

**December 21, 2017**

Identifier Locator Addressing Mapping Protocol  
[draft-herbert-ila-ilamp-00](#)

Abstract

Identifier-locator protocols rely on a mapping system that is able to map identifiers to locators. ILA nodes that perform ILA translations need to access the mapping system via a protocol. This document specifies the ILA Mapping Protocol that is used by ILA forwarding nodes and hosts to populate and maintain a cache of ILA mappings.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2017 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

(<http://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

<a href="#">1</a>	Introduction . . . . .	<a href="#">4</a>
<a href="#">2</a>	Reference topology . . . . .	<a href="#">4</a>
<a href="#">2.1</a>	Functional components . . . . .	<a href="#">5</a>
<a href="#">2.2</a>	ILA forwarding nodes and hosts . . . . .	<a href="#">5</a>
<a href="#">2.2.1</a>	ILA forwarding nodes . . . . .	<a href="#">5</a>
<a href="#">2.2.2</a>	ILA hosts . . . . .	<a href="#">5</a>
<a href="#">2.2.3</a>	ILA address to SIR address translation . . . . .	<a href="#">5</a>
<a href="#">2.2.4</a>	ILA Forwarding . . . . .	<a href="#">5</a>
<a href="#">2.3</a>	ILA routers . . . . .	<a href="#">6</a>
<a href="#">2.3.1</a>	Forwarding routers . . . . .	<a href="#">6</a>
<a href="#">2.3.2</a>	Mapping routers . . . . .	<a href="#">6</a>
<a href="#">2.3.3</a>	ILA router synchronization . . . . .	<a href="#">7</a>
<a href="#">3</a>	ILA Mapping Protocol . . . . .	<a href="#">7</a>
<a href="#">3.1</a>	Common header format . . . . .	<a href="#">8</a>
<a href="#">3.2</a>	Hello messages . . . . .	<a href="#">8</a>
<a href="#">4</a>	ILAMP Version 0 . . . . .	<a href="#">9</a>
<a href="#">4.1</a>	Map request . . . . .	<a href="#">9</a>
<a href="#">4.2</a>	Map information . . . . .	<a href="#">10</a>
<a href="#">4.3</a>	Extended map information . . . . .	<a href="#">11</a>
<a href="#">4.4</a>	Locator unreachable . . . . .	<a href="#">12</a>
<a href="#">4.5</a>	Identifier and locator types . . . . .	<a href="#">13</a>
<a href="#">4.5.1</a>	Identifier types . . . . .	<a href="#">13</a>
<a href="#">4.5.2</a>	Locator types . . . . .	<a href="#">13</a>
<a href="#">5</a>	Operation . . . . .	<a href="#">14</a>
<a href="#">5.1</a>	Version negotiation . . . . .	<a href="#">14</a>
<a href="#">5.2</a>	Populating an ILA cache . . . . .	<a href="#">14</a>
<a href="#">5.2.1</a>	ILA Redirects . . . . .	<a href="#">15</a>
<a href="#">5.2.1.1</a>	Proactive push with redirect . . . . .	<a href="#">15</a>
<a href="#">5.2.1.2</a>	Redirect rate limiting . . . . .	<a href="#">15</a>
<a href="#">5.2.2</a>	Map request/reply . . . . .	<a href="#">15</a>
<a href="#">5.2.3</a>	Push mappings . . . . .	<a href="#">16</a>
<a href="#">5.3</a>	Cache maintenance . . . . .	<a href="#">16</a>
<a href="#">5.3.1</a>	Timeouts . . . . .	<a href="#">16</a>
<a href="#">5.3.2</a>	Cache refresh . . . . .	<a href="#">16</a>
<a href="#">5.3.3</a>	Cache timeout values . . . . .	<a href="#">17</a>
<a href="#">5.4</a>	ILA forwarding node and host receive processing . . . . .	<a href="#">17</a>
<a href="#">5.5</a>	Locator unreachable handling . . . . .	<a href="#">17</a>

