Internet Engineering Task Force (IETF)
Request for Comments: 8005
Obsoletes: 5205
Category: Standards Track
ISSN: 2070-1721

J. Laganier
Luminate Wireless, Inc.

October 2016

Host Identity Protocol (HIP) Domain Name System (DNS) Extension

Abstract

This document specifies a resource record (RR) for the Domain Name System (DNS) and how to use it with the Host Identity Protocol (HIP). This RR allows a HIP node to store in the DNS its Host Identity (HI), the public component of the node public-private key pair; its Host Identity Tag (HIT), a truncated hash of its public key (PK); and the domain names of its rendezvous servers (RVSs). This document obsoletes RFC 5205.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc8005.

Table of Contents

1. Introduction

This document specifies a resource record (RR) for the Domain Name System (DNS) [RFC1034] and how to use it with the Host Identity Protocol (HIP) [RFC7401]. This RR allows a HIP node to store in the DNS its Host Identity (HI), the public component of the node public-private key pair; its Host Identity Tag (HIT), a truncated hash of its HI; and the domain names of its rendezvous servers (RVSs) [RFC8004].

Currently, most of the Internet applications that need to communicate with a remote host first translate a domain name (often obtained via user input) into one or more IP addresses. This step occurs prior to communication with the remote host and relies on a DNS lookup.

With HIP, IP addresses are intended to be used mostly for on-the-wire communication between end hosts, while most Upper Layer Protocols (ULPs) and applications use HIs or HITs instead (ICMP might be an example of a ULP not using them). Consequently, we need a means to translate a domain name into an HI. Using the DNS for this translation is pretty straightforward: We define a HIP RR. Upon query by an application or ULP for a name-to-IP-address lookup, the resolver would then additionally perform a name-to-HI lookup and use it to construct the resulting HI-to-IP-address mapping (which is internal to the HIP layer). The HIP layer uses the HI-to-IP-address mapping to translate HIs and HITs into IP addresses, and vice versa.

The HIP specification [RFC7401] specifies the HIP base exchange between a HIP Initiator and a HIP Responder based on a four-way handshake involving a total of four HIP packets (I1, R1, I2, and R2). Since the HIP packets contain both the Initiator and the Responder HIT, the Initiator needs to have knowledge of the Responder's HI and HIT prior to initiating the base exchange by sending an I1 packet.

The HIP Rendezvous Extension [RFC8004] allows a HIP node to be reached via the IP address(es) of a third party, the node's RVS. An Initiator willing to establish a HIP association with a Responder served by an RVS would typically initiate a HIP base exchange by sending the I1 packet initiating the exchange towards the RVS IP address rather than towards the Responder IP address. Consequently, we need a means to find the name of an RVS for a given host name.

This document introduces the HIP DNS RR to store the RVS, HI, and HIT information.

2. Conventions Used in This Document

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].

3. Usage Scenarios

In this section, we briefly introduce a number of usage scenarios where the DNS is useful with HIP.

With HIP, most applications and ULPs are unaware of the IP addresses used to carry packets on the wire. Consequently, a HIP node could take advantage of having multiple IP addresses for failover, redundancy, mobility, or renumbering, in a manner that is transparent to most ULPs and applications (because they are bound to HIs; hence, they are agnostic to these IP address changes).

In these situations, for a node to be reachable by reference to its Fully Qualified Domain Name (FQDN), the following information should be stored in the DNS:

- A set of IP addresses via A [RFC1035] and AAAA [RFC3596] Resource Record Sets (RRSets) [RFC2181].

- An HI, a HIT, and possibly a set of RVSs through HIP RRs.

The HIP RR is class independent.

When a HIP node wants to initiate communication with another HIP node, it first needs to perform a HIP base exchange to set up a HIP association towards its peer. Although such an exchange can be initiated opportunistically, i.e., without prior knowledge of the Responder's HI, by doing so both nodes knowingly risk man-in-the-middle (MitM) attacks on the HIP exchange. To prevent these attacks, it is recommended that the Initiator first obtains the HI of the Responder and then initiates the exchange. This can be done, for example, through manual configuration or DNS lookups. Hence, a HIP RR is introduced.

