

Internet Engineering Task Force (IETF)
Request for Comments: 6835
Category: Informational
ISSN: 2070-1721

D. Farinacci
D. Meyer
Cisco Systems
January 2013

The Locator/ID Separation Protocol Internet Groper (LIG)

Abstract

A simple tool called the Locator/ID Separation Protocol (LISP) Internet Groper or 'lig' can be used to query the LISP mapping database. This document describes how it works.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

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/rfc6835>.

Copyright Notice

Copyright (c) 2013 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

- 1. Introduction 2
- 2. Definition of Terms 3
- 3. Basic Overview 5
- 4. Implementation Details 6
 - 4.1. LISP Router Implementation 6
 - 4.2. Public Domain Host Implementation 8
- 5. Testing the ALT 9
- 6. Future Enhancements 10
- 7. Deployed Network Diagnostic Tools 10
- 8. Security Considerations 10
- 9. References 11
 - 9.1. Normative References 11
 - 9.2. Informative References 11
- Appendix A. Acknowledgments 12

1. Introduction

The Locator/ID Separation Protocol [RFC6830] specifies an architecture and mechanism for replacing the addresses currently used by IP with two separate namespaces: Endpoint IDs (EIDs), used within sites, and Routing Locators (RLOCs), used on the transit networks that make up the Internet infrastructure. To achieve this separation, LISP defines protocol mechanisms for mapping from EIDs to RLOCs. In addition, LISP assumes the existence of a database to store and propagate those mappings globally. Several such databases have been proposed, among them: a Content distribution Overlay Network Service for LISP [LISP-CONS], a Not-so-novel EID-to-RLOC Database (LISP-NERD) [RFC6837], and LISP Alternative Topology (LISP+ALT) [RFC6836], with LISP+ALT being the system that is currently being implemented and deployed on the pilot LISP network.

In conjunction with the various mapping systems, there exists a network-based API called LISP Map-Server [RFC6833]. Using Map-Resolvers and Map-Servers allows LISP sites to query and register into the database in a uniform way independent of the mapping system used. Sending Map-Requests to Map-Resolvers provides a secure mechanism to obtain a Map-Reply containing the authoritative EID-to-RLOC mapping for a destination LISP site.

The 'lig' is a manual management tool to query the mapping database. It can be run by all devices that implement LISP, including Ingress Tunnel Routers (ITRs), Egress Tunnel Routers (ETRs), Proxy-ITRs, Proxy-ETRs, Map-Resolvers, Map-Servers, and LISP-ALT Routers, as well as by a host system at either a LISP-capable or non-LISP-capable site.

The mapping database system is typically a public database used for wide-range connectivity across Internet sites. The information in the public database is purposely not kept private so it can be generally accessible for public use.

2. Definition of Terms

Map-Server: a network infrastructure component that learns EID-to-RLOC mapping entries from an authoritative source (typically, an ETR, though static configuration or another out-of-band mechanism may be used). A Map-Server advertises these mappings in the distributed mapping database.

Map-Resolver: a network infrastructure component that accepts LISP Encapsulated Map-Requests, typically from an ITR, quickly determines whether or not the destination IP address is part of the EID namespace; if it is not, a Negative Map-Reply is immediately returned. Otherwise, the Map-Resolver finds the appropriate EID-to-RLOC mapping by consulting the distributed mapping database system.

Routing Locator (RLOC): the IPv4 or IPv6 address of an Egress Tunnel Router (ETR). It is the output of an EID-to-RLOC mapping lookup. An EID maps to one or more RLOCs. Typically, RLOCs are numbered from topologically aggregatable blocks that are assigned to a site at each point to which it attaches to the global Internet. Thus, the topology is defined by the connectivity of provider networks, and RLOCs can be thought of as Provider-Assigned (PA) addresses. Multiple RLOCs can be assigned to the same ETR device or to multiple ETR devices at a site.