T. Herbert

Expires June 24, 2018

[Page 2]

<a href="#">5.6</a>	Control Connections . . . . .	<a href="#">18</a>
<a href="#">5.7</a>	Protocol errors . . . . .	<a href="#">18</a>
<a href="#">6</a>	Security Considerations . . . . .	<a href="#">19</a>
<a href="#">7</a>	IANA Considerations . . . . .	<a href="#">20</a>
<a href="#">8</a>	References . . . . .	<a href="#">20</a>
<a href="#">8.1</a>	Normative References . . . . .	<a href="#">20</a>
<a href="#">8.2</a>	Informative References . . . . .	<a href="#">20</a>
	Author's Address . . . . .	<a href="#">20</a>

## 1 Introduction

The ILA Mapping Protocol (ILAMP) is a control plane protocol that provides ILA nodes mapping information. ILA [[ILA](#)] nodes that perform ILA translation rely on a mapping system to provide identifier to locator mappings. There are two levels of mapping protocols to be defined: one used by ILA routers that require the full set of ILA mappings for a domain, and one used by ILA nodes that maintain a caches of mappings.

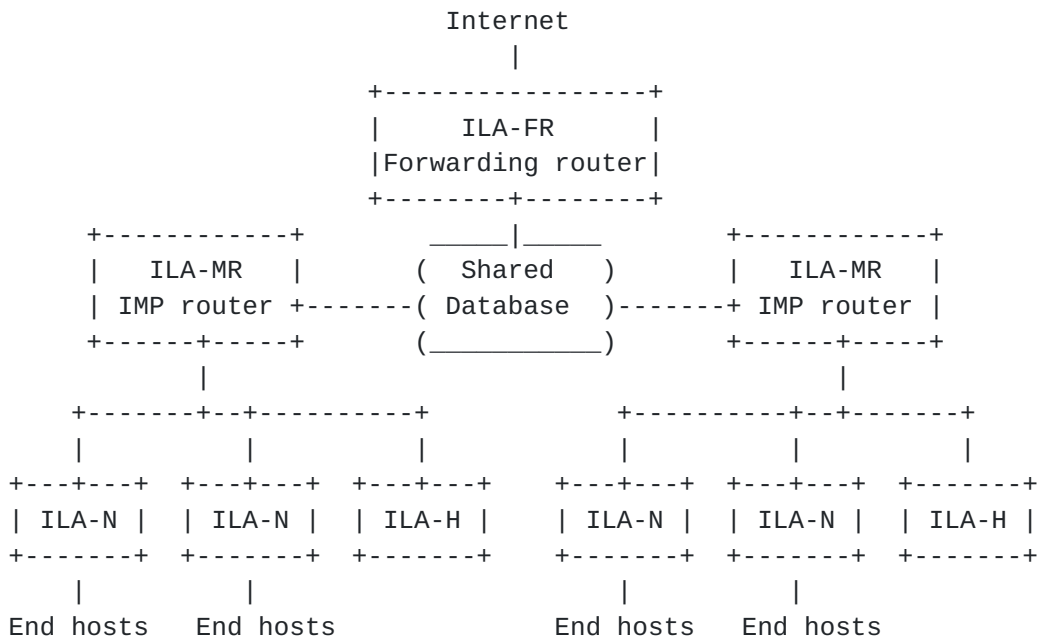
The ILA mapping system is effectively a key/value database that maps identifiers to locators. The protocol for sharing mapping information amongst ILA routers, nodes that maintain a full table of mappings, can thus be a database. A database schema for the ILA mapping system will be described in a separate document.

ILA separates the control plane from the data plane, so alternative control plane protocols may be used with a common data plane. For instance, BGP could be used as a mapping system protocol [[ILABGP](#)].

## 2 Reference topology

This section provides a reference topology for ILA.

The topology is general however can be adapted to specific ILA use cases such as data center virtualization and mobility networks.