When a HIP node is frequently changing its IP address(es), the natural DNS latency for propagating changes may prevent it from publishing its new IP address(es) in the DNS. For solving this problem, the HIP Architecture [RFC4423] introduces RVSs [RFC8004]. A HIP host uses an RVS as a rendezvous point to maintain reachability with possible HIP Initiators while moving [RFC5206]. Such a HIP node would publish in the DNS its RVS domain name(s) in a HIP RR, while keeping its RVS up-to-date with its current set of IP addresses.

When a HIP node wants to initiate a HIP exchange with a Responder, it will perform a number of DNS lookups. Depending on the type of implementation, the order in which those lookups will be issued may vary. For instance, implementations using HIT in Application Programming Interfaces (APIs) may typically first query for HIP RRs at the Responder FQDN, while those using an IP address in APIs may typically first query for A and/or AAAA RRs.

In the following, we assume that the Initiator first queries for HIP RRs at the Responder FQDN.

If the query for the HIP type was responded to with a DNS answer with RCODE=3 (Name Error), then the Responder's information is not present in the DNS, and further queries for the same owner name SHOULD NOT be made.

In case the query for the HIP records returned a DNS answer with RCODE=0 (No Error) and an empty answer section, it means that no HIP information is available at the Responder name. In such a case, if the Initiator has been configured with a policy to fall back to opportunistic HIP (initiating without knowing the Responder's HI) or plain IP, it would send out more queries for A and AAAA types at the Responder's FQDN.

Depending on the combinations of answers, the situations described in Sections 3.1 and 3.2 can occur.

Note that storing HIP RR information in the DNS at an FQDN that is assigned to a non-HIP node might have ill effects on its reachability by HIP nodes.

3.1. Simple Static Single-Homed End Host

In addition to its IP address or addresses (IP-R), a HIP node (R) with a single static network attachment that wishes to be reachable by reference to its FQDN (www.example.com) to act as a Responder would store in the DNS a HIP RR containing its Host Identity (HI-R) and Host Identity Tag (HIT-R).

An Initiator willing to associate with a node would typically issue the following queries:

- Query #1: QNAME=www.example.com, QTYPE=HIP

(QCLASS=IN is assumed and omitted from the examples)

Which returns a DNS packet with RCODE=0 and one or more HIP RRs with the HIT and HI (e.g., HIT-R and HI-R) of the Responder in the answer section, but no RVS.

- Query #2: QNAME=www.example.com, QTYPE=A

- Query #3: QNAME=www.example.com, QTYPE=AAAA

Which would return DNS packets with RCODE=0 and, respectively, one or more A or AAAA RRs containing the IP address(es) of the Responder (e.g., IP-R) in their answer sections.

Caption: In the remainder of this document, for the sake of keeping diagrams simple and concise, several DNS queries and answers are represented as one single transaction, while in fact there are several queries and answers flowing back and forth, as described in the textual examples.

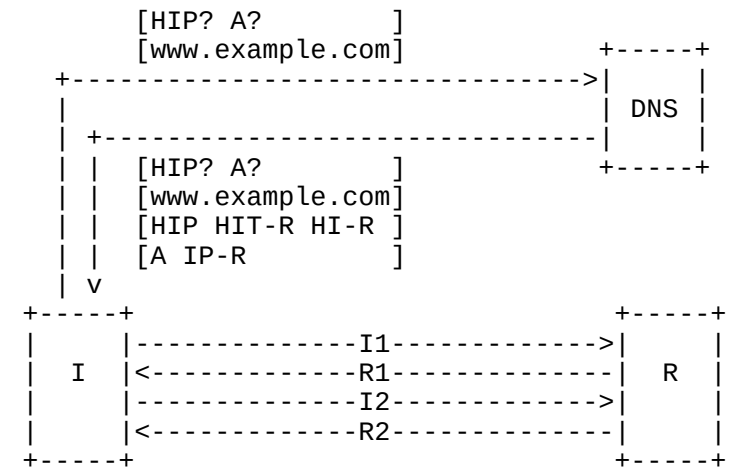

Static Single-Homed Host

The Initiator would then send an I1 to the Responder's IP addresses (IP-R).

3.2. Mobile End Host

A mobile HIP node (R) wishing to be reachable by reference to its FQDN (www.example.com) would store in the DNS, possibly in addition to its IP address or addresses (IP-R), its HI (HI-R), its HIT (HIT-R), and the domain name or names of its RVS or servers (e.g., rvs.example.com) in a HIP RR or records. The mobile HIP node also needs to notify its RVSs of any change in its set of IP addresses.

