Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: September 18, 2014

D. Farinacci
lispers.net
A. Jain
Juniper Networks
I. Kouvelas
D. Lewis
cisco Systems
March 17, 2014

LISP-DDT Referral Internet Groper (RIG)
draft-farinacci-lisp-rig-03

Abstract

A simple tool called the LISP-DDT Referral Internet Groper or 'rig' can be used to query the LISP-DDT hierarchy. This draft describes how the 'rig' tool works.

Table of Contents

1. Requirements Language 3
2. Introduction 4
3. Definition of Terms 5
4. Basic Overview 8
5. Implementation Details 10
6. Security Considerations 12
7. IANA Considerations 13
8. References 14
8.1. Normative References 14
8.2. Informative References 14
Appendix A. Acknowledgments 15
Appendix B. Document Change Log 16
B.1. Changes to draft-farinacci-lisp-rig-03.txt 16
B.2. Changes to draft-farinacci-lisp-rig-02.txt 16
B.3. Changes to draft-farinacci-lisp-rig-01.txt 16
B.4. Changes to draft-farinacci-lisp-rig-00.txt 16
Authors' Addresses 17

1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Introduction

LISP [RFC6830] 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. This document focuses on the LISP-DDT [LISP-DDT] mapping database system.

The 'rig' is a manual management tool to query LISP-DDT mapping database hierarchy. It can be run by all devices which implement LISP, including ITRs, ETRs, PITRs, PETRs, Map-Resolvers, Map-Servers, and LISP-DDT nodes, as well as by a host system at either a LISP-capable or non-LISP-capable site.

The LISP-DDT 'rig' tool is similar to the 'lig' [RFC6835] tool in that they are both diagnostic tools to query a distributed database. However, 'rig' is used to find Map-Servers serving an EID-prefix, specifically within a LISP-DDT mapping database framework. And 'lig' can be used on top of any mapping database system to retrieve locators used for packet encapsulation.

3. Definition of Terms

Endpoint ID (EID): An EID is 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 Domain Name System (DNS) [RFC1034] lookup or Session Invitation Protocol (SIP) [RFC3261] exchange. The source EID is obtained via existing mechanisms used to set a host's "local" IP address. An EID used on the public Internet must have the same properties as any other IP address used in that manner; this means, among other things, that it must be globally unique. 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. In theory, the bit string that represents an EID for one device can represent an RLOC for a different device. As the architecture is realized, if a given bit string is both an RLOC and an EID, it must refer to the same entity in both cases. When used in discussions with other Locator/ID separation proposals, a LISP EID will be called a "LEID". Throughout this document, any references to "EID" refers to an LEID.

Extended EID (XEID): A LISP EID, optionally extended with a non-zero Instance ID (IID) if the EID is intended for use in a context where it may not be a unique value, such as on a Virtual Private Network where "private" address space is used. See "Using Virtualization and Segmentation with LISP" in [RFC6830] for more discussion of Instance IDs.

Routing Locator (RLOC): A RLOC is an IPv4 [RFC0791] or IPv6 [RFC2460] address of an egress tunnel router (ETR). A RLOC 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; 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.

DDT Node: A network infrastructure component responsible for specific XEID-prefix and for delegation of more-specific sub-prefixes to other DDT nodes.

DDT Client: A network infrastructure component that sends Map-Request messages and implements the iterative following of Map-Referral results. Typically, a DDT client will be a Map Resolver but it is also possible for an ITR to implement DDT client functionality. A DDT client can be a device that is originating 'rig' requests.

DDT Map Server: A DDT node that also implements Map Server functionality (forwarding Map-Requests and/or returning Map-Replies if offering proxy-mode service) for a subset of its delegated prefixes.

DDT Map Resolver: A network infrastructure element that accepts a Map-Request, adds the XEID to its lookup queue, then queries one or more DDT nodes for the requested EID, following returned referrals until it receives one with the the ms-ack action. This indicates that the Map-Request has been sent to a Map-Server that will forward it to an ETR that, in turn, will provide a Map-Reply to the original sender. A DDT Map Resolver maintains both a cache of Map-Referral message results containing RLOCs for DDT nodes responsible for XEID-prefixes of interest (termed the "referral cache") plus a lookup queue of XEIDs that are being resolved through iterative querying of DDT nodes.

Encapsulated Map-Request: A LISP Map-Request carried within an Encapsulated Control Message, which has an additional LISP header prepended. Sent to UDP destination port 4342. The "outer" addresses are globally-routable IP addresses, also known as RLOCs. Used by an ITR when sending to a Map-Resolver and by a Map-Server when forwarding a Map-Request to an ETR as documented in [RFC6833].