## **[2.1](#) Functional components**

There are four types of functional nodes in the ILA architecture:

ILA-N: ILA forwarding nodes

ILA-H: ILA hosts

ILA-FR: ILA forwarding routers

ILA-MR: ILA mapping routers

## **[2.2](#) ILA forwarding nodes and hosts**

### **[2.2.1](#) ILA forwarding nodes**

ILA forwarding nodes (ILA-N) are deployed in the network infrastructure towards the edges to provide caches for ILA forwarding. ILA forwarding nodes have two functions: ILA address to SIR address translation and ILA forwarding. As indicated in the reference topology, forwarding nodes may be deployed near the point of device attachment (e.g. base station, eNodeB) of user devices (e.g. UEs).

### **[2.2.2](#) ILA hosts**

ILA hosts (ILA-H) are end hosts that participate in ILA. These may be servers that provide ILA to work with virtualization techniques such as VMs or containers. ILA hosts perform the same functions as ILA forwarding nodes, however the source of packets is local to the same host so there are some optimizations that may be applied.

### **[2.2.3](#) ILA address to SIR address translation**

ILA forwarding nodes and hosts perform ILA address to SIR address translation. This is a reverse ILA translation in order to restore the original addresses in a packet for delivery. Forwarding the packet on to the destination is done based on the SIR address. For instance, a forwarding node may map a SIR address to a layer 2 address of a directly attached device that has the SIR address. Note that this functionality is required somewhere in the path between the node that writes a locator into an address and the ultimate destination device.

### **[2.2.4](#) ILA Forwarding**

An ILA forwarding node or host may perform ILA translation and forward packets directly to peer ILA nodes in the same domain. The





mappings for this are maintained in a working set cache. As a cache there must be methods to populate, evict, and timeout entries. A cache is considered an optimization so the system should be functional without its use (e.g. if the cache has no entries).

### **2.3 ILA routers**

ILA routers (denoted by ILA-R\*) are deployed within the network infrastructure and collectively contain a database of all identifier to locator mappings in the domain, as well as all identifier group information for the domain. The database may be sharded across some number of ILA routers for scalability. ILA routers that maintain the database or a shard may be replicated for scalability and availability.

ILA routers provide two main functions: ILA forwarding and mapping resolution. An ILA router may perform one or the other of these functions or may provide both at the same time.

#### **2.3.1 Forwarding routers**

Forwarding routers (ILA-FR) perform ILA translation when packets enter a domain. A destination address of a packet that is a SIR address is translated to an ILA address. The process is that the router performs a lookup on the destination address in the mapping table and a locator is returned. The locator is written into the destination address of the packet (that is the high order sixty-four bits are overwritten with a locator).

In the case of a sharded database being used, the high order bits of the identifier indicate the shard number. This is included in a routing prefix so that the packet is routed to an ILA router that contains the database for the relevant shard.

#### **2.3.2 Mapping routers**

A mapping router (ILA-MR) provides ILA forwarding nodes and hosts mapping information. A mapping router may also perform ILA translations and forwarding. A mapping router implements ILAMP to communicate with ILA forwarding nodes and hosts. Mapping information is provided by a request/reply protocol, ILA redirects, or by a push mechanism.

An ILA router that is performing mapping resolution will respond to mapping requests from ILA forwarding nodes or ILA hosts. The mapping request protocol allows a node to request the locator information for an identifier address.



A mapping router can send ILA redirects to ILA forwarding nodes and ILA hosts in order to inform them of a direct ILA path. A redirect is sent to the upstream ILA forwarding node or host of the source which is determined by an ILA lookup the source address.

### **2.3.3 ILA router synchronization**

ILA routers, both those that are forwarding and those that provide mapping resolution, must synchronize the contents of the database. This synchronization is done for each shard. When a change occurs to an identifier-locator mapping, for instance the locator for an identifier changes, the shard that contains the identifier must be synchronized in as little time to converge as possible.

There are a number of options to use for implementing the ILA mapping system and router protocol. One option is to use a key/value database (such as a NoSQL database like Redis).