An Initiator willing to associate with such a mobile node would typically issue the following queries:

- Query #1: QNAME=www.example.com, QTYPE=HIP

Which returns a DNS packet with RCODE=0 and one or more HIP RRs with the HIT, HI, and RVS domain name or names (e.g., HIT-R, HI-R, and rvs.example.com) of the Responder in the answer section.

- Query #2: QNAME=rvs.example.com, QTYPE=A

- Query #3: QNAME=rvs.example.com, QTYPE=AAAA

Which return DNS packets with RCODE=0 and, respectively, one or more A or AAAA RRs containing an IP address(es) of the Responder's RVS (e.g., IP-RVS) in their answer sections.

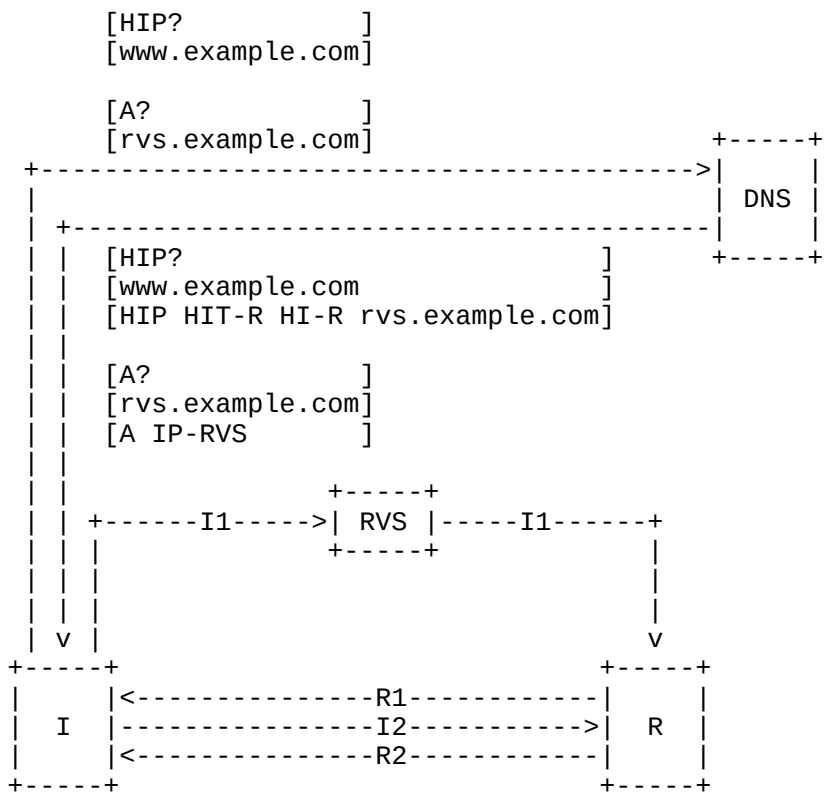

Mobile End Host

The Initiator would then send an I1 to the RVS IP address (IP-RVS). Following, the RVS will relay the I1 up to the mobile node's IP address (IP-R), which will complete the HIP exchange.

4. Overview of Using the DNS with HIP

4.1. Storing HI, HIT, and RVS in the DNS

For any HIP node, its HI, the associated HIT, and the FQDN of its possible RVSs can be stored in a DNS HIP RR. Any conforming implementation may store an HI and its associated HIT in a DNS HIP RDATA format. HI and HIT are defined in Section 3 of the HIP specification [RFC7401].

Upon return of a HIP RR, a host MUST always calculate the HI-derivative HIT to be used in the HIP exchange, as specified in Section 3 of the HIP specification [RFC7401], while the HIT included in the HIP RR SHOULD only be used as an optimization (e.g., table lookup).

The HIP RR may also contain one or more domain names of one or more RVSs towards which HIP I1 packets might be sent to trigger the establishment of an association with the entity named by this RR [RFC8004].

The Rendezvous Server field of the HIP RR stored at a given owner name MAY include the owner name itself. A semantically equivalent situation occurs if no RVS is present in the HIP RR stored at that owner name. Such situations occur in two cases:

- The host is mobile, and the A and/or AAAA RR(s) stored at its host name contain the IP address(es) of its RVS rather than its own one.
- The host is stationary and can be reached directly at the IP address(es) contained in the A and/or AAAA RR(s) stored at its host name. This is a degenerate case of rendezvous service where the host somewhat acts as an RVS for itself.