Endpoint ID (EID): a 32-bit (for IPv4) or 128-bit (for IPv6) value used in the source and destination address fields of the first (most inner) LISP header of a packet. The host obtains a destination EID the same way it obtains a destination address today, for example, through a DNS lookup. The source EID is obtained via existing mechanisms used to set a host's "local" IP address. An EID is allocated to a host from an EID-prefix block associated with the site where the host is located. An EID can be used by a host to refer to other hosts. EIDs must not be used as LISP RLOCs. Note that EID blocks may be assigned in a hierarchical manner, independent of the network topology, to facilitate scaling of the mapping database. In addition, an EID block assigned to a site may have site-local structure (subnetting) for routing within the site; this structure is not visible to the global routing system.

EID-to-RLOC Cache: a short-lived, on-demand table in an ITR that stores, tracks, and is responsible for timing-out and otherwise validating EID-to-RLOC mappings. This cache is distinct from the full "database" of EID-to-RLOC mappings; the cache is dynamic, local to the ITR(s), and relatively small, while the database is distributed, relatively static, and much more global in scope.

EID-to-RLOC Database: a global distributed database that contains all known EID-prefix to RLOC mappings. Each potential ETR typically contains a small piece of the database: the EID-to-RLOC mappings for the EID prefixes "behind" the router. These map to one of the router's own, globally-visible, IP addresses.

Encapsulated Map-Request (EMR): an EMR is a Map-Request message that is encapsulated with another LISP header using UDP destination port number 4342. It is used so an ITR, PITR, or a system initiating a 'lig' command can get the Map-Request to a Map-Resolver by using locator addresses. When the Map-Request is decapsulated by the Map-Resolver, it will be forwarded on the ALT network to the Map-Server that has injected the EID-prefix for a registered site. The Map-Server will then encapsulate the Map-Request in a LISP packet and send it to an ETR at the site. The ETR will then return an authoritative reply to the system that initiated the request. See [RFC6830] for packet format details.

Ingress Tunnel Router (ITR): An ITR is a router that accepts an IP packet with a single IP header (more precisely, an IP packet that does not contain a LISP header). The router treats this "inner" IP destination address as an EID and performs an EID-to-RLOC mapping lookup. The router then prepends an "outer" IP header with one of its globally routable RLOCs in the source address field and the result of the mapping lookup in the destination address field. Note that this destination RLOC may be an intermediate, proxy device that has better knowledge of the EID-to-RLOC mapping closer to the destination EID. In general, an ITR receives IP packets from site end-systems on one side and sends LISP-encapsulated IP packets toward the Internet on the other side.

Egress Tunnel Router (ETR): An ETR is a router that accepts an IP packet where the destination address in the "outer" IP header is one of its own RLOCs. The router strips the "outer" header and forwards the packet based on the next IP header found. In general, an ETR receives LISP-encapsulated IP packets from the Internet on one side and sends decapsulated IP packets to site end-systems on the other side. ETR functionality does not have to be limited to a router device. A server host can be the endpoint of a LISP tunnel as well.

Proxy-ITR (PITR): A PITR, also known as a PTR, is defined and described in [RFC6832]. A PITR acts like an ITR but does so on behalf of non-LISP sites that send packets to destinations at LISP sites.

Proxy-ETR (PETR): A PETR is defined and described in [RFC6832]. A PETR acts like an ETR but does so on behalf of LISP sites that send packets to destinations at non-LISP sites.

xTR: An xTR is a reference to an ITR or ETR when direction of data flow is not part of the context description. xTR refers to the router that is the tunnel endpoint; it is used synonymously with the term "tunnel router". For example, "an xTR can be located at the Customer Edge (CE) router" means that both ITR and ETR functionality is at the CE router.