The idea of the database is that each shard is a distributed database instance with some number of replicas. When a write is done in the database, the change is propagated throughout all of the replicas for the shard using the standard database replication mechanisms. Mapping information is written to the database using common database API and requires authenticated write permissions. Each ILA router can read the database for the associated shard to perform its function.

The specifics of applying a database and a database schema for ILA will be provided in other documents.

## **3 ILA Mapping Protocol**

The ILA Mapping Protocol (ILAMP) is used between ILA forwarding nodes and ILA mapping routers. The purpose of the protocol is to populate and maintain the ILA mapping cache in forwarding nodes.

ILAMP defines redirects, a request/response protocol, and a push mechanism to populate the mapping table. Unlike traditional routing protocols that run over UDP, this protocol is intended to be run over TCP. TCP provides reliability, statefulness implied by established connections, ordering, and security in the form of TLS. Secure redirects are facilitated by the use of TCP.

ILAMP is used to send message between ILA routers and ILA forwarding nodes or hosts. The messages are sent over the TCP stream and must be delineated by a receiver. Different versions of ILAMP are allowed and the version used for communication is negotiated by Hello messages.



### 3.1 Common header format

All ILAMP messages begin with a two octet common header:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Type |          Length          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The contents of the common header are:

- o Type: Indicates the type of message. A type 0 message is a hello message. Types greater than zero are interpreted per the negotiated version.
- o Length: Length of the message in octets. This includes the common header. The minimal length of a message is 2 octets and the maximum length is 1,048,575 octets.

Following the two octet common header is variable length data that is specific to the version and type the message.

### 3.2 Hello messages

Hello messages indicate the versions of ILAMP that a node supports. Hello message MUST be sent by each side as the first message in the connection.

The format of an ILAMP Hello message is:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| 0 |          4          |R| Rsvd   | MinV | MaxV |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The contents of the Hello message are:

- o Type = 0. This indicates the type is a Hello message.
- o Router bit: Indicates the sender is an ILA router. If the sender is an ILA forwarding node or host this bit is cleared.
- o Rsvd: Reserved bits. Must be set to zero on transmit.
- o MinV: Minimum version number supported by the sending node.
- o MaxV: Maximum version number supported by the sending node.

Version numbers are from 0 to 15. This document describes version 0 of ILAMP.



## 4 ILAMP Version 0

The message types in version 0 of IMP are:

- o Map request (Type = 1)
- o Map information (Type = 2)
- o Extended map information (Type = 3)
- o Locator unreachable (Type = 4)

### 4.1 Map request

A map request is sent by an ILA forwarding node or host to an ILA mapping router to request mapping information for a list of identifiers. The format of a map request is:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| 0   |           Length           |           Rsvd           | IDType|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                           | |
~                               Identifier ~ ent
|                                           | |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The contents of the map request message are:

- o Type = 1. This indicates the type is map request.
- o Length: Set to 4 plus the size of the identifier times the number of identifiers in the list.
- o Rsvd: Reserved bits. Must be set to zero when sending.
- o IDType: Identifier type. Specifies the identifier type. This also implies the length of each identifier in the request list. Identifier types are defined below.
- o Identifier: An identifier of type indicated by IDType. The size of an identifier is specified by the type.

The Identifier field is repeated for each identifier in the list. The number of identifiers being requested is  $(\text{message length} - 4) / (\text{identifier size})$ .





## 4.2 Map information

Map information messages are sent by an ILA router to an ILA forwarding node or host. This message provides a list of identifier to locator mappings. The format of a map information message is:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  2  |           Length           | Rsvd  |SubType|LocType| IDType|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                           | |
~                               Identifier ~ e
|                                           | n
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ t
|                                           | |
~                               Locator ~ |
|                                           | /
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The contents of the map information message are:

- o Type = 2. This indicates the type is map information.
- o Length: Set to 4 plus the size of an identifier and locator times the number of identifier/locator pairs in the list.
- o Rsvd: Reserved bits. Must be set to zero when sending.
- o SubType: Specifies the reason that ILA-R sent this message. Sub types are
  - o 0: Redirect
  - o 1: Map reply to a map request
  - o 2: Push map information
- o LocType: Locator type. Specifies the locator type. This also implies the length of each locator in the list. Locator types are defined below.
- o IDType: Identifier type. Specifies the identifier type. This also implies the length of each identifier in the list. Identifier types are defined below.
- o Identifier: An identifier of type indicated by IDType. The size of an identifier is specified by the type.
- o Locator: A locator of type indicated by LocType. The size of a



locator is specified by the type.