An RVS receiving such an I1 would then relay it to the appropriate Responder (the owner of the I1 receiver HIT). The Responder will then complete the exchange with the Initiator, typically without ongoing help from the RVS.

4.2. Initiating Connections Based on DNS Names

On a HIP node, a HIP exchange SHOULD be initiated whenever a ULP attempts to communicate with an entity, and the DNS lookup returns HIP RRs.

HIP RRs have a Time To Live (TTL) associated with them. When the number of seconds that passed since the record was retrieved exceeds the record’s TTL, the record MUST be considered no longer valid and deleted by the entity that retrieved it. If access to the record is necessary to initiate communication with the entity to which the record corresponds, a new query MUST be made to retrieve a fresh copy of the record.

There may be multiple HIP RRs associated with a single name. It is outside the scope of this specification as to how a host chooses between multiple RRs when more than one is returned. The RVS

information may be copied and aligned across multiple RRs, or may be different for each one; a host MUST check that the RVS used is associated with the HI being used, when multiple choices are present.

5. HIP RR Storage Format

The RDATA for a HIP RR consists of a PK Algorithm Type, the HIT length, a HIT, a PK (i.e., an HI), and optionally one or more RVSs.

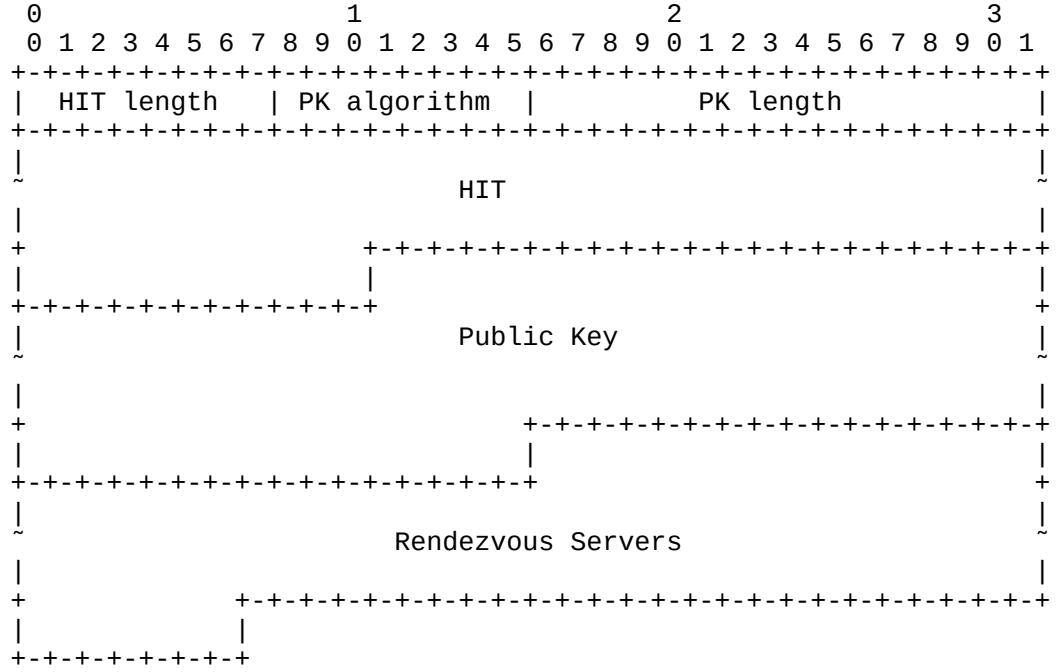

The HIT length, PK algorithm, PK length, HIT, and Public Key fields are REQUIRED. The Rendezvous Server field is OPTIONAL.

5.1. HIT Length Format

The HIT length indicates the length in bytes of the HIT field. This is an 8-bit unsigned integer.

5.2. PK Algorithm Format

The PK algorithm field indicates the PK cryptographic algorithm and the implied Public Key field format. This is an 8-bit unsigned integer. This document reuses the values defined for the 'Algorithm Type' of the IPSECKEY RR [RFC4025].

Presently defined values are listed in Section 9 for reference.

5.3. PK Length Format

The PK length indicates the length in bytes of the Public Key field. This is a 16-bit unsigned integer in network byte order.

5.4. HIT Format

The HIT is stored as a binary value in network byte order.