Provider-Assigned (PA) Addresses: PA addresses are an address block assigned to a site by each service provider to which a site connects. Typically, each block is a sub-block of a service-provider Classless Inter-Domain Routing (CIDR) [RFC4632] block and is aggregated into the larger block before being advertised into the global Internet. Traditionally, IP multihoming has been implemented by each multihomed site acquiring its own globally visible prefix. LISP uses only topologically assigned and aggregatable address blocks for RLOCs, eliminating this demonstrably non-scalable practice.

3. Basic Overview

When the 'lig' command is run, a Map-Request is sent for a destination EID. When a Map-Reply is returned, the contents are displayed to the user. The information displayed includes:

- o The EID-prefix for the site that the queried destination EID matches.
- o The locator address of the Map Replier.
- o The Locator-Set for the mapping entry, which includes the locator address, up/down status, priority, and weight of each Locator.
- o A round-trip-time estimate for the Map-Request/Map-Reply exchange.

A possible syntax for a 'lig' command could be:

```
lig <destination> [source <source>] [to <map-resolver>]
```

Parameter description:

<destination>: is either a Fully Qualified Domain Name (FQDN) or a destination EID for a remote LISP site.

source <source>: is an optional source EID to be inserted in the 'Source EID' field of the Map-Request.

to <map-resolver>: is an optional FQDN or RLOC address for a Map-Resolver.

The 'lig' utility has two use cases. The first is a way to query the mapping database for a particular EID. The other is to verify if a site has registered successfully with a Map-Server.

The first usage has already been described. Verifying registration is called "ligging yourself"; it happens as follows. In the 'lig' initiator, a Map-Request is sent for one of the EIDs for the 'lig' initiator's site. The Map-Request is then returned to one of the ETRs for the 'lig'-initiating site. In response to the Map-Request, a Map-Reply is sent back to the locator address of the 'lig' initiator (note the Map-Reply could be sent by the 'lig' initiator). That Map-Reply is processed, and the mapping data for the 'lig'-initiating site is displayed for the user. Refer to the syntax in Section 4.1 for an implementation of "ligging yourself". However, for host-based implementations within a LISP site, "lig self" is less useful since the host may not have an RLOC with which to receive a Map-Reply. But, 'lig' can be used in a non-LISP site, as well as from infrastructure hosts, to get mapping information.

4. Implementation Details

4.1. LISP Router Implementation

The Cisco LISP prototype implementation has support for 'lig' for IPv4 and IPv6. The command line description is:

```
lig <dest-eid> [source <source-eid>] [to <mr>] [count <1-5>]
```

This command initiates the LISP Internet Groper. It is similar to the DNS analogue 'dig' but works on the LISP mapping database. When this command is invoked, the local system will send a Map-Request to the configured Map-Resolver. When a Map-Reply is returned, its contents will be displayed to the user. By default, up to three Map-Requests are sent if no Map-Reply is returned, but, once a Map-Reply is returned, no other Map-Requests are sent. The destination can take a DNS name, or an IPv4 or IPv6 EID address. The <source-eid> can be one of the EID addresses assigned to the site in the default

Virtual Routing and Forwarding (VRF) table. When <mr> is specified, then the Map-Request is sent to the address. Otherwise, the Map-Request is sent to a configured Map-Resolver. When a Map-Resolver is not configured, then the Map-Request is sent on the ALT network if the local router is attached to the ALT. When "count <1-5>" is specified, 1, 2, 3, 4, or 5 Map-Requests are sent.

Some sample output:

```
router# lig abc.example.com
Send map-request to 10.0.0.1 for 192.168.1.1 ...
Received map-reply from 10.0.0.2 with rtt 0.081468 secs

Map-Cache entry for abc.example.com EID 192.168.1.1:
192.168.1.0/24, uptime: 13:59:59, expires: 23:59:58,
via map-reply, auth
Locator      Uptime      State  Priority/Weight  Packets In/Out
10.0.0.2     13:59:59   up     1/100            0/14
```

Using 'lig' to "lig yourself" is accomplished with the following syntax:

```
lig {self | self6} [source <source-eid>] [to <mr>] [count <1-5>]
```

Use this command for a simple way to see if the site is registered with the mapping database system. The destination-EID address for the Map-Request will be the first configured EID-prefix for the site (with the host bits set to 0). For example, if the site's EID-prefix is 192.168.1.0/24, the destination-EID for the Map-Request is 192.168.1.0. The source-EID address for the Map-Request will also be 192.168.1.0 (in this example), and the Map-Request is sent to the configured Map-Resolver. If the Map-Resolver and Map-Server are the same LISP system, then the "lig self" is testing if the Map-Resolver can "turn back a Map-Request to the site". If another Map-Resolver is used, it can test that the site's EID-prefix has been injected into the ALT infrastructure, in which case the 'lig' Map-Request is processed by the Map-Resolver and propagated through each ALT-Router hop to the site's registered Map-Server. Then, the Map-Server returns the Map-Request to the originating site. In that case, an xTR at the originating site sends a Map-Reply to the source of the Map-Request (could be itself or another xTR for the site). All other command parameters are described above. Using "lig self6" tests for registering of IPv6 EID-prefixes.

Some sample output for "ligging yourself":

```
router# lig self
Send loopback map-request to 10.0.0.1 for 192.168.2.0 ...
Received map-reply from 10.0.0.3 with rtt 0.001592 secs

Map-Cache entry for EID 192.168.2.0:
192.168.2.0/24, uptime: 00:00:02, expires: 23:59:57
via map-reply, self
  Locator      Uptime      State  Priority/Weight  Packets In/Out
  10.0.0.3     00:00:02   up     1/100            0/0

router# lig self6
Send loopback map-request to 10.0.0.1 for 2001:db8:1:: ...
Received map-reply from 10::1 with rtt 0.044372 secs

Map-Cache entry for EID 192:168:1:::
2001:db8:1::/48, uptime: 00:00:01, expires: 23:59:58
via map-reply, self
  Locator      Uptime      State  Priority/Weight  Packets In/Out
  10.0.0.3     00:00:01   up     1/100            0/0
  2001:db8:ffff::1 00:00:01   up     2/0              0/0
```

4.2. Public Domain Host Implementation

There is a public domain implementation that can run on any x86-based system. The only requirement is that the system that initiates 'lig' must have an address assigned from the locator namespace.

```
lig [-d] <eid> -m <map-resolver> [-c <count>] [-t <timeout>]
```

Parameter description:

- d: prints additional protocol debug output.
- <eid>: the destination EID or FQDN of a LISP host.
- m <map-resolver>: the RLOC address or FQDN of a Map-Resolver.
- c <count>: the number of Map-Requests to send before the first Map-Reply is returned. The default value is 3. The range is from 1 to 5.
- t <timeout>: the amount of time, in seconds, before another Map-Request is sent when no Map-Reply is returned. The default value is 2 seconds. The range is from 1 to 5.

Some sample output:

```
% lig xyz.example.com -m 10.0.0.1
Send map-request to 10.0.0.1 for 192.168.1.1 ...
Received map-reply from 10.0.0.2 with rtt 0.04000 sec
```

```
Mapping entry for EID 192.168.1.1:
192.168.1.0/24, record ttl: 60
Locator          State      Priority/Weight
10.0.0.1         up        1/25
10.0.0.2         up        1/25
10.0.0.3         up        1/25
10.0.0.4         up        2/25
```

The public domain implementation of 'lig' is available at
<<http://github.com/LISPMob/lig>>.

5. Testing the ALT

There are cases where a Map-Reply is returned from a 'lig' request, but the user doesn't really know how much of the mapping infrastructure was tested. There are two cases to consider -- avoiding the ALT and traversing the ALT.

When an ITR sends a 'lig' request to its Map-Resolver for a destination-EID, the Map-Resolver could also be configured as a Map-Server. And if the destination-EID is for a site that registers with this Map-Server, the Map-Request is sent to the site directly without testing the ALT. This occurs because the Map-Server is the source of the advertisement for the site's EID-prefix. So, if the map-reply is returned to the 'lig'-requesting site, you cannot be sure that other sites can reach the same destination-EID.

