INTERNET-DRAFT
Intended Status: Informational
Expires: September 22, 2016

T. Herbert
Facebook

March 21, 2016

ILA Control Messages
draft-herbert-ila-messages-00

Abstract

This specification defines control messages for Identifier Locator Addressing. These control messages are sent between ILA hosts or from ILA routers to ILA hosts to notify a sending host to change its entry for an identifier in its ILA mapping table. Messages indicate that no mapping was found for a destination identifier or indicate a redirect to use a different locator for an identifier.

Table of Contents

1 Introduction

Identifier locator addressing is described in [I-D.herbert-nvo3-ila] and an architecture for ILA deployment is described in [I-D.lapukhov-ila-deployment]. In this specification we describe ILA control messages. These are messages concerning the state of ILA mappings and are sent between ILA hosts or between ILA routers and ILA hosts. These messages are sent in response to data plane packets in order to notify a sender about status of the ILA mapping for the identifier in a destination address. There are four different types of notifications:

- "Host No Identifier" is sent by an ILA host if it receives a packet addressed with its local locator but cannot match the identifier.

- "Router No Mapping" is sent by an ILA router when it receives a SIR addressed packet however does not have a mapping for the identifier. In this case the destination identifier is unreachable

- "Host redirect" is sent by an ILA host when it receives a packet addressed with its local locator and the identifier has a forwarding mapping entry. This is used to redirect senders to a new location after an identifier has been moved to a different host.

- "Router redirect" is sent by an ILA router when it receives a SIR addressed packet and the locator for the identifier is in the mapping table. Upon receiving this message a host is able to send packets directly using the locator instead of the SIR address.

2 ILA control message format

ILA messages are sent in the payload of UDPv6 packets. The general format is:

