HIP R. Moskowitz
Internet-Draft X. Xu
Intended status: Standards Track Huawei
Expires: February 4, 2017 August 3, 2016

# Hierarchical HITs for HIPv2
## draft-moskowitz-hierarchical-hip-00.txt

Abstract

This document describes using a hierarchical HIT to facilitate large deployments in mobile networks.

Table of Contents



## 1. Introduction

This document expands on HIPv2 [RFC7401] to describe the structure of a hierarchical HIT, the Registry services to support this hierarchy, and given a hierarchical HIT, how a peer is found in the network.

A separate document will further expand on the registry service and how a device can advertise its availability and services provided.

## 2. Terms and Definitions

### 2.1. Requirements Terminology

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].

### 2.2. Definitions

HDA (Hierarchical HIT Domain Authority): The 14 bit field identifying the HIT Domain Authority under a RAA.

HID (Hierarchy ID): The 32 bit field providing the HIT Hierarchy ID.

RAA (Registered Assigning Authority): The 18 bit field identifying the Hierarchical HIT Assigning Authority.

## 3. Problem Space

### 3.1. Managing a large flat address space

For HIP to be successfully used in mobile networks, it must support an Identity per person, or upwards to 10 billion Identities. Perhaps a Distributed Hash Table [I-D.irtf-hiprg-dht] can scale this large. There is still the operational challenges in establishing such a world-wide DHT implementation and how RVS [I-D.ietf-hip-rfc5204-bis] works with such a large population.

Even though the probability of collisions with 7B HITs in a 96 bit flat address space is 3.9E-10, it is still real. How are collisions managed? It is also possible with weak key uniqueness, as has been shown in deployed TLS certificates, results in a much greater probability of collisions. Thus resolution of collisions needs to be a feature in a globally mobile network.

### 3.2. Desire for administrative control by RVS providers

An RVS provider may only what to provide discovery services to HIP clients it knows and trusts. A flat HIT space does not provide any intrinsic functionality to support this. A hierarchical HIT space can be mapped to the RVS provider.

## 4. The Hierarchical Host Identity Tag (HIT)

The Hierarchical HIT is a small but important enhancement over the flat HIT space. It represents the HI in only a 64 bit hash and uses the other 32 bits to create a hierarchical administration

organization for HIT domains. Hierarchical HITs are ORCHIDs [RFC7343]. The change in construction rules are in Section 4.1.4.

A Hierarchical HIT is built from the following fields:

- o 28 bit IANA prefix
- o 4 bit HIT Suite ID
- o 32 bit Hierarchy ID (HID)
- o 64 bit ORCHID hash

## 4.1. The Hierarchy ID (HID)

The Hierarchy ID (HID) provides the structure to organize HITs into administrative domains. HIDs are further divided into 2 fields:

- o 14 bit Registered Assigning Authority (RAA)
- o 18 bit Hierarchical HIT Domain Authority (HDA)

### 4.1.1. The Registered Assigning Authority (RAA)

The RAA is a 14 bit field (16,384 RAAs) assigned sequentially by a numbers management organization, perhaps ICANN. An RAA must provide a set of services to allocate HDAs to organizations. It must have a public policy on what is necessary to obtain an HDA. The RAA need not maintain any HIP related services. It must maintain a DNS zone for discovering HID RVS servers.

This DNS zone may be a reverse PTR for its RAA. Assume that the RAA is 100. The PTR record is constructed at a 2 bit grouping:

```
1.3.1.0.0.0.0.arpa   IN PTR      raa.bar.com.
```

### 4.1.2. The Hierarchical HIT Domain Authority (HDA)

The HDA is an 18 bit field (262,144 HDAs per RAA) assigned sequentially by an RAA. An HDA should maintain a set of RVS servers that its client HIP-enabled customers use. How this is done and scales to the potentially millions of customers is outside the scope of this document. This service should be discoverable through the DNS zone maintained by the HDA's RAA.

An RAA may assign a block of values to an individual organization. This is completely up to the individual RAA's published policy for delegation.

#### 4.1.3. Example of the HID DNS

HID related services should be discoverable via DNS. For example the RVS for a HID could be found via the following. Assume that the RAA is 100 and the HDA is 50. The PTR record is constructed at a 2 bit grouping:

```
2.0.3.0.0.0.0.0.0.1.3.1.0.0.0.0.arpa   IN PTR      rvs.foo.com.
```

#### 4.1.4. Changes to ORCHIDv2 to support Hierarchical HITs

ORCHIDv2 has a number of inputs including a context, some header bits, the hash algorithm, and the public key. The output is a 96 bit value. Hierarchical HIT makes the following changes. The HID is added as part of the header bits and the output is a 64 bit value, derived the same way as the 96 bit hash.

Hierarchical HIT uses the same context as all other HIPv2 HIT Suites as they are clearly separated by the distinct HIT Suite ID.

#### 4.1.5. Collision risks with Hierarchical HITs

The 64 bit hash size does have an increased risk of collisions over the 96 bit hash size used for the other HIT Suites. There is a 0.01% probability of a collision in a population of 66 million. The probability goes up to 1% for a population of 663 million. See Appendix A for the collision probability formula.