5.5. Public Key Format

Two of the PK types defined in this document (RSA and Digital Signature Algorithm (DSA)) reuse the PK formats defined for the IPSECKEY RR [RFC4025].

The DSA key format is defined in RFC 2536 [RFC2536].

The RSA key format is defined in RFC 3110 [RFC3110], and the RSA key size limit (4096 bits) is relaxed in the IPSECKEY RR [RFC4025] specification.

In addition, this document similarly defines the PK format of type Elliptic Curve Digital Signature Algorithm (ECDSA) as the algorithm-specific portion of the DNSKEY RR RDATA for ECDSA [RFC6605], i.e, all of the DNSKEY RR DATA after the first four octets, corresponding to the same portion of the DNSKEY RR that must be specified by documents that define a DNSSEC algorithm.

5.6. Rendezvous Servers Format

The Rendezvous Server field indicates one or more variable length wire-encoded domain names of one or more RVSs, concatenated and encoded as described in Section 3.3 of RFC 1035 [RFC1035]: "<domain-name> is a domain name represented as a series of labels, and terminated by a label with zero length". Since the wire-encoded format is self-describing, the length of each domain name is implicit: The zero length label termination serves as a separator between multiple RVS domain names concatenated in the Rendezvous Server field of a same HIP RR. Since the length of the other portion of the RR's RRDATA is known, and the overall length of the RR's RDATA is also known (RDLENGTH), all the length information necessary to parse the HIP RR is available.

The domain names MUST NOT be compressed. The RVS or servers are listed in order of preference (i.e., the first RVS or servers are preferred), defining an implicit order amongst RVSs of a single RR.

When multiple HIP RRs are present at the same owner name, this implicit order of RVSs within an RR MUST NOT be used to infer a preference order between RVSs stored in different RRs.

6. HIP RR Presentation Format

This section specifies the representation of the HIP RR in a zone master file.

The HIT length field is not represented, as it is implicitly known thanks to the HIT field representation.

The PK algorithm field is represented as unsigned integers.

The HIT field is represented as the Base16 encoding [RFC4648] (a.k.a. hex or hexadecimal) of the HIT. The encoding MUST NOT contain whitespaces to distinguish it from the Public Key field.

The Public Key field is represented as the Base64 encoding of the PK, as defined in Section 4 of [RFC4648]. The encoding MUST NOT contain whitespace(s) to distinguish it from the Rendezvous Server field.

The PK length field is not represented, as it is implicitly known thanks to the Public Key field representation containing no whitespaces.

The Rendezvous Server field is represented by one or more domain names separated by whitespace(s). Such whitespace is only used in the HIP RR representation format and is not part of the HIP RR wire format.

The complete representation of the HIP record is:

```
IN  HIP   ( pk-algorithm
            base16-encoded-hit
            base64-encoded-public-key
            rendezvous-server[1]
                    ...
            rendezvous-server[n] )
```

When no RVSs are present, the representation of the HIP record is:

```
IN  HIP   ( pk-algorithm
            base16-encoded-hit
            base64-encoded-public-key )
```

7. Examples

In the examples below, the Public Key field containing no whitespace is wrapped, since it does not fit in a single line of this document.

Example of a node with an HI and a HIT but no RVS:

```
www.example.com.      IN  HIP ( 2 200100107B1A74DF365639CC39F1D578
                                AwEAAbdxyhNuSutc5EMzxTs9LBPCIkOFH8cI
vM4p9+LrV4e19WzK00+CI6zBCQTdtWsuxKbWIy87UOoJTwkUs7lBu+Upr1gsNrut79ry
ra+bSRGQb1slImA8YVJyuIDsj7kwzG7jnERNqnWxZ48AWkskmdHaVDP4BcelrTI3rMXd
XF5D )
```

Example of a node with an HI, a HIT, and one RVS:

```
www.example.com.      IN  HIP ( 2 200100107B1A74DF365639CC39F1D578
                                AwEAAbdxyhNuSutc5EMzxTs9LBPCIkOFH8cI
vM4p9+LrV4e19WzK00+CI6zBCQTdtWsuxKbWIy87UOoJTwkUs7lBu+Upr1gsNrut79ry
ra+bSRGQb1slImA8YVJyuIDsj7kwzG7jnERNqnWxZ48AWkskmdHaVDP4BcelrTI3rMXd
XF5D
                                rvs.example.com. )
```