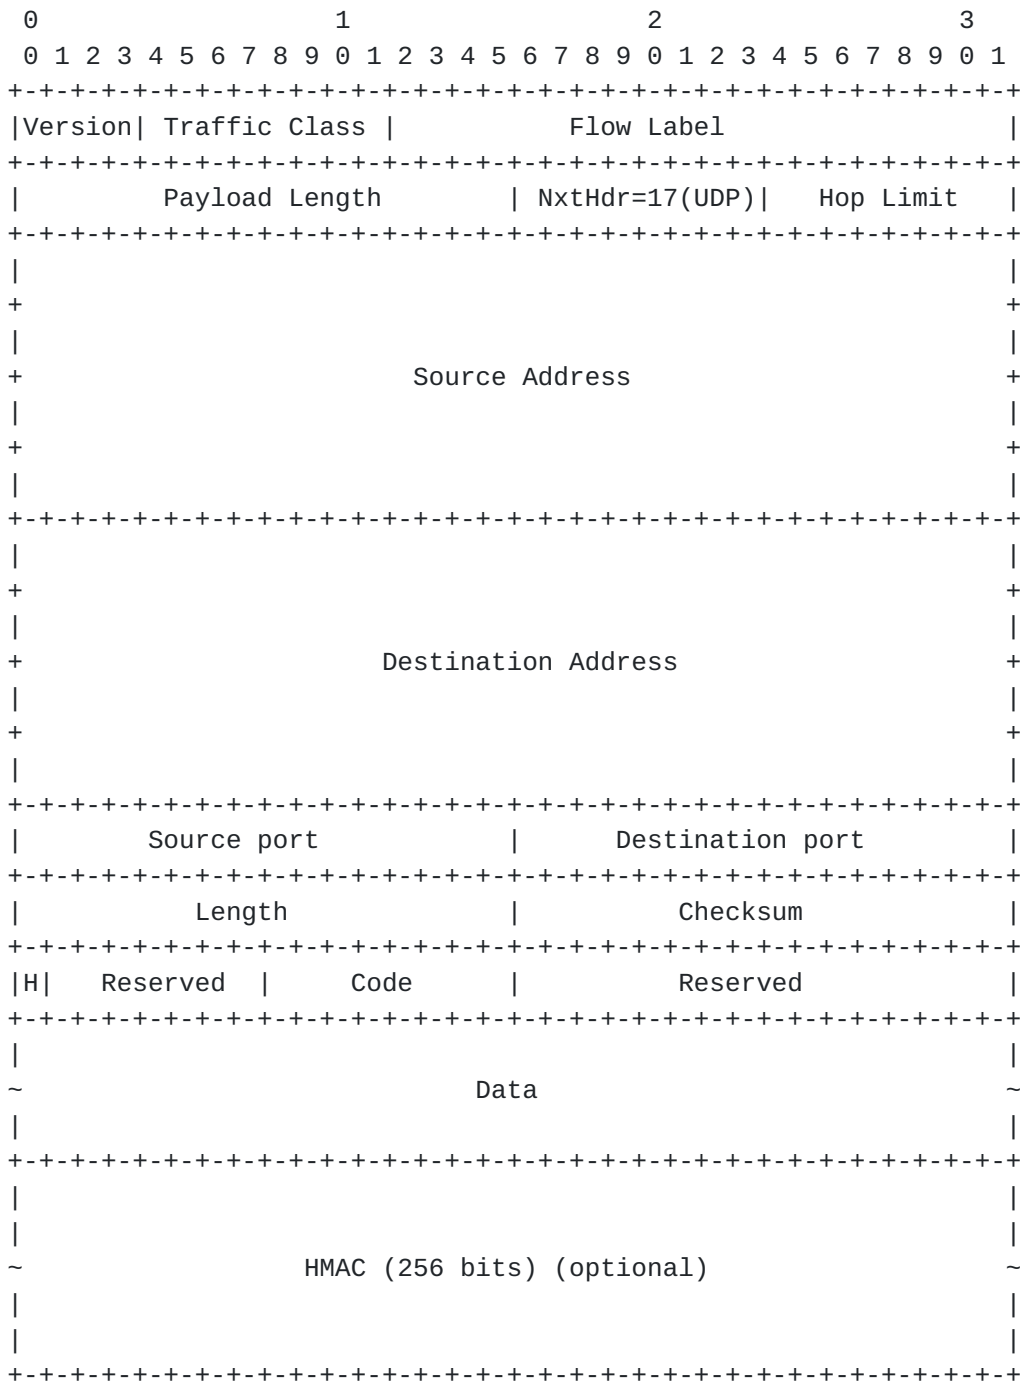

o Source address is the common locator address (CLA) of the device

sending the message. (the common locator address has the form: <locator>::1).

- Destination address is the common locator address derived from the source address of the received packet. The source address of the received packet is a SIR address, so to send an ILA control message the SIR address must be reverse mapped to a locator which is set in the CLA.

- Source port: set to ILA control message port number (assignment TDB).

- Destination UDP port: set to ILA control message port number (assignment TBD).

- H: Indicates the HMAC field is present.

- Code indicates the type of the ILA message. Defined codes are:

  - 0x1: Host No Identifier

  - 0x2: Router No Mapping

  - 0x3: Host Redirect

  - 0x4: Router Redirect

- Data: The data portion of the control messages.

- HMAC: HMAC Key ID and HMAC field, and their use are defined in Section 4.

2.1 Host No Identifier

The Host No Identifier message is sent by an ILA host if it receives a packet addressed with its local locator but cannot match the identifier. The format of this message is:

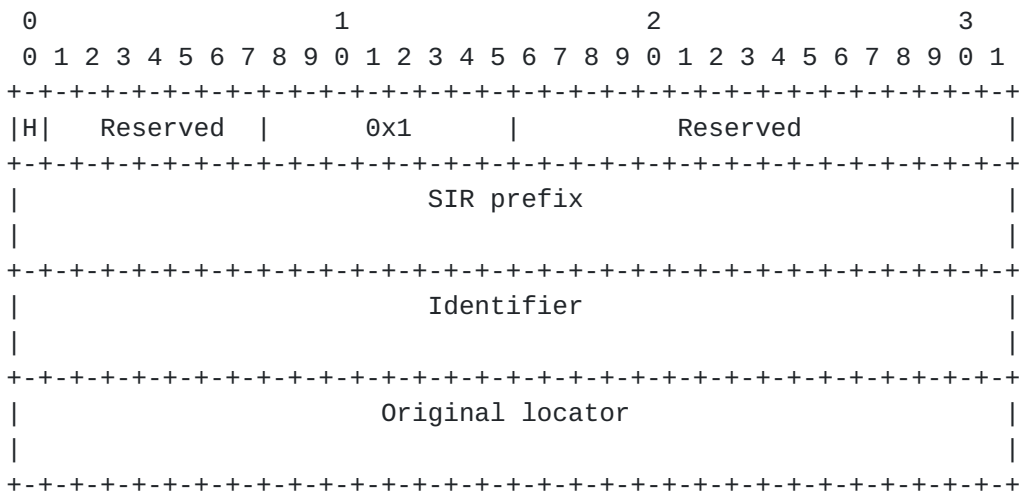

o SIR prefix: the SIR prefix associated with the locator of the sending host (i.e. the locator in the original packet)

o Identifier: The identifier received in the packet. The host sending the ILA message has no mapping for this identifier.