Map-Referral: A LISP message sent by a DDT node when it receives a DDT Map-Request for an XEID that matches a configured XEID-prefix delegation. The Map-Referral message includes a "referral", a set of RLOCs for DDT nodes that have more information about the sub-prefix; a DDT client "follows the referral" by sending another DDT Map-Request to one of those RLOCs to obtain either an answer or another referral to DDT nodes responsible for a more-specific XEID-prefix.

Authoritative XEID-prefix: An XEID-prefix delegated to a DDT node and for which the DDT node may provide further delegations of more-specific sub-prefixes.

4. Basic Overview

When the rig command is run, an Encapsulated Control Message Map-Request is sent for a destination EID. When a LISP-DDT Map-Referral is returned, the contents are displayed to the user. The information displayed includes:

- A delegated EID-prefix configured in a DDT-node or a configured site EID-prefix in a DDT Map-Server that matches the requested EID.
- The type of DDT-node which sent the Map-Referral.
- The action code and ttl set by the sender of the Map-Referral.
- The referral RLOC addresses from the Map-Referral message.
- An round-trip-time estimate for the ECM-Map-Request/Map-Referral message exchange.

A possible syntax for a rig command could be:

```
rig [instance-id <iid>] <eid> to <ddt-node> [follow-all-referrals]
```

Parameter description:

[instance-id <iid>]:   is the instance-ID portion of the Extended EID used as VPN identifer. When the DDT hierarchy is not configured with instance-IDs, this argument is omitted from the command line.

<eid>:   is either a Fully Qualified Domain Name or a destination EID that is being queried in the LISP-DDT mapping database.

<ddt-node>:   is the RLOC address of any DDT-node in the DDT hierarchy. This can be the DDT root node, a DDT transit node, or a DDT Map-Server.

[follow-all-referrals]:   when this keyword is used each referral RLOC is queried so rig can descend the entire DDT hierarchy starting from the node <ddt-node>. When this keyword is not used, one of the referral RLOCs will be selected to descend a branch of the DDT hierarchy.

The rig utility not only shows you branches of the delegation hierarchy but can also report:

- When a DDT Map-Server would forward a Map-Request to the ETRs at a registered LISP site. This is known as an 'ms-acknowledgement' action.

- When a DDT Map-Server sends a negative referral indicating a requested EID is configured but not registered to the mapping database system. This is known as an 'ms-not-registered' action.

- When a DDT node is sending referrals for a transit or leaf node in the hierarchy. These are known as 'node-referral' and 'ms-referral' actions, respectively.

- When a DDT-node finds a hole in the address space that has not been allocated or configured in the delegation hierarchy. This is typically associated with a hole in a DDT node's configured authoritative prefix. This is known as a 'delegation-hole' action.

- When a DDT-node finds a hole in the address space that has not been allocated or configured in the delegation hierarchy at all. This is typically associated with a hole that is outside of a DDT node's authoritative prefix. This is known as 'non-authoritative' action.

Refer to [LISP-DDT] for more detail about Map-Referral actions.

5. Implementation Details

The cisco LISP prototype implementations on IOS and NX-OS has rig support for IPv4 and IPv6 EIDs in either the default instance or a non-zero instance-ID.

The IOS syntax is:

```
rig [instance-id <iid>] <eid> to <ddt-node> [follow-all-referrals]
```

The NX-OS syntax is:

```
rig [instance-id <iid>] <hostname> | {<eid> | <eid6>}}
                        to {<ddt-hostname> | {<ddt> | <ddt6>}}
```