Example of a node with an HI, a HIT, and two RVSs:

```
www.example.com.      IN  HIP ( 2 200100107B1A74DF365639CC39F1D578
                                AwEAAbdxyhNuSutc5EMzxTs9LBPCIkOFH8cI
vM4p9+LrV4e19WzK00+CI6zBCQTdtWsuxKbWIy87UOoJTwkUs7lBu+Upr1gsNrut79ry
ra+bSRGQb1slImA8YVJyuIDsj7kwzG7jnERNqnWxZ48AWkskmdHaVDP4BcelrTI3rMXd
XF5D
                                rvs1.example.com.
                                rvs2.example.com. )
```

8. Security Considerations

This section contains a description of the known threats involved with the usage of the HIP DNS Extension.

In a manner similar to the IPSECKEY RR [RFC4025], the HIP DNS Extension allows for the provision of two HIP nodes with the public keying material (HI) of their peer. These HIs will be subsequently used in a key exchange between the peers. Hence, the HIP DNS Extension is subject, as the IPSECKEY RR, to threats stemming from attacks against unsecured HIP RRs, as described in the remainder of this section.

A HIP node SHOULD obtain HIP RRs from a trusted party through a secure channel ensuring data integrity and authenticity of the RRs. DNSSEC [RFC4033] [RFC4034] [RFC4035] provides such a secure channel. However, it should be emphasized that DNSSEC only offers data integrity and authenticity guarantees to the channel between the DNS server publishing a zone and the HIP node. DNSSEC does not ensure that the entity publishing the zone is trusted. Therefore, the RRSIG of the HIP RRSet MUST NOT be misinterpreted as a certificate binding the HI and/or the HIT to the owner name.

In the absence of a proper secure channel, both parties are vulnerable to MitM and Denial-of-Service (DoS) attacks, and unrelated parties might be subject to DoS attacks as well. These threats are described in the following sections.

8.1. Attacker Tampering with an Insecure HIP RR

The HIP RR contains public keying material in the form of the named peer's PK (the HI) and its secure hash (the HIT). Both of these are not sensitive to attacks where an adversary gains knowledge of them. However, an attacker that is able to mount an active attack on the DNS, i.e., tampers with this HIP RR (e.g., using DNS spoofing), is able to mount MitM attacks on the cryptographic core of the eventual HIP exchange (Responder's HIP RR rewritten by the attacker).

The HIP RR may contain an RVS domain name resolved into a destination IP address where the named peer is reachable by an I1, as per the HIP Rendezvous Extension [RFC8004]. Thus, an attacker that is able to tamper with this RR is able to redirect I1 packets sent to the named peer to a chosen IP address for DoS or MitM attacks. Note that this kind of attack is not specific to HIP and exists independently of whether or not HIP and the HIP RR are used. Such an attacker might tamper with A and AAAA RRs as well.

An attacker might obviously use these two attacks in conjunction: It will replace the Responder's HI and RVS IP address by its own in a spoofed DNS packet sent to the Initiator HI, and then redirect all exchanged packets to him and mount a MitM on HIP. In this case, HIP won't provide confidentiality nor Initiator HI protection from eavesdroppers.

8.2. Hash and HITs Collisions

As with many cryptographic algorithms, some secure hashes (e.g., SHA1, used by HIP to generate a HIT from an HI) eventually become insecure, because an exploit has been found in which an attacker with reasonable computation power breaks one of the security features of the hash (e.g., its supposed collision resistance). This is why a

HIP end-node implementation SHOULD NOT authenticate its HIP peers based solely on a HIT retrieved from the DNS, but rather SHOULD use HI-based authentication.

8.3. DNSSEC

In the absence of DNSSEC, the HIP RR is subject to the threats described in RFC 3833 [RFC3833].

9. IANA Considerations

[RFC5205], obsoleted by this document, made the following definition and reservation in the "Resource Record (RR) TYPEs" subregistry under "Domain Name System (DNS) Parameters":

Value	Type
55	HIP

In the "Resource Record (RR) TYPEs" subregistry under "Domain Name System (DNS) Parameters", references to [RFC5205] have been replaced by references to this document.

As [RFC5205], this document reuses the Algorithm Types defined by [RFC4025] for the IPSEC KEY RR. Presently defined values are shown here for reference only:

Value	Description
1	A DSA key is present, in the format defined in [RFC2536]
2	A RSA key is present, in the format defined in [RFC3110]

IANA has made the following assignment in the "Algorithm Type Field" subregistry under "IPSECKEY Resource Record Parameters" [RFC4025]:

Value	Description
3	An ECDSA key is present, in the format defined in [RFC6605]

10. References

10.1. Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <http://www.rfc-editor.org/info/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <http://www.rfc-editor.org/info/rfc1035>.

