

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: November 6, 2009

D. Farinacci
D. Meyer
Cisco Systems
May 5, 2009

**LISP Internet Groper (LIG)
draft-farinacci-lisp-lig-01.txt**

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/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on November 6, 2009.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

A simple tool called the LISP Internet Groper or 'lig' can be used to query the LISP mapping database. This draft describes how it works.

Table of Contents

1.	Requirements Notation	3
2.	Introduction	4
3.	Definition of Terms	5
4.	Basic Overview	7
5.	Implementation Details	9
5.1.	LISP Router Implementation	9
5.2.	Public Domain Host Implementation	10
6.	Testing the ALT	12
7.	Future Enhancements	13
8.	Security Considerations	14
9.	References	15
9.1.	Normative References	15
9.2.	Informative References	15
Appendix A.	Acknowledgments	16
	Authors' Addresses	17

1. Requirements Notation

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 [[RFC2119](#)].

2. Introduction

LISP [[LISP](#)] specifies an architecture and mechanism for replacing the addresses currently used by IP with two separate name spaces: 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, the Locator/ID Separation Protocol (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: LISP-CONS [[CONS](#)], LISP-NERD [[NERD](#)], and LISP+ALT [[ALT](#)], 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 [[LISP-MS](#)]. 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 which implement LISP, including ITRs, ETRs, PTR, 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.

3. Definition of Terms

Map-Server: a network infrastructure component which 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 which 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 a 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; where the topology is defined by the connectivity of provider networks, RLOCs can be thought of as 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 an destination address today, for example through a DNS lookup or SIP exchange. 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, it 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 which is encapsulated with another LISP header using UDP destination port number 4341. It is used so an ITR, PTR, 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.

4. 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 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 An 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 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 Fully Qualified Domain Name or RLOC address for a Map-Resolver.

The `lig` utility has two usage cases. The first being a way to query the mapping database for a particular EID. And the other 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". What occurs is 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 `lig` initiating site is displayed for the user. Refer to the syntax in section [Section 5.1](#) 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 to receive a Map-Reply with. But, `lig` can be used in a non-LISP

site as well as from infrastructure hosts to get mapping information.

5. Implementation Details

5.1. 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 3 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 VRF. 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:

```
titanium-dino# lig titanium-dmm.lisp4.net
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 titanium-dmm.lisp4.net 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, propagated through each ALT router hop to the site's registered Map-Server. Then the Map-Server returns the Map-Request to originating site. In which 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:

```

rutile# 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

titanium-simlo# lig self6
Send loopback map-request to 10.0.0.1 for 192:168:1:: ...
Received map-reply from 10::1 with rtt 0.044372 secs

Map-cache entry for EID 192:168:1:::
192:168:1::/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
  10::1         00:00:01  up     2/0             0/0

```

5.2. 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>: is the destination EID or FQDN of a LISP host.

-m <map-resolver>: is 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 titanium-test.lisp4.net -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 sourceforge.net.

6. 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 2 lig requests, each which 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.

7. 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 the Map-Request initiated by lig request is being rejected due to rate-limiting on the replier.

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 [\[LISP\]](#), [\[ALT\]](#), and [\[LISP-MS\]](#) for descriptions of the security properties of the LISP infrastructure.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

9.2. Informative References

- [ALT] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "LISP Alternative Topology (LISP-ALT)", [draft-fuller-lisp-alt-03.txt](#) (work in progress), February 2009.
- [CONS] Farinacci, D., Fuller, V., and D. Meyer, "LISP-CONS: A Content distribution Overlay Network Service for LISP", [draft-meyer-lisp-cons-03.txt](#) (work in progress), November 2007.
- [LISP] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", [draft-farinacci-lisp-12.txt](#) (work in progress), March 2009.
- [LISP-MS] Farinacci, D. and V. Fuller, "LISP Map Server", [draft-fuller-lisp-ms-00.txt](#) (work in progress), March 2009.
- [NERD] Lear, E., "NERD: A Not-so-novel EID to RLOC Database", [draft-lear-lisp-nerd-04.txt](#) (work in progress), April 2008.

[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. 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.

Authors' Addresses

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

Email: dino@cisco.com

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

Email: dmm@cisco.com