o Original locator: The locator that was in the destination address of the received packet. If there is a only one locator assigned per host this will be the same as the locator used in the source address of the ILA message (see above).

2.2 Router No Mapping

Router No Mapping is sent by an ILA router when it receives a SIR addressed packet however does not have a mapping for the identifier. The format of this message is:

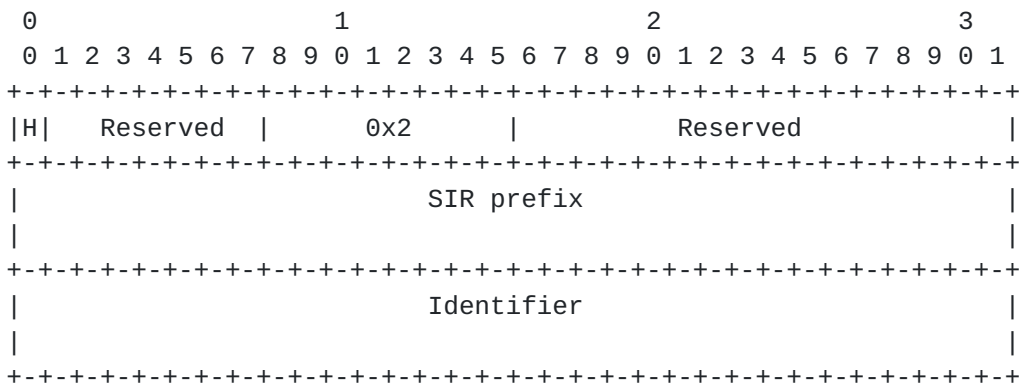

o SIR prefix: the SIR prefix associated with the locator of the sending host (i.e. the locator in the original packet)

o Identifier: The identifier received in the packet. The ILA

router sending the ILA message has no mapping for this identifier.

2.3 Host redirect

A Host Redirect is sent by an ILA host when receives a packet addressed with its local locator and the identifier has a forwarding mapping entry. The format of this message is:

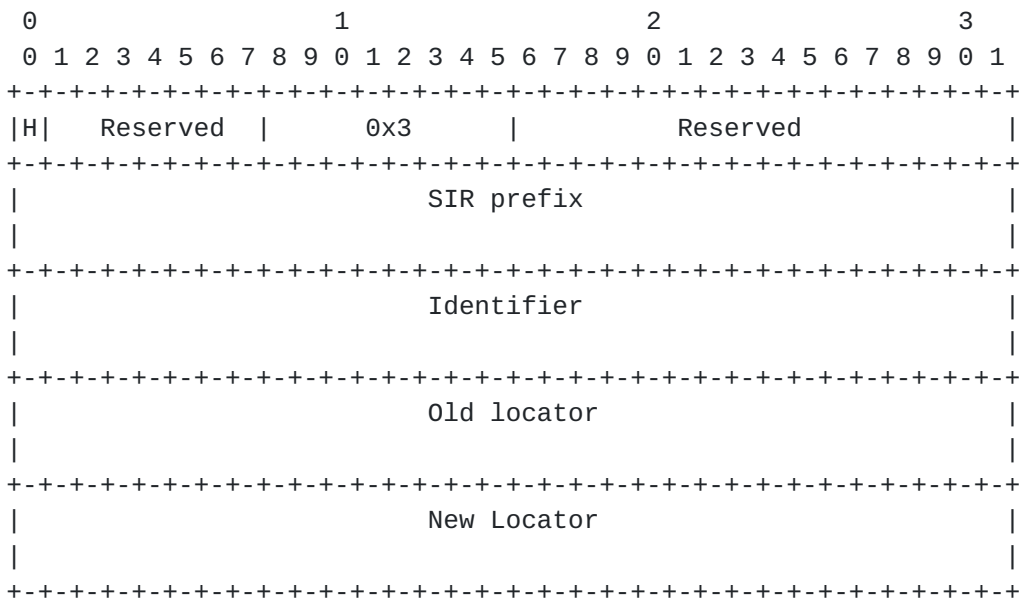

- SIR prefix: the SIR prefix associated with the locator of the sending host (i.e. the locator in the original packet)

- Identifier: The identifier received in the packet. The host sending the ILA message is suggesting a new locator for this identifier.