Here is some sample IOS output:

```
Router# rig 12.0.1.1 to 1.1.1.1

Send Map-Request to DDT-node 1.1.1.1 ... node referral, rtt: 0 ms
EID-prefix: [0] 12.0.0.0/16, ttl: 1440
referrals: 2.2.2.2

Send Map-Request to DDT-node 2.2.2.2 ... node referral, rtt: 0 ms
EID-prefix: [0] 12.0.1.0/24, ttl: 1440
referrals: 4.4.4.4, 5.5.5.5

Send Map-Request to DDT-node 4.4.4.4 ... map-server acknowledgement,
                                         rtt: 0 ms
EID-prefix: [0] 12.0.1.0/28, ttl: 1440
referrals: 4.4.4.4, 5.5.5.5
```

```
Router# rig 12.0.1.1 to 1.1.1.1 follow-all-referrals

Send Map-Request to DDT-node 1.1.1.1 ... node referral, rtt: 4 ms
EID-prefix: [0] 12.0.0.0/16, ttl: 1440
referrals: 2.2.2.2

Send Map-Request to DDT-node 2.2.2.2 ... node referral, rtt: 0 ms
EID-prefix: [0] 12.0.1.0/24, ttl: 1440
referrals: 4.4.4.4, 5.5.5.5

Send Map-Request to DDT-node 4.4.4.4 ... map-server acknowledgement,
                                         rtt: 0 ms
EID-prefix: [0] 12.0.1.0/28, ttl: 1440
referrals: 4.4.4.4, 5.5.5.5

Send Map-Request to DDT-node 5.5.5.5 ... map-server acknowledgement,
                                         rtt: 0 ms
EID-prefix: [0] 12.0.1.0/28, ttl: 1440
referrals: 4.4.4.4, 5.5.5.5

No more referrals to pursue.
```

Here is some sample NX-OS output:

```
Router# rig 12.0.1.1 to 1.1.1.1

rig LISP-DDT hierarchy for EID [0] 12.0.1.1
Send Map-Request to DDT-node 1.1.1.1 ... replied, rtt: 0.003509 secs
EID-prefix [0] *, ttl: 1440, action: node-referral, referrals:
  2.2.2.2, priority/weight: 0/0

Send Map-Request to DDT-node 2.2.2.2 ... replied, rtt: 0.003173 secs
EID-prefix [0] 12.0.0.0/20, ttl: 1440, action: node-referral,
  referrals:
  3.3.3.3, priority/weight: 0/0

Send Map-Request to DDT-node 3.3.3.3 ... replied, rtt: 0.004145 secs
EID-prefix [0] 12.0.1.0/24, ttl: 1440, action: node-referral,
  referrals:
  5.5.5.5, priority/weight: 0/0
  6.6.6.6, priority/weight: 0/0

Send Map-Request to DDT-node 6.6.6.6 ... replied, rtt: 0.005800 secs
EID-prefix [0] 12.0.1.0/28, ttl: 1440, action: ms-ack, referrals:
  5.5.5.5, priority/weight: 0/0
  6.6.6.6, priority/weight: 0/0
```

6. Security Considerations

The use of rig does not affect the security of the LISP infrastructure as it is simply a tool that facilities diagnostic querying. See [RFC6830], [LISP-DDT], and [RFC6833] for descriptions of the security properties of the LISP infrastructure.

LISP rig 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.

7. IANA Considerations

This document makes no request of the IANA.

8. References

8.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

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

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, January 2013.

[RFC6833] Fuller, V. and D. Farinacci, "Locator/ID Separation Protocol (LISP) Map-Server Interface", RFC 6833, January 2013.

[RFC6835] Farinacci, D. and D. Meyer, "The Locator/ID Separation Protocol Internet Groper (LIG)", RFC 6835, January 2013.

8.2. Informative References

[LISP-DDT]
Fuller, V., Lewis, D., Ermagan, V., and A. Jain, "LISP Delegated Database Tree", draft-ietf-lisp-ddt-01.txt (work in progress).

Appendix A. Acknowledgments

The authors would like to thank the following people for their ideas and comments. They are TBD.

Appendix B. Document Change Log

B.1. Changes to draft-farinacci-lisp-rig-03.txt

- Draft posted March 2014.

- Resetting timer for expired draft.

B.2. Changes to draft-farinacci-lisp-rig-02.txt

- Draft posted September 2013.

- Resetting timer for expired draft.

- Update author affliations and reference RFCs.

B.3. Changes to draft-farinacci-lisp-rig-01.txt

- Draft posted March 2012.

- Updated reference to LISP-DDT".

B.4. Changes to draft-farinacci-lisp-rig-00.txt

- Initial draft posted March 2012.

Authors' Addresses

Dino Farinacci
lispers.net
San Jose, California
USA

Phone: 408-718-2001
Email: farinacci@gmail.com

Amit Jain
Juniper Networks
San Jose, California
USA

Phone:
Email: atjain@juniper.net

Isidor Kouvelas
cisco Systems
Tasman Ave.
San Jose, California
USA

Phone:
Email: kouvelas@cisco.com

Darrel Lewis
cisco Systems
Tasman Ave.
San Jose, California
USA

Phone:
Email: darlewis@cisco.com