This risk, however, is within a single HDA. Further, all HDAs are expected to provide a registration process for reverse lookup validation. This registration process would reject a collision, forcing the client to generate a new HI and thus hierarchical HIT and reapplying to the registration process.

## 5. HIP Parameters

The HIP parameters carry information that is necessary for establishing and maintaining a HIP association. For example, the peer's public keys as well as the signaling for negotiating ciphers and payload handling are encapsulated in HIP parameters. Additional information, meaningful for end hosts or middleboxes, may also be included in HIP parameters. The specification of the HIP parameters and their mapping to HIP packets and packet types is flexible to allow HIP extensions to define new parameters and new protocol behavior.

## 5.1. HIT_SUITE_LIST

The HIT_SUITE_LIST parameter contains a list of the supported HIT suite IDs of the Responder. Based on the HIT_SUITE_LIST, the Initiator can determine which source HIT Suite IDs are supported by the Responder. The HIT_SUITE_LIST parameter is defined in Section 5.2.10 of [RFC7401].

The following HIT Suite IDs are defined for Hierarchical HITs, and the relationship between the four-bit ID value used in the OGA ID field and the eight-bit encoding within the HIT_SUITE_LIST ID field is clarified:

| HIT Suite | Four-bit ID | Eight-bit encoding |
|---|---|---|
| ECDSA/hier/SHA-256 | 4 | 0x40 |

Note that the Hierarchical HIP HIT Suite ID allows the peers to use the hierarchical RVS discovery and authentication services to validate the peer and discover available services. The Responder SHOULD respond with a HIP hierarchical HIT suite ID when the HIT of the Initiator is a HIP hierarchical HIT.

## 5.2. CLIENT_INFO

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
/                      Client Information                       /
/                                                               /
/                               +-------------------------------+
/                               |            Padding            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

| | |
|---|---|
| Type | [TBD-IANA] |
| Length | length in octets, excluding Type, Length, and Padding |
| Client Information | The information required by the HDA in the format required by the HDA. |

This parameter contains client information to include in the HIT registration. The specific content and format is set by the HDA.

## 6. HIT Registry services to support hierarchical HITs

Hierarchical HIT registration SHOULD be performed using the HIP Registration Extension [I-D.ietf-hip-rfc5203-bis]. The client either uses an X.509 certificate [I-D.ietf-hip-rfc6253-bis], or use a PSK, as defined in Appendix A of HIP-DEX [I-D.ietf-hip-dex], to validate the registration.

The Registration should include additional client information. This information may be contained within the X.509 certificate and/or is carried in the CLIENT_INFO parameter, see Section 5.2. The Registrar can include its requirements in the R1 packet, and the client provide its information in the I2 packet. This parameter may be encrypted within the ENCRYPTED parameter. If the CLIENT_INFO contains Personal Identifying Information (PII), then it MUST be placed into the ENCRYPTED parameter.

The content and internal format of the CLIENT_INFO parameter is set by the HDA's policy and is outside the scope of this document. Examples of client information can by phone number, IMEI, and ICCID.

### 6.1. Hierarchical HIT Registration using X.509 Certificates

This requires the HIP client to possess a client certificate trusted by the HDA/Registrar. Registration will either succeed or fail.

### 6.2. Hierarchical HIT Registration using a PSK

This requires the HIP client and the HDA/Registrar to share a PSK. The PSK may already exist prior to starting the registration and just be used within the registration. A PSK out-of-band exchange may be triggered by performing the registration without any authentication.

If no client authentication is included in the I2 packet, the registration fails with "No Authentication provided". If the I2 packet included the proper HDA required client information, the HDA can use it to set up a side channel for an out-of-band delivery of a PSK. And example of this would be to send an SMS message with the PSK. Once the client possesses the PSK, it can rerun the registration at which point the HI and HIT duplicate checks are performed.

### 6.3. Hierarchical HIT Registration Type

The Registration Type used in the REG_REQUEST is:

| Number | Registration Type |
| --- | --- |
| 2 | HIT Registration |

### 6.4. Hierarchical HIT Registration Failure Type

The Registration may fail. In fact, with PSK, this may be the response to expect an SMS message with the PSK to use in a second registration request. Failure Types used in the REG_FAIL are:

| Failure Type | Reason |
| --- | --- |
| [TBD-IANA] | Hierarchical HIT Already Registered |
| [TBD-IANA] | HI Already Registered |
| [TBD-IANA] | Previously Registered HI with different device information |
| [TBD-IANA] | No Authentication provided |
| [TBD-IANA] | Invalid Authentication |
| [TBD-IANA] | Invalid Authentication, new PSK sent via SMS |

### 6.5. Registration failure behavior

If the failure type is "Hierarchical HIT Already Registered", the client's HI is hashing to an existing HIT and must generate a new HI and hierarchical HIT and reregister. If the failure is "HI Already Registered", the client should assume it is registered. If the failure is "Previously Registered HI with different device information", either the client managed to generate a duplicate HI, probably indicating a weak key generation algorithm, or the client was previously registered on a different device. Resolving this conflict will be left to the HDA's policy.

## 7. Using hierarchical HITs

All HIP clients with hierarchical HITs maintain an RVS connection with their HDA's RVS server(s). How the HDA scales this service up to a potential population in the millions is out of scope of this document. Lifetime management of these connections is also out of scope.

### 7.1. Contacting a HIP client

A service Initiator uses some service to discover the HIT of the service Responder. The Initiator uses the hierarchical information in the HIT to find the Responder's RVS. An I1 is sent to that RVS which forwards it to the Responder.

The potential Responder uses the HIT in the I1 to query the Initiator's RVS about the Initiator. The nature of information, and

method of communication are determined by the Initiator's HDA and the Responder's (and or HDA's) relationship with it. Based on the Responder's local policy, this information will be used to determine if the contact is to be accepted. If accepted, the Responder may proceed sending an R1 to the Initiator. It may alternatively initiate some non-HIP process.

It should be noted that this R1 may contain a REG_INFO list for the Initiator to validate that the Responder does offer the desired service.

## 8. IANA Considerations

The following change to the "Host Identity Protocol (HIP) Parameters" registries has been made:

HIT Suite ID: This document defines the new HIT Suite "Hierarchy with ECDSA/SHA256" (see Section 5.1).

CLIENT_INFO: This document defines the new CLIENT_INFO parameter (see Section 5.2). The parameter value will be assigned by IANA.

Reg Type: This document defines the new Registration Type for the REG_REQUEST parameter "HIT Registration" (see Section 6.3).

Reg Fail: This document defines the new Failure Types for the REG_FAIL parameter (see Section 6.4).

## 9. RAA Management Organization Considerations

Introducing the RAA management organization may be the largest hurdle for hierarchical HITs. Thus it would be best if this were adopted by an organization already in the business of allocating numbers within either the Internet or the Mobile, cellular, infrastructure.

One consideration would be to reserve the first N RAA values to map to the existing DNS TLDs. For example, these TLDs can be organized in an ascending order and numbered accordingly. Thus the 2 character TLDs will be a lower number than the 3 character TLDs. After that, it could be a first come, first numbered assignment process.

## 10. Security Considerations

There are potential risks with the hierarchical HIT, the Registry service, and the discovery of potential peer using its hierarchical HIT.

The two risks with hierarchical HITs is the use of an invalid HID and forced HIT collisions. Use of the hhit.arpa. DNS zone is a strong protection against invalid HIDs. Querying an HDA's RVS for a HIT under the HDA protects against talking to unregistered clients. The Registry service has direct protection against forced or accidental HIT hash collisions.

By using the HIP Registration Extension, the Registry service is protected from direct attacks. This service does rely on either the integrity of a PKI service or an out-of-band PSK delivery process. Thus the risk to the Registry service is highly related to the trust in these authentication setup services. Further, the duplicate HI resolution process may require human interaction with related social engineering risks.

Finally the peer discovery process relies on trusting the finding the proper HDA for the peer and its forwarding the I1 to the proper Responder. A rogue RVS, impersonating the RVS for the HIT, could redirect the I1 to a client that has forced a collision with the HIT and the Initiator would be none the wiser. The only defense against this is if the Initiator has some other source for the Responder HI and validate the HI in the R1.

## 11. References

### 11.1. Normative References

[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>.

### 11.2. Informative References

[I-D.ietf-hip-dex]
Moskowitz, R. and R. Hummen, "HIP Diet EXchange (DEX)", draft-ietf-hip-dex-03 (work in progress), June 2016.

[I-D.ietf-hip-rfc5203-bis]
Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Registration Extension", draft-ietf-hip-rfc5203-bis-10 (work in progress), January 2016.

[I-D.ietf-hip-rfc5204-bis]
Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extension", draft-ietf-hip-rfc5204-bis-07 (work in progress), December 2015.

[I-D.ietf-hip-rfc6253-bis]
Heer, T. and S. Varjonen, "Host Identity Protocol Certificates", draft-ietf-hip-rfc6253-bis-09 (work in progress), July 2016.

[I-D.irtf-hiprg-dht]
Ahrenholz, J., "Host Identity Protocol Distributed Hash Table Interface", draft-irtf-hiprg-dht-05 (work in progress), December 2011.

[RFC7343] Laganier, J. and F. Dupont, "An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers Version 2 (ORCHIDv2)", RFC 7343, DOI 10.17487/RFC7343, September 2014, <http://www.rfc-editor.org/info/rfc7343>.

[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>.

## Appendix A. Calculating Collision Probabilities

The accepted formula for calculating the probability of a collision is:

```
p = 1 - e^{-k^2/(2n)}


P       Collision Probability
n       Total possible population
k       Actual population
```

## Authors' Addresses

Robert Moskowitz
Huawei
Oak Park, MI 48237
USA

Email: rgm@htt-consult.com

Xiaohu Xu
Huawei
Huawei Bld, No.156 Beiqing Rd.
Beijing, Hai-Dian District 100095
China

Email: xuxiaohu@huawei.com