- Old locator: The locator that was set in destination of the received packet.

- New Locator: The locator that the a sending host is being redirected to use.

2.4 Router redirect

A Router Redirect is sent by an ILA router when it receives a SIR addressed packet and the locator for the identifier in the destination address is in the mapping table. The format of this message is:

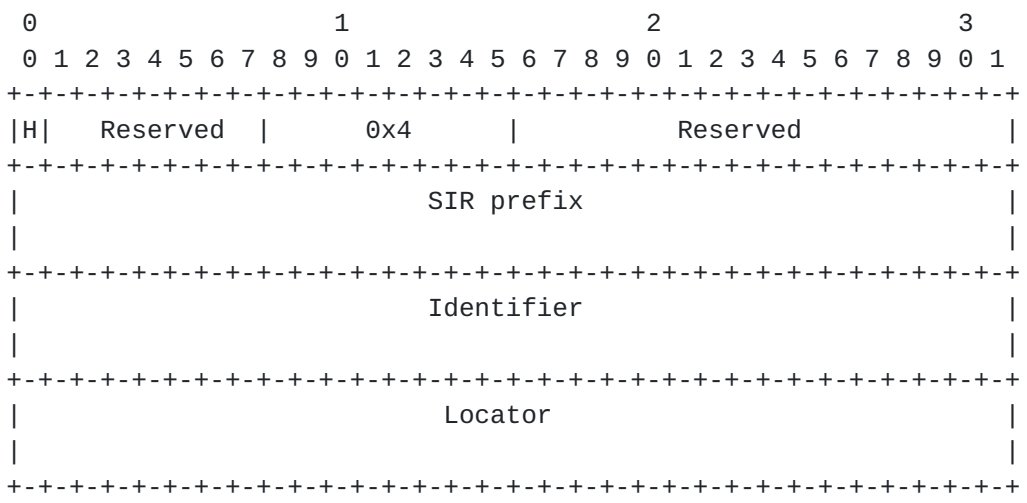

- o SIR prefix: the SIR prefix associated with the locator of the sending host (i.e. the locator in the original packet)

- o Identifier: The identifier received in the packet. The router sending the ILA message is suggesting a new locator for this identifier.

- o Locator: The locator that the a sending host is being redirected to use.

3. Operation

ILA messages can only be sent in response to packets whose source address is a SIR address in the same domain as the host or router sending the control messages. ILA messages must not be sent to non-ILA hosts or ILA hosts in a different domain (i.e. use a different SIR prefix).

ILA messages must be verified for authenticity (see Section 4).

3.1 Host control message generation

An ILA host may send control messages in response to packets received that have its local locator in the destination address.

The steps in receive processing of an ILA host are:

1) Receive packets where the destination locator matches the local locator of the host.

2) Lookup the identifier in the destination address in the ILA locator mapping database.

3) If the identifier is found and it is marked as a local identifier (i.e. this is the location for the identifier), then overwrite the locator in the destination address with the SIR prefix and receive the packet in the local stack as per ILA processing.

4) Else, if the source address is not a SIR address in the same ILA domain as the host, then drop the packet and take no further action.

5) Lookup the identifier in the source address in the ILA mapping table. If no entry is found or there is no locator associated with the entry, then drop the packet and take no further action.

6) If the identifier in the destination address was found in the mapping database and there is a forwarding address set in the mapping entry, then send a Host Redirect message:

   a) The destination address of the control message is set to the common locator address which is comprised of the locator associated with the identifier in the source address and ::1 in the low order 64 bits.
   b) The source address of the control message packet is set to the common locator address of the host sending the message.
   c) The SIR prefix in the control message is set to that associated with the locator the packet was received on.
   d) The identifier in the control message is set to the identifier in the destination address of the received packet.
   e) The Old Locator in the control message is set to the locator in the destination address of the received packet (should be a local locator).
   f) The New locator in the control message data is set to the corresponding value in the mapping table.