If a Map-Resolver is used that is not a Map-Server for the EID-prefix being sought, then the ALT infrastructure can be tested. This test case is testing the functionality of the Map-Resolver, traversal of the ALT (testing BGP-over-GRE), and the Map-Server.

It is recommended that users issue two 'lig' requests; they send Map-Requests to different Map-Resolvers.

The network can have a LISP-ALT Router deployed as a "ALT looking-glass" node. This type of router has BGP peering sessions with other ALT Routers where it does not inject any EID-prefixes into the ALT but just learns ones advertised by other ALT Routers and Map-Servers. This router is configured as a Map-Resolver. 'lig' users can point to the ALT looking-glass router for Map-Resolver services via the "to <map-resolver>" parameter on the 'lig' command. The ALT looking-

glass node can be used to 'lig' other sites as well as your own site. When the ALT looking-glass is used as a Map-Resolver, you can be assured the ALT network is being tested.

6. Future Enhancements

When Negative Map-Replies have been further developed and implemented, 'lig' should be modified appropriately to process and clearly indicate how and why a Negative Map-Reply was received. Negative Map-Replies could be sent in the following cases: the 'lig' request was initiated for a non-EID address or there was rate-limiting on the replier.

7. Deployed Network Diagnostic Tools

There is a web-based interface to do auto-polling with 'lig' on the back-end for most of the LISP sites on the LISP test network. The web page can be accessed at <http://www.lisp4.net/status>.

There is a LISP site monitoring web-based interface that can be found at <http://lispmon.net>.

At <http://baldomar.ccaba.upc.edu/lispmon>, written by the folks at Universitat Politecnica de Catalunya (UPC), shows a geographical map indicating where each LISP site resides.

8. Security Considerations

The use of 'lig' does not affect the security of the LISP infrastructure as it is simply a tool that facilitates diagnostic querying. See [RFC6830], [RFC6836], and [RFC6833] for descriptions of the security properties of the LISP infrastructure.

'lig' provides easy access to the information in the public mapping database. Therefore, it is important to protect the mapping information for private use. This can be provided by disallowing access to specific mapping entries or to place such entries in a private mapping database system.

9. References

9.1. Normative References

- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, January 2013.
- [RFC6832] Lewis, D., Meyer, D., Farinacci, D., and V. Fuller, "Interworking between Locator/ID Separation Protocol (LISP) and Non-LISP Sites", RFC 6832, January 2013.
- [RFC6833] Farinacci, D. and V. Fuller, "Locator/ID Separation Protocol (LISP) Map Server Interface", RFC 6833, January 2013.

9.2. Informative References

- [LISP-CONS] Farinacci, D., Fuller, V., and D. Meyer, "LISP-CONS: A Content distribution Overlay Network Service for LISP", Work in Progress, April 2008.
- [RFC6836] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", RFC 6836, January 2013.
- [RFC6837] Lear, E., "NERD: A Not-so-novel Endpoint ID (EID) to Routing Locator (RLOC) Database", RFC 6837, January 2013.

Appendix A. Acknowledgments

Thanks and kudos to John Zwiebel, Andrew Partan, Darrel Lewis, and Vince Fuller for providing critical feedback on the 'lig' design and prototype implementations. To these folks, as well as all the people on `lisp-beta@external.cisco.com` who tested 'lig' functionality and continue to do so, we extend our sincere thanks.

This document is based on an individual contribution.

Authors' Addresses

Dino Farinacci
Cisco Systems
Tasman Drive
San Jose, CA 95134
USA

E-Mail: farinacci@gmail.com

Dave Meyer
Cisco Systems
170 Tasman Drive
San Jose, CA
USA

E-Mail: dmm@cisco.com