed in Section 3.2 of [\[I-D.arkko-farrell-arch-model-t\]](#). Hosting encrypted DNS server even in case of split-VPN configuration minimizes the attack vector (e.g., a compromised network device cannot monitor/modify DNS traffic). This specification describes a mechanism to restrict access to the DNS messages to only the parties that need to know.

The initiator may trust the encrypted DNS servers supplied by means of IKEv2 from a trusted responder more than the locally provided DNS

servers, especially in the case of connecting to unknown or untrusted networks (e.g., coffee shops or hotel networks).

If IKEv2 responder has used NULL Authentication method [RFC7619] to authenticate itself, the initiator MUST NOT use returned ENCDNS_IP* servers configuration unless it is pre-configured in the OS or the browser.

This specification does not extend the scope of accepting DNSSEC trust anchors beyond the usage guidelines defined in [Section 6 of \[RFC8598\]](#).

6. IANA Considerations

6.1. Configuration Payload Attribute Types

This document requests IANA to assign the following new IKEv2 Configuration Payload Attribute Types from the "IKEv2 Configuration Payload Attribute Types" namespace available at <https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml#ikev2-parameters-21>.

Value	Attribute Type	Multi-Valued	Length	Reference
-----	-----	-----	-----	-----
TBA1	ENCDNS_IP4	YES	0 or more	RFC XXXX
TBA2	ENCDNS_IP6	YES	0 or more	RFC XXXX
TBA3	ENCDNS_ENCDNS_DIGEST_INFO	YES	0 or more	RFC XXXX

7. Acknowledgements

Many thanks to Yoav Nir, Christian Jacquenet, Paul Wouters, and Tommy Pauly for the review and comments.

Yoav and Paul suggested the use of one single attribute carrying both the name and an IP address instead of depending on the existing INTERNAL_IP6_DNS and INTERNAL_IP4_DNS attributes.

Christian Huitema suggested to return a port number in the attributes.

8. References

8.1. Normative References

[Hash] "IKEv2 Hash Algorithms",
<<https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml#hash-algorithms>>.

[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)", [draft-ietf-dnsop-svcb-https-08](#) (work in progress), October 2021.

[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>>.

[RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", [RFC 5890](#), DOI 10.17487/RFC5890, August 2010, <<https://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, <<https://www.rfc-editor.org/info/rfc7296>>.

[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>>.

[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>>.

8.2. Informative References

[I-D.arkko-farrell-arch-model-t]

Arkko, J. and S. Farrell, "Challenges and Changes in the Internet Threat Model", [draft-arkko-farrell-arch-model-t-04](#) (work in progress), July 2020.

[I-D.ietf-dprive-dnsquic]

Huitema, C., Dickinson, S., and A. Mankin, "DNS over Dedicated QUIC Connections", [draft-ietf-dprive-dnsquic-06](#) (work in progress), October 2021.

[RFC7619] Smyslov, V. and P. Wouters, "The NULL Authentication Method in the Internet Key Exchange Protocol Version 2 (IKEv2)", [RFC 7619](#), DOI 10.17487/RFC7619, August 2015, <<https://www.rfc-editor.org/info/rfc7619>>.

- [RFC7671] Dukhovni, V. and W. Hardaker, "The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance", [RFC 7671](#), DOI 10.17487/RFC7671, October 2015, <<https://www.rfc-editor.org/info/rfc7671>>.
- [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>>.
- [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>>.
- [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>>.

[Appendix A](#). Sample Deployment Scenarios

[A.1](#). Roaming Enterprise Users

In this Enterprise scenario ([Section 1.1.3 of \[RFC7296\]](#)), a roaming user connects to the Enterprise network through an IPsec tunnel. The split-tunnel Virtual Private Network (VPN) configuration allows the endpoint to access hosts that resides in the Enterprise network [\[RFC8598\]](#) using that tunnel; other traffic not destined to the Enterprise does not traverse the tunnel. In contrast, a non-split-tunnel VPN configuration causes all traffic to traverse the tunnel into the enterprise.

For both split- and non-split-tunnel configurations, the use of encrypted DNS instead of Do53 provides privacy and integrity protection along the entire path (rather than just to the VPN termination device) and can communicate the encrypted DNS server policies.

For split-tunnel VPN configurations, the endpoint uses the Enterprise-provided encrypted DNS server to resolve internal-only domain names.

For non-split-tunnel VPN configurations, the endpoint uses the Enterprise-provided encrypted DNS server to resolve both internal and external domain names.

Enterprise networks are susceptible to internal and external attacks. To minimize that risk all enterprise traffic is encrypted (Section 2.1 of [[I-D.arkko-farrell-arch-model-t](#)]).

[A.2.](#) VPN Service Provider

Legacy VPN service providers usually preserve end-users' data confidentiality by sending all communication traffic through an encrypted tunnel. A VPN service provider can also provide guarantees about the security of the VPN network by filtering malware and phishing domains.

Browsers and OSes support DoH/DoT; VPN providers may no longer expect DNS clients to fallback to Do53 just because it is a closed network.

The encrypted DNS server hosted by the VPN service provider can be securely discovered by the endpoint using the IKEv2 Configuration Payload Attribute Type.

[A.3.](#) DNS Offload

VPN service providers typically allow split-tunnel VPN configuration in which users can choose applications that can be excluded from the tunnel. For example, users may exclude applications that restrict VPN access.

The encrypted DNS server hosted by the VPN service provider can be securely discovered by the endpoint using the IKEv2 Configuration Payload Attribute Type.

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tirumaleswar Reddy
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

Valery Smyslov
ELVIS-PLUS
RU

Email: svan@elvis.ru