The Identifier/Locator pair is repeated for each mapping being reported in the list. The number of identifiers being requested is  $(\text{message length} - 4) / (\text{identifier size} + \text{locator size})$ .

#### 4.3 Extended map information

An extended locator map information message is sent by an ILA router to associate more than one locator with an identifier as well as providing an expiration time for an identifier locator mapping and additional locator specific attributes. The format of an extended map information is:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  3  |           Length           | Rsvd | SubType|LocType|IDType |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ <--+
|                                           | |
~                               Identifier                               ~ |
|                                           | r
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ e
| Num locator |           Record timeout           | c
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ \ o
| Prio  | Rsvd  | Weight          |           Rsvd          | e r
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ t d
|                                           | t |
~                               Locator                               ~ | |
|                                           | / |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ <--+

```

The contents of the map reply message are:

- o Type = 3. This indicates an extended map information message
- o Length: Set to 4 plus the sum of sizes for each identifier record in the list.
- o Rsvd: Reserved bits. Must be set to zero when sending.
- o SubType: Specifies the reason that an ILA router sent this messages. Sub types are:
  - o 0: Map reply to a map request
  - o 1: Redirect
  - o 2: Push map information



- o LocType: Locator type. Specifies the locator type. This also implies the length of each locator in the list. Locator types are defined below.
- o IDType: Identifier type. Specifies the identifier type. This also implies the length of each identifier in the list. Identifier types are defined below.
- o Num locator: Number of locators being reported for an identifier.
- o Record timeout: The time to live for the identifier information in seconds. A value of zero indicates the default is used.
- o Priority: Relative priority of a locator. Locators with higher priority values have preference to be used. Locators that have the same priority may be used for load balancing.
- o Weight: Relative weights assigned to each locator. In the case that locators have the same priority the weights are used to control how traffic is distributed. A weight of zero indicates no weight and the mapping is not used unless all locators for the same priority have a weight of zero.
- o Locator: A locator of type specified in LocType.

The identifier record is repeated for each mapping being reported and the locator entry is repeated for each locator being reported for an identifier. The total number of identifiers being reported is determined by parsing the message.

#### [4.4](#) Locator unreachable

A locator unreachable message is sent by an ILA router to ILA forwarding node or host in the event that a locator or locators are known to no longer be reachable. The format of a locator unreachable message is:

```

+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  4  |          Length          |   Rsvd   |LocType| Rsvd  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
~                               Locator                               ~ ent
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+/

```

The contents of the locator unreachable message are:

- o Type = 4. This indicates the type is a locator unreachable



message.

- o Length: Set to 4 plus the size of the locator times the number of locators in the list.
- o Rsvd: Reserved bits. Must be set to zero when sending.
- o LocType: Locator type. Specifies the locator type. This also implies the length of each locator in the list. Locator types are defined below.
- o Locator: A locator of type indicated by LocType. The size of a locator is specified by the type.

The Locator field is repeated for each locator in the list. The number of locators being reported is  $(\text{message length} - 4) / (\text{locator size})$ .

## **[4.5 Identifier and locator types](#)**

### **[4.5.1 Identifier types](#)**

Identifier types used in IDType fields of ILAMP messages are:

- o IPv6 address (IDType = 1): 128 bit IPv6 address
- o ILA Identifier (IDType = 2): 64 bit ILA identifier
- o 32 bit index (IDType = 3): 32 bit index into an identifier table
- o 64 bit index (IDType = 4): 64 bit index into an identifier table

For the table index types it is assumed that a table mapping index to and identifier is shared with ILA forwarding nodes and hosts.

### **[4.5.2 Locator types](#)**

Locator types used in LocType fields of ILAMP messages are:

- o IPv6 address (LocType = 1): 128 bit IPv6 address
- o ILA Identifier (LocType = 2): 64 bit ILA locator
- o 32 bit index (LocType = 3): 32 bit index into a locator table
- o 64 bit index (LocType = 4): 64 bit index into a locator table

For the table index types it is assumed that a table mapping index to





and locator is shared with ILA forwarding nodes and hosts.

## **5 Operation**

### **5.1 Version negotiation**

The first message sent by each side of an ILAMP connection is a Hello message. Hello messages contain the minimum and maximum versions of ILAMP supported. The minimum and maximum values form an inclusive range.

When a host receives an ILAMP Hello it determines which version is negotiated. The negotiated version is the maximum version number supported by both sides. For instance, if a node advertises a minimum version of 0 and maximum of 1 and receives a peer Hello message with a minimum version of 0 and maximum of 2; then the negotiated version is 1 since that is the greatest version supported by both sides. The peer host will also determine that 1 is the negotiated version.

If there is no common version supported between the peers, that is their supported version ranges are disjoint, then version negotiation fails. The connection **MUST** be terminated and error message **SHOULD** be logged.

If both sides set the router bit or both clear the router bit in a Hello message, then this is an error and the connection **MUST** be terminated and error message **SHOULD** be logged. Both sides cannot have the same role in an ILAMP session.

### **5.2 Populating an ILA cache**

ILA forwarding nodes and ILA hosts maintain a cache of identifier to locator mappings. There are three means that this cache can be populated by ILAMP:

- o ILA redirect
- o Mapping request/reply
- o Push mappings

ILA redirects are **RECOMMENDED** to be the primary means of obtaining mapping information. Request/reply and push mappings may be used in limited circumstances, however generally these techniques don't scale.

Note that forwarding nodes and hosts do not hold packets that are pending mapping resolution. If a node does not have a mapping for a



destination in its cache then packet is forwarded in the network. The packet should be translated by an ILA router and sent to the proper destination node.

### **5.2.1 ILA Redirects**

A mapping router can send ILA redirects in conjunction with forwarding packets. Redirects are sent to ILA forwarding nodes and ILA hosts in order to inform them of a direct ILA path. A redirect is sent to the upstream ILA forwarding node or host of the source which is determined by an ILA lookup on the source address of the packet being forwarded. The found locator is used to infer an address of the ILA forwarding node or host. For instance, the address of the forwarding node might be <locator>::1 where <locator> is a sixty-four bit prefix and 0:0:0:1 is reserved as a special identifier. Note that this technique assumes a symmetric path towards the source. If a redirect is sent then the received packet that motivated the redirect MUST be ILA translated and forwarded by the router.