[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>.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <http://www.rfc-editor.org/info/rfc2181>.

[RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, DOI 10.17487/RFC3596, October 2003, <http://www.rfc-editor.org/info/rfc3596>.

[RFC4025] Richardson, M., "A Method for Storing IPsec Keying Material in DNS", RFC 4025, DOI 10.17487/RFC4025, March 2005, <http://www.rfc-editor.org/info/rfc4025>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <http://www.rfc-editor.org/info/rfc4033>.

[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>.

[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, <http://www.rfc-editor.org/info/rfc4035>.

[RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <http://www.rfc-editor.org/info/rfc4648>.

[RFC6605] Hoffman, P. and W. Wijngaards, "Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC", RFC 6605, DOI 10.17487/RFC6605, April 2012, <http://www.rfc-editor.org/info/rfc6605>.

[RFC7401] Moskowitz, R., Ed., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, DOI 10.17487/RFC7401, April 2015, <http://www.rfc-editor.org/info/rfc7401>.

[RFC8004] Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extension", RFC 8004, DOI 10.17487/RFC8004, October 2016, <http://www.rfc-editor.org/info/rfc8004>.

10.2. Informative References

[RFC2536] Eastlake 3rd, D., "DSA KEYs and SIGs in the Domain Name System (DNS)", RFC 2536, DOI 10.17487/RFC2536, March 1999, <http://www.rfc-editor.org/info/rfc2536>.

[RFC3110] Eastlake 3rd, D., "RSA/SHA-1 SIGs and RSA KEYs in the Domain Name System (DNS)", RFC 3110, DOI 10.17487/RFC3110, May 2001, <http://www.rfc-editor.org/info/rfc3110>.

[RFC3833] Atkins, D. and R. Austein, "Threat Analysis of the Domain Name System (DNS)", RFC 3833, DOI 10.17487/RFC3833, August 2004, <http://www.rfc-editor.org/info/rfc3833>.

[RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, DOI 10.17487/RFC4423, May 2006, <http://www.rfc-editor.org/info/rfc4423>.

[RFC5205] Nikander, P. and J. Laganier, "Host Identity Protocol (HIP) Domain Name System (DNS) Extensions", RFC 5205, DOI 10.17487/RFC5205, April 2008, <http://www.rfc-editor.org/info/rfc5205>.

[RFC5206] Nikander, P., Henderson, T., Ed., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol", RFC 5206, DOI 10.17487/RFC5206, April 2008, <http://www.rfc-editor.org/info/rfc5206>.

Appendix A. Changes from RFC 5205

- Updated HIP references to revised HIP specifications.
- Extended DNS HIP RR to support for Host Identities based on ECDSA.
- Clarified that new query must be made when the time that passed since an RR was retrieved exceeds the TTL of the RR.
- Added considerations related to multiple HIP RRs being associated with a single name.
- Clarified that the Base64 encoding in use is as per Section 4 of [RFC4648].
- Clarified the wire format when more than one RVS is defined in one RR.
- Clarified that "whitespace" is used as the delimiter in the human-readable representation of the RR but is not part of the wire format.

Acknowledgments

As usual in the IETF, this document is the result of a collaboration between many people. The authors would like to thank the author (Michael Richardson), contributors, and reviewers of the IPSECKEY RR [RFC4025] specification, after which this document was framed. The authors would also like to thank the following people, who have provided thoughtful and helpful discussions and/or suggestions, that have helped improve this document: Jeff Ahrenholz, Rob Austein, Hannu Flinck, Olafur Gudmundsson, Tom Henderson, Peter Koch, Olaf Kolkman, Miika Komu, Andrew McGregor, Gabriel Montenegro, and Erik Nordmark. Some parts of this document stem from the HIP specification [RFC7401]. Finally, thanks to Sheng Jiang for performing the Internet Area Directorate review of this document in the course of the publication process.

Contributors

Pekka Nikander coauthored an earlier, experimental version of this specification [RFC5205].

Author's Address

Julien Laganier
Luminate Wireless, Inc.
Cupertino, CA
United States of America

Email: julien.ietf@gmail.com