

A Method for Storing IPsec Keying Material in DNS

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document describes a new resource record for the Domain Name System (DNS). This record may be used to store public keys for use in IP security (IPsec) systems. The record also includes provisions for indicating what system should be contacted when an IPsec tunnel is established with the entity in question.

This record replaces the functionality of the sub-type #4 of the KEY Resource Record, which has been obsoleted by RFC 3445.

Table of Contents

1.	Introduction	2
1.1.	Overview	2
1.2.	Use of DNS Address-to-Name Maps (IN-ADDR.ARPA and IP6.ARPA)	3
1.3.	Usage Criteria	3
2.	Storage Formats	3
2.1.	IPSECKEY RDATA Format	3
2.2.	RDATA Format - Precedence	4
2.3.	RDATA Format - Gateway Type	4
2.4.	RDATA Format - Algorithm Type	4
2.5.	RDATA Format - Gateway	5
2.6.	RDATA Format - Public Keys	5
3.	Presentation Formats	6
3.1.	Representation of IPSECKEY RRs	6
3.2.	Examples	6
4.	Security Considerations	7

4.1.	Active Attacks Against Unsecured IPSECKEY Resource Records	8
4.1.1.	Active Attacks Against IPSECKEY Keying Materials.	8
4.1.2.	Active Attacks Against IPSECKEY Gateway Material.	8
5.	IANA Considerations	9
6.	Acknowledgements	10
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	10
	Author's Address	11
	Full Copyright Statement	12

1. Introduction

Suppose a host wishes (or is required by policy) to establish an IPsec tunnel with some remote entity on the network prior to allowing normal communication to take place. In many cases, this end system will be able to determine the DNS name for the remote entity (either by having the DNS name given explicitly, by performing a DNS PTR query for a particular IP address, or through some other means, e.g., by extracting the DNS portion of a "user@FQDN" name for a remote entity). In these cases, the host will need to obtain a public key to authenticate the remote entity, and may also need some guidance about whether it should contact the entity directly or use another node as a gateway to the target entity. The IPSECKEY RR provides a mechanism for storing such information.

The type number for the IPSECKEY RR is 45.

This record replaces the functionality of the sub-type #4 of the KEY Resource Record, which has been obsoleted by RFC 3445 [11].

1.1. Overview

The IPSECKEY resource record (RR) is used to publish a public key that is to be associated with a Domain Name System (DNS) [1] name for use with the IPsec protocol suite. This can be the public key of a host, network, or application (in the case of per-port keying).

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 [3].

1.2. Use of DNS Address-to-Name Maps (IN-ADDR.ARPA and IP6.ARPA)

Often a security gateway will only have access to the IP address of the node with which communication is desired and will not know any other name for the target node. Because of this, frequently the best way of looking up IPSECKEY RRs will be by using the IP address as an index into one of the reverse mapping trees (IN-ADDR.ARPA for IPv4 or IP6.ARPA for IPv6).

The lookup is done in the fashion usual for PTR records. The IP address' octets (IPv4) or nibbles (IPv6) are reversed and looked up with the appropriate suffix. Any CNAMEs or DNAMEs found MUST be followed.

Note: even when the IPsec function is contained in the end-host, often only the application will know the forward name used. Although the case where the application knows the forward name is common, the user could easily have typed in a literal IP address. This storage mechanism does not preclude using the forward name when it is available but does not require it.

1.3. Usage Criteria

An IPSECKEY resource record SHOULD be used in combination with DNSSEC [8] unless some other means of authenticating the IPSECKEY resource record is available.


It is expected that there will often be multiple IPSECKEY resource records at the same name. This will be due to the presence of multiple gateways and a need to roll over keys.

This resource record is class independent.

2. Storage Formats

2.1. IPSECKEY RDATA Format

The RDATA for an IPSECKEY RR consists of a precedence value, a gateway type, a public key, algorithm type, and an optional gateway address.



2.2. RDATA Format - Precedence

This is an 8-bit precedence for this record. It is interpreted in the same way as the PREFERENCE field described in section 3.3.9 of RFC 1035 [2].

Gateways listed in IPSECKEY records with lower precedence are to be attempted first. Where there is a tie in precedence, the order should be non-deterministic.

2.3. RDATA Format - Gateway Type

The gateway type field indicates the format of the information that is stored in the gateway field.

The following values are defined:

- 0 No gateway is present.
- 1 A 4-byte IPv4 address is present.
- 2 A 16-byte IPv6 address is present.
- 3 A wire-encoded domain name is present. The wire-encoded format is self-describing, so the length is implicit. The domain name MUST NOT be compressed. (See Section 3.3 of RFC 1035 [2].)

2.4. RDATA Format - Algorithm Type

The algorithm type field identifies the public key's cryptographic algorithm and determines the format of the public key field.

A value of 0 indicates that no key is present.

The following values are defined:

- 1 A DSA key is present, in the format defined in RFC 2536 [9].
- 2 A RSA key is present, in the format defined in RFC 3110 [10].

2.5. RDATA Format - Gateway

The gateway field indicates a gateway to which an IPsec tunnel may be created in order to reach the entity named by this resource record.

There are three formats:

A 32-bit IPv4 address is present in the gateway field. The data portion is an IPv4 address as described in section 3.4.1 of RFC 1035 [2]. This is a 32-bit number in network byte order.

A 128-bit IPv6 address is present in the gateway field. The data portion is an IPv6 address as described in section 2.2 of RFC 3596 [12]. This is a 128-bit number in network byte order.

The gateway field is a normal wire-encoded domain name, as described in section 3.3 of RFC 1035 [2]. Compression MUST NOT be used.

2.6. RDATA Format - Public Keys

Both the public key types defined in this document (RSA and DSA) inherit their public key formats from the corresponding KEY RR formats. Specifically, the public key field contains the algorithm-specific portion of the KEY RR RDATA, which is all the KEY RR DATA after the first four octets. This is the same portion of the KEY RR that must be specified by documents that define a DNSSEC algorithm. Those documents also specify a message digest to be used for generation of SIG RRs; that specification is not relevant for IPSECKEY RRs.

Future algorithms, if they are to be used by both DNSSEC (in the KEY RR) and IPSECKEY, are likely to use the same public key encodings in both records. Unless otherwise specified, the IPSECKEY public key field will contain the algorithm-specific portion of the KEY RR RDATA for the corresponding algorithm. The algorithm must still be designated for use by IPSECKEY, and an IPSECKEY algorithm type number (which might be different from the DNSSEC algorithm number) must be assigned to it.

The DSA key format is defined in RFC 2536 [9]

The RSA key format is defined in RFC 3110 [10], with the following changes:

The earlier definition of RSA/MD5 in RFC 2065 [4] limited the exponent and modulus to 2552 bits in length. RFC 3110 extended that limit to 4096 bits for RSA/SHA1 keys. The IPSECKEY RR imposes no length limit on RSA public keys, other than the 65535 octet limit

imposed by the two-octet length encoding. This length extension is applicable only to IPSECKEY; it is not applicable to KEY RRs.

3. Presentation Formats

3.1. Representation of IPSECKEY RRs

IPSECKEY RRs may appear in a zone data master file. The precedence, gateway type, algorithm, and gateway fields are REQUIRED. The base64 encoded public key block is OPTIONAL; if it is not present, the public key field of the resource record MUST be construed to be zero octets in length.

The algorithm field is an unsigned integer. No mnemonics are defined.

If no gateway is to be indicated, then the gateway type field MUST be zero, and the gateway field MUST be "."

The Public Key field is represented as a Base64 encoding of the Public Key. Whitespace is allowed within the Base64 text. For a definition of Base64 encoding, see RFC 3548 [6], Section 5.2.

The general presentation for the record is as follows:

```
IN      IPSECKEY ( precedence gateway-type algorithm
                  gateway base64-encoded-public-key )
```

3.2. Examples

An example of a node, 192.0.2.38, that will accept IPsec tunnels on its own behalf.

```
38.2.0.192.in-addr.arpa. 7200 IN      IPSECKEY ( 10 1 2
                  192.0.2.38
                  AQNRU3mG7TVTO2BkR47usntb102uFJtugbo6BSGvgqt4AQ== )
```

An example of a node, 192.0.2.38, that has published its key only.

```
38.2.0.192.in-addr.arpa. 7200 IN      IPSECKEY ( 10 0 2
                  .
                  AQNRU3mG7TVTO2BkR47usntb102uFJtugbo6BSGvgqt4AQ== )
```

An example of a node, 192.0.2.38, that has delegated authority to the node 192.0.2.3.

```
38.2.0.192.in-addr.arpa. 7200 IN      IPSECKEY ( 10 1 2
                               192.0.2.3
                               AQNRU3mG7TVTO2BkR47usntb102uFJtugbo6BSGvgqt4AQ== )
```

An example of a node, 192.0.1.38 that has delegated authority to the node with the identity "mygateway.example.com".

```
38.1.0.192.in-addr.arpa. 7200 IN      IPSECKEY ( 10 3 2
                               mygateway.example.com.
                               AQNRU3mG7TVTO2BkR47usntb102uFJtugbo6BSGvgqt4AQ== )
```

An example of a node, 2001:0DB8:0200:1:210:f3ff:fe03:4d0, that has delegated authority to the node 2001:0DB8:c000:0200:2::1

```
$ORIGIN 1.0.0.0.0.0.2.8.B.D.0.1.0.0.2.ip6.arpa.
0.d.4.0.3.0.e.f.f.f.3.f.0.1.2.0 7200 IN      IPSECKEY ( 10 2 2
                               2001:0DB8:0:8002::2000:1
                               AQNRU3mG7TVTO2BkR47usntb102uFJtugbo6BSGvgqt4AQ== )
```

4. Security Considerations

This entire memo pertains to the provision of public keying material for use by key management protocols such as ISAKMP/IKE (RFC 2407) [7].

The IPSECKEY resource record contains information that SHOULD be communicated to the end client in an integral fashion; i.e., free from modification. The form of this channel is up to the consumer of the data; there must be a trust relationship between the end consumer of this resource record and the server. This relationship may be end-to-end DNSSEC validation, a TSIG or SIG(0) channel to another secure source, a secure local channel on the host, or some combination of the above.

The keying material provided by the IPSECKEY resource record is not sensitive to passive attacks. The keying material may be freely disclosed to any party without any impact on the security properties of the resulting IPsec session. IPsec and IKE provide defense against both active and passive attacks.

Any derivative specification that makes use of this resource record MUST carefully document its trust model and why the trust model of DNSSEC is appropriate, if that is the secure channel used.

An active attack on the DNS that caused the wrong IP address to be retrieved (via forged address), and therefore the wrong QNAME to be queried, would also result in a man-in-the-middle attack. This situation is independent of whether the IPSECKEY RR is used.

4.1. Active Attacks Against Unsecured IPSECKEY Resource Records

This section deals with active attacks against the DNS. These attacks require that DNS requests and responses be intercepted and changed. DNSSEC is designed to defend against attacks of this kind. This section deals with the situation in which DNSSEC is not available. This is not the recommended deployment scenario.

4.1.1. Active Attacks Against IPSECKEY Keying Materials

The first kind of active attack is when the attacker replaces the keying material with either a key under its control or with garbage.

The gateway field is either untouched or is null. The IKE negotiation will therefore occur with the original end-system. For this attack to succeed, the attacker must perform a man-in-the-middle attack on the IKE negotiation. This attack requires that the attacker be able to intercept and modify packets on the forwarding path for the IKE and data packets.

If the attacker is not able to perform this man-in-the-middle attack on the IKE negotiation, then a denial of service will result, as the IKE negotiation will fail.

If the attacker is not only able to mount active attacks against DNS but also in a position to perform a man-in-the-middle attack on IKE and IPsec negotiations, then the attacker will be able to compromise the resulting IPsec channel. Note that an attacker must be able to perform active DNS attacks on both sides of the IKE negotiation for this to succeed.

4.1.2. Active Attacks Against IPSECKEY Gateway Material

The second kind of active attack is one in which the attacker replaces the gateway address to point to a node under the attacker's control. The attacker then either replaces the public key or removes it. If the public key were removed, then the attacker could provide an accurate public key of its own in a second record.

This second form creates a simple man-in-the-middle attacks since the attacker can then create a second tunnel to the real destination. Note that, as before, this requires that the attacker also mount an active attack against the responder.

Note that the man-in-the-middle cannot just forward cleartext packets to the original destination. While the destination may be willing to speak in the clear, replying to the original sender, the sender will already have created a policy expecting ciphertext. Thus, the attacker will need to intercept traffic in both directions. In some cases, the attacker may be able to accomplish the full intercept by use of Network Address/Port Translation (NAT/NAPT) technology.

This attack is easier than the first one because the attacker does NOT need to be on the end-to-end forwarding path. The attacker need only be able to modify DNS replies. This can be done by packet modification, by various kinds of race attacks, or through methods that pollute DNS caches.

If the end-to-end integrity of the IPSECKEY RR is suspect, the end client MUST restrict its use of the IPSECKEY RR to cases where the RR owner name matches the content of the gateway field. As the RR owner name is assumed when the gateway field is null, a null gateway field is considered a match.

Thus, any records obtained under unverified conditions (e.g., no DNSSEC or trusted path to source) that have a non-null gateway field MUST be ignored.

This restriction eliminates attacks against the gateway field, which are considered much easier, as the attack does not need to be on the forwarding path.

In the case of an IPSECKEY RR with a value of three in its gateway type field, the gateway field contains a domain name. The subsequent query required to translate that name into an IP address or IPSECKEY RR will also be subject to man-in-the-middle attacks. If the end-to-end integrity of this second query is suspect, then the provisions above also apply. The IPSECKEY RR MUST be ignored whenever the resulting gateway does not match the QNAME of the original IPSECKEY RR query.

5. IANA Considerations

This document updates the IANA Registry for DNS Resource Record Types by assigning type 45 to the IPSECKEY record.

This document creates two new IANA registries, both specific to the IPSECKEY Resource Record:

This document creates an IANA registry for the algorithm type field.

Values 0, 1, and 2 are defined in Section 2.4. Algorithm numbers 3 through 255 can be assigned by IETF Consensus (see RFC 2434 [5]).

This document creates an IANA registry for the gateway type field.

Values 0, 1, 2, and 3 are defined in Section 2.3. Gateway type numbers 4 through 255 can be assigned by Standards Action (see RFC 2434 [5]).

6. Acknowledgements

My thanks to Paul Hoffman, Sam Weiler, Jean-Jacques Puig, Rob Austein, and Olafur Gudmundsson, who reviewed this document carefully. Additional thanks to Olafur Gurmundsson for a reference implementation.

7. References

7.1. Normative References

- [1] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [2] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [4] Eastlake 3rd, D. and C. Kaufman, "Domain Name System Security Extensions", RFC 2065, January 1997.
- [5] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [6] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 3548, July 2003.

7.2. Informative References

- [7] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.
- [8] Eastlake 3rd, D., "Domain Name System Security Extensions", RFC 2535, March 1999.
- [9] Eastlake 3rd, D., "DSA KEYS and SIGs in the Domain Name System (DNS)", RFC 2536, March 1999.

- [10] Eastlake 3rd, D., "RSA/SHA-1 SIGs and RSA KEYS in the Domain Name System (DNS)", RFC 3110, May 2001.
- [11] Massey, D. and S. Rose, "Limiting the Scope of the KEY Resource Record (RR)", RFC 3445, December 2002.
- [12] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.

Author's Address

Michael C. Richardson
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z 5V7
CA

EMail: mcr@sandelman.ottawa.on.ca
URI: <http://www.sandelman.ottawa.on.ca/>

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the IETF's procedures with respect to rights in IETF Documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