#### **5.2.1.1 Proactive push with redirect**

In addition to sending an ILA redirect to the ILA forwarding node or host, a mapping router MAY send an ILA push to the ILA forwarding node or host of the destination to inform it of the identifier to locator mapping for the source address in a packet. This is an optimization to push the ILA translation that will be used in the reverse direction of the communications. In order to do this, the mapping router performs an ILA lookup on the source address (which should already be done to perform the redirect). An ILA push message is then sent to the forwarding node or host based on its locator.

#### **5.2.1.2 Redirect rate limiting**

A mapping router SHOULD rate limit the number of redirects it sends to a forwarding node or host for each redirected address. The rate limit SHOULD be configurable. The default SHOULD be that no more than one redirect is sent every one half of the minimum identifier timeout being used. The minimum rate limit SHOULD be to send no more than one redirect per second per redirected identifier. If a mapping change is detected the rate limiting SHOULD be reset so that redirects for a new mapping can be sent immediately.

### **5.2.2 Map request/reply**

A forwarding node or host may send a map request message to obtain mapping information for a locator. If the receiving mapping router has the mapping information it responds with a map information message. If the mapping router does not have a mapping entry for the



requested identifier it MAY reply with in all zeros locator (LocType = 2 and 64 bit locator is all zeroes).

Map requests are NOT RECOMMENDED to be used to populate entries in the cache table that are not present. The problem with this technique is that an ILA forwarding node or host may generate a map request for each new destination that it gets from a downstream end host. A downstream end host could launch a Denial of Service (DOS) attack whereby it sends packets with random destination addresses that requires a mapping looking. In the worse case scenario the mapping router would send a map request for every packet received. Rate limiting the sending of map requests does not mitigate the problem since that would prevent the cache from getting mappings for legitimate destinations.

### **[5.2.3](#) Push mappings**

A mapping router may push mappings to an ILA forwarding node or host without being requested to do so. This mechanism could be used to pre-populate an ILA cache. Pre-populating the cache might be done if the network has a very small number of identifiers or there are a set of identifiers that are likely to be used for forwarding in most ILA forwarding nodes and hosts (identifiers for common services in the network for instance). When a mapping router detects a changed mapping, the locator changes for instance, the new mapping can be pushed to the ILA nodes and hosts.

The push model is NOT RECOMMENDED as a primary means to populate an ILA cache since this does not scale. Conceivably, one could keep track of all ILA mappings and to which nodes the mapping information was provided. When a mapping changes, mapping information could be sent to those nodes that expressed interest. Such a scheme will not scale in deployments that have many mappings.

## **[5.3](#) Cache maintenance**

### **[5.3.1](#) Timeouts**

A node SHOULD apply a timeout for the mapping entry using either the default timeout or record timeout if one was received in an extended map information message. If the timeout fires then the mapping entry is removed. Subsequent packets may cause a mapping router to send a redirect so that the mapping entry gets repopulated in the cache.

### **[5.3.2](#) Cache refresh**

In order to avoid cycling a mapping entry with a redirect for a mapping that times out, a node MAY try to refresh the mapping before



timeout. This should only be done if the cache entry has been used to forward a packet during the timeout interval.

A cache refresh is performed by sending a map request for an identifier before its cache entry expires. If a map information messages is received for the identifier then the timeout can be reset and there are no other side effects.

### **[5.3.3](#) Cache timeout values**

The RECOMMENDED default timeout for identifiers is one minute. If a node sends a map request to refresh a mapping, the RECOMMENDED default is to send the request ten seconds before the the mapping expires.

## **[5.4](#) ILA forwarding node and host receive processing**

If an ILA forwarding node or host receives an ILA addressed packet with its locator it will check its local mapping database to determine if the identifier is local. If the identifier is local, a forwarding node will forward the packet to its destination after ILA to SIR address translation has been done on the packet's destination address. Similarly, an ILA host will receive the packet into it's local stack after ILA to SIR address translation.

If the identifier is not local then the ILA forwarding node or host will perform ILA to SIR address translation on the destination address and forward the packet into the network. This may happen if an end node has moved to be attached to a different ILA forwarding node in host and the new locator has not yet been propagated to all ILA nodes. The packet should traverse a mapping router which can send an ILA redirect back the source's ILA forwarding node or host as described above.

When a node migrates its point of attachment from one forwarding node or host to another, the local mapping on the old node is removed so that any packets that are received and destined to the migrated identifier are re-injected with a SIR address as described above. A "negative" mapping with timeout may also be set ensure that the node is able to infer the SIR address from a destination address (e.g. would be needed with foreign identifiers).

## **[5.5](#) Locator unreachable handling**

When connectivity to a locator is loss the mapping system should detect this. A locator unreachable message MAY be sent by an ILA router to ILA forwarding nodes or hosts informing them that a locator is no longer reachable. Each forwarding node or host SHOULD remove



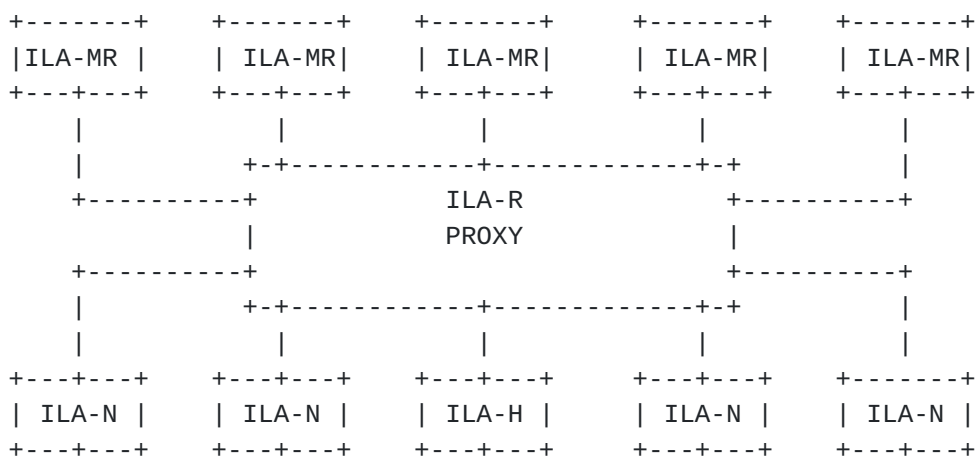


any cache entries using that locator and MAY send a map request for the affected identifiers.

## 5.6 Control Connections

ILA nodes and routers must create ILAMP connections to all the mapping routers that might provide routing information. In a simple network there may be just one mapping router to connect to. In a more complex network with ILA routers for sharded and replicated mapping system database there may be many. A list of ILA routers to connect to is provided to each ILA forwarding node and host. This list could be provided by configuration, a shared database, or an external protocol to ILAMP.

Conceivably, the number of mapping routers in a network that might report mapping information to a host could be quite large (into the thousands). If managing a large number of connections at the ILA forwarding nodes or hosts is problematic, ILA mapping router proxies could be used that consolidate connections as illustrated below:



In the above diagram a single ILA mapping router proxy serves five ILA routers and five ILA forwarding nodes and nodes. The proxy creates one connection to each router and each ILA forwarding node and host creates one connection to the proxy.

## 5.7 Protocol errors

If a protocol error is encountered in processing ILAMP messages a peer MUST terminate the connection. It SHOULD log the error and MAY attempt to restart the connection. There are no error messages defined in ILAMP.

Protocol errors include mismatch of length for given data, reserved bit not set to zero, unknown identifier type or locator types,



unknown type, unknown sub-type, and loss of message synchronization in a TCP stream. Note that if the end of a message does not end on field or record boundary this also considered a protocol error.

## **6 Security Considerations**

ILAMP must have protection against message forgery. In particular secure redirects and mapping information message are required to prevent and attacked from spoofing messages and illegitimately redirecting packets. This security is provided by using TCP connections so that origin of the messages is never ambiguous.

Transport Layer Security (TLS) [[RFC5246](#)] MAY be used to provide secrecy, authentication, and integrity check for ILAMP messages.

The TCP Authentication Option [[RFC5925](#)] MAY be used to provide authentication for ILAMP messages.



## [7](#) IANA Considerations

## [8](#) References

### [8.1](#) Normative References

- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [ILA] Herbert, T., and Lapukhov, P., "Identifier Locator Addressing for IPv6" [draft-herbert-intarea-ila-00](#)

### [8.2](#) Informative References

- [[RFC5246](#)] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", [RFC 5925](#), DOI 10.17487/RFC5925, June 2010, <<https://www.rfc-editor.org/info/rfc5925>>.
- [BGPILA] Lapukhov, P., "Use of BGP for dissemination of ILA mapping information" [draft-lapukhov-bgp-ila-afi-02](#)

#### Author's Address

Tom Herbert  
Quantonium  
Santa Clara, CA  
USA

Email: [tom@quantonium.net](mailto:tom@quantonium.net)