7) Else, send a Host No Identifier. The host has no information regarding the received identifier:

   a) The destination address of the control message is set to the common locator address which is comprised of the locator associated with the identifier in the source address and ::1 in the low order 64 bits.
   b) The source address of the control message packet is set to the common locator address of the host sending the message.
   c) The SIR prefix in the control message is set to that associated with the locator the packet was received on.
   d) The identifier in the control message is set to the identifier in the destination address of the received packet.

e) The Original Locator in the control message is set to the locator in the destination address of the received packet (should be a local locator).

3.2 ILA router processing

An ILA router should send control messages in response to packets received whose destination address is a SIR address.

The steps in receive processing of an ILA router are:

1) Receive anycast SIR addressed packets that are routed to the router.

2) Lookup the identifier in the destionation address in the ILA locator database.

3) If an entry for the identifier is found and there is an associated locator, then overwrite the SIR prefix in the destination address with the found locator and forward the packet per ILA processing.

4) Else, if the source address is not a SIR address in the same ILA domain as the host, then drop the packet and take no further action.

5) Lookup the identifier in the source address in the ILA mapping table. If no entry is found or there is no locator associated with the entry, then take no further action.

6) If the locator was found in the mapping table then send a Router Redirect message:

   a) The destination address of the control message is set to the common locator address which is comprised of the locator associated with the identifier in the source address and ::1 in the low order 64 bits.
   b) The source address of the control message packet is set to the common locator address of the ILA router sending the message.
   c) The SIR prefix in the control message is set to that associated with the locator the packet was received on.
   d) The identifier in the control message is set to the identifier in the destination address of the received packet.
   f) The Locator in the control message data is set to the corresponding value in the mapping table.

7) Else, send a Router No Mapping messages

a) The destination address of the control message is set to the common locator address which is comprised of the locator associated with the identifier in the source address and ::1 in the low order 64 bits.
b) The source address of the control message packet is set to the common locator address of the ILA router sending the message.
c) The SIR prefix in the control message is set to that associated with the locator the packet was received on.
d) The identifier in the control message is set to the identifier in the destination address of the received packet.

3.3 Host processing of received ILA messages

Hosts listen on the appropriate UDP port for ILA control messages. The steps in processing control messages are:

1) Control messages are received on the common locator address for the host.

2) If the message code is Host No Mapping

a) Lookup the specified identifier in the ILA mapping cache. If an entry is found, there is a locator in the mapping, and the locator matches that in the Original Locator in the control message-- then invalidate the locator. The next packet that is sent to that identifier will not be translated and so will be sent with a destination SIR address.
b) If the mapping entry has no locator set or a different locator than what is reported in the Host No mapping, then disregard the message

3) If the message code is Router No Mapping

a) Lookup the specified identifier in the ILA mapping cache. If a match is found, then mark the entry is unreachable (subject to a timer)
b) If an entry is not found then disregard the message.

4) If the message code is Host redirect

a) Lookup the specified identifier in the ILA mapping cache. If an entry is found, there is a locator in the mapping, and the locator matches that in the Old Locator in the control message-- then set the locator to the value of New Locator. The next packet that is sent to that identifier will use the new locator.
b) If the mapping entry has no locator set or a different

locator than what is reported in the Host Redirect then disregard the message

5) If the message is Router Redirect

a) Lookup the specified identifier in the ILA mapping cache. If an entry is found then set the locator to the value of Locator in the control message. The next packet that is sent to that identifier will use the new locator.

b) If no mapping entry is found for the identifier then disregard the message or, optionally, create a new cache entry for the redirected identifier (this permits a mechanism for ILA routers to push mappings).

3.4 Other properties

Host No Mapping and Host Redirect are considered optional messages for hosts to send. If they are sent, then rate limiting should be a applied to avoid sending more than 100 control messages per second. If control messages are not generated, then a sender that has an out of date mapping should eventually timeout and stop sending packets to an incorrect locator. A host should not forward ILA addressed packets even in the case that it has the forwarding address for an identifier.

Router No Mapping and Router Redirect must be sent in order to keep hosts synchronized with mapping state. The number of messages sent to particular host should be rate limited so that no more than 100 messages per second are sent. An ILA router may proactively send Router No Mapping and Router Redirect messages upon a mapping state change to those hosts that are known to possess a mapping (the list of hosts that have sent packets through the router using a mapping could be kept in the mapping database). Conceivably, an ILA router could also multicast messages to inform hosts of a state (the details for this are outside the scope of this document).

4 Security Considerations

ILA control messages convey sensitive control information so they need to be protected against spoofing. ILA control messages should be authenticated and may be encrypted.

4.1 HMAC authentication and integrity

ILA control messages may optionally use HMAC authentication. The use of HMAC in a message is indicated by the H bit being set.

4.1.1 Security fields in ILA control messages

This section summarizes the use of specific fields in the ILA messages. They are based on a key-hashed message authentication code (HMAC).

The security-related fields in ILA messages are:

- HMAC Key-id, 8 bits wide;
- HMAC, 256 bits wide (optional, exists only if HMAC Key-id is not 0).

The HMAC field is the output of the HMAC computation (per RFC 2104 [RFC2104]) using a pre-shared key and hashing algorithm identified by HMAC Key-id and of the text which consists of the concatenation of:

- the source IPv6 address;
- the ILA control message header and data

The purpose of the HMAC field is to verify the validity, the integrity and the authorization of the ILA control messages itself. If an outsider of the ILA domain does not have access to a current pre-shared secret, then it cannot compute the right HMAC so when the receiver of a control messages checks the validity of the HMAC it will simply reject the packet.

The HMAC Key-id field allows for the simultaneous existence of several hash algorithms (SHA-256, SHA3-256 ... or future ones) as well as pre-shared keys. This allows for pre-shared key roll-over when two pre-shared keys are supported for a while when all ILA nodes converged to a fresher pre-shared key. The HMAC Key-id field is opaque, i.e., it has neither syntax not semantic except as an index to the right combination of pre-shared key and hash algorithm and except that a value of 0 means that there is no HMAC field. It could also allow for interoperation among different ILA domains if allowed by local policy and assuming a collision-free Key Id allocation which is out of scope of this memo.

4.1.2 Selecting a hash algorithm

The HMAC field in the ILA control messages is 256 bits wide. Therefore, the HMAC MUST be based on a hash function whose output is at least 256 bits. If the output of the hash function is 256, then this output is simply inserted in the HMAC field. If the output of the hash function is larger than 256 bits, then the output value is truncated to 256 by taking the least-significant 256 bits and

inserting them in the HMAC field.

ILA implementations can support multiple hash functions but MUST implement SHA-2 [FIPS180-4] in its SHA-256 variant.

4.1.3 Pre-shared key management

The field HMAC Key-id allows for:

- key roll-over: when there is a need to change the key (the hash pre-shared secret), then multiple pre-shared keys can be used simultaneously. The validating routing can have a table of <HMAC Key-id, pre-shared secret, hash algorithm> for the currently active and future keys.

- different algorithm: by extending the previous table to <HMAC Key-id, hash function, pre-shared secret>, the validating router can also support simultaneously several hash algorithms (see section Section 4.1.2)

The pre-shared secret distribution can be done:

- in the configuration of the validating nodes, either by static configuration or any SDN oriented approach;

- dynamically using a trusted key distribution such as [RFC6407]

  The intent of this document is NOT to define yet-another-key-distribution-protocol.

4.2 DTLS

For stronger security, including encryption of ILA control messages, DTLS may be used. The use of DTLS necessitates a separate UDP port.

5 IANA Considerations

IANA will be requested to assign one UDP port number for ILA control messages, and one for ILA control messages with DTLS.

6 References

6.1 Normative References

[I-D.herbert-nvo3-ila] Herbert, T., "Identifier-locator addressing for network virtualization", draft-herbert-nvo3-ila-02 (work in progress), March 2016.

6.2 Informative References

[I-D.lapukhov-ila-deployment] Lapukhov, P., "Deploying Identifier-Locator Addressing (ILA) in datacenter", draft-lapukhov-ila-deployment-00 (work in progress), March 2016.

[RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, <http://www.rfc-editor.org/info/rfc2104>.

[RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, DOI 10.17487/RFC6407, October 2011, <http://www.rfc-editor.org/info/rfc6407>.

Authors' Addresses

Tom Herbert
Facebook
1 Hacker Way
Menlo Park, CA
US

EMail: tom@herbertland